

©ISTOCKPHOTO.COM/GEORAL

[Yan Chen, Chunxiao Jiang, Chih-Yu Wang, Yang Gao, and K.J. Ray Liu]

Decision Learning

[Data analytic learning with strategic decision making]

With the increasing ubiquity and power of mobile devices as well as the prevalence of social systems, more activities in our daily life are being recorded, tracked, and shared, creating the notion of social media. Such abundant and still growing real-life data, known as *big data*, provide a tremendous research opportunity in many fields. To analyze, learn, and understand such user-generated data, machine learning has been an important tool, and various machine-learning algorithms have been developed. However, since the user-generated data are the outcome of users' decisions, actions, and

socioeconomic interactions, which are highly dynamic, without considering users' local behaviors and interests, existing learning approaches tend to focus on optimizing a global objective function at the macroeconomic level, while totally ignoring users' local interactions at the microeconomic level. As such, there is a growing need to combine learning with strategic decision making, which are two traditionally distinct research disciplines, to be able to jointly consider both global phenomena and local effects to better understand, model, and analyze the newly arising issues in the emerging social media with user-generated data. In this article, we present an overview of the emerging notion of decision learning, i.e., learning with strategic decision making, which involves users' behaviors and interactions by combining learning with strategic decision making. We will discuss some examples from social

Digital Object Identifier 10.1109/MSP.2015.2479895
Date of publication: 29 December 2015

media with real data to show how decision learning can be used to better analyze users' optimal decision from a user's perspective, as well as design a mechanism from the system designer's perspective to achieve a desirable outcome.

INTRODUCTION

With the rapid development of communication and information technologies, the last decade has witnessed a proliferation of emerging social systems that help to promote the connectivity of people to an unprecedentedly high level. Examples of these emerging systems can be found in a wide range of domains, from online social networks like Facebook or Twitter to crowd-sourcing sites like Amazon Mechanical Turk or Topcoder where people solve various tasks by assigning them among a large pool of online workers to online question-and-answer (Q&A)

sites like Quora or Stack Overflow where people ask all kinds of questions, and all the way to new paradigms of power systems like smart grid. In a social system, there are two key characters: the system designer, who designs the social system, and the users, who participate in the social system and generate the data.

Because of the increasing ubiquity and power of mobile devices, the prevalence of social systems, and the rise of cloud data storage, our daily activities are increasingly being recorded, tracked, and shared. Big data provides a tremendous research opportunity in many fields, for example, behavior and sentiment analysis, epidemics and diseases propagation modeling, grid and network traffic management, and financial market trends tracking, just to name a few.

Machine learning has been an important tool to understand user-generated data, [1], [2]. Learning aims to use reasoning to find new, relevant information, given some background knowledge through representation, evaluation, and optimization. However, there are some limitations and constraints. For example, the generalization assumption that the training set is statistically

consistent with the testing set is often not true because users behave differently at different times and under different settings. Moreover, the single-objective function cannot address all interests of users since interests vary and, thus, there will be different objective functions. Besides, users are rational and naturally selfish, so they want to optimize their own objective functions [3], [4]. In addition, the data are the outcome of users' interactions, while learning algorithms cannot naturally involve users' individual local interests. Therefore, the knowledge contained in the data is


difficult to fully exploit from such a macroscopic view.

Existing learning approaches tend to focus on optimizing a global objective function at the macroeconomic level but totally ignore users' local decisions and interactions at the microeconomic level. Indeed, user-generated data are the outcome of users' decisions, actions, and their

social-economic interactions, which are highly dynamic, and, thus, the interactions of users and their decision-making process should be taken into consideration. As such, there is a growing need to bridge learning and strategic decision making to be more effective in mining, reasoning, and extracting knowledge and information from big data.

Yet there is a missing link. Traditionally, both learning and decision making are two distinct research disciplines. Successfully bridging them allows us to jointly consider both global phenomena and local effects to better understand, model, and analyze user-generated data from social media. Besides, learning is for making optimal decisions. In essence, learning and decision making are destined to couple due to the network externality, i.e., the influence of other users' behaviors on one user's reward [5]. In this article, we describe the emerging research field of decision learning, i.e., learning with strategic decision making, which involves users' behaviors and interactions by combining learning with strategic decision making, as illustrated in Figure 1. In decision learning, there are two major elements of data-driven issues:

LEARNING APPROACHES TEND TO FOCUS ON OPTIMIZING A GLOBAL OBJECTIVE FUNCTION AT THE MACROECONOMIC LEVEL BUT TOTALLY IGNORE USERS' LOCAL DECISIONS AND INTERACTIONS AT THE MICROECONOMIC LEVEL.



[FIG1] Decision learning: bridging learning and strategic decision making. Learning and decision making are traditionally two distinct research disciplines. Bridging them allows us to jointly consider both global and local effects to better understand, model, and analyze user-generated data from social media.

one is the modeling, analysis, and understanding of user behaviors and their interactions, and the other is the design of mechanisms to achieve the desired outcomes. The former considers the issues from user perspectives, while the latter is motivated by the system's point of view.

In contrast to traditional networks and systems, where users are constrained by fixed and predetermined rules, user interactions in social networks are generally self-enforcing [6], [7]. On the one hand, users in these systems have great flexibility in their actions and have the ability to observe, learn, and make intelligent decisions. On the other hand, because of their selfish nature, users will act to pursue their own interests, which often conflict with other users' objectives and the system designer's goal. These new features call for new theoretical and practical solutions to the designs of social networks. How can system designers design their systems to resolve the conflicting interests among users? And given the various and conflicting interests among users, how is the desired system-wide performance achieved?

The aforementioned questions motivate the study of user behaviors and incentive mechanisms in social systems. Incentive mechanisms refer to schemes that aim to steer user behaviors through the allocation of various forms of rewards, such as monetary rewards, virtual points, and reputation status. Plenty of empirical evidence can be found in the social psychology literature that demonstrates user behaviors in social networks are indeed highly influenced by these rewards [8]–[13]. Although we can use the social psychology literature to learn what factors influence user

behaviors and, thus, can be used as rewards, how to allocate the rewards to achieve the desired user behavior is still not well understood, which leads to ad hoc or poor designs of incentive mechanisms in many social networks in practice.

How can we fundamentally understand user behavior under the presence of rewards in social networks? Moreover, based on such understanding, how should a system developer design incentive mechanisms to achieve various objectives in a systematic way?

The focus of this article is to open a discussion in a tutorial way of an

emerging field, termed *decision learning*, which jointly combines learning with decision making toward a better fundamental understanding of user behaviors embedded under the tsunami of user-generated big data. In this article, we present three game-theoretic frameworks (readers who are interested in the preliminaries of game theory are referred to [3], [4], [6], [7], and [14]) to formally model user participation and interactions under various scenarios in social networks: decision learning with evolutionary user behavior, decision learning with sequential user behavior, and decision learning with mechanism design. A high-level comparison of the three frameworks is shown in Figure 2. On the evolutionary behavior, how information diffuses over online social networks using a graphical evolutionary game will be presented [15], [16], and the focus will be on how to learn users' utility function from Twitter and MemeTracker data for understanding and modeling strategic decision making; on the sequential behavior, how customers learn and choose the "best" deals using the Chinese-restaurant-game framework will be considered [17], [18]. In addition, Groupon deals

LEARNING AIMS TO USE REASONING TO FIND NEW, RELEVANT INFORMATION, GIVEN SOME BACKGROUND KNOWLEDGE THROUGH REPRESENTATION, EVALUATION, AND OPTIMIZATION.

	Sequential Behavior	Evolutionary Behavior	Mechanism Design
Perspective	User	User	System
Approach	Bottom Up	Bottom Up	Top Down
Rationality	Fully Rational	Limited	Fully Rational
User Type or Population	Static	Dynamic	Probabilistic
Solution Concept	Subgame-Perfect Nash Equilibrium	Evolutionarily Stable Strategy	Symmetric Bayesian Equilibrium
Solution Method	Backward Induction	Replicator Dynamic	Punishment Design
Performance Metrics and Desired Properties	Existence and Uniqueness of Equilibrium, Price of Anarchy	Evolutionarily Stable, Price of Anarchy	Strategy-Proof, Individual Rationality, Budget Balance
Application	Information Sharing	Deal Selection	Crowdsourcing

[FIG2] A high-level comparison of decision learning with evolutionary user behavior, decision learning with sequential user behavior, and decision learning with mechanism design.

and Yelp rating data will be used to discuss how users can learn from each other's interactions for better strategic decision making. On the mechanism design, how to design the mechanism to collect high-quality data with low cost from crowdsourcing will be illustrated [19]. Using these frameworks, we can theoretically analyze and predict user behaviors through equilibrium analysis. And, based on the analysis, one can optimize in a systematic way the design of incentive mechanisms for social networks to achieve a wide range of system objectives and analyze their performances accordingly. Finally, recent related works on the intersection of learning and strategic decision making will be surveyed and discussed.

INCENTIVE MECHANISMS REFER TO SCHEMES THAT AIM TO STEER USER BEHAVIORS THROUGH THE ALLOCATION OF VARIOUS FORMS OF REWARDS SUCH AS MONETARY REWARDS, VIRTUAL POINTS, AND REPUTATION STATUS.

USER BEHAVIOR MODELING AND ANALYSIS IN DECISION LEARNING

In this section, we will address decision learning from the user's point of view. Both the evolutionary and sequential user behaviors are commonly exhibited in social systems. How learning with strategic decision making may arise from both settings will be illustrated, first with information diffusion over online social networks using the graphical evolutionary game framework from Twitter and MemeTracker data, and then with the optimal restaurant strategy using the Chinese-restaurant-game framework from both Groupon deals and Yelp ratings, respectively.

EVOLUTIONARY USER BEHAVIOR: GRAPHICAL EVOLUTIONARY GAME FRAMEWORK

One typical user behavior in social systems is the repetitive and evolutionary decision making. A good example is that users repetitively decide whether to post information or not on online social networks. Figure 3 shows the top 50 threads in the news cycle with highest volume for the period of 1 August–31 October 2008, where each thread consists of all new articles and blog posts containing a textual variant of a particular quoted phrase. The five large peaks between late August and late September corresponding to the Democratic and Republican National Conventions illustrate the spread of comments and phrasing by candidates. Notice that the information forwarding is often not unconditional. One has to make a decision on whether or not to do so based on many factors, such as if the information is exciting or if friends are interested in it, etc. Other examples include repetitive online purchasing and review posting.

We find that, in essence, the repetitive and evolutionary decision-making process on social systems follows the evolution process in natural ecological systems [21]. It is a process that evolves from one state at a particular instance to another when information is shared and a decision is made. Thus, the evolutionary game is an ideal tool to model and analyze the social system users' repetitive and evolutionary behavior. Evolutionary game theory (EGT) is an application of the mathematical theory of games to the interaction-dependent strategy evolution in populations [21]. Arising from the realization

that frequency-dependent fitness introduces a strategic aspect to evolution, EGT becomes an essential component of a mathematical and computational approach to biological contexts, such as genes, viruses, cells, and humans. Recently, EGT has also become of increased interest to economists, sociologists, anthropologists, and social scientists. Here, we show how the evolutionary game theory is deployed to study users' repetitive and evolutionary behavior in social systems.

In the setting of our consideration, the social system user topology can be treated as a graph structure, and the user with a new decision can be regarded as the mutant. By considering the decision-making process as the mutant-spreading process [to forward or not to forward when an event (mutation) takes place], the graphical evolutionary game provides us with an analytical means to find the evolutionary dynamics and equilibrium of user behavior.

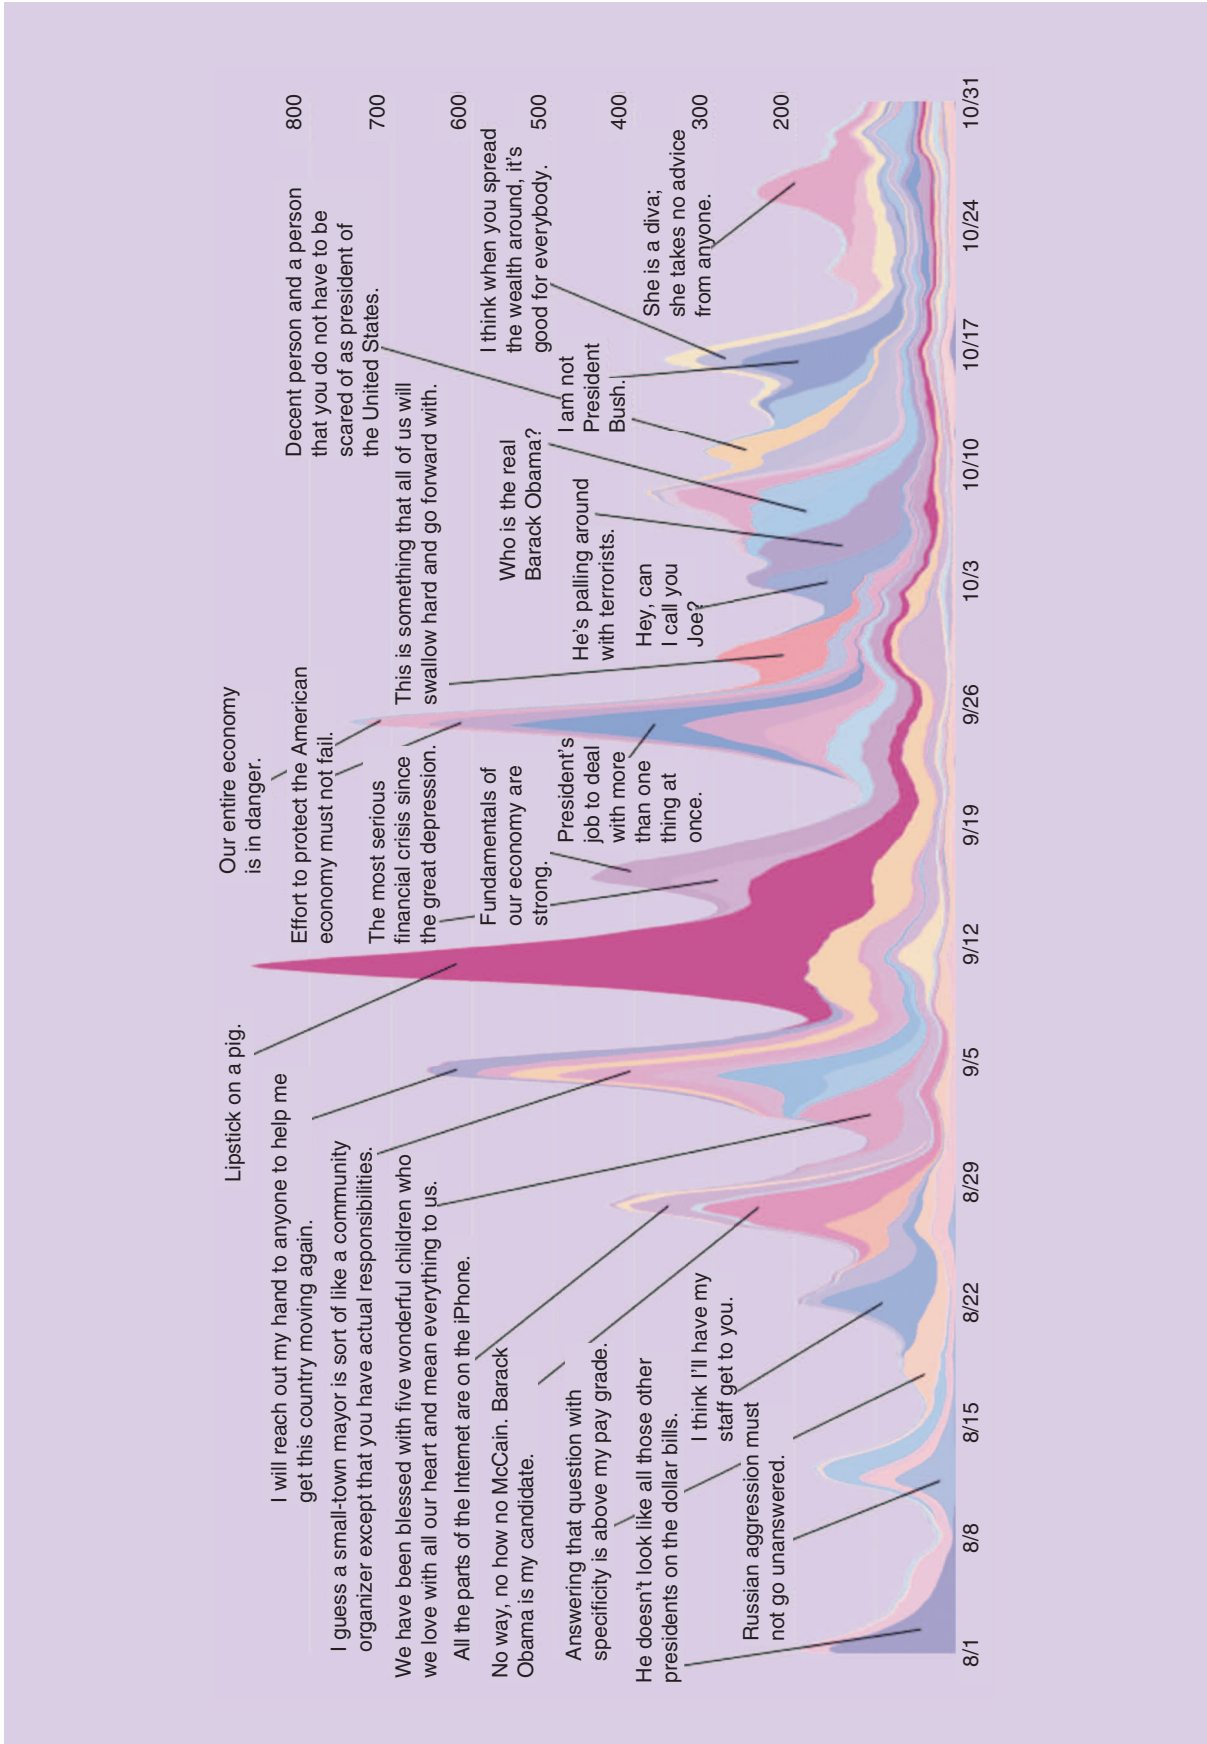
GRAPHICAL EVOLUTIONARY GAME FRAMEWORK

In EGT, the utility of a player is referred to as *fitness* [22]. Specifically, the fitness Φ is a linear combination of the baseline fitness (B) representing the player's inherent property and the player's payoff (U), which is determined by the predefined payoff matrix and the player's interactions with others as follows:


$$\Phi = (1 - \alpha)B + \alpha U, \quad (1)$$

where the combining weight α is called the *selection intensity*. One can interpret that one's fitness is not only determined by one's own strength, but also from one's environment affecting with a selection intensity α . The case that $\alpha \rightarrow 0$ represents the limit of weak selection [23], while $\alpha \rightarrow 1$ denotes strong selection. The selection intensity can also be time varying, e.g., $\alpha = \beta e^{-\alpha t}$, which means that the contribution of game interaction decreases with time.

With the fitness function, the EGT studies and characterizes how a group of players converge to a stable equilibrium after a period of strategic interactions. Such a final equilibrium state is called the *evolutionarily stable state (ESS)*, which is "a strategy such that, if all members of the population adopt it, then no mutant strategy could invade the population under the influence of natural selection" [21]. In other words, even if a small fraction of players may not be rational and take out-of-equilibrium strategies, ESS is still a locally stable state. How to find the ESSs is an important issue in EGT. One common approach is to find the stable points of the system state dynamic, which is known as *replicator dynamics*. The corresponding underlying physical meaning is that, if adopting a certain strategy can lead to a higher fitness than the average level, the proportion of population adopting this strategy will increase, and the increasing rate is proportional to the difference between the average fitness with this strategy and the average fitness of the whole population. Note that when the total population is sufficiently large




[FIG3] MemeTracker phrase cluster during the 2008 U.S. presidential election: spreading of comments and phrasing by candidates [20].



[FIG4] The imitation strategy updating rule.

and homogeneous, the proportion of players adopting a certain strategy is equivalent to the probability of one individual player adopting such a strategy, i.e., the strategy distribution over the whole population can be interpreted as each player's mixed strategy, and the replicator dynamics can be interpreted as each player's mixed strategy update.

Graphical EGT studies the strategies' evolution in such a structured population [24]. In the EGT, in addition to the entities of players, strategy, and fitness matrix, each game model is associated with a graph structure, where the vertices represent players and the edges determine which player to interact with. Since the players only have limited connections with others, each player's fitness is locally determined from interactions with all adjacent players.



[FIG5] A Facebook subnetwork [16], [26].

The commonly used strategy updating rules [25] are originated from the evolutionary biology field and used to model the mutant evolution process. Figure 4 illustrates the detailed evolution procedures of the imitation strategy update rule. In the first step, a user is randomly chosen from the population for imitation. Then, the fitness of the chosen user and all corresponding neighbors is computed. Finally, the user will, in all probability, either be imitated by one of the neighbors or remain with the current strategy, with the probability being proportional to fitness. There are also other rules, such as the birth–death strategy update rule and death–birth strategy update rule, but through theoretical analysis [15], we find that these rules are equivalent when the network degree is sufficiently large.

INFORMATION-DIFFUSION FORMULATION AND ANALYSIS

A social network is usually illustrated by a graph, e.g., a Facebook subnetwork is shown in Figure 5, where each node represents a user, and lines represent the relationships between users. When some new information is originated from one user, the information may be propagated over the network, depending on other users' actions to forward the information or not. For each user, whether he or she forwards the information is determined by several factors, including the user's own interest in the information and the neighbor's actions, in the sense that, if all the neighbors forward the information, the user may also forward the information with a relatively high probability. In such a case, the users' actions are coupled with each other through their social interactions. This is very similar to the player's strategy update in the graphical evolutionary game, where players' strategies are also influenced with each other through the graph structure. In the graphical evolutionary game, a user's strategy can influence one of the neighbors when the fitness of adopting this strategy is high. Similarly, in the information-diffusion process, when forwarding the information can bring a user more utility, the user's neighbors may also be influenced to forward the information in the near future. Therefore, the information-diffusion process can be well modeled by the graphical evolutionary game, as illustrated in Figure 6.

There are two possible actions for each user, i.e., to forward (S_f) or not forward (S_n), and the corresponding users' payoff matrix can be written as

$$\begin{pmatrix} u_{ff} & u_{fn} \\ u_{fn} & u_{nn} \end{pmatrix}, \quad (2)$$

where a symmetric payoff structure is considered, i.e., when a user with strategy S_f meets a user with strategy S_n , each of them receives the same payoff u_{fn} . Note that the payoff matrix is related to the fitness in the graphical evolutionary game according to (1). The physical meaning of

the payoff can be either the popularity of a user in a social network or the hit rate of a website. These three parameters will be learned from the data and then used for decision making. Under different application scenarios, the values of the payoff matrix may be different. For example, if the information is related to recent hot topics and forwarding of the information can attract more attentions from other users or websites, the payoff matrix should have the following characteristic: $u_{ff} \geq u_{fn} \geq u_{nn}$. According to (1), the fitness of forwarding is larger, and, thus, the probability of forwarding will be higher. On the other hand, if the information is about useless advertisements, the payoff matrix would exhibit $u_{nn} \geq u_{fn} \geq u_{ff}$, i.e., the fitness of not forwarding is higher, and, thus, users tend not to forward the information. Furthermore, if the information is supposed to be shared only within a circle, i.e., a small group with the same interest, the payoff matrix could exhibit $u_{fn} \geq u_{ff} \geq u_{nn}$.

Since the player's payoff is determined by both his or her own strategy and the opponent's strategy, to characterize the global population dynamics, we need to first derive the local influence dynamics as well as the corresponding influence equilibria. We find in [15] that the local network states, i.e., the neighbors' strategy distribution given a player's strategy, evolve with a rate of order 1, while the global network state, i.e., the strategy distribution of the whole population, evolves with a rate at the order of the selection intensity α , which is much smaller than one due to the weak selection [23]. In such a case, the local network states will converge to equilibria at a much faster rate than the global network state. This is because the dynamics of local network states are only in terms of a local area, which contains only the neighbors. At such a small scale, the local dynamics can change and converge quite fast. On the other hand, if the dynamics of the global network state are associated with all users, i.e., the whole network, the dynamics would be much slower. Therefore, the global network state can be regarded as constant during the convergence of influence dynamics. By doing so, the equilibria of the local influence dynamics can be obtained, which are found to be linear functions of the global network state.

With the equilibria of the local influence dynamics, the global population dynamics can be derived through analyzing the strategy updating rules specified in the graphical evolutionary game [25]. It is found that the global population dynamics can be represented as a two-parameter, third-order polynomial function of the global network state [15]

$$\dot{p}_f(t) = \frac{\alpha(\bar{k}-1)(\bar{k}^2-2\bar{k})}{(\bar{k}^2-\bar{k})^2} p_f(t)[1-p_f(t)][ap_f(t)+b], \quad (3)$$

where $p_f(t)$ is the proportion of the population forwarding the information, $\dot{p}_f(t)$ is the corresponding dynamics, $\bar{k} = E[k]$ is the average degree of the network, $\bar{k}^2 = E[k^2]$ is the second moment of the degree of the network, and a and b are two parameters determined by the payoff matrix shown in (2).

From (3), we can see that, given the characteristic of the network, i.e., the average degree \bar{k} and the second moment of the degree \bar{k}^2 , the evolution dynamics of the information diffusion can be modeled by a simple two-parameter, third-order

polynomial function, where the two parameters a and b are determined by the payoff in the payoff matrix, i.e., u_{ff} , u_{fn} and u_{nn} . Therefore, by learning the payoff from the data, we are able to characterize the evolution dynamics of information diffusion using the evolutionary game-theoretic framework.

By evaluating the global population dynamics at the steady state, the global population equilibria can be found [16], which is zero (no user shares the information with the neighbors), one (all users share the information with their neighbors), or only a portion of users share the information with their neighbors where the amount of such users is purely determined by the payoff matrices as follows:

$$p_f^* = \begin{cases} 0, & \text{if } u_{nn} > u_{fn} > u_{ff}; \\ 1, & \text{if } u_{ff} > u_{fn} > u_{nn}; \\ \frac{(\bar{k}^2/\bar{k}-2)(u_{fn}-u_{nn})+(u_{ff}-u_{nn})}{(\bar{k}^2/\bar{k}-2)(2u_{fn}-u_{ff}-u_{nn})}, & \text{else.} \end{cases} \quad (4)$$

From (4), we can see that neither user forwarding the information can gain the most payoff, while both forwarding gains the least payoff, $p_f^* = 0$. This corresponds to the scenario where the released information is useless or a negative advertisement, forwarding that can only incur unnecessary cost. On the contrary, both users forwarding the information can gain the most payoff, while not forwarding gains the least payoff, $p_f^* = 1$. This corresponds to the scenario where the released information is an extremely hot topic, and forwarding it can attract more attention. For other cases, p_f^* lies between zero and one. For this third ESS, some approximations can be made as follows:

$$p_f^* = \frac{(\bar{k}^2/\bar{k}-2)(u_{fn}-u_{nn})+(u_{ff}-u_{nn})}{(\bar{k}^2/\bar{k}-2)(2u_{fn}-u_{ff}-u_{nn})}, \\ \doteq \frac{1}{1 + \frac{u_{fn}-u_{ff}}{u_{fn}-u_{nn}}}, \quad (5)$$

where the last approximation is due to $\bar{k}^2/\bar{k} \geq \bar{k}$ and the assumption that the average network degree $\bar{k} \gg 2$ in real social

Graphical EGT	Social Network
Graph Structure	Social Network Topology
Players	Users in the Social Network
Strategy	S_f : Forward the Information S_n : Not Forward the Information $\begin{matrix} S_f & S_n \\ S_f & \begin{pmatrix} u_{ff} & u_{fn} \\ u_{fn} & u_{nn} \end{pmatrix} \end{matrix}$
Fitness	Utility from Forwarding or Not
ESS	Stable Information-Diffusion State

[FIG6] Information diffusion as a graphical evolutionary game.

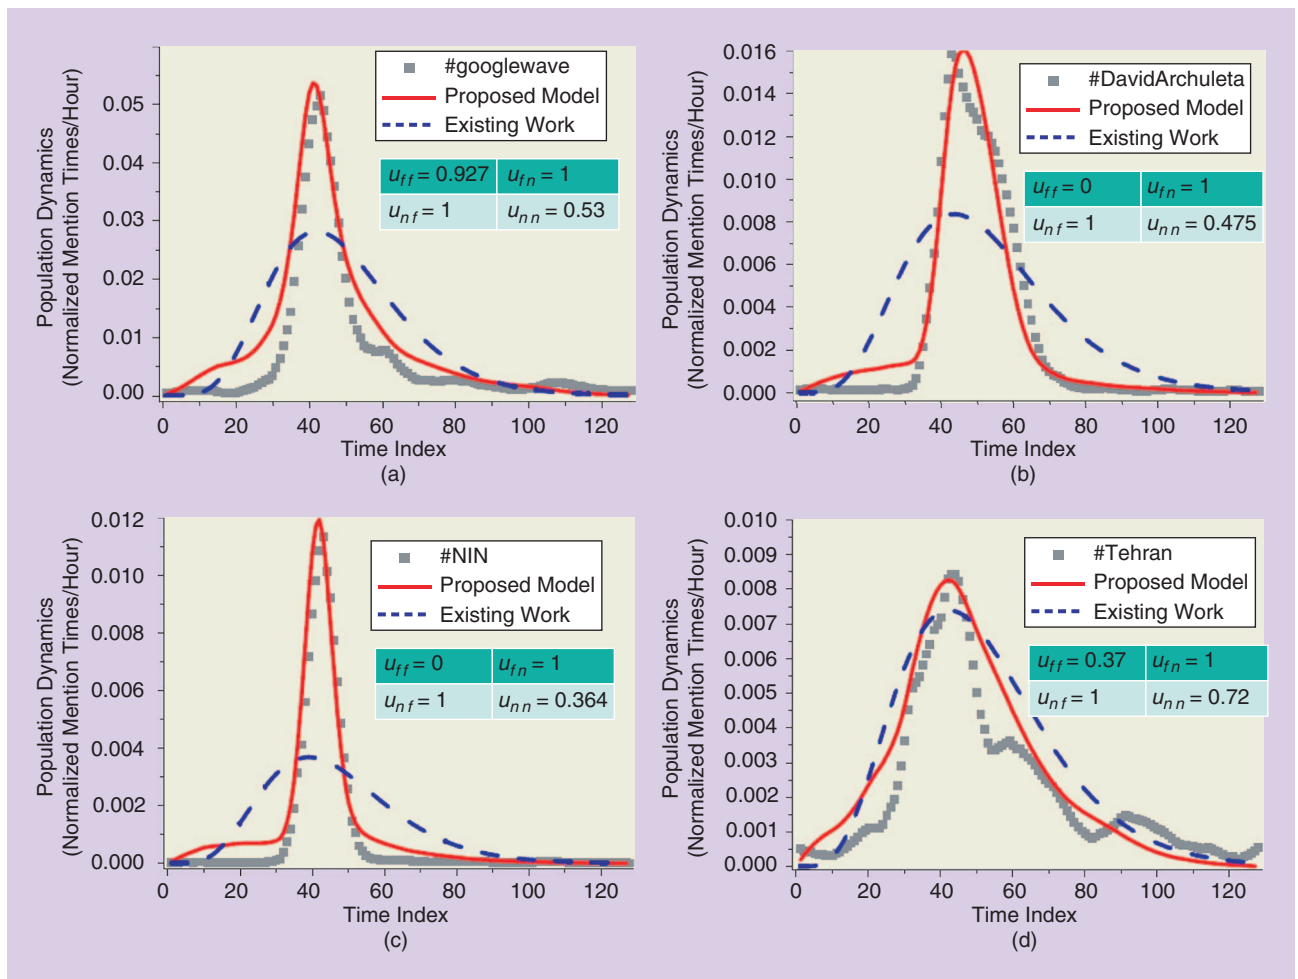
networks. We can see that, when the average network degree \bar{k} is sufficiently large, the information-diffusion result is independent of the network scale, i.e., there is a scale-free phenomenon for the information-diffusion equilibrium.

EXPERIMENTS WITH REAL-WORLD DATA SETS

The real-world data sets are used to validate the proposed model. We first use the Twitter hashtag data set to validate the evolutionary population dynamics [15]. Specifically, we learn the payoff matrix in (2) by fitting the real temporal dynamics with the evolution dynamics in (3) and generate the corresponding evolution dynamics based on the estimated payoff matrix. The Twitter hashtag data set contains the number of mentions per hour of 1,000 Twitter hashtags with the corresponding time series, which are the 1,000 hashtags with highest total mention times among 6 million hashtags from June to December 2009 [26]. We compare our results with one of the most related existing works using a data mining method [20]. Figure 7 shows the comparison results, where the vertical axis is the dynamics and the mention times of different hashtags per hour in the Twitter data set are normalized within interval [0, 1]

and denoted by solid gray square. From the figure, we can see that the game-theoretic model can fit very well the real-world information-diffusion dynamics, better than the data mining method in [20] since the users' interactions and decision-making behaviors are taken into account.

We then use the MemeTracker data set to validate the ESS [16]. The data set contains more than 172 million news articles and blog posts from 1 million online sources [20]. When a site publishes a new post, it will put hyperlinks to related posts in some other sites published earlier as its sources. Later, the site will also be cited by other newer posts as well. An example is shown in Figure 8. In such a case, the hyperlinks between articles and posts can be used to represent the spreading of information from one site to another. We extract five groups of sites, where each group includes 500 sites. Each group is regarded as a complete graph, and each site is considered a user. We divide the data set into two halves, where the first half is used to train the payoff matrix and the second half is used for testing. Figure 9 shows the results using the proposed model and the results from the real-world data set, from which we can see they match well with each other. We also depict the variances



[FIG7] Experimental results of the evolutionary population dynamics: (a) #googlewave, (b) #DavidArchuleta, (c) #NIN, and (d) #Tehran.

2 8 we're not commenting on that story i'm afraid 2131865
 3 3 we're not commenting on that 489007
 2008-08-18 14:23:05 1 M <http://business.theage.com.au/business/bb-chief-set-to-walk-plank-20080818-3xp7.html>
 2008-11-26 01:27:13 1 B <http://sfweekly.com/2008-11-26/news/buy-line>
 2008-11-27 18:55:30 1 B <http://aconstantineblacklist.blogspot.com/2008/11/re-researcher-matt-janovic.html>
 5 2 we're not commenting on that story 2131864
 2008-12-08 14:50:18 3 B <http://videogaming247.com/2008/12/08/home-in-10-days-were-not-commenting-on-that-story-says-scee>
 2008-12-08 19:35:31 2 B <http://jplaystation.com/2008/12/08/home-in-10-days-were-not-commenting-on-that-story-says-scee>

[FIG8] An example of a MemeTracker phrase-cluster data set [20].

of the estimated results in Figure 9, which shows that the simulated results are always in the variance interval of the corresponding estimated results. Figure 9 also reveals the cohesiveness of different groups. We can see that the sites in Group 5 behave cohesively or share major common interests, while the sites in Group 1 share relatively few common interests. This is particularly interesting in advertisement or advocacy scenarios where certain cohesive focus groups need to be mined to target high return value.

**SEQUENTIAL USER BEHAVIOR:
 THE CHINESE-RESTAURANT-GAME FRAMEWORK**

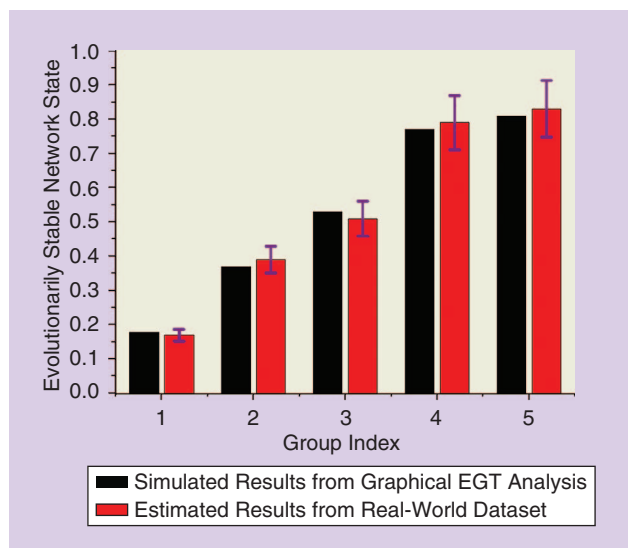
Another distinguishing feature of social systems is that users often participate sequentially in their own time and space. For example, users sequentially Q&A sites like Yahoo! Answers and Stack Overflow and decide whether to provide an answer, to vote on an existing answer, or not to participate. Other examples include online reviews, where customers write reviews for the product they purchase, and social news sites where online users post and promote stories under various categories.

The existence of network externality [5] in a social group dictates that users' actions and decisions influence each other. The network externality can be either positive or negative. When it is positive, users will have higher utilities when making the same decisions. On the contrary, when negative, users tend to make different decisions from others to achieve higher utilities. To achieve better performance, users should take into account the effect of network externality when making decisions.

On the other hand, users' decisions are also influenced by their knowledge of the system. In general, a user's knowledge of the system may be very limited because of the uncertainty in observations. This limitation reduces the accuracy of the user's decision and, thus, the overall system performance. The phenomenon of limited knowledge can be overcome through learning [27]–[30]. Users can learn from their previous experiences through machine-learning techniques and/or from other users' decisions and observations through social learning. All such information can help users to construct a belief, which can be probabilistic, on the

unknown system states. In most cases, the accuracy of users' decisions can be greatly enhanced by taking the belief into account.

Therefore, to achieve the best utilities, users need to consider the effects of both learning and network externality when making decisions. While there are some existing works on combining positive network externality with learning [31]–[33], few works have been done on combining negative network externality with learning in the literature, mainly due to the difficulty of the problem, where a user has to consider the previous users' decisions and predict those of the subsequent users' decisions. Furthermore, the information leaked by a user's decision may eventually impair the utility the user can obtain. However, in practice, negative network externality commonly exists in social systems where users share and/or compete for resources and contents. To address this issue, we have developed a joint learning-decision-making framework, called the *Chinese-restaurant game* [17], [18], to study users' sequential learning and decision-making behavior in social systems.




[FIG9] Experimental results of the evolutionarily stable strategy.

CHINESE-RESTAURANT-GAME FRAMEWORK

The well-known Chinese-restaurant process has been used in various fields, including machine learning, speech recognition, text modeling, and object detection in images and biological data clustering [34]. It offers an ideal structure to jointly formulate the decision-making problems with negative network externality. The Chinese-restaurant process is a nonparametric learning method for an unbounded number of objects in machine learning. In a Chinese-restaurant process, a restaurant has an infinite number of tables and customers arrive the restaurant sequentially. When a customer enters, he/she either joins one of the existing tables or requests a new table with a predetermined probability. However, there is not yet any notion of strategic decision making in the Chinese-restaurant process.

By introducing the strategic behavior into the nonstrategic Chinese-restaurant process, we proposed a new framework, the Chinese-restaurant game [17], [18], to study the learning and decision-making problem with negative network externality. To illustrate the framework, let us start with a Chinese restaurant that has a fixed number of tables, and customers sequentially come in requesting seats at these tables. Each customer may request a table. Since tables are available to all customers, there may be multiple customers requesting to sit at the same table, which, thereafter, incurs the negative network externality. We can imagine that the more personal space a customer has, the more comfortable the dining experience. Moreover, when the table sizes are unknown to the customers (before arriving at the restaurant), each of them may resort to some signals (e.g., through advertisements or discussions with previous customers) about the table sizes. By observing previous actions or signals, a user can exercise a learning process to make up the shortcoming of limited knowledge. With the proposed Chinese-restaurant game, we are able to develop an analytical framework involving the learning and decision making with negative network externality.

As shown in Figure 10, in the Chinese-restaurant game, there is a Chinese restaurant with K tables numbered $1, 2, \dots, K$ and N customers labeled $1, 2, \dots, N$. The table sizes are determined by the restaurant state $\theta \in \Theta$ and the table size functions $\{R_1(\theta), R_2(\theta), \dots, R_K(\theta)\}$. When customer i arrives, he or she receives a signal s_i about the state θ and makes a decision about which table to choose such that he or she can maximize his or her utility, based on what was observed and the prediction of future customers' decisions. The prior distribution of the state



[FIG10] A system model of the Chinese-restaurant game.

information is assumed to be known by all customers. The signal is generated from a predefined distribution. Since there are uncertainties on the table sizes, customers who arrive first may not choose the right tables, and, consequently, their utilities may be lower. On the other hand, customers who arrive later may eventually have better chances to get the better tables since they can collect more information to make the right decisions. In other words, when signals are not perfect, learning can help to result in higher utilities for customers choosing later. Therefore, there is a tradeoff between more choices when playing first and more accurate signals when playing later. To study this tradeoff, some questions need to be answered: How can customers learn from their own signals and the information revealed by other customers? How can customers predict the decisions of future customers? And what are the best strategies for the customers?

To study how customers learn from the revealed information from others and their own signals, we first introduce the concept of belief to describe customers' uncertainty about the system state. One customer's belief on the system state is the conditional probability of the system state, given all the information observed by the customer, as follows:

$$\mathbf{g}_i = \{g_{i,l} | g_{i,l} = P(\theta = l | \mathbf{h}_i, s_i, \mathbf{g}_0), \forall l \in \Theta\}, \quad (6)$$

where $\mathbf{h}_i = \{s_1, s_2, \dots, s_{i-1}\}$ is the signals observed by customer i , and $\mathbf{g}_0 = \{g_{0,l} | g_{0,l} = P(\theta = l), \forall l \in \Theta\}$ is the prior distribution.

With Bayesian learning [29], rational customers use Bayes' rule to find the optimal estimate about the system state and update their belief on the system state as follows:

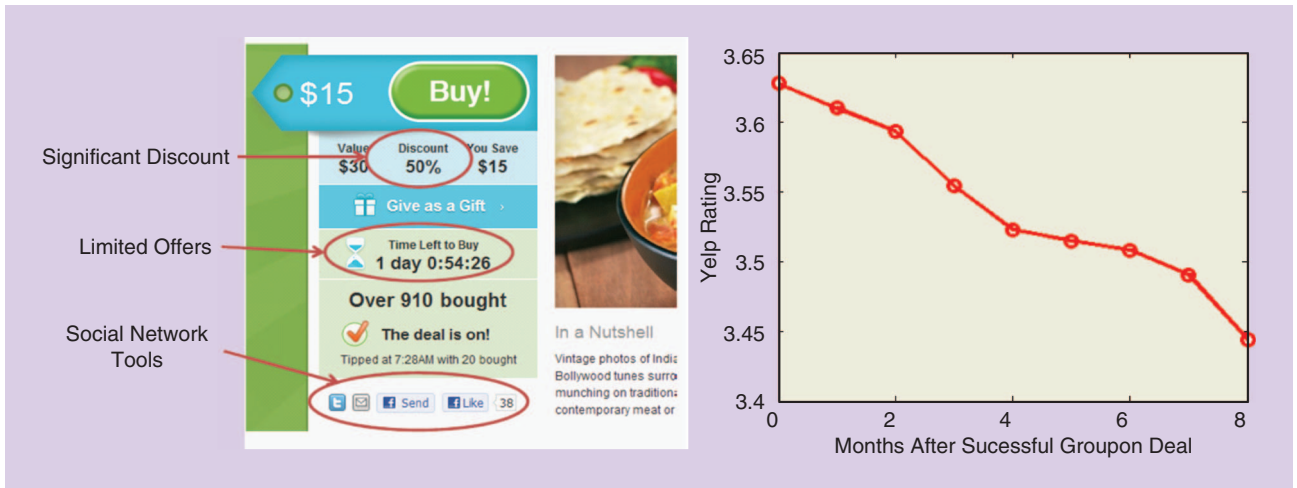
$$g_{i,l} = \frac{g_{0,l} P(\mathbf{h}_i, s_i | \theta = l)}{\sum_{l' \in \Theta} g_{0,l'} P(\mathbf{h}_i, s_i | \theta = l')}. \quad (7)$$

Because of their rationality and selfish nature, customers will choose their strategies to maximize their own utilities. In such a case, considering the incomplete information about the future customers, the best response of a customer is to maximize his or her expected utility based on all the observed information as follows:

$$\text{BE}_i(\mathbf{n}_i, s_i, \mathbf{h}_i) = \arg \max_j E[U_i(R_j(\theta), n_j^*) | \mathbf{n}_i, s_i, \mathbf{h}_i, x_i = j], \quad (8)$$

where n_j^* is the final number of customers choosing table j ; $U_i(R_j(\theta), n_j^*)$ is the utility of customer i choosing table j ; $\mathbf{n}_i = \{n_{i1}, n_{i2}, \dots, n_{iK}\}$ is the grouping observed by customer i , with n_{ik} being the number of customers choosing table k before customer i ; and x_i is the action of customer i .

Note that the best response is determined by the final grouping, which depends on the subsequent customers' decisions. Since the decisions of subsequent customers are unknown to a customer when the customer is making the decision, a closed-form solution to the best response function is generally impossible and impractical. To find the best response for each customer, a recursive method based on backward induction is designed [18]. The key idea is to use the next customer's best response $\text{BE}_{i+1}(\mathbf{n}_{i+1}, s_{i+1}, \mathbf{h}_{i+1})$ to derive the current customer's best response $\text{BE}_i(\mathbf{n}_i, s_i, \mathbf{h}_i)$.




[FIG11] Yelp star rating declines after a successful Groupon deal.

With the best-response function, a customer's optimal decision is purely determined by the received signal given the grouping, i.e., the number of customers at each table and the information revealed by other customers. Therefore, we can partition the signal space into subspaces where within each subspace the customer will choose a specific table. By integrating the signal over each subspace, we can derive a recursive form of the probability mass function for the final grouping, i.e., the final number of customers at each table. With the recursive form of the final grouping, the expected utility of each customer can be computed, and the best response of all customers using backward induction can be derived.

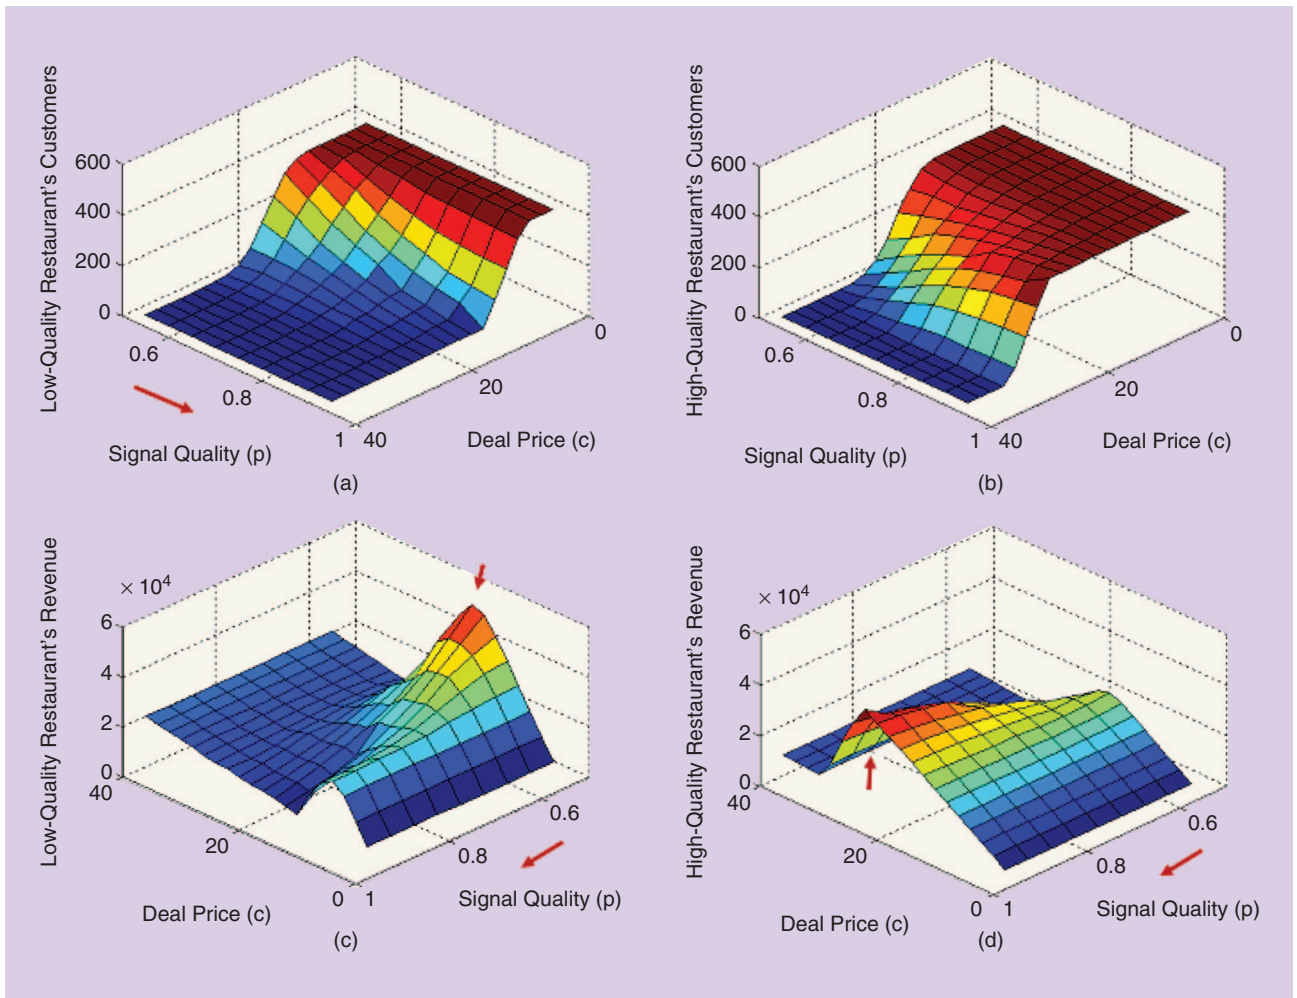
From previous discussions, we can see that the learning and decision making in the Chinese-restaurant-game framework are interweaved. On one hand, customers learn the system state from the information revealed by previous customers for better decision making. On the other hand, the decisions and information revealed by the customers will affect subsequent customers' learning and decision-making processes. Moreover, before any decision making, the utility function in (8) needs to be learned from the real data.

EXPERIMENTS WITH REAL DATA FROM SOCIAL SYSTEMS

We use deal selection on Groupon as an example to illustrate the Chinese-restaurant-game framework. Many have the experiences that some deals on Groupon look good but eventually turn out to have poor quality because of the overwhelming number of customers showing up at the same time, i.e., the negative network externality is at work. By collecting the data on Groupon and Yelp around the Washington, D.C., area for eight months, we indeed observe the decline of Yelp review ratings after some successful Groupon deals, as depicted in Figure 11. One can see a nonlinear decline function in review ratings. Let us use the real Yelp rating data to train the utility function of customers by approximating it as a linear model, as shown in Figure 12(a). Then, based on (8), we evaluate the



[FIG12] (a) Utility function modeling using real data from Groupon and Yelp and (b) the performance comparison of our method with the learning method without negative network externality.



[FIG13] A new restaurant's strategy: the number of customers choosing the restaurant with (a) low quality and (b) high quality and the result of the revenue of the (a) low-quality and (b) high-quality restaurants.

average utility, which is the average review rating customers can obtain. A comparison was made between the decision-learning method, denoted as *learning with negative network externality*, and that which does not consider negative network externality, denoted as *learning without negative network externality*. Note that the former considers the interplay between the learning and decision making, while the latter only considers the learning of system state but totally ignores the influence among customers' decision making. The results are shown in Figure 12(b). One can see that by combining learning with negative network externality, the proposed method can achieve much better utility for customers.

We further study the best pricing and promotion strategy of a new restaurant under the Chinese-restaurant-game framework. Let us consider two restaurants. One is always of high quality, and the other is a new restaurant, which could be of low or high quality. The same utility function trained from the real data in the above experiment is used to infer the strategy. The results are shown in Figure 13. One can see that if the new restaurant is of low quality, then the number of customers

choosing the new restaurant decreases as signal quality increases and vice versa. One can also see that the optimal deal price of the high-quality restaurant is higher than that of the low-quality restaurant. Therefore, the high-quality restaurant should make every effort to increase the signal quality, while the low-quality restaurant should hide the quality information and use a low deal price to attract customers to increase the revenue. This offers a vivid example of using data to learn and come out with an optimal strategy.

EXTENSION TO THE CHINESE-RESTAURANT-GAME FAMILY

We have discussed the Chinese-restaurant game under a fixed population setting, i.e., there are a finite number of customers choosing the tables sequentially. However, in some applications, customers may arrive and leave the restaurant at any time, which results in the dynamic population setting. Examples include cloud storage service selection, deal selection on Groupon, and Wi-Fi access point selection in a conference hall [35]. In such a case, the utilities of customers will change from time to time because of the dynamic number of customers on each

table. To tackle this challenge, we have extended the Chinese-restaurant game to the dynamic population setting [36], [37], where we consider the scenario that customers may arrive and leave the restaurant with, e.g., a Poisson process. With such a dynamic population setting, each newly arriving customer not only learns the system state according to the information received and revealed by former customers, but also predicts the future customers' decisions to maximize the utility.

The Chinese-restaurant game is proposed by introducing strategic decision-making into the Chinese-restaurant process, where each customer can choose one table to maximize the utility. However, in some applications, users may want to simultaneously choose multiple resources. For example, mobile terminals may access multiple channels, cloud users may have multiple cloud storage services, and students may take multiple online courses. To further generalize the setting, we have introduced strategic decision making into another well-known random process, the Indian-buffet process [38], and develop a new framework, called the *Indian-buffet game*, to study the learning and decision-making problem with negative network externality under the scenario that customers can have multiple choices [39]. In the Indian-buffet game framework, we also consider multislot interactions where customers can interact and make decisions repeatedly, and only partial information is revealed, i.e., customers only reveal beliefs instead of full signals to others. We use the non-Bayesian social learning to learn from each other to improve the knowledge of the system and thus make better decisions. Similar extension can be applied to multiarmed-bandit problems by introducing decision-making processing into their formulation.

ONE KEY FACTOR FOR THE SUCCESS OF SUPERVISED AND SEMISUPERVISED LEARNING IS THE LARGE-SCALE LABELED DATA SET.

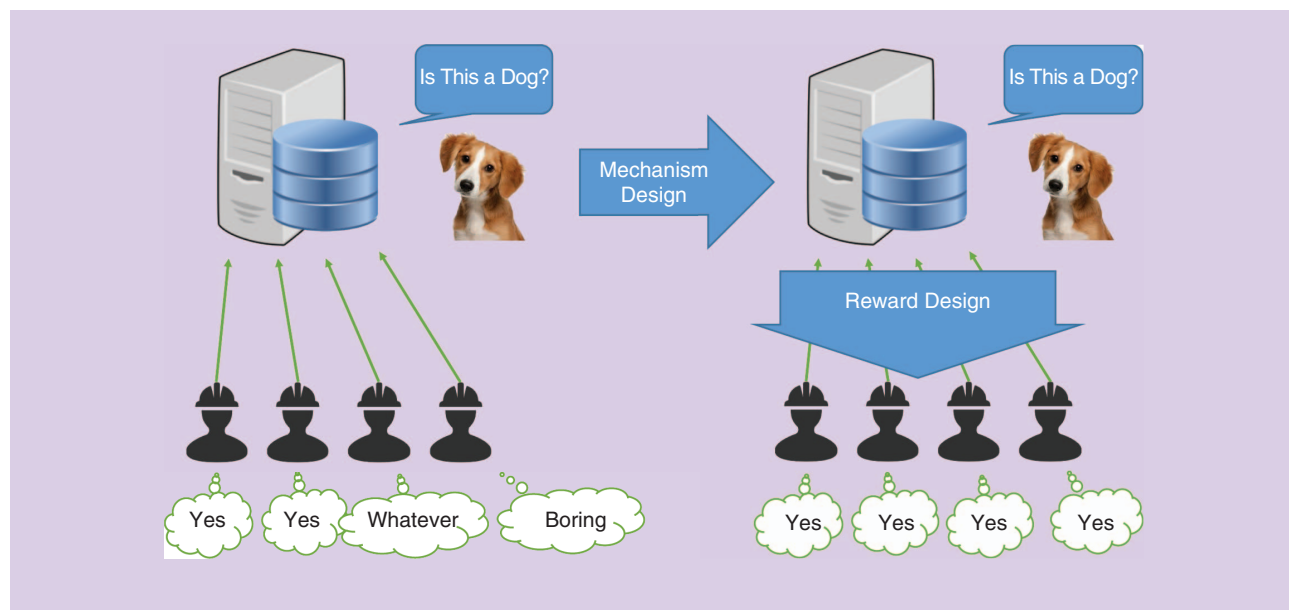
MECHANISM DESIGN IN DECISION LEARNING

In this section, we will address decision learning from the system point of view. Can we design mechanisms for users to learn the desired behavior

and thus achieve the goals of the system designer, as shown in Figure 14? In the following, we will use microtask crowdsourcing to illustrate how to design a mechanism to obtain high-quality data for data analytics.

One key factor for the success of supervised and semisupervised learning is a large-scale labeled data set [1], [2]. In general, a larger-scale data set will lead to a more accurate model and, thus, better performance. However, large-scale annotation is very expensive, which often becomes one of the bottlenecks of supervised and semisupervised learning. To address this challenge, microtask crowdsourcing, with access to a large and relatively cheap online labor pool, is a promising method since it can generate large volume of labeled data in a short time at a much lower price compared with traditional in-house solutions. An example of microtask crowdsourcing is illustrated in Figure 15.

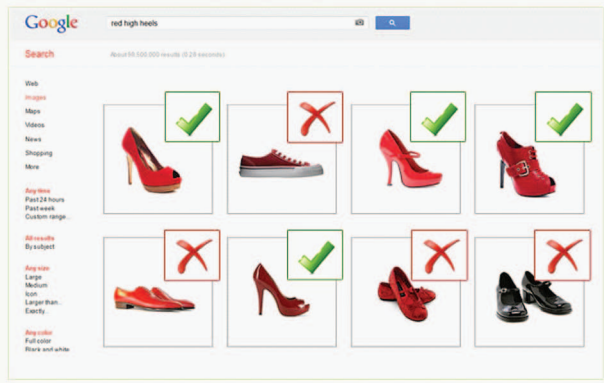
On the other hand, because of the lack of proper incentives, microtask crowdsourcing suffers from quality issues. Since workers are paid a fixed amount of money per task they complete, it is profitable for them to provide random or bad-quality solutions to increase the number of submissions within a certain amount of time or with the least effort. It has been reported that most workers on Mturk, a leading marketplace for microtask crowdsourcing, do not contribute high-quality work [40]. To address this issue, a common machine-learning solution is to either add a data-curation



[FIG14] An illustration of mechanism design. Users may behave randomly without a well-designed mechanism, which can greatly degrade the system performance. Through mechanism design, the system designer can guide users to learn the desired behavior and thus achieve the expected goals at the system level.

Example: Search Relevance – Image Results

Identify Relevant Product Results for "Red High Heels"

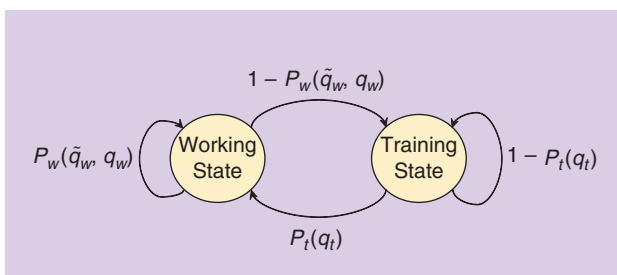


[FIG15] An example of microtask crowdsourcing.

phase to filter out low-quality data or to modify the learning algorithm to accept noisy labels [41]–[45].

In contrast to existing machine-learning solutions, we tackle such a problem by incentivizing the high-quality data in the first place [19], e.g., from the workers. This problem is challenging because of the inherent conflict between incentivizing high-quality solutions from workers and maintaining the low-cost advantage of microtask crowdsourcing for requesters. On one hand, requesters typically have a very low budget for each task in microtask crowdsourcing. On the other hand, the implementation of incentive mechanisms is costly as the operation of verifying the quality of submitted solutions is expensive [41]. Such a conflict makes it extremely challenging to design proper incentives for microtask crowdsourcing and motivates us to ask the following question: What incentive mechanisms should requesters employ to collect high-quality solutions in a cost-effective way? In a general sense, the core problem is how to design mechanism for obtaining good data.

To answer the question, we first study and model the behavior of workers. Specifically, let us consider a model with strategic workers, where the action of a worker is the quality of the solution $q \in [0, 1]$ and the primary objective of a worker is to maximize his or her own utility, defined as the reward he or she will receive minus the cost of producing solutions of a certain quality $c(q)$. Based on this model, we analyze two basic mechanisms that are widely adopted in existing microtask crowdsourcing applications: reward consensus mechanism M_c and reward accuracy mechanism M_a [19].



[FIG16] The state transition diagram of the incentive mechanism M_t .

REWARD CONSENSUS MECHANISM M_c

With this mechanism, a task is assigned to multiple workers. Only the same answer that is submitted by the majority of workers will be chosen as the correct solution and the workers whose solution agrees with the correct one will receive a positive reward. Through analyzing this mechanism, we find that there exists a minimum mechanism cost per task to obtain high-quality solutions [19]

$$C_{M_c}^* = 3c'(1), \quad (9)$$

where $c'(1)$ is the first-order derivative of the cost function $c(q)$ evaluating at the desired solution quality $q = 1$.

REWARD ACCURACY MECHANISM M_a

This mechanism assigns each task only to one worker. The requester evaluates with a certain probability the quality of submitted solutions directly, where each validation incurs a constant cost d . The validation can be erroneous with a probability of ϵ . The workers whose solutions are not evaluated or evaluated and confirmed as correct will receive a positive reward. Through analyzing this mechanism, we again find that there exists a minimum mechanism cost per task to obtain high-quality solutions [19]

$$C_{M_a}^* = \begin{cases} 2\sqrt{\frac{c'(1)d}{1-2\epsilon}} - \epsilon \frac{c'(1)}{1-2\epsilon}, & \text{if } d \geq \frac{c'(1)}{1-2\epsilon}; \\ \frac{c'(1)(1-\epsilon)}{1-2\epsilon} + d, & \text{otherwise.} \end{cases} \quad (10)$$

From (9) and (10), we can see that to obtain high-quality solutions using the two basic mechanisms (M_c and M_a), the unit cost incurred by requesters per task is subject to a lower bound constraint, which is beyond the control of requesters. In case that the budget of the requester is lower than the minimum cost constraint, it becomes impossible for the requester to achieve the desired quality solutions with these two basic mechanisms. In other words, neither of these two basic mechanisms is cost effective.

INCENTIVE MECHANISM VIA TRAINING M_t

To tackle this challenge, we design a cost-effective mechanism by employing quality-aware worker training as a tool to stimulate workers to provide high-quality solutions [19]. Different from current microtask crowdsourcing applications where training tasks are usually assigned to workers at the very beginning and are irrelevant to the quality of submitted solutions, we use the training tasks in a more effective way by assigning them to workers when they perform poorly. That is, when a worker performs poorly, he/she will be required to enter a training session without a reward to regain accreditation to be able to go back to perform in the regular session with a reward.

With the introduction of quality-aware training tasks, there will be two system states in our proposed mechanism: the working state and the training state. The working state is for production purposes, where workers work on standard tasks in return for rewards, while the training state is an auxiliary state, where workers do a set of training tasks to gain qualifications for the

working state. The state transition diagram is shown in Figure 16, where $P_w(\tilde{q}_w, q_w)$ represents the probability of a solution with quality q_w being accepted in the working state when other submitted solutions from the working state are of quality \tilde{q}_w , and $P_t(q_t)$ is the probability of a worker who produces solutions of quality q_t at the training state being allowed to enter the working state next time.

From Figure 16, we can see that the current action of a worker will affect the future system state of the worker. In other words, the quality of a worker's solution to one task will affect not only the worker's immediate utility but also his future utility due to the possible change of the system state. Such a dependence provides requesters with an extra degree of freedom in designing incentive mechanisms and thus enables them to collect high-quality solutions while still having control over their incurred costs.

To find the optimal action, each worker must solve a Markov decision process (MDP), according to the state transition diagram shown in Figure 16; the MDP faced by each worker also depends on other workers' actions. In essence, this is a challenging game-theoretic MDP problem [19]. Through analyzing the incentive mechanism M_t , we find that, as long as the number of training tasks is large enough, there always exists a desirable equilibrium where workers submit high-quality solutions at the working state. In other words, given any parameters in the working state, one can always guarantee the existence of a desirable equilibrium through the design of the training state. When the desirable equilibrium is adopted by all workers by following a certain design procedure, the minimum mechanism cost is theoretically proved to be zero [19], i.e.,

$$C_{M_t}^* = 0, \quad (11)$$

A COMMON MACHINE-LEARNING SOLUTION IS TO EITHER ADD A DATA-CURATION PHASE TO FILTER OUT LOW-QUALITY DATA OR TO MODIFY THE LEARNING ALGORITHM TO ACCEPT NOISY LABELS.

which means that one can collect high-quality solutions with an arbitrarily low cost. In other words, given any predetermined budget, the incentive mechanism M_t enables the requester to collect high-quality solutions while still staying within the budget.

Notice that one can easily achieve better learning purposes with the high-quality data collected by the


incentive mechanism M_t . Therefore, through modeling and analyzing users' strategic decision-making processes, one can design mechanisms from the system point of view to steer users' strategic behaviors to obtain better-quality data for better learning.

REAL BEHAVIORAL EXPERIMENTS

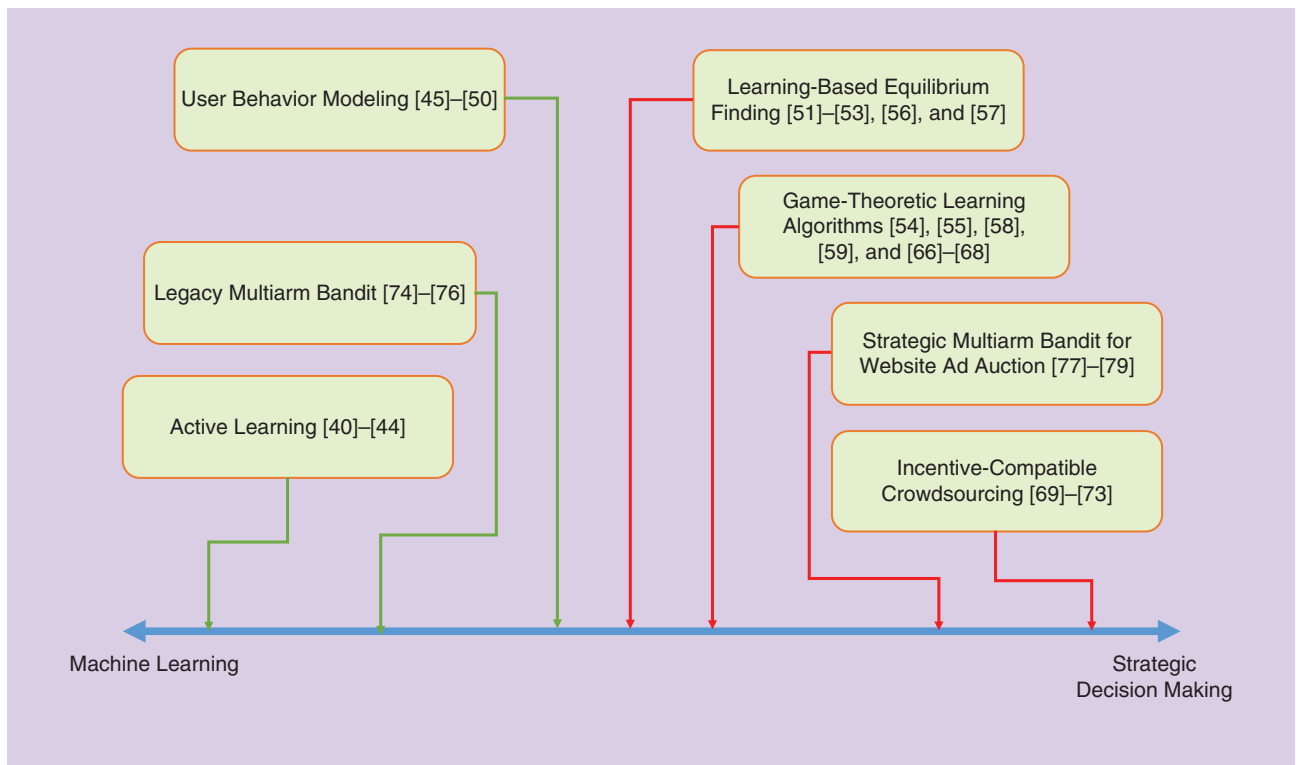
A set of behavioral experiments are conducted to test the incentive mechanism M_t in practice. We evaluate the performance of participants on a set of simple computational tasks under different incentive mechanisms. We compare the incentive mechanism M_t with the reward accuracy mechanism M_a , where the quality of submitted solutions is evaluated with a certain probability.

There are 41 participants in our experiments, most of whom are engineering graduate students. We use the accuracy of each participant as an indicator to the effectiveness of the incentive mechanisms, and the results are shown in Figure 17, where (a) shows the results of the reward accuracy mechanism M_a with sampling probability 1, i.e., every submitted solution is evaluated; (b) shows the results of the reward accuracy mechanism M_a with sampling probability 0.3, i.e., 30% of the submitted solutions are evaluated; and (c) shows the results using the incentive mechanism M_t with the sampling probability as that in (b).

As shown in Figure 17(a), with the highest sampling probability, most participants respond positively by submitting solutions with very high qualities. There is only one participant who had relatively low accuracy compared with others in that he/she



[FIG17] Histogram of accuracy: (a) the results using the reward accuracy mechanism M_a with sampling probability 1; (b) the results using the reward accuracy mechanism M_a with sampling probability 0.3; and (c) the results using the incentive mechanism M_t with sampling probability 0.3.



[FIG18] A summary of related works.

was playing the strategy of avoiding difficult tasks according to our exit survey. When a much lower sampling probability of 0.3 is used, it becomes profitable to increase the number of submissions by submitting lower-quality solutions, as most errors will simply not be detected. This explains why the majority of participants had very low accuracies, as shown in Figure 17(b). Notably, a few workers, five of 41, still exhibited very high accuracies in this case. Our exit survey suggests that their behaviors are influenced by a sense of work ethics, which prevents them from playing strategically to exploit the mechanism vulnerability. With the incentive mechanism M_t , as the introduction of training tasks makes it more costly to submit wrong solutions, participants need to reevaluate their strategies to achieve a good tradeoff between accuracy and the number of submitted tasks. From Figure 17(c), we can see that the accuracy of the participants with the incentive mechanism M_t has a very similar distribution to that of the group using the reward accuracy mechanism M_a with the highest sampling probability. Therefore, through the use of quality-aware worker training, the incentive mechanism M_t can greatly improve the effectiveness of the basic reward accuracy mechanism M_a with a low sampling probability to a level that is comparable to the one that has the highest sampling probability.

RELATED WORKS

Although not referred to specifically as decision learning, there has been a growing body of literature in recent years on the intersection of learning and strategic decision making, as

summarized in Figure 18. One class of related works is learning to understand how human beings make strategic decisions from real data. For example, classical machine-learning techniques are used in [46] to predict how people make and respond to offers during negotiations and how they reveal information and their response to potential revelation actions by others. Their results showed that the strategies derived from machine-learning algorithms, even when not optimal, can beat real human beings [46]. The year-long study of empirical data shows that an experienced human being in a repeated game will be more cooperative but turn the tables more definitely when he or she is betrayed by the opponent [47]. Additionally, the study in [48] shows that human beings have very limited memory space and computation capability, which limits the optimality of their decisions. It has also been shown in [49] that a dynamic belief model, by ignoring the older signals in constructing the belief, works best in predicting human decisions. Through empirically analyzing the purchase history on Taobao, a large-scale online shopping social network, Guo et al. revealed that a real human values purchase experiences shared by his or her friend and would be willing to pay a higher price for trustworthy vendors [50]. Nevertheless, in such a complicated system, it is still difficult to predict the purchase decisions with more than 50% accuracy using traditional machine-learning algorithms [50]. In [51], how users make decisions on social computing systems is learned from real data and used to guide the design of mechanisms for the systems. The aforementioned works establish a solid foundation of decision learning. These

works reveal the real nature of people in making decisions and suggest that people usually behave in a rational way with several limitations such as short memory. These evidences become a strong support to the rationality assumption in decision learning.

Another class of related works is finding equilibrium through learning. Finding Nash equilibrium is critical yet challenging in most game models since the difficulty has been shown to be polynomial parity arguments on directed graphs in general settings or even nondeterministic polynomial time-complete in specific problems [52]. Given that a general and exact solution is intractable, it is a natural choice to design proper learning algorithms to find the solutions. No-regret learning, for instance, has been shown to be a practical candidate. It has been applied in extensive-form games to reduce the number of subgame trees to explore [53]. The sufficient conditions for such type of learning algorithms to converge in the selfish routing problem [54] are also theoretically studied. Reinforcement learning is another candidate since its action-reward structure naturally forms the best response dynamic in game theory. Since traditional Q-learning may fail to converge if directly applied in a game, especially when the Nash equilibrium is not unique, maxmin Q-learning is proposed in [55] to find the Nash equilibrium in a two-player zero-sum game. The objective of maxmin Q-learning is modified from pure reward maximization into a max–min problem with an opponent's actions in mind. Nash Q-learning, a more general Q-learning algorithm, is proposed in [56] to handle multiplayer game with non-zero sum. The objective of Nash Q-learning is replaced with equilibrium conditions defined in game theory. The experiment results show that Nash Q-learning can help identify better Nash equilibrium than the traditional Q-learning algorithm. Learning has also been used to reduce the complexity in finding the subgame-perfect Nash equilibrium in a sequential game [57], which is PSPACE-hard in general. In [58], MDP and Monte Carlo simulation is used to reduce the complexity in identifying the optimal bidding strategy in sequential auctions. Finding the equilibrium point is also a critical part in decision learning since the performance of the decision-learning method relies on the equilibrium point. The aforementioned works contribute useful methods to help find the equilibrium point efficiently.

There have also been some related works that formulate the training problem in machine learning as a game. For instance, it was shown in [59] that a class of online learning algorithms can be modeled as a drifting game with both the trainer and the system as players. The learning algorithm in such a formation becomes the best response of the trainer to the system's reply to each training problem. Another application is maintaining fairness in a multiagent sequential decision problem. Given that the objective of the system is max–min fairness, one may model the learning model as a two-player game, where the first player

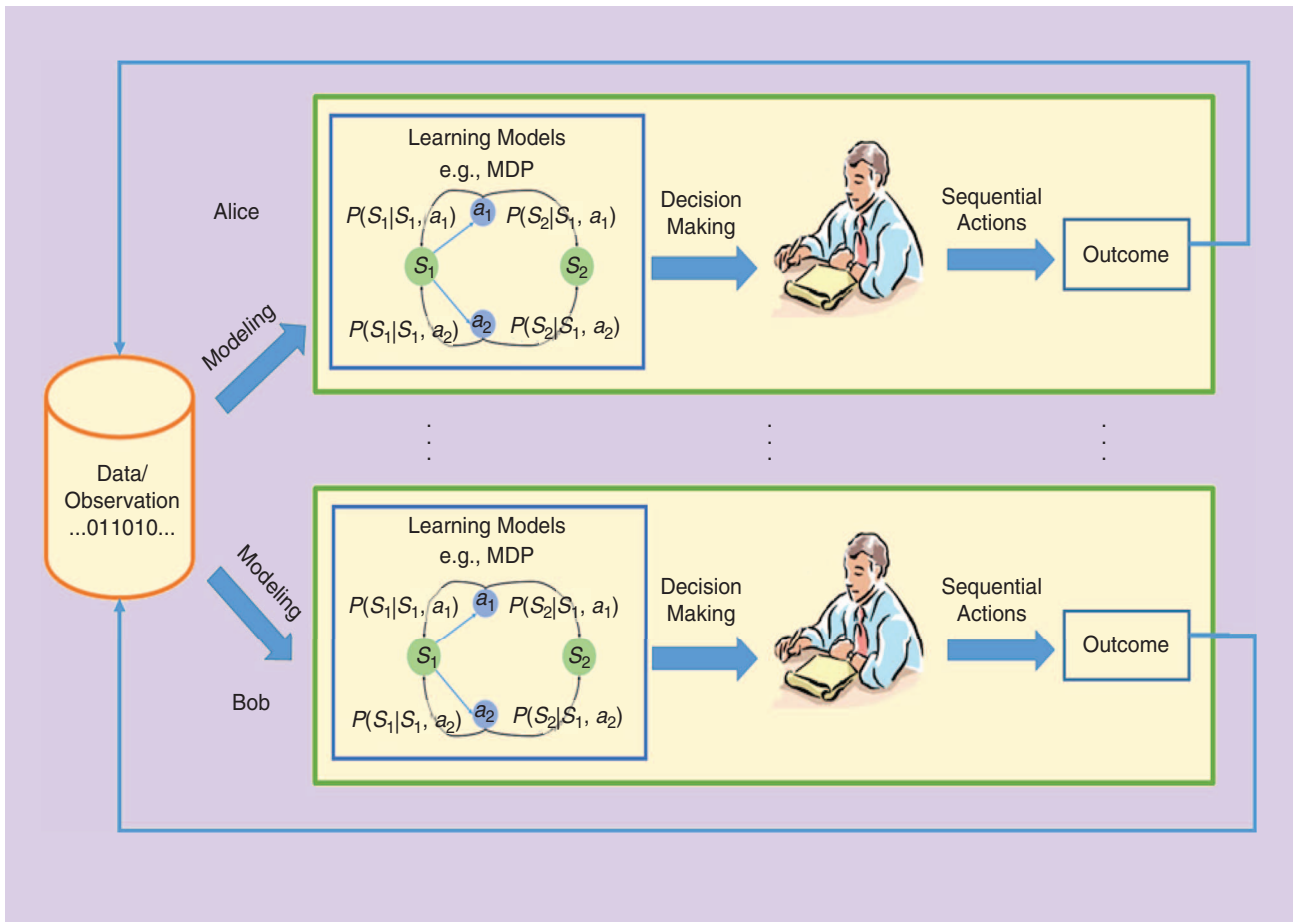
DECISION LEARNING IS AN EMERGING RESEARCH AREA TO BRIDGE LEARNING FROM LARGE VOLUMES OF DATA WITH STRATEGIC DECISION MAKING THAT UNDERSTANDS AND MODELS USER BEHAVIORS.

aims at maximizing the utility of the target agent who is chosen by second player, while the second player chooses the agent with lowest utility as the target agent [60]. Decision learning not only inherits the spirit of the previously mentioned works but also makes one critical extension to address the competition effect in real world.

That is, in decision learning, we introduce the game-theoretic competition effect into not only the system-versus-user relation but also user-versus-user relations.

Active learning is another related field [61]. Through actively choosing which data to learn from, active learning algorithms have the potential to greatly reduce the amount of labeling effort in machine-learning algorithms. Active learning with an explicit labeling cost has been widely studied [62]–[65]. It is known that active learning algorithms degrade quickly as the noise rate of labels increases [66]. To address the quality issue in labeled data collection, a variety of approaches have been proposed to filter low-quality labels and to increase the robustness of machine-learning algorithms [41]–[45]. In [67], a game-theoretic dynamics was proposed to approximately denoise the data to exploit the power of active learning. Incentive mechanisms have also been used to improve the quality of collected data. In [11], [68], and [69], all-pay auctions are applied to incentivize high-quality user contributions. In [70], Shaw et al. conducted an experiment to compare the effectiveness of a collection of social and financial incentive mechanisms. A reputation-based incentive mechanism was proposed and analyzed for microtask crowdsourcing in [71]. In [72] and [73], Singer and Mittal proposed an online mechanism for microtask crowdsourcing where tasks are dynamically priced and allocated to workers based on their bids. In [74], Singla and Krause proposed a posted price scheme where workers are offered a take-it-or-leave-it price offer. The earlier works are closely related to decision learning since they also focus on the rationality and incentives of users in the learning process. Crowdsourcing is one of the main applications of decision learning, of which we also gave a detailed illustration previously.

A nonstrategic but related problem is the multiarmed-bandit problem [75]. In the multiarmed-bandit problem, a bandit with multiple arms is provided to a gambler. The gambler may have different levels of rewards by playing different arms each time. Thus, the gambler may try and learn in each play to maximize his collected rewards. Liu and Zhao extended this model by considering multiple agents and including the network externality in [76]. They studied how agents learn the expected payoff and the other agent's choice by estimating the regrets after choosing different arms based on his or her current belief. A multiarmed-bandit problem with costs in observations is discussed in [77]. Nevertheless, traditional studies in the multiarmed-bandit problem are generally nonstrategic. They assume agents will always follow the



[FIG19] In the coming big data tsunami, when a large volume of data is available, users can learn better models to improve their own decision making. On the other hand, their actions result in changes of the data pool, which consequently affects the models learned by the users. In summary, users' decisions and actions affect each other in an ever-changing fashion for user-generated data applications. Decision learning is an emerging research area to bridge learning from large volumes of data with strategic decision making that understands and models user behaviors.

learning rule designed by the system designer. Combining strategic thinking with the multiarmed-bandit problem has gained more attention recently mainly because of a popular and practical application: website ad auctions. The ad slot on a website is usually sold through auction. The value of the ad depends on two factors: the value of the product in the ad and the expected number of clicks on the ad. The former is known by the advertiser and can be collected through truthful auction such as a Vickery auction. Nevertheless, the expected number of clicks, or the click-through-rate (CTR) of the ad, is unknown to both the website owner and the advertiser. The no-regret algorithm in the multiarmed-bandit problem can be used to learn CTR while maintaining the truthfulness of the auction [78], [79]. In [80], it is shown that increasing number of explore stages will push the buyers to reveal their true valuation more, with fewer exploit stages for sellers to gain extra revenue from the learned valuation in return. The strategic multiarm-bandit problem is the one receiving the most attention in the transition from traditional machine learning to decision learning because of its structure and potentials in real world applications such as ad auctions. The great success also

shows the potential of decision learning in real-world applications, especially those involving competitions.

CONCLUSIONS AND FINAL THOUGHTS

Decision learning is learning with strategic decision making that can analyze users' optimal behaviors from users' perspectives while designing optimal mechanisms from system designers' perspectives. In this article, we have used three social media examples to highlight the concepts of decision learning. Specifically, information diffusion over online social networks was used to illustrate how to learn users' utility functions from real data for understanding and modeling strategic decision making. Deal selection on Groupon with Yelp data was used to discuss how users learn from each other's interactions for better strategic decision making. Microtask crowdsourcing was used to discuss how to design mechanisms to steer users' strategic behaviors to obtain better-quality data for better learning. Besides the three examples discussed in this article, there can be other forms of joint learning with strategic decision making, including those discussed previously. In essence, in the coming big data tsunami, when a large

volume of data is available, users can learn better models to improve their own decision making, as depicted in Figure 19. On the other hand, their actions result in changes of the data pool, which consequently affects the models learned by the users.

In summary, users' decisions and actions affect each other in an ever-changing fashion for user-generated data applications. Decision learning is an emerging research area to bridge learning from large volumes of data with strategic decision making that understands and models user behaviors.

AUTHORS

Yan Chen (yan@umd.edu) received the Ph.D. degree from the University of Maryland, College Park, in 2011. A founding member and a principal technologist, he joined Origin Wireless, Inc., in 2013. He is currently a professor at the University of Electronic Science and Technology of China. His research interests are in data science, game theory, signal processing, and wireless communications. He is the recipient of the Best Paper Award from IEEE GLOBECOM in 2013, the Future Faculty Fellowship and Distinguished Dissertation Fellowship Honorable Mention in 2010 and 2011, respectively, and the Chinese Government Award for outstanding students abroad in 2011.

Chunxiao Jiang (chx.jiang@gmail.com) received the B.S. degree from Beihang University in 2008 and the Ph.D. degree from Tsinghua University (THU), Beijing, in 2013. From 2011 to 2014, he was a visiting researcher with the Signals and Information Group at the University of Maryland, College Park, with Prof. K.J. Ray Liu. He currently holds a postdoctoral position with the Department of Electrical Engineering, THU. His research interests include wireless communication and networking and social networks. He was a recipient of the Best Paper Award from IEEE GLOBECOM in 2013, the Beijing Distinguished Graduate Student Award, the Chinese National Fellowship, and the Tsinghua Outstanding Distinguished Doctoral Dissertation award in 2013.

Chih-Yu Wang (cywang@citi.sinica.edu.tw) received the B.S. degree in electrical engineering from National Taiwan University in 2007. He received the Ph.D. degree from the Graduate Institute of Communication Engineering, National Taiwan University. He was a visiting student at the University of Maryland, College Park, in 2011. He is currently an assistant research fellow at the Research Center for Information Technology Innovation, Academia Sinica. His research interests include game theory, wireless communications, social networks, and data science.

Yang Gao (yanggao.umd@gmail.com) received the B.S. degree from Tsinghua University, Beijing, China, in 2009 and the Ph.D. degree from the University of Maryland, College Park, in 2014, both in electrical engineering. Currently, he is a research scientist with Facebook, Inc. His research interests are in strategic behavior analysis and incentive mechanism design for crowdsourcing and social computing systems. He is a recipient of the silver medal of the 21st National Chinese Physics Olympiad, the honor of Excellent Graduate of Tsinghua University in 2009, the University of Maryland Future Faculty Fellowship in 2012, and the IEEE GLOBECOM 2013 Best Paper Award.

K.J. Ray Liu (kjrlu@umd.edu) was named a Distinguished Scholar-Teacher of the University of Maryland, College Park, in 2007, where he is Christine Kim Eminent Professor of Information Technology. He was a recipient of the 2016 IEEE Leon K. Kirchmayer Technical Field Award on graduate teaching and mentoring, IEEE Signal Processing Society 2014 Society Award, and IEEE Signal Processing Society 2009 Technical Achievement Award. Recognized by Thomson Reuters as a highly cited researcher, he is a Fellow of the IEEE and AAAS. He is on the IEEE Board of Directors. He was the 2012–2013 president of the IEEE Signal Processing Society.

REFERENCES

- [1] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2006.
- [2] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning*. Cambridge, U.K.: MIT Press, 2012.
- [3] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, U.K.: MIT Press, 1991.
- [4] M. J. Osborne, *An Introduction to Game Theory*. Oxford, U.K.: Oxford Univ. Press, 2004.
- [5] M. Katz and C. Shapiro, "Technology adoption in the presence of network externalities," *J. Political Econ.*, vol. 94, no. 4, pp.822–841, 1986.
- [6] Y. Chen and K. J. R. Liu, "Understanding microeconomic behaviors in social networking: An engineering view," *IEEE Signal Processing Mag.*, vol. 29, no. 2, pp. 53–64, 2012.
- [7] H. V. Zhao, W. Lin, and K. J. R. Liu, *Behavior Dynamics in Media-Sharing Social Networks*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [8] M. Burke, C. Marlow, and T. Lento, "Feed me: Motivating newcomer contribution in social network sites," in *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2009.
- [9] K. Ling, G. Beenen, X. Wang, K. Chang, D. Frankowski, P. Resnick, and R. Kraut, "Using social psychology to motivate contributions to online communities," *J. Comput. Mediate. Commun.*, vol. 10, no. 4, 2005.
- [10] J. Yang, L. A. Adamic, and M. S. Ackerman, "Crowdsourcing and knowledge sharing: Strategic user behavior on task CN," in *Proc. 9th ACM Conf. Electronic Commerce (EC)*, 2008.
- [11] D. DiPalantino and M. Vojnovic, "Crowdsourcing and all-pay auctions," in *Proc. 10th ACM Conf. Electronic Commerce (EC)*, 2009.
- [12] Y. R. Tausczik and J. W. Pennebaker, "Participation in an online mathematics community: Differentiating motivations to add," in *Proc. ACM 2012 Conf. Computer Supported Cooperative Work (CSCW)*, 2012.
- [13] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec, "Steering user behavior with badges," in *Proc. 22nd Int. Conf. World Wide Web (WWW)*, 2013.
- [14] B. Wang, Y. Wu, and K. J. R. Liu, "Game theory for cognitive radio networks: An overview," *Comput. Netw.*, vol. 54, no. 14, pp. 2537–2561, 2010.
- [15] C. Jiang, Y. Chen, and K. J. R. Liu, "Evolutionary dynamics of information diffusion over social networks," *IEEE Trans. Signal Processing*, vol. 62, no. 17, pp. 4573–4586, Sept. 2014.
- [16] C. Jiang, Y. Chen, and K. J. R. Liu, "Graphical evolutionary game for information diffusion over social networks," *IEEE J. Sel. Top. Signal Processing.*, vol. 8, no. 4, pp. 524–536, Aug. 2014.
- [17] C. Y. Wang, Y. Chen, and K. J. R. Liu, "Chinese restaurant game," *IEEE Signal Processing. Lett.*, vol. 19, no. 12, pp. 898–901, 2012.
- [18] C. Y. Wang, Y. Chen, and K. J. R. Liu, "Sequential Chinese restaurant game," *IEEE Trans. Signal Processing*, vol. 61, no. 3, pp. 571–584, 2013.
- [19] Y. Gao, Y. Chen, and K. J. R. Liu, "On cost-effective incentive mechanisms in microtask crowdsourcing," *IEEE Trans. Computat. Intell. AI Games*, vol. 7, no. 1, pp. 3–15, 2015.
- [20] J. Leskovec, L. Backstrom, and J. Kleinberg, "Meme-tracking and the dynamics of the news cycle," in *Proc. Int. Conf. Knowledge Discovery and Data Mining (KDD)*, 2009, pp. 497–506.
- [21] J. M. Smith, *Evolutionary and the Theory of Games*. Cambridge, U.K.: Cambridge Univ. Press, 1982.
- [22] M. A. Nowak and K. Sigmund, "Evolutionary dynamics of biological games," *Science*, vol. 303, no. 5659, pp. 793–799, 2004.
- [23] H. Ohtsuki, M. A. Nowak, and J. M. Pacheco, "Breaking the symmetry between interaction and replacement in evolutionary dynamics on graphs," *Phys. Rev. Lett.*, vol. 98, no. 10, p. 108106, 2007.

- [24] P. Shakarian, P. Roos, and A. Johnson, "A review of evolutionary graph theory with applications to game theory," *Biosystems*, vol. 107, no. 2, pp. 66–80, 2012.
- [25] H. Ohtsukia and M. A. Nowak, "The replicator equation on graphs," *J. Theor. Biol.*, vol. 243, no. 1, pp. 86–97, 7 Nov. 2006.
- [26] J. Leskovec, Stanford large network data set collection. [Online]. Available: <http://snap.stanford.edu/data>
- [27] V. Bala and S. Goyal, "Learning from neighbours," *Rev. Econ. Stud.*, vol. 65, no. 3, pp. 595–621, 1999.
- [28] B. Golub and M. O. Jackson, "Naive learning in social networks and the wisdom of crowds," *Amer. Econom. J.: Microecon.*, vol. 2, no. 1, pp. 112–149, 2010.
- [29] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *Rev. Econ. Stud.*, vol. 78, no. 4, pp. 1201–1236, 2011.
- [30] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, 2011.
- [31] J. Wit, "Social learning in a common interest voting game," *Games Econ. Behav.*, vol. 26, no. 1, pp. 131–156, 1999.
- [32] M. Battaglini, "Sequential voting with abstention," *Games Econ. Behav.*, vol. 51, no. 2, pp. 445–463, 2005.
- [33] A. Dasgupta, "Coordination and delay in global games," *J. Econ. Theory*, vol. 134, no. 1, pp. 195–225, 2007.
- [34] J. Pitman, *Combinatorial Stochastic Processes*. Lecture Notes for St. Flour Summer School. New Springer-Verlag, 2002.
- [35] Y. H. Yang, Y. Chen, C. X. Jiang, C. Y. Wang, and K. J. R. Liu, "Wireless access network selection game with negative network externality," *IEEE Trans. Wireless Commun.*, vol. 12, no. 10, pp. 5048–5060, 2013.
- [36] C. Jiang, Y. Chen, Y.-H. Yang, and K. J. R. Liu, "Dynamic chinese restaurant game: Theory and application to cognitive radio networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 4, pp. 1960–1973, Apr. 2014.
- [37] C. Y. Wang, Y. Chen, H. Y. Wei, and K. J. R. Liu, "Optimal pricing in stochastic scalable video coding multicasting system," in *Proc. IEEE INFOCOM Mini-Conf.*, Apr. 2013.
- [38] T. L. Griffiths and Z. Ghahramani, "The Indian buffet process: An introduction and review," *J. Mach. Learn. Res.*, vol. 12, pp. 1185–1224, Feb. 2011.
- [39] C. Jiang, Y. Chen, Y. Gao, and K. J. R. Liu, "Indian buffet game with negative network externality and non-Bayesian social learning," *IEEE Trans. Syst., Man Cyber.: Syst.*, vol. 45, no. 4, pp. 609–623, Apr. 2015.
- [40] P. Wais, S. Lingamnei, D. Cook, J. Fennell, B. Goldenberg, D. Lubarov, D. Marin, and H. Simons, "Towards building a high-quality workforce with mechanical turk," in *NIPS Workshop on Computational Social Science and the Wisdom of Crowds*, 2010.
- [41] P. G. Ipeirotis, F. Provost, and J. Wang, "Quality management on Amazon mechanical turk," in *Proc. ACM SIGKDD Workshop on Human Computation, HCOMP 10*, 2010, pp. 64–67.
- [42] J. Whitehill, P. Ruvolo, T. Wu, J. Bergsma, and J. Movellan, "Whose vote should count more: Optimal integration of labels from labelers of unknown expertise," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2009, pp. 2035–2043.
- [43] P. Welinder, S. Branson, S. Belongie, and P. Perona, "The multidimensional wisdom of crowds," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [44] V. C. Raykar, S. Yu, L. H. Zhao, G. H. Valadez, C. Florin, L. Bogoni, and L. Moy, "Learning from crowds," *J. Mach. Learn. Res.*, vol. 11, pp. 1297–1322, Apr. 2010.
- [45] V. C. Raykar and S. Yu, "Eliminating spammers and ranking annotators for crowdsourced labeling tasks," *J. Mach. Learn. Res.*, vol. 13, pp. 491–518, Feb. 2012.
- [46] N. Peled, Y. K. Gal, and S. Kraus, "An agent design for repeated negotiation and information revelation with people," in *Proc. 27th AAAI Conf. Artificial Intelligence*, 2013.
- [47] W. Mason, S. Suri, and D. J. Watts, "Long-run learning in games of cooperation," in *Proc. 15th ACM Conf. Economics and Computation (EC)*, June 2014, pp. 821–838.
- [48] D. Fudenberg and A. Peysakhovich, "Recency, records and recaps: Learning and non-equilibrium behavior in a simple decision problem," in *Proc. 15th ACM Conf. Economics and Computation (EC)*, June 2014, pp. 971–986.
- [49] S. Zhang and A. J. Yu, "Forgetful Bayes and myopic planning: Human learning and decision-making in a bandit setting," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2013, pp. 2607–2615.
- [50] S. Guo, M. Wang, and J. Leskovec, "The role of social networks in online shopping: Information passing, price of trust, and consumer choice," in *Proc. 12th ACM Conf. Electronic Commerce (EC)*, 2011, pp. 157–166.
- [51] Y. Gao, Y. Chen, and K. J. R. Liu, "Understanding sequential user behavior in social computing: To answer or to vote?," *IEEE Trans. Netw. Sci. Eng.*, vol. 2, no. 3, pp. 112–126, Aug. 2015.
- [52] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, "The complexity of computing a Nash equilibrium," *SIAM J. Comput.*, vol. 39, no. 1, pp. 195–259, 2009.
- [53] N. Burch, M. Lanctot, D. Szafron, and R. G. Gibson, "Efficient Monte Carlo counterfactual regret minimization in games with many player actions," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [54] W. Krichene, B. Drighes, and A. Bayen, "On the convergence of no-regret learning in selfish routing," in *Proc. 31st Int. Conf. Machine Learning (ICML)*, 2014.
- [55] M. L. Littman, "Markov games as a framework for multi-agent reinforcement learning," in *Proc. 11th Int. Conf. Machine Learning (ICML)*, 1994, pp. 157–163.
- [56] J. Hu and M. P. Wellman, "Nash Q-learning for general-sum stochastic games," *J. Mach. Learn. Res.*, vol. 4, pp. 1039–1069, Dec. 2003.
- [57] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia, "Strategic behavior when allocating indivisible goods sequentially," in *Proc. 27th AAAI Conf. Artificial Intelligence*, 2013.
- [58] A. Greenwald, J. Li, and E. Sodomka, "Approximating equilibria in sequential auctions with incomplete information and multi-unit demand," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [59] H. Luo and R. E. Schapire, "A drifting-games analysis for online learning and applications to boosting," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [60] C. Zhang and J. A. Shah, "Fairness in multi-agent sequential decision-making," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [61] B. Settles, "Active learning literature survey," Tech. Rep., Univ. Wisconsin-Madison, 2010.
- [62] D. D. Margineantu, "Active cost-sensitive learning," in *Proc. IJCAI*, 2007.
- [63] B. Settles, M. Craven, and L. Friedlan, "Active learning with real annotation costs," in *Proc. NIPS Workshop on Cost-Sensitive Learning*, 2008.
- [64] S. Sabato, A. D. Sarwate, and N. Srebro, "Auditing: Active learning with outcome-dependent query costs," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2013.
- [65] L. Yang and J. Carbonell, "Buy-in-bulk active learning," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2013.
- [66] A. Beygelzimer, S. Dasgupta, and J. Langford, "Importance weighted active learning," in *Proc. ICML*, 2009.
- [67] M. F. Balcan, C. Berline, A. Blum, E. Cohen, K. Patnaik, and L. Song, "Active learning and best-response dynamics," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [68] S. Chawla, J. D. Hartline, and B. Sivan, "Optimal crowdsourcing contests," in *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2012.
- [69] R. Cavallo and S. Jain, "Efficient crowdsourcing contests," in *Proc. Int. Conf. Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.
- [70] A. Shaw, J. Horton, and D. Chen, "Designing incentives for inexpert human raters," in *Proc. ACM Conf. Computer Supported Cooperative Work (CSCW)*, 2011.
- [71] Y. Zhang and M. van der Schaar, "Reputation-based incentive protocols in crowdsourcing applications," in *Proc. IEEE INFOCOM*, 2012.
- [72] Y. Singer and M. Mittal, "Pricing tasks in online labor markets," in *Proc. 3rd Workshop on Human Computation, (HCOMP) 11*, 2011.
- [73] Y. Singer and M. Mittal, "Pricing mechanisms for crowdsourcing markets," in *Proc. 22nd Int. Conf. World Wide Web (WWW)*, 2013.
- [74] A. Singla and A. Krause, "Truthful incentives in crowdsourcing tasks using regret minimization mechanisms," in *Proc. 22nd Int. Conf. World Wide Web (WWW)*, 2013.
- [75] D. Berry and B. Fristedt, *Bandit Problems: Sequential Allocation of Experiments*. London: Chapman and Hall, 1985.
- [76] K. Liu and Q. Zhao, "Distributed learning in multi-armed bandit with multiple players," *IEEE Trans. Signal Processing*, vol. 58, no. 11, pp. 5667–5681, 2010.
- [77] Y. Seldin, P. Bartlett, K. Crammer, and Y. Abbasi-Yadkori, "Prediction with limited advice and multiarmed bandits with paid observations," in *Proc. 31st Int. Conf. Machine Learning (ICML)*, 2014.
- [78] N. R. Devanur and S. M. Kakade, "The price of truthfulness for pay-per-click auctions," *Proc. 10th ACM Conf. Electronic Commerce (EC)*, 2009.
- [79] N. Gatti, A. Lazaric, and F. Tróvò, "A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities," in *Proc. 13th ACM Conf. Electronic Commerce (EC)*, 2012, pp. 605–622.
- [80] K. Amin, A. Rostamizadeh, and U. Syed, "Learning prices for repeated auctions with strategic buyers," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2013, pp. 1169–1177.