Learning Quadratic Games on Networks

Yan Leng [1]   **Xiaowen Dong** [2]   **Junfeng Wu** [3]   **Alex Pentland** [4]

Abstract

Individuals, or organizations, cooperate with or compete against one another in a wide range of practical situations. Such strategic interactions are often modeled as games played on networks, where an individual's payoff depends not only on her action but also on that of her neighbors. The current literature has largely focused on analyzing the characteristics of network games in the scenario where the structure of the network, which is represented by a graph, is known beforehand. It is often the case, however, that the actions of the players are readily observable while the underlying interaction network remains hidden. In this paper, we propose two novel frameworks for learning, from the observations on individual actions, network games with linear-quadratic payoffs, and in particular the structure of the interaction network. Our frameworks are based on the Nash equilibrium of such games and involve solving a joint optimization problem for the graph structure and the individual marginal benefits. Both synthetic and real-world experiments demonstrate the effectiveness of the proposed frameworks, which have theoretical as well as practical implications for understanding strategic interactions in a network environment.

1. Introduction

We live in an increasingly connected society. First studied by the American sociologist Stanley Milgram via his 1960s experiments and later popularized by John Guare's 1990 eponymous play, the theory of "six degrees of separation" has been recently re-analyzed on the social networking site Facebook, only to find out that any pair of Facebook users can actually be connected via approximately three and a half other ones (Backstrom et al., 2012). Individuals, unsurprisingly, are not merely connected; their decisions and actions often influence the ones around them. Indeed, Christakis and Fowler (Christakis & Fowler, 2009) have found in a series of studies that, one's emotion, health habit, and political opinion can affect individuals who are as far as three degrees of separation in her social circle. Furthermore, such influence on the decision-making process may take place via either explicit (Aral et al., 2009; Leng et al., 2018) or implicit interactions (Bandura & McClelland, 1977; Dong et al., 2018).

To study the decision-making of a group of interacting agents, recent literature in economics has increasingly focused on the modeling of such interactions as games played on networks (Jackson & Zenou, 2014; Bramoullé & Kranton, 2016). The underlying assumption in this setting is that, in a game played by a group of players who form a social network, the payoff of a player depends on her action, e.g., an effort made to achieve a specific task, as well as that of her neighbors in the network. Two types of actions have been studied in the literature, i.e., strategic complements and strategic substitutes. In the former case, one's action increases her neighbors' incentives for action, e.g., students putting an effort together into a joint assignment or firms working on a collaborative research project (Goyal & Moraga-Gonzaléz, 2001). In the latter case, however, the situation is the opposite, such as the scenarios of firms competing on market prices or individuals on local public goods (Bramoullé & Kranton, 2007).

In a network game, the underlying structure of the network carries critical information and dictates the behavior and actions of the players. Typically, graphs are used as mathematical tools to represent the structure of these networks, and the current literature in this area has predominately focused on studying the characteristics of games on known or predefined graphs (Ballester et al., 2006; Bramoullé et al., 2014; Galeotti et al., 2017). However, it is increasingly common that while ample observations on the actions of the agents are available, the underlying complex relationships among them, which may be captured by an interaction network, remain mostly hidden due to cost in observation or privacy concern. In this case, the network needs to be

[1]McCombs School of Business, The University of Texas at Austin, Austin, TX, USA [2]Department of Engineering Science, University of Oxford, Oxford, UK [3]College of Control Science and Engineering, Zhejiang University, Hangzhou, China [4]Media Lab, Massachusetts Institute of Technology, Cambridge, MA, USA. Correspondence to: Yan Leng <yan.leng@mccombs.utexas.edu>.

estimated to better understand the present and predict the future actions of these agents. The primary goal of this paper is, therefore, to study the problem of learning, given the observations on the actions of the agents, a graph structure that best explains the observed actions in the setting of a network game.

Such a problem, generally speaking, may be thought of as an instance of the ones of learning relationships, often in the form of graph structures, from observations made on a set of data entities. Classical approaches from the machine learning and signal processing communities tackle this problem by building statistical models (e.g., probabilistic graphical models (Koller & Friedman, 2009; Friedman et al., 2008)), physically-motivated models (e.g., diffusion processes on networks (Gomez-Rodriguez et al., 2010; Gomez-Rodriguez et al., 2011)), or more recently signal processing models (Dong et al., 2019; Mateos et al., 2019). These approaches, however, do not take into account the game-theoretic aspect of the decision-making of players in a network environment.

In the computer science literature, network games are known as graphical games (Kearns et al., 2001), and there have been a few studies recently on learning the games from observed action data. For example, the works in Irfan & Ortiz (2014); Honorio & Ortiz (2015); Ghoshal & Honorio (2016; 2017) have proposed to learn graphical games by observing actions from linear influence games with linear influence functions, where Ghoshal & Honorio (2018) has considered polymatrix games with pairwise matrix payoff functions. The work in Garg & Jaakkola (2016) has proposed to learn potential games on tree-structured networks of symmetric interactions. These conditions have been relaxed in Garg & Jaakkola (2017) where the authors have studied aggregative games where a player's payoff is convex and Lipschitz in an aggregate of their neighbors' actions defined via a local aggregator function. All these works, however, either consider a binary or a finite discrete action space, which may be restrictive in certain practical scenarios where actions take continuous values. Very recently, Barik & Honorio (2019) has considered learning continuous-action graphical games, which is similar in spirit to our study albeit under a slightly different action (which is budgeted) and payoff setting.

In this study, we focus on learning games with linear-quadratic payoffs (Ballester et al., 2006; Bramoullé et al., 2014; Acemoglu et al., 2015; Galeotti et al., 2017). We propose a learning framework where, given the Nash equilibrium action of the games, we jointly infer the graph structure that represents the interaction network as well as the individual marginal benefits. We further develop a second framework by considering the homophilous effect of individual marginal benefits in the interaction network. The first framework involves solving a convex optimization problem, while the second leads to a non-convex one for which we

develop an algorithm based on alternating minimization. We test the performance of the proposed algorithms in inferring graph structures for network games and show that it is superior to the baseline approaches of sample correlation and regularized graphical Lasso (Lake & Tenenbaum, 2010), albeit developed for slightly different learning settings.

The main contributions of the paper are as follows. First, the proposed learning frameworks, to the best of our knowledge, are the first to address the problem of learning the graph structure of the broad class of network games with linear-quadratic payoffs and continuous actions. Second, our framework also allows for the inference of marginal benefits of the players which permits a range of applications such as target interventions. Third, we analyze several factors in the quadratic games that affect the learning performance, such as the strength of strategic complements or substitutes, the topological characteristics of the networks, and the homophilous effect of individual marginal benefits. Overall, our paper constitutes a theoretical contribution to the studies of network games and may shed light on the understanding of strategic interactions in a wide range of practical scenarios, including business, education, governance, and technology adoption.

2. Network Games of Linear-Quadratic Payoffs

Consider a weighted network of N individuals represented by a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ and $\mathcal{E}$ denote the node and edge sets, respectively. For any pair of individuals i and j, $G_{ij} = G_{ji} > 0$ if $(i, j) \in \mathcal{E}$ and $G_{ij} = G_{ji} = 0$ otherwise, where G_{ij} is the ij-th entry of the adjacency matrix $\mathbf{G}$. In a network game of linear-quadratic payoffs, an individual i chooses her action a_i to maximize her payoff, u_i, which has the following form (Ballester et al., 2006; Bramoullé et al., 2014; Galeotti et al., 2017; Acemoglu et al., 2015):

$$u_i = b_i a_i - \frac{1}{2} a_i^2 + \beta a_i \sum_{j \in \mathcal{V}} G_{ij} a_j. \tag{1}$$

In Eq. (1), the first term is contributed by i's own action where the parameter b_i is called the marginal benefit, and the third term comes from the peer effect weighted by the actions of her neighbors. The parameter β captures the nature and the strength of such peer effect: if $\beta > 0$, actions are called strategic complements; and if $\beta < 0$, actions are called strategic substitutes.

The quadratic game with payoff function in Eq. (1) represents a broad class of games that have been extensively studied in the literature, and has a number of desirable properties. First, it naturally allows for continuous actions (i.e., a_i is considered to be continuous); second, it can be used for modelling games of both strategic complements and substitutes, i.e., positive and negative spillover effect; third, it may

also be used to approximate games with complex non-linear payoffs. For these reasons, games of linear-quadratic payoffs have been used to analyse crime activity, educational outcome, firm cooperation, and urban dynamics just to name a few (Jackson & Zenou, 2014).

One important advantage of the game in Eq. (1) is that it allows for an explicit solution for equilibrium behavior as a function of the network. To see this, let us define the vectorial forms $\mathbf{a} = [a_1, a_2, \cdots, a_N]^T$, $\mathbf{b} = [b_1, b_2, \cdots, b_N]^T$, and $\mathbf{u} = [u_1, u_2, \cdots, u_N]^T$, where we use the convention that the subscript i indicates the i-th entry of the vector. Taking the first-order derivative of the payoff u_i with respect to the action a_i in Eq. (1), we have:

$$\frac{\partial u_i}{\partial a_i} = b_i - a_i + \beta(\mathbf{Ga})_i. \tag{2}$$

Combining Eq. (2) for all i, it is clear that the following relationship holds, as pointed out in Ballester et al. (2006), for any (pure strategy) Nash equilibrium action $\mathbf{a}$:

$$(\mathbf{I} - \beta\mathbf{G})\,\mathbf{a} = \mathbf{b}, \tag{3}$$

hence

$$\mathbf{a} = (\mathbf{I} - \beta\mathbf{G})^{-1}\,\mathbf{b}, \tag{4}$$

where $\mathbf{I} \in \mathbb{R}^{N \times N}$ is the identity matrix. We adopt the critical assumption that the spectral radius of the matrix $\beta\mathbf{G}$, denoted by $\rho(\beta\mathbf{G})$, is less than 1, which guarantees the inversion of Eq. (4). Furthermore, as proved in Ballester et al. (2006), this assumption also ensures the uniqueness and stability of the Nash equilibrium action $\mathbf{a}$.

The equilibrium action $\mathbf{a}$ can be rewritten as $\mathbf{a} = \sum_{p=0}^{\infty} \beta^p \mathbf{G}^p \mathbf{b}$, and therefore has the following interpretations. If $\mathbf{b}$ is the all-one vector, then each entry of $\mathbf{a}$ is the Katz-Bonacich centrality (Katz, 1953; Bonacich, 1987) of the corresponding node, i.e., the number of walks of any length p originated from that node discounted exponentially by β. As pointed out in Jackson & Zenou (2014), interestingly, this means despite the local neighborhood relationship in Eq. (1) the payoff interdependency actually spreads indirectly throughout the network. On the other hand, the formulation of Eq. (4) can also be interpreted as computing steady-state opinions in studying opinion dynamics under a linear DeGroot model (DeGroot, 1974) and has been used in works on social network sensing (Wai et al., 2016).

From a different perspective, notice that $\mathbf{G}$ is a real and symmetric matrix hence has the following eigendecomposition: $\mathbf{G} = \boldsymbol{\chi}\boldsymbol{\Lambda}\boldsymbol{\chi}^T$. Plugging this into Eq. (4), the equilibrium action $\mathbf{a}$ can then be rewritten as $\mathbf{a} = \boldsymbol{\chi}(\mathbf{I} - \beta\boldsymbol{\Lambda})^{-1}\boldsymbol{\chi}^T\mathbf{b}$. Treating the marginal benefit $\mathbf{b}$ as a signal defined on the node set of the graph, the operation $\boldsymbol{\chi}^T\mathbf{b}$ can be interpreted as a Fourier-like transform for $\mathbf{b}$ according to the graph signal processing literature (Ortega et al., 2018). Given that the

eigenvector associated with the largest/smallest eigenvalue of $\mathbf{G}$ is the most smooth/non-smooth hence corresponds to low-/high-frequency signal on the graph, the action $\mathbf{a}$ can thus be interpreted as a low-pass filtered version of $\mathbf{b}$ for $\beta > 0$, and a high-pass filtered version of it for $\beta < 0$. This matches our intuition that equilibrium action tends to be smooth on the interaction network for the case of strategic complements, and non-smooth for strategic substitutes.

3. Learning Games with Independent Marginal Benefits

Given the graph with an adjacency matrix $\mathbf{G}$, the marginal benefits $\mathbf{b}$, and the parameter β, Eq. (4) provides a way of computing $\mathbf{a}$, the Nash equilibrium action of the players. The graph structure, in many cases, can be naturally chosen from the application domain, such as a social or business network. However, these natural choices of graphs may not necessarily describe well the strategic interactions between the players, and a natural graph might not be easy to define at all in some applications. Compared to the underlying relationships captured by $\mathbf{G}$, it is often easier to observe the individual actions $\mathbf{a}$, such as the amount of effort committed by students in a joint course project, or the strategic moves made by firms in an industrial setting. In these cases, given the actions and the dependencies described in Eq. (1), it is therefore of considerable interest to infer the structure of the graph on which the game is played, hence revealing the hidden relationships between the players.

3.1. Learning Framework

We consider N players, connected by a fixed interaction network $\mathbf{G}$, playing K different and independent games in each of which their payoffs depend not only on their own actions but also that of their neighbors[1]. Let us define the marginal benefits for these K games as $\mathbf{B} = [\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \cdots, \mathbf{b}^{(k)}] \in \mathbb{R}^{N \times K}$, where each column of $\mathbf{B}$ is the marginal benefit vector for one game, and the corresponding actions of the players as $\mathbf{A} = [\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \cdots, \mathbf{a}^{(K)}] \in \mathbb{R}^{N \times K}$. We first consider in this section the case where, for each game, the marginal benefits of individual players follow independent and identical Gaussian distributions, and then address in Section 4 the dependent case. The parameter that captures the strength of the network effect, β, can be either positive or negative, corresponding to strategic complements and strategic substitutes, respectively. Notice that we assume a fixed beta for all games in the current setting, which implies that: 1) the relationship between a player and her neighbors remains the same for all players (either complements or sub-

---

[1]The setting of K games of scalar actions could also be interpreted as one game with a K-dimensional action, in which case the assumption is that each dimension of the player's action satisfies the equilibrium condition.

stitutes); 2) the contribution to the individual payoff from the neighbors is scaled by a beta for all players. A more flexible setting for the β parameter could be an interesting future direction.

Given the observed actions $\mathbf{A}$ and the parameter β, the goal is to infer a graph structure $\mathbf{G}$ as well as the marginal benefits $\mathbf{B}$, which best explain $\mathbf{A}$ in terms of the relationship in Eq. (3). To this end, we propose the following joint optimization problem of $\mathbf{G}$ and $\mathbf{B}$:

$$\min_{\mathbf{G},\mathbf{B}} \ f(\mathbf{G},\mathbf{B})$$
$$= ||(\mathbf{I} - \beta\mathbf{G})\mathbf{A} - \mathbf{B}||_F^2 + \theta_1||\mathbf{G}||_F^2 + \theta_2||\mathbf{B}||_F^2, \quad (5)$$
$$\text{s.t. } G_{ij} = G_{ji}, \ G_{ij} \geq 0, \ G_{ii} = 0 \text{ for } \forall i,j \in \mathcal{V},$$
$$||\mathbf{G}||_1 = N,$$

where $||\cdot||_F$ and $||\cdot||_1$ denote the Frobenius norm and element-wise L^1-norm of a matrix, respectively, and θ_1 and θ_2 are two non-negative regularization parameters. The first line of constraints ensures that $\mathbf{G}$ is a valid adjacency matrix, and the second constraint (the constraint on the L^1-norm) fixes the volume of the graph and permits to avoid trivial solutions. Without loss of generality the volume is chosen to be N. It is clear that, in the problem of Eq. (5), we aim at a joint inference of the graph structure $\mathbf{G}$ and the marginal benefits $\mathbf{B}$, such that the observed actions $\mathbf{A}$ are close to the Nash Equilibria of the K games played on the graph. The Frobenius norm on $\mathbf{G}$ is added as a penalty term to control the distribution of the edge weights of the learned graph (the off-diagonal entries of $\mathbf{G}$)[2], which, together with the L^1-norm constraint, bears similarity to the linear combination of L^1 and L^2 penalties in an elastic net regularization (Zou & Hastie, 2005).

The effectiveness of the formulation in Eq. (5) depends on $\rho(\beta G)$. To see this, notice that under the assumption that $\mathbf{b}$ is IID Gaussian, the equilibrium action $\mathbf{a}$ follows a Gaussian distribution with covariance $(\mathbf{I} - \beta\mathbf{G})^{-2}$. If $\rho(\beta G)$ is close to zero, then $\mathbf{a}$ is almost independent from G, and it would be difficult to infer $\mathbf{G}$ from $\mathbf{a}$ in this scenario. On the other hand, if $\rho(\beta G)$ is close to one, the covariance is dominated by the eigenvector associated with the largest (when $\beta > 0$) or smallest (when $\beta < 0$) eigenvalue of $\mathbf{G}$. In this case, the action $\mathbf{a}$ clearly contains information about $\mathbf{G}$ which facilitates learning.

It can be seen from Eq. (5) that, comparing to the work in Barik & Honorio (2019) which learns the neighbors for each player separately via a regression framework, our approach learns the network at once by solving a single optimization problem. This difference is analogous to that between the neighborhood selection of Meinshausen et al. (2006) and the graphical Lasso of Friedman et al. (2008) in the context

---

[2]Similar constraints have been adopted in Hu et al. (2015); Dong et al. (2016) for graph inference.

of covariance estimation.

3.2. Learning Algorithm

Given the non-negativity of G_{ij}, we can re-write the constraint: $||\mathbf{G}||_1 = \mathbf{1}^T \mathbf{G} \mathbf{1} = N$, where $\mathbf{1} \in \mathbb{R}^N$ is the all-one vector. The constraints in Eq. (5) therefore form a convex set. The problem of Eq. (5) is thus a quadratic program jointly convex in $\mathbf{B}$ and $\mathbf{G}$, and can be solved efficiently via the interior point methods (Boyd & Vandenberghe, 2004). In our experiments, we solve the problem of Eq. (5) using the Python software package CVXOPT (Andersen et al., Version 1.2.0. Available at cvxopt.org, 2018). In case of graphs of very large number of vertices, we can instead use operator splitting methods, e.g., alternating direction method of multipliers (ADMM) (Boyd et al., 2011), to find a solution. The algorithm is summarized in Algorithm 1.

Algorithm 1 Learning Games with Independent Marginal Benefits

Input: Actions $\mathbf{A} \in \mathbb{R}^{N \times K}$ for K games, β, θ_1, θ_2
Output: Network $\mathbf{G} \in \mathbb{R}^{N \times N}$, marginal benefits $\mathbf{B} \in \mathbb{R}^{N \times K}$ for K games
Solve for $\mathbf{G}$ and $\mathbf{B}$ in Eq. (5)
return: $\mathbf{G}$, $\mathbf{B}$

---

4. Learning Games with Homophilous Marginal Benefits

A large number of studies in the literature of social sciences and economics have analyzed the phenomenon of homophily in social networks, which describes that individuals tend to associate and form ties with those that are similar to themselves (McPherson et al., 2001; Jackson, 2010). Since the marginal benefit vector $\mathbf{b}$ in each game can be thought of as the individual preferences toward a particular action, they may contribute, in the presence of the homophily effect, to the formation of the interaction network on which the game is played. The second formulation in our paper is, therefore, to address the problem of learning games with such homophilous marginal benefits.

4.1. Learning Framework

The homophily effect that is present in the marginal benefit vector $\mathbf{b}$ implies that $\mathbf{b}$ as a signal defined on the graph is relatively smooth, in the sense that nodes that are connected share similar marginal benefits. This may be quantified by the so-called Laplacian quadratic form on the graph:

$$\mathbf{b}^T \mathbf{L} \mathbf{b} = \frac{1}{2} \sum_{i,j \in \mathcal{V}} G_{ij} (b_i - b_j)^2, \quad (6)$$

where $\mathbf{L} = \text{diag}(\sum_{j \in \mathcal{V}} G_{ij}) - \mathbf{G}$ is the unnormalized (combinatorial) graph Laplacian matrix (Chung, 1997). We there-

fore propose to replace the norm on $\mathbf{B}$ with this measure in the objective function of Eq. (5) to promote homophilous marginal benefits. This essentially assumes that the marginal benefits follow a multivariate Gaussian distribution with the precision matrix being the graph Laplacian. This leads to the following optimization problem:

$$
\begin{aligned}
\min_{\mathbf{G},\mathbf{B}} \ & h(\mathbf{G},\mathbf{B}) \\
= & ||(\mathbf{I}-\beta\mathbf{G})\mathbf{A}-\mathbf{B}||_F^2 + \theta_1||\mathbf{G}||_F^2 + \theta_2\,\mathrm{tr}(\mathbf{B}^T\mathbf{L}\mathbf{B}), \\
\text{s.t. } & G_{ij}=G_{ji},\ G_{ij}\geq 0,\ G_{ii}=0 \ \text{ for } \forall i,j\in\mathcal{V}, \\
& ||\mathbf{G}||_1 = N, \\
& \mathbf{L}=\mathrm{diag}(\sum_{j\in\mathcal{V}}G_{ij})-\mathbf{G},
\end{aligned}
\tag{7}
$$

where $\mathrm{tr}(\cdot)$ denotes the trace operator. The third term in the objective is the sum of the Laplacian quadratic form for all the columns in $\mathbf{B}$, and the third constraint comes from the definition of the graph Laplacian $\mathbf{L}$. Like in Eq. (5), θ_1 and θ_2 are two non-negative regularization parameters. The problem of Eq. (7) is similar to that of Eq. (5), but with a different assumption that there exists the effect of homophily in the marginal benefits $\mathbf{b}$, whose strength is controlled by the regularization parameter θ_2, i.e., a larger θ_2 favors a stronger homophily effect, and vice versa.

4.2. Learning Algorithm

Unlike the problem of Eq. (5), the problem of Eq. (7) is not jointly convex in $\mathbf{G}$ and $\mathbf{B}$ due to the third term in the objective function. We, therefore, adopt an alternating minimization scheme to optimize for the graph structure $\mathbf{G}$ and the marginal benefits $\mathbf{B}$ where, at each step, we solve for one variable while fixing the other.

Given $\mathbf{B}$, we first solve for $\mathbf{G}$ in Eq. (7). The constraints on $\mathbf{G}$ in Eq. (7) are the same as that in Eq. (5) and thus convex. Since $\theta_1\geq 0$ and $\theta_2\geq 0$, fixing $\mathbf{B}$ and solving for $\mathbf{G}$ results in a strongly convex objective, and consequently the problem admits a unique solution. We again solve this convex quadratic program using the package CVXOPT. Next, we fix $\mathbf{G}$ and solve for $\mathbf{B}$ in Eq. (7). By fixing $\mathbf{G}$, Eq. (7) becomes an unconstrained convex quadratic program, and thus admits a closed-form solution which can be obtained by setting the derivative to zero:

$$
\frac{\partial h(\mathbf{G},\mathbf{B})}{\partial \mathbf{B}} = -2\big((\mathbf{I}-\beta\mathbf{G})\mathbf{A}-\mathbf{B}\big)+2\theta_2\mathbf{L}\mathbf{B}=\mathbf{0}, \tag{8}
$$

hence

$$
\mathbf{B}=(\mathbf{I}+\theta_2\mathbf{L})^{-1}(\mathbf{I}-\beta\mathbf{G})\mathbf{A}. \tag{9}
$$

We iterative between the two steps until either the change in the objective $h(\mathbf{G},\mathbf{B})$ is smaller than 10^{-4}, or a maximum number of iterations has been reached. This strategy is called block coordinate descent (BCD) and, since both subproblems are strongly convex, is guaranteed to converge to a local minimum (see Proposition 2.7.1 in Bertsekas (1995)). The complete algorithm is summarized in Algorithm 2.

Algorithm 2 Learning Games with Homophilous Marginal Benefits

---

   Input: Actions $\mathbf{A}\in\mathbb{R}^{N\times K}$ for K games, β,θ_1,θ_2
   Output: Network $\mathbf{G}\in\mathbb{R}^{N\times N}$, marginal benefits $\mathbf{B}\in\mathbb{R}^{N\times K}$ for K games
   Initialize: $\mathbf{B}_0(:,k)\sim\mathcal{N}(\mathbf{0},\mathbf{I})$ for $k=1,\cdots,K$, $t=1$, $\Delta=1$
   if $\Delta\geq 10^{-4}$ and $t\leq$ # iterations **then**
      Solve for $\mathbf{G}_t$ in Eq. (7) given $\mathbf{B}_{t-1}$
      Compute $\mathbf{L}_t$ using $\mathbf{G}_t$
      $\mathbf{B}_t=(\mathbf{I}+\theta_2\mathbf{L}_t)^{-1}(\mathbf{I}-\beta\mathbf{G}_t)\mathbf{A}$
      $\Delta=|h(\mathbf{G}_t,\mathbf{B}_t)-h(\mathbf{G}_{t-1},\mathbf{B}_{t-1})|$ (for $t>1$)
      $t=t+1$
   end if
   return: $\mathbf{G}=\mathbf{G}_t, \mathbf{B}=\mathbf{B}_t$.

---

5. Experiments on Synthetic Data

In this section, we evaluate the performance of the proposed learning frameworks on synthetic networks that follow three types of random graph models, i.e., the Erdős–Rényi (ER), the Watts-Strogatz (WS), and the Barabási-Albert (BA) models. In the ER graph, an edge is created with a probability of $p=0.2$ independently from all other possible edges. In the WS graph, we set the average degree of the vertices to be $k=\lfloor\log_2(N)\rfloor$, with a probability of $p=0.2$ for the random rewiring process. Finally, in the BA graph, we add $m=1$ new node at each time by connecting it to an existing node in the graph via preferential attachment. All the graphs have $N=20$ vertices in our experiments. Once the graphs are constructed, we compute $\beta>0$ such that the spectral radius, $\rho(\beta\mathbf{G})$, varies between 0 and 1 hence satisfying the assumption in Section 2.

We adopt two different settings, one for generating the independent marginal benefits $\mathbf{b}$ and the other for the homophilous $\mathbf{b}$. In the independent case of Section 3, for each game, we generate realizations by considering $\mathbf{b}\sim\mathcal{N}(\mathbf{0},\mathbf{I})$. In the homophilous setting of Section 4, we generate realizations by considering $\mathbf{b}\sim\mathcal{N}(\mathbf{0},\mathbf{L}^\dagger)$, where $\mathbf{L}^\dagger$ is the Moore-Penrose pseudoinverse of the groundtruth graph Laplacian $\mathbf{L}$. In both cases we further add Gaussian noise $\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\frac{1}{10}\mathbf{I})$ to the simulated marginal benefits. Now, given $\mathbf{b}$ and β, we compute the players' Nash equilibrium action $\mathbf{a}$ according to Eq. (4). We consider $K=50$ games for each of which we generate the action $\mathbf{a}$.

We apply Algorithm 1 and Algorithm 2 to the respective settings to infer graph structures and compare against the

groundtruth ones in a scenario of binary classification, i.e., either there exists an edge between i and j (positive case), or not (negative case). Since the ratio of positive cases is small for all the three types of graphs, we use the area under the curve (AUC) for the evaluation of the learning performance. We compare our algorithms with two baseline methods for inferring graph structures given data observations: the sample correlation and the regularized graphical Lasso in Lake & Tenenbaum (2010). In the former case we consider the correlations between each pair of variables as "edge weights" in a learned graph, while in the latter case a graph adjacency matrix is computed as in our algorithms.

Notice that in the synthetic experiments, we focus on the case of strategic complements, i.e., $\beta > 0$, to facilitate a fair comparison with the two baselines that only apply to this case. Our methods therefore also have the unique advantage of dealing with the case of strategic substitutes, i.e., $\beta < 0$.

5.1. Comparison of Learning Performance

The performance of the three methods in comparison is shown in Fig. 1 (top) for the case of independent marginal benefits. For Algorithm 1 and regularized graphical Lasso, we report the results using the parameter values that give the best average performance over 20 randomly generated graph instances[3]. First, we see that the performance of all the three methods increases with the spectral radius $\rho(\beta\mathbf{G})$ for the majority of the cases. This pattern indicates that stronger strategic dependencies between actions of potential neighbors reveal more information about the existence of the corresponding links. Indeed, as $\rho(\beta\mathbf{G})$ increases, the action matrix $\mathbf{A}$ contains more information about the graph structure as explained in Section 3.1. Second, the performance of the proposed Algorithm 1 generally outperforms the two baselines in terms of recovering the locations of the edges of the groundtruth. Notice that for regularized graphical Lasso, the performance drops with a larger value for $\rho(\beta\mathbf{G})$. One possible explanation is that, as $\rho(\beta\mathbf{G})$ becomes close to 1, the smallest eigenvalue of $\mathbf{I} - \beta\mathbf{G}$ approaches 0 resulting in a large ratio between the smallest and the largest eigenvalues of the empirical covariance of $\mathbf{a}$, which may lead to inaccurate estimation of the precision matrix in the graphical Lasso. In comparison, our method does not seem to be affected by such phenomenon. Finally, the performance of all the methods for the WS and BA graphs is generally better than that of the ER graphs, possibly because there exists more structural information in the former models than the latter.

The same results for the case of homophilous marginal benefits are shown in Fig. 1 (bottom). We observe the same increase in performance as $\rho(\beta\mathbf{G})$ increases for all the

---

[3]Analysis of robustness of performance against regularization parameters is presented in Supplementary Material.

three methods, as well as the drop in performance towards large $\rho(\beta\mathbf{G})$ for regularized graphical Lasso. The proposed Algorithm 2 generally achieves superior performance in this scenario, which is expected due to the way the observations $\mathbf{A}$ are generated taking into account the regularization term in the objective in Eq. (7) that enforces homophily.

5.2. Learning Performance with Respect to Different Factors in Network Games

We now examine the performance of Algorithm 2 with respect to a number of factors, including the number of games, the noise intensity, the structure of the groundtruth network, and the strength of the homophily effect in marginal benefits (in Supplementary Material). The same results for Algorithm 1 are presented in Supplementary Material.

Number of games. We are first interested in understanding the influence of the number of games K on the learning performance. In the following and all subsequent analyses, we choose $\rho(\beta\mathbf{G}) = 0.6$, and fix the parameters in Algorithm 2 to be the ones that lead to the best learning performance. In Fig. S3 (top), we vary the number of games and evaluate its effect on the performance. We see that in general, the performance of the algorithm increases, as more observations become available. The benefit is least obvious for the ER graph, suggesting that adding more observations does not help as much in improving the performance in this case when the edges in the graph appear more randomly.

Noise intensity in the marginal benefits. We now analyze the robustness of the result against noise intensity in the marginal benefits. With more noise in the marginal benefits, the observed actions $\mathbf{A}$ becomes noisier as well, hence possibly affecting the learning performance. As shown in Fig. S3 (bottom), the learning performance generally decays as the intensity of noise increases, which is expected. The performance of the model is relatively stable until the standard deviation of the noise becomes larger than 1.

Network structure. The random graphs used in our experiments have parameters that may affect the performance of the proposed algorithms. We, therefore, analyze the effect of p in the ER, k in the WS, and m in the BA graphs on the learning performance of the proposed algorithm. The larger these parameters, the higher the edge density in these random graph models. As shown in Fig. 2, the density of edges has a substantial effect on the learning performance for all the networks, i.e., the denser the edges, the worse the performance. One possible explanation is that, in a sparse network the correlations between individuals' actions might contain more accurate information about the existence of dependencies hence edges between them, while in a dense network the influence from one neighbor is often mingled with that from another, which makes it more challenging to uncover pairwise dependencies.
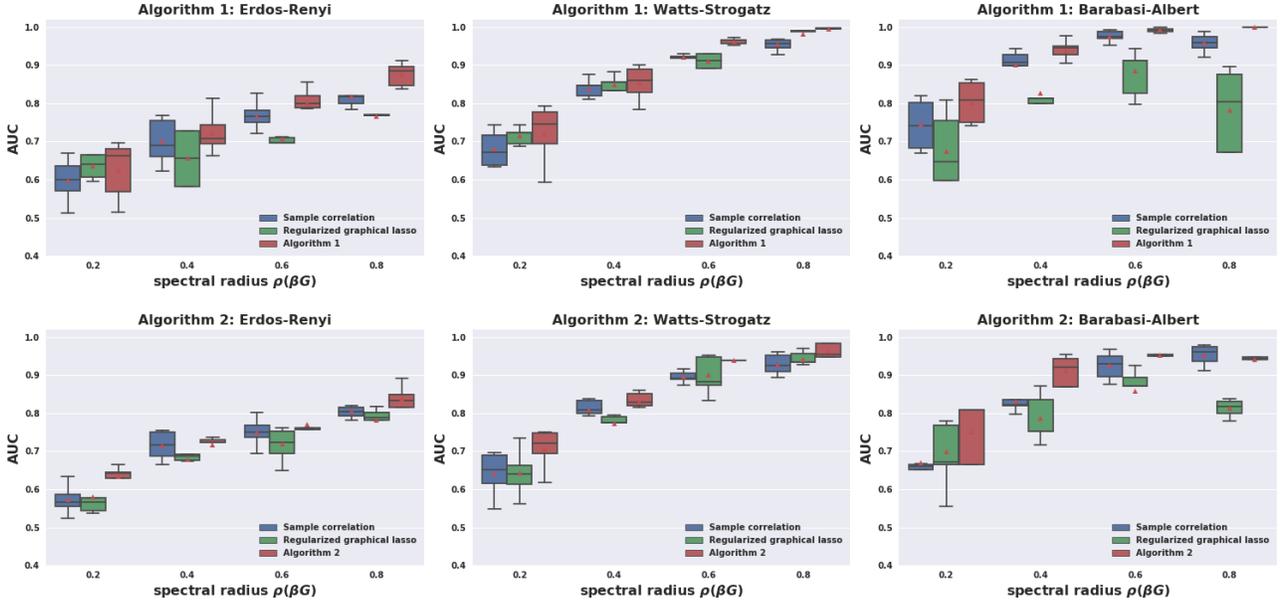
Figure 1: Performance of the proposed algorithm and baselines in the setting of independent (top) and homophilous (bottom) marginal benefits. The red triangle, the middle line, lower and upper boundaries of the box (interquartile range or IQR) correspond to mean, median, and 25/75 percentile of the data, respectively. The lower and upper whiskers extend maximally 1.5 times of IQR from 25 percentile downwards and 75 percentile upwards, respectively.
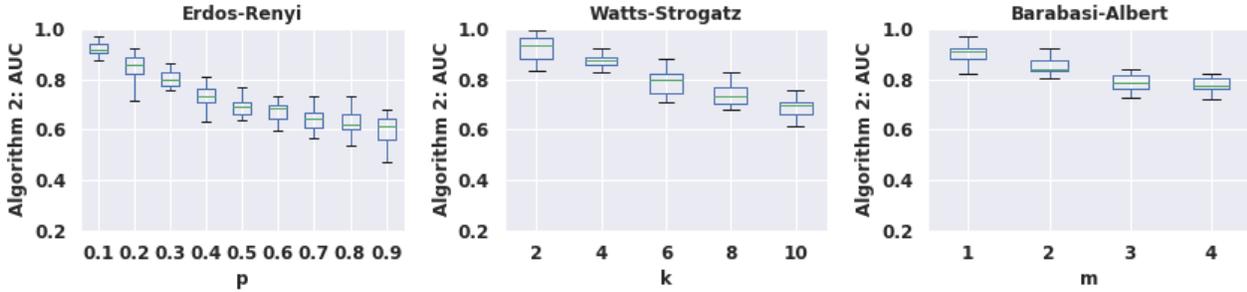
Figure 2: Performance of Algorithm 2 versus structural properties of the network.

5.3. Learning the Marginal Benefits

In our framework, we jointly infer the graph structure and the marginal benefits of the players. This is one of the main advantages of our algorithms, since the inference of marginal benefits can be critical for targeting strategies and interventions (Galeotti et al., 2017). To test the performance of learning marginal benefits, for each random graph model, we generate a network with 20 nodes and simulate 50 games with $\rho(\beta\mathbf{G}) = 0.6$, for both independent and homophilous marginal benefits. We repeat this process for 30 times, and report the average performance of learning the marginal benefits in Table 1. The performance is measured in terms of the coefficients of determination (R^2), by treating the groundtruth and learned marginal benefits (both in vectorized form) as dependent and independent variables, respec-

tively. As we can see, in both cases the R^2 values are above 0.9, which indicates that the learned marginal benefits are very similar to the groundtrith ones.

Table 1: Performance (R^2) of learning marginal benefits.

	Algorithm 1		Algorithm 2	
	mean	std	mean	std
ER graph	0.959	0.005	0.982	0.002
WS graph	0.955	0.007	0.921	0.010
BA graph	0.937	0.008	0.909	0.010

6. Experiments on Real-World Data

The strategic interactions between players in real-world situations may follow the formulation of network games. Given this broad assumption, we present three examples

of inferring the network structure in quadratic games in practical scenarios, e.g., the inference of social, trade, and political networks.

6.1. Social Network

We first consider inferring a social network between households in a village in rural India (Banerjee et al., 2013). In particular, following the setting in Banerjee et al. (2013), we consider the actions of each household as choosing the number of rooms, beds, and other facilities in their houses. The assumption is that there may exist strategic interactions between these households regarding constructing such facilities. In particular, when deciding to adopt new technologies or innovations, people have an incentive to conform to the social norms they perceive (Young, 2009; Montanari & Saberi, 2010), which are formed by the decisions made by their neighbors. For example, if neighbors adopt a specific facility, villagers tend to gain higher payoff after adopting the same facility by complying with social norms (i.e., strategic complements).

We consider each action as a strategy in a quadratic game, and we have 31 games with discrete actions made by 182 households. We then apply the proposed algorithms to infer the relationships between these households, and compare against a groundtruth network of self-reported friendship. Since we do not observe β, we treat it as a hyperparameter, and tune it within the range of $\beta \in [-3, 3]$. It can be seen from Table 2 that both of the proposed methods outperform regularized graphical Lasso by about 2.5% and sample correlation by about 10.7%[4], indicating that they can recover a social network structure closer to the groundtruth.

6.2. Trade Network

We next consider inferring a global trade network. Specifically, we consider the overall trading activities of 235 countries on 96 export products and 96 import products in the year 2008 as our observed actions[5]. This leads to 192 games (for both import and export actions) played by 235 agents (countries). By applying the proposed algorithms, we infer the relationships among nations regarding their strategic trading decisions and compare against a groundtruth, which

is the trading network in year 2002[6]. In constructing the groundtruth, we consider the edge weight between each pair of nations as the logarithmic of the total amount of trades (import plus export) between the two nations.

In the groundtruth trade network, each nation is connected with the ones with which it traded in 2002. This implies that the nation has different demand and supply compared to its neighbors, and their import and export actions tend to be different in the near future. Therefore, we expect a strategic substitute relationship between the nations when looking at their import and export activities in 2008.

We tune β within the range of $\beta \in [-1, 1]$. Table 2 shows that Algorithm 1 and Algorithms 2 outperform regularized graphical Lasso by 12.09% and 24.85%, respectively[7]. The larger performance gain in this case is due to the fact that both sample correlation and regularized graphical Lasso are suitable only for strategic complement and not strategic substitute relationships. Furthermore, Algorithm 2 performs better than Algorithm 1 in this example, which implies a homophilous distribution of marginal benefits across neighboring nations.

Table 2: Performance (AUC) of learning the structure of the social network and the trade network.

	Social network	Trade network
Sample correlation	0.525	0.523
Regularized graphical Lasso	0.564	0.570
Algorithm 1	0.575	0.622
Algorithm 2	0.576	0.677

6.3. Political Network

The third real-world example we considered is the inference of the relationship between the cantons in Switzerland in terms of their political preference. To this end, we consider voting statistics from the national referendums for 37 federal initiatives in Switzerland between 2008 and 2012[8]. Specifically, we consider the percentage of voters supporting each initiative in the 26 Swiss cantons as the observed actions. This leads to 37 games (initiatives) played by 26 agents (cantons). By applying the proposed algorithms, we infer a network that captures the strategic political relationship

---

[4]The improvement is calculated by the absolute improvement in AUC normalized by the room for improvement. The best performance of Algorithm 1 is obtained with $\beta = 0.1$, $\theta_1 = 2^{-8.5}$, and $\theta_2 = 2^1$, while that of Algorithm 2 is obtained with $\beta = 2.6$, $\theta_1 = 2^7$, and $\theta_2 = 2^{-5.5}$. The positive sign of β in both cases indicates a strategic complement relationship between the households, which is consistent with our hypothesis.

[5]Data can be accessed via https://atlas.media.mit.edu/en/resources/data/. The trading activities are classified by the 2002 edition of the HS (Harmonized System).

[6]The trading network from previous years provides a foundation for nations to make decisions and thus can be thought of as a groundtruth. The year 2002 is the latest year before 2008 for which trading data are available.

[7]The best performance of Algorithm 1 is obtained with $\beta = -0.6$, $\theta_1 = 2^1$, and $\theta_2 = 2^{-10}$, and that of Algorithm 2 is obtained with $\beta = -0.7$, $\theta_1 = 2^{11.5}$, and $\theta_2 = 2^{-15.5}$. The negative sign of β in both cases indicates a strategic substitute relationship between the nations, which is consistent with our hypothesis.

[8]The voting statistics were obtained via http://www.swissvotes.ch.
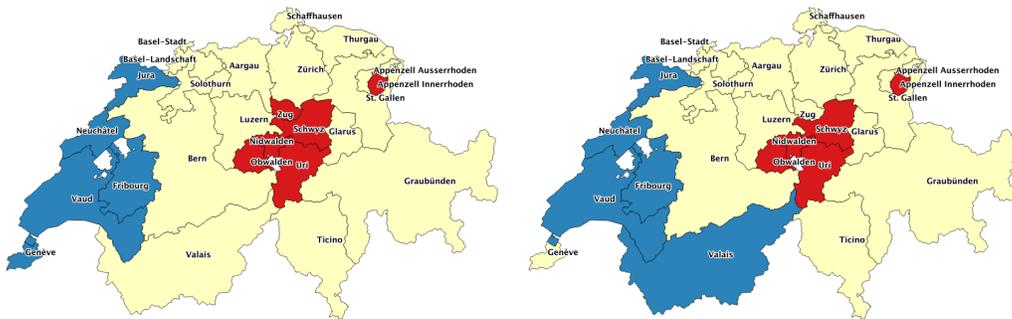
Figure 3: Clustering of Swiss cantons based on the political network learned by Algorithm 1 (left) and Algorithm 2 (right).

between these cantons reflected by their votes in the national referendums[9].

Unlike the previous examples, it is more difficult to define a groundtruth network in this case. Instead, we apply spectral clustering (von Luxburg, 2007) to the learned network and interpret the obtained clusters of cantons. The three-cluster partition of the networks learned by Algorithm 1 and Algorithm 2 are presented in Fig. 3. As we can see, the clusters obtained in the two cases are largely consistent, with the blue and yellow clusters generally corresponding to the French-speaking and German-speaking cantons, respectively. The red cluster, in both cases, contains the five cantons of Uri, Schwyz, Nidwalden, Obwalden and Appenzell Innerrhoden, which are all considered among the most conservative ones in Switzerland. This demonstrates that the learned networks are able to capture the strategic dependence between cantons within the same cluster, which tend to vote similarly in national referendums.

7. Discussion

In this paper, we have proposed two novel learning frameworks for joint inference of graph structure and individual marginal benefits for a broad class of network games, i.e., games with linear-quadratic payoffs. We believe that the present paper may shed light on the understanding of network games (in particular those with linear-quadratic payoffs), and contribute to the vibrant literature of learning hidden relationships from data observations.

The proposed approaches can benefit a wide range of prac-

tical scenarios. For instance, the learned graph, which captures the strategic interactions between the players, may be used for detecting communities formed by the players (Fortunato, 2010). This can, in turn, be used for purposes such as stratification. Another use case is to compute centrality measures of the nodes in the network, which may help in designing efficient targeting strategies in marketing scenarios (Leng et al., 2020). Finally, the joint inference of the graph and the marginal benefits can help a central planner who wishes to design intervention mechanisms achieve specific planning objectives. One such objective could be the maximization of the total payoffs of all players, which can be done by adjusting, according to the network topology, the marginal benefits via incentivization (Galeotti et al., 2017). Another objective could be the reduction of inequality between the players in terms of their payoffs, which can be done by adjusting network topology via encouraging the formation of certain new relationships.

There remain many interesting directions to explore. For example, building upon the promising empirical results presented in this paper, it would be important to study the theoretical guarantees of the proposed algorithms in recovering the graph structure. It would also be interesting to consider graph inference given partial or incomplete observations of the actions, especially in the case where it is costly to observe the actions of all the network players, or consider a setting where the underlying relationships between the players may evolve over time, which can be modeled by dynamic graph topologies. Finally, the inference framework might need to be adapted accordingly for network games of different payoff functions. In this sense, it would be very interesting to investigate the possibility of inferring the graph structure in a purely date-driven fashion via a neural network, without the explicit knowledge of the payoff function. We leave these studies as future work.

---

[9]We tune β within the range of $[-1, 1]$. For Algorithm 1 we report results with $\beta = 0.6$, $\theta_1 = 2^{-6.2}$, and $\theta_3 = 2^{-1.65}$. For Algorithm 2 we report results with $\beta = 0.67$, $\theta_1 = 2^2$, and $\theta_2 = 2^3$. The positive sign of β in both cases indicates a strategic complement relationship between the cantons.

Acknowledgement

References

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. Networks, shocks, and systemic risk. Working Paper 20931, National Bureau of Economic Research, 2015. URL http://www.nber.org/papers/w20931.

Andersen, M. S., Dahl, J., and Vandenberghe, L. CVXOPT: A Python package for convex optimization. Version 1.2.0. Available at cvxopt.org, 2018.

Aral, S., Muchnik, L., and Sundararajan, A. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51):21544–21549, 2009.

Backstrom, L., Boldi, P., Rosa, M., Ugander, J., and Vigna, S. Four degrees of separation. In *Proceedings of the 4th Annual ACM Web Science Conference*, pp. 33–42, 2012.

Ballester, C., Calvó-Armengol, A., and Zenou, Y. Who's who in networks. Wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.

Bandura, A. and McClelland, D. C. *Social learning theory*. Prentice-Hall, 1977.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. The diffusion of microfinance. *Science*, 341(6144): 1236498, 2013.

Barik, A. and Honorio, J. Provable computational and statistical guarantees for efficient learning of continuous-action graphical games. *arXiv:1911.04225*, 2019.

Bertsekas, D. P. *Nonlinear programming*. Athena Scientific, 1995.

Bonacich, P. Power and centrality: A family of measures. *American journal of sociology*, 92(5):1170–1182, 1987.

Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge University Press, 2004.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

Bramoullé, Y. and Kranton, R. Public goods in networks. *Journal of Economic Theory*, 135(1):478–494, 2007.

Bramoullé, Y. and Kranton, R. Games played on networks. In *The Oxford Handbook of the Economics of Networks, Yann Bramoullé, Andrea Galeotti, and Brian Rogers (Eds.)*, 2016.

Bramoullé, Y., Kranton, R., and D'amours, M. Strategic interaction and networks. *American Economic Review*, 104(3):898–930, 2014.

Christakis, N. A. and Fowler, J. H. *Connected: The surprising power of our social networks and how they shape our lives*. Little, Brown, 2009.

Chung, F. R. K. *Spectral graph theory*. American Mathematical Society, 1997.

DeGroot, M. Reaching a consensus. *Journal of the American Statistical Association*, 69:118–121, 1974.

Dong, X., Thanou, D., Frossard, P., and Vandergheynst, P. Learning laplacian matrix in smooth graph signal representations. *IEEE Transactions on Signal Processing*, 64(23):6160–6173, 2016.

Dong, X., Suhara, Y., Bozkaya, B., Singh, V. K., Lepri, B., and Pentland, A. Social bridges in urban purchase behavior. *ACM Transactions on Intelligent Systems and Technology*, 9(3):33:1–33:29, 2018.

Dong, X., Thanou, D., Rabbat, M., and Frossard, P. Learning graphs from data: A signal representation perspective. *IEEE Signal Processing Magazine*, 36(3):44–63, 2019.

Fortunato, S. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.

Friedman, J., Hastie, T., and Tibshirani, R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.

Galeotti, A., Golub, B., and Goyal, S. Targeting interventions in networks. *arXiv:1710.06026*, 2017.

Garg, V. and Jaakkola, T. Learning tree structured potential games. In *Advances in Neural Information Processing Systems*, pp. 1552–1560, 2016.

Garg, V. and Jaakkola, T. Local aggregative games. In *Advances in Neural Information Processing Systems*, pp. 5341–5351, 2017.

Ghoshal, A. and Honorio, J. From behavior to sparse graphical games: Efficient recovery of equilibria. In *Proceedings of the IEEE Allerton Conference on Communication, Control, and Computing*, pp. 1220–1227, 2016.

Ghoshal, A. and Honorio, J. Learning graphical games from behavioral data: Sufficient and necessary conditions. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pp. 1532–1540, 2017.

Ghoshal, A. and Honorio, J. Learning sparse polymatrix games in polynomial time and sample complexity. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, pp. 1486–1494, 2018.

Gomez-Rodriguez, M., Leskovec, J., and Krause, A. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1019–1028, 2010.

Gomez-Rodriguez, M., Balduzzi, D., and Schölkopf, B. Uncovering the temporal dynamics of diffusion networks. In *Proceedings of the 28th International Conference on Machine Learning*, pp. 561–568, 2011.

Goyal, S. and Moraga-Gonzaléz, J. L. R&D networks. *Rand Journal of Economics*, 32(4):686–707, 2001.

Honorio, J. and Ortiz, L. E. Learning the structure and parameters of large-population graphical games from behavioral data. *Journal of Machine Learning Research*, 16 (1):1157–1210, 2015.

Hu, C., Cheng, L., Sepulcre, J., Johnson, K. A., Fakhri, G. E., Lu, Y. M., and Li, Q. A spectral graph regression model for learning brain connectivity of alzheimer's disease. *PLOS ONE*, 10(5):e0128136, 2015.

Irfan, M. and Ortiz, L. On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach. *Artificial Intelligence*, 215:79–119, 2014.

Jackson, M. O. *Social and economic networks*. Princeton University Press, 2010.

Jackson, M. O. and Zenou, Y. Games on networks. In *Handbook of Game Theory, Vol. 4, Peyton Young and Shmuel Zamir (Eds.)*, pp. 95–163, 2014.

Katz, L. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953.

Kearns, M., Littman, M., and Singh, S. Graphical models for game theory. In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, pp. 253–260, 2001.

Koller, D. and Friedman, N. *Probabilistic graphical models: Principles and techniques*. MIT Press, 2009.

Lake, B. and Tenenbaum, J. Discovering structure by learning sparse graph. In *Proceedings of the Annual Cognitive Science Conference*, 2010.

Leng, Y., Dong, X., Moro, E., and Pentland, A. The rippling effect of social influence via phone communication network. In *Complex Spreading Phenomena in Social Systems, S. Lehmann and Y.-Y. Ahn (Eds.)*, pp. 323–333, 2018.

Leng, Y., Sella, Y., Ruiz, R., and Pentland, A. Contextual centrality: Going beyond network structures. *Scientific Reports*, 10(9401), 2020.

Mateos, G., Segarra, S., Marques, A. G., and Ribeiro, A. Connecting the dots: Identifying network structure via graph signal processing. *IEEE Signal Processing Magazine*, 36(3):16–43, 2019.

McPherson, M., Smith-Lovin, L., and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001.

Meinshausen, N., Bühlmann, P., et al. High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, 34(3):1436–1462, 2006.

Montanari, A. and Saberi, A. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.

Ortega, A., Frossard, P., Kovačević, J., Moura, J. M. F., and Vandergheynst, P. Graph signal processing: Overview, challenges and applications. *Proceedings of the IEEE*, 106(5):808–828, 2018.

von Luxburg, U. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.

Wai, H.-T., Scaglione, A., and Leshem, A. Active sensing of social networks. *IEEE Transactions on Signal and Information Processing over Networks*, 2(3):406–419, 2016.

Young, H. P. Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American economic review*, 99(5):1899–1924, 2009.

Zou, H. and Hastie, T. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67(2):301–320, 2005.