
Learning Quadratic Games on Networks

Yan Leng¹, Xiaowen Dong^{1,2} and Alex ‘Sandy’ Pentland¹

¹ Media Lab, Massachusetts Institute of Technology, Cambridge, MA, USA

² Department of Engineering Science, University of Oxford, Oxford, UK

Abstract

Individuals, or organizations, cooperate with or compete against one another in a wide range of practical situations. In the economics literature, such strategic interactions are often modeled as games played on networks, where an individual’s payoff depends not only on her action but also 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. We test the proposed frameworks in synthetic settings and further study several factors that affect their learning performance. Moreover, with experiments on three real world examples, we show that our methods can effectively and more accurately learn the games than the baselines. The proposed approach is among the first of its kind for learning quadratic games, and have both theoretical and 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 [1]. Individuals, unsurprisingly, are not merely connected; their decisions and actions often influence the ones around them. Indeed, Christakis and Fowler [2] 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 [3, 4] or implicit interactions [5, 6].

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 [7, 8]. 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, players are incentivized to follow similar actions to maximize their payoffs, e.g., students putting an effort together into a joint assignment or firms working on a collaborative research project [9]. In the latter case, however, one’s action reduces her neighbors’ incentives for action, such as the scenarios of firms competing on market prices or individuals on local public goods [10].

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 largely focused on studying the characteristics of games on known or predefined graphs [11, 12, 13]. 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, remains mostly unknown and needs to be estimated to understand the present better 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 [14, 15]), physically-motivated models (e.g., diffusion processes on networks [16, 17], or more recently signal processing models [18, 19]). 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 [20] and there has been a few studies recently proposing to learn the games from observed action data. For example, the works in [21, 22, 23] have proposed to learn graphical games by observing actions from linear influence games with linear influence functions, where the authors of [24] has considered polymatrix games with pairwise matrix payoff functions. The work in [25] has proposed to learn potential games on tree-structured networks of symmetric interactions. These conditions have been relaxed in [26] 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 (or strategy) space, which may be restrictive in certain practical scenarios where actions take continuous values.

In this study, we focus on learning quadratic games, i.e., games with linear-quadratic payoffs [11, 12, 13, 27]. Quadratic games are a broad class of games that have been extensively studied in the literature, not only because they allow for continuous actions and can be used to model actions of both strategic complements and substitutes, but also because they can be used to approximate games with complex non-linear payoffs. We propose a learning framework where, given the Nash equilibrium actions 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 [28], 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 network games with linear-quadratic payoffs. Second, 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. Third, we carry out three real world experiments, where we show that the proposed frameworks can infer the underlying social, trading, and political relationships between players under the assumption of quadratic games. 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 network of N individuals represented by a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where \mathcal{V} and \mathcal{E} denote the vertex and edge sets, respectively. For any pair of individuals i and j , $G_{ij} = G_{ji} = 1$ 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} . Following

the literature, in a network game of linear-quadratic payoffs, an individual i chooses her action a_i to maximize her utility, u_i , which has the following form [11, 12, 13, 27]:

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

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.

Let us define the vectorial forms $\mathbf{a} = [a_1, a_2, \dots, a_N]^T$, $\mathbf{b} = [b_1, b_2, \dots, b_N]^T$, and $\mathbf{u} = [u_1, u_2, \dots, u_N]^T$, where we use the convention that the subscript i indicates the i -th entry of the vector. Taking the first-order derivative of the utility u_i with respect to the action a_i in Eq. (1), we have:

$$\frac{\partial u_i}{\partial a_i} = b_i - a_i + \beta(\mathbf{G}\mathbf{a})_i. \quad (2)$$

Letting the right hand side of Eq. (2) to be 0, and combining the results for all i , we have:

$$\mathbf{b} - \mathbf{a} + \beta\mathbf{G}\mathbf{a} = \mathbf{0}. \quad (3)$$

It is clear from Eq. (3) that the following relationship holds, as pointed out in [11], for any (pure strategy) Nash equilibrium action \mathbf{a} :

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

hence

$$\mathbf{a} = (\mathbf{I} - \beta\mathbf{G})^{-1} \mathbf{b}, \quad (5)$$

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. (5). Furthermore, as proved in [11], this assumption also ensures the uniqueness and stability of the Nash equilibrium action \mathbf{a} . Notice that the formulation of Eq. (5) can be interpreted as computing steady state opinions in studying opinion dynamics under a linear DeGroot model [29] and has been used in works on social network sensing [30].

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. (5) provides a way of computing \mathbf{a} , the Nash equilibrium actions 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 a financial market. 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 Problem formulation

In the setting of this paper, we consider N players, connected by a fixed interaction network \mathbf{G} , playing K different games in each of which their payoffs depend not only on their own actions but also that of their neighbors. Let us define the marginal benefits for these K games as $\mathbf{B} = [\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \dots, \mathbf{b}^{(k)}] \in \mathbb{R}^{N \times K}$, where each column of \mathbf{B} is the marginal benefit vector for one game, and the corresponding actions of the players as $\mathbf{A} = [\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \dots, \mathbf{a}^{(K)}] \in \mathbb{R}^{N \times K}$. We first consider in this section the case where, for each game, the marginal benefits of individual players follow independent Gaussian distributions, and then move to the dependent case in Section 4. In our setting, 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. 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. (4).

To this end, we propose the following joint optimization problem of the graph structure \mathbf{G} and the marginal benefits \mathbf{B} :

$$\begin{aligned} & \underset{\mathbf{G}, \mathbf{B}}{\text{minimize}} && f(\mathbf{G}, \mathbf{B}) = \|(\mathbf{I} - \beta \mathbf{G})\mathbf{A} - \mathbf{B}\|_F^2 + \theta_1 \|\mathbf{G}\|_F^2 + \theta_2 \|\mathbf{B}\|_F^2, \\ & \text{subject to} && G_{ij} = G_{ji}, G_{ij} \geq 0, G_{ii} = 0 \text{ for } \forall i, j \in \mathcal{V}, \\ & && \|\mathbf{G}\|_1 = N, \end{aligned} \tag{6}$$

where $\text{tr}(\cdot)$, $\|\cdot\|_F$, and $\|\cdot\|_1$ denote operators that compute the trace, Frobenius norm, and entry-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. It is clear that, in the problem of Eq. (6), 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})¹, 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 [33].

3.2 Learning algorithm

Given the non-negativity of G_{ij} , we can re-write the constraint:

$$\|\mathbf{G}\|_1 = \mathbf{1}^T \mathbf{G} \mathbf{1} = N, \tag{7}$$

where $\mathbf{1} \in \mathbb{R}^N$ is the all-one vector. The constraints in Eq. (6) therefore form a convex set. The problem of Eq. (6) is thus a quadratic program jointly convex in \mathbf{B} and \mathbf{G} , and can be solved efficiently via the interior point methods [34]. In our experiments, we solve the problem of Eq. (6) using the Python software package CVXOPT [35]. 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) [36]) to find a solution. The algorithm is summarized in Algorithm 1.

Algorithm 1 Learning games with independent marginal benefits

- 1: **Input** Observed actions $\mathbf{A} \in \mathbb{R}^{N \times K}$ for K games, $\beta, \theta_1, \theta_2$
 - 2: **Output** Network $\mathbf{G} \in \mathbb{R}^{N \times N}$, marginal benefits $\mathbf{B} \in \mathbb{R}^{N \times K}$ for K games
 - 3: Solve for \mathbf{G} and \mathbf{B} in Eq. (6)
 - 4: **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 [37, 38]. Since the marginal benefit vector \mathbf{b} can be thought of as the individual preferences (toward a particular action), they may contribute in the presence of homophily 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 Problem formulation

The effect of homophily present in the marginal benefit vector \mathbf{b} may be quantified by the so-called Laplacian quadratic form on the graph:

$$\mathcal{Q}(\mathbf{b}) = \mathbf{b}^T \mathbf{L} \mathbf{b} = \frac{1}{2} \sum_{i,j \in \mathcal{V}} G_{ij} (b_i - b_j)^2, \tag{8}$$

¹Similar constraints have been adopted in [31, 32] for graph inference.

where $\mathbf{L} = \text{diag}(\sum_{j \in \mathcal{V}} G_{ij}) - \mathbf{G}$ is the unnormalized (combinatorial) graph Laplacian matrix [39]. We therefore propose to add this measure in the objective function of the problem of Eq. (6) to promote homophilous marginal benefits, which essentially assumes that the marginal benefits follow a multivariate Gaussian distribution. This leads to the following optimization problem:

$$\begin{aligned} & \underset{\mathbf{G}, \mathbf{B}}{\text{minimize}} \quad h(\mathbf{G}, \mathbf{B}) = \|(\mathbf{I} - \beta \mathbf{G})\mathbf{A} - \mathbf{B}\|_F^2 + \theta_1 \|\mathbf{G}\|_F^2 + \theta_2 \text{tr}(\mathbf{B}^T \mathbf{L} \mathbf{B}), \\ & \text{subject to} \quad G_{ij} = G_{ji}, G_{ij} \geq 0, G_{ii} = 0 \text{ for } \forall i, j \in \mathcal{V}, \\ & \quad \|\mathbf{G}\|_1 = N, \\ & \quad \mathbf{L} = \text{diag}(\sum_{j \in \mathcal{V}} G_{ij}) - \mathbf{G}, \end{aligned} \tag{9}$$

where 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. (6), θ_1 and θ_2 are two non-negative regularization parameters. The problem of Eq. (9) is similar to that of Eq. (6), 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. (6), the problem of Eq. (9) 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. (9). Let us take a closer look at Eq. (9). The constraint set is the same as that in Eq. (6), 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 optimization problem using the package CVXOPT.

Next, we fix \mathbf{G} and solve for \mathbf{B} in Eq. (9). By fixing \mathbf{G} , Eq. (9) becomes an unconstrained convex quadratic program, and thus admits a closed-form solution which can be obtained by setting the derivative to zero:

$$\frac{\partial h(\mathbf{G}, \mathbf{B})}{\partial \mathbf{B}} = -2((\mathbf{I} - \beta \mathbf{G})\mathbf{A} - \mathbf{B}) + 2\theta_2 \mathbf{L} \mathbf{B} = \mathbf{0}, \tag{10}$$

hence

$$\mathbf{B} = (\mathbf{I} + \theta_2 \mathbf{L})^{-1} (\mathbf{I} - \beta \mathbf{G}) \mathbf{A}. \tag{11}$$

We iterate 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 and, since both subproblems are strongly convex, is guaranteed to converge to a local minimum (see Proposition 2.7.1 in [40]). The complete algorithm is summarized in Algorithm 2.

Algorithm 2 Learning games with homophilous marginal benefits

- 1: **Input** Observed actions $\mathbf{A} \in \mathbb{R}^{N \times K}$ for K games, $\beta, \theta_1, \theta_2$
 - 2: **Output** Network $\mathbf{G} \in \mathbb{R}^{N \times N}$, marginal benefits $\mathbf{B} \in \mathbb{R}^{N \times K}$ for K games
 - 3: **Initialize** $\mathbf{B}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $t = 1$, $\Delta = 1$
 - 4: **while** $\Delta \geq 10^{-4}$ and $t \leq \#$ iterations **do**
 - 5: Solve for \mathbf{G}_t in Eq. (9) given \mathbf{B}_{t-1}
 - 6: Compute \mathbf{L}_t using \mathbf{G}_t
 - 7: $\mathbf{B}_t = (\mathbf{I} + \theta_2 \mathbf{L}_t)^{-1} (\mathbf{I} - \beta \mathbf{G}_t) \mathbf{A}$
 - 8: $\Delta = |h(\mathbf{G}_t, \mathbf{B}_t) - h(\mathbf{G}_{t-1}, \mathbf{B}_{t-1})|$ (for $t > 1$)
 - 9: $t = t + 1$
 - 10: **return** $\mathbf{G} = \mathbf{G}_{\text{iter}}$, $\mathbf{B} = \mathbf{B}_{\text{iter}}$.
-

5 Experiments on synthetic data

5.1 Experimental setting

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 vertex at each time by connecting it to an existing vertex in the graph via preferential attachment. Once the graphs are constructed, we compute β 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 and the other for the homophilous marginal benefits \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 $\epsilon \sim \mathcal{N}(\mathbf{0}, \frac{1}{10}\mathbf{I})$ to the simulated marginal benefits. Now, given \mathbf{b} and β , we compute the players’ Nash equilibrium actions according to Eq. (5). We consider $K = 50$ games for each of which we generate the actions \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, which is widely adopted in case of classification with imbalanced class labels. We compare our algorithms with two baseline methods for inferring graph structures given data observations: the sample correlation and the regularized graphical Lasso in [28]. 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.

5.2 Comparison of learning performance

The performance of the three methods in comparison is shown in the top row of Fig. 1 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. 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} contain more information about the graph structure. 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 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, 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.

Next, the same results for the case of homophilous marginal benefits are shown in the bottom row of Fig. 1. We observe the same increase in performance as $\rho(\beta\mathbf{G})$ increases for all the 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 unsurprising due to the way the observations \mathbf{A} are generated taking into account the regularization term in the objective in Eq. (9) that enforces homophily.

Robustness against regularization parameters. We next analyze the robustness of the performance of Algorithm 1 against the regularization parameter θ_1 in Eq. (6), and the results averaged

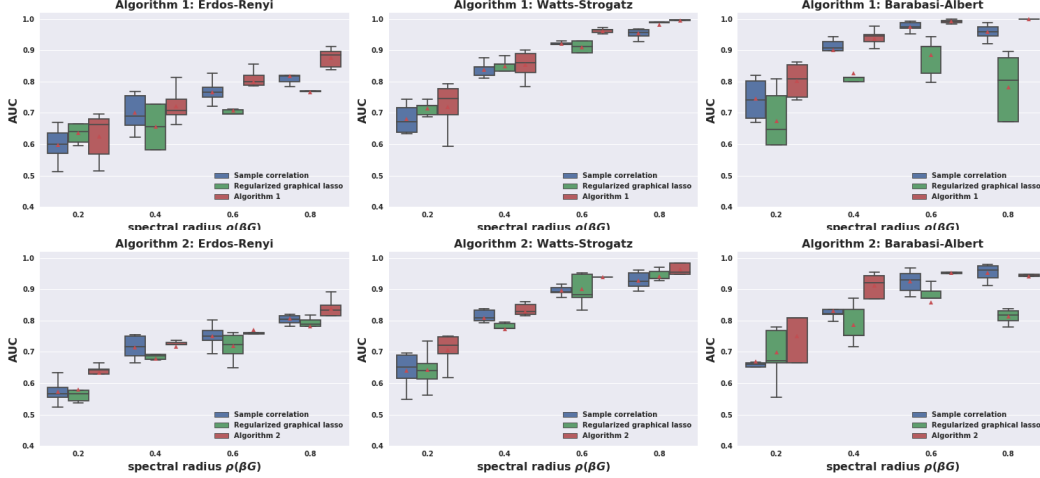


Figure 1: Comparison of learning performance between the proposed algorithm and baselines in the setting of independent marginal benefits (top row) and homophilous marginal benefits (bottom row). The red triangle, the middle line, lower and upper boundaries of the box (interquartile range or IQR) correspond 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.

over 20 random graph instances are presented in Fig. 2. In general, in addition to the effect of $\rho(\beta\mathbf{G})$ discussed above, we see a consistent pattern across the three graph models that link the values of θ_1 and θ_2 to the learning performance. Specifically, when θ_1 is smaller than around 10^2 , there is a region where a certain ratio of θ_1 to θ_2 leads to optimal performance, suggesting that in this case, the second and third terms are the dominating factors in the optimization of Eq. (6). A phase transition takes place when θ_1 is larger than 10^2 , where the performance becomes largely constant. The reason behind this behavior is as follows. When θ_1 increases, the Frobenius norm of \mathbf{G} in the objective function of Eq. (6) tends to be small. Given a fixed entry-wise L^1 -norm of \mathbf{G} , this leads to a more uniform distribution of the off-diagonal entries. When θ_1 is large enough, the edge weights become almost the same, leading to a constant AUC measure.

Similarly, we present in Fig. 3 the performance of Algorithm 2 with respect to different values of θ_1 and θ_2 in Eq. (9). We see that the patterns are generally consistent with that in Fig. 2, with one noticeable difference being that there also seems to be a phase transition taking place around the value of 10^{-1} for θ_2 . One possible explanation for this behavior is that, when θ_2 is large enough, the trace term in the objective function of Eq. (9) tends to be small, making the resulting graph with fewer edges but with larger weights. This contributes to an AUC score that is mostly constant.

5.3 Learning performance with respect to different factors in network games

In this section we examine the learning performance of the proposed algorithms 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.

Learning performance versus 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 1 and Algorithm 2 to be the ones that lead to the best learning performance in Fig. 2 and Fig. 3, respectively. In Fig. 4, we vary the number of games and evaluate its effect on the performance. We see that in general, the performance of both algorithms increase, as more observations become available. The benefit is least obvious for the ER graph with independent marginal benefits, 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.




Figure 2: Learning performance (in terms of AUC) of Algorithm 1 with respect to $\rho(\beta\mathbf{G})$, θ_1 , and θ_2 .




Figure 3: Learning performance (in terms of AUC) of Algorithm 2 with respect to $\rho(\beta\mathbf{G})$, θ_1 , and θ_2 .

Learning performance versus noise intensity in the marginal benefits. We now analyze the robustness of the result against noise intensity in the marginal benefits. It is clear that 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. 5, the learning performance generally decays as the intensity of noise increases, which is expected for both algorithms. The performance of the model is relatively stable until the standard deviation of the noise becomes larger than 1.

Learning performance versus 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, p and k in the WS, and m in the BA graphs on the learning performance of the two proposed algorithms. As shown in Fig. 6, increasing the randomness of edges in the WS graph via a higher rewiring probability decreases the performance of learning. This pattern also

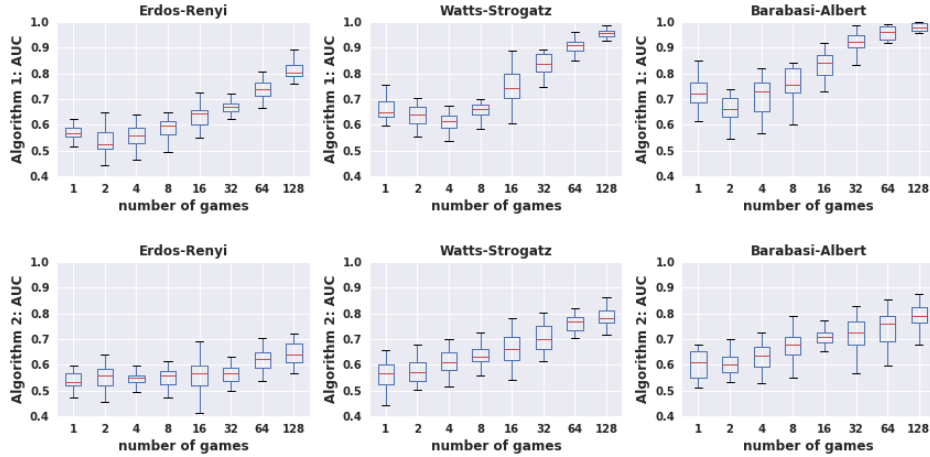


Figure 4: Performance measure versus number of games for Algorithm 1 (top row) and Algorithm 2 (bottom row).

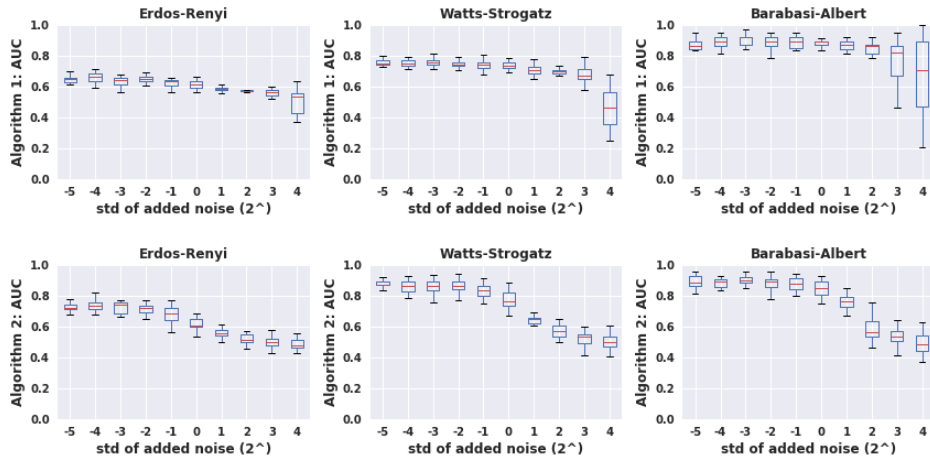


Figure 5: Performance measure versus noise intensity in the marginal benefits for Algorithm 1 (top row) and Algorithm 2 (bottom row).

explains why the performance on the ER graph (in which edges appear more randomly) is generally the worst among the three types of networks.

The remaining plots in Figure. 6 show that 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.

Learning performance versus strength of homophily. Finally, we analyze the influence of the strength of homophily on the learning performance of Algorithm 2. We consider three scenarios, i.e., weak, medium and strong homophily effect. To this end, we generate the marginal benefits \mathbf{b} as linear combinations of the eigenvectors corresponding to the 1st-5th, 6th-10th, and 11th-15th smallest eigenvalues of the graph Laplacian matrix. Due to the properties of the eigenvectors, these three sets lead to different quantities for the Laplacian quadratic form $\mathcal{Q}(\mathbf{b})$, hence corresponding to weak, medium and strong homophily effect, respectively. Notice that the presence of the homophily effect in \mathbf{B} tends to imply homophily in \mathbf{A} for the following reason. Regardless of the characteristics of

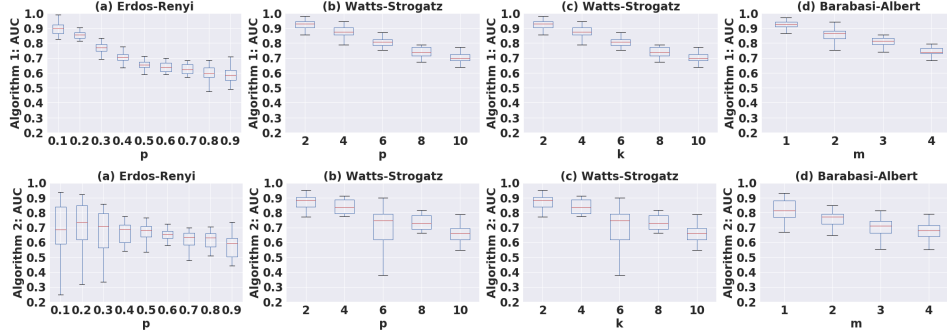


Figure 6: Learning performance versus structural properties of the network for Algorithm 1 (top row) and Algorithm 2 (bottom row).




Figure 7: Learning performance versus strength of homophily in the marginal benefits for Algorithm 2.

the game, a higher marginal benefit \mathbf{b} is more likely to incentivize higher activity level \mathbf{a} due to the first term of the utility function in Eq. (1). Therefore, homophily in \mathbf{B} tends to lead to homophily in \mathbf{A} , hence revealing more information about the graph structure. As shown in Fig. 7, for all the three types of networks, the stronger the homophily in the marginal benefits, the better the learning performance.

5.4 Learning the marginal benefits

In learning quadratic games, 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 [13]. To this end, 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, respectively. 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 groundtruth ones.

Table 1: Performance (in terms of 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 the network games. Given this broad assumption, we present three examples of inferring the network structure in quadratic games, hence demonstrating the effectiveness of the proposed algorithms in

practical scenarios. Our experiments cover the inference of three real world networks of social relationship, trading behavior, and political preference.

6.1 Social network

The first example is the inference of the structure of a social network between households in a village in rural India [41]. In particular, following the setting in [41], 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 [42, 43], which are formed by the decisions made by their neighbors. For example, if neighbors adopt a specific facility, villagers tend to gain higher utility after adopting the same facility by complying with social norms.

We consider each action as a strategy in a quadratic game, and we have in total 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%², indicating that they can recover a social network structure closer to the groundtruth.

Table 2: Performance (in terms of 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

6.2 Trading relationship

The second example is the inference of the structure of the global trade network. Specifically, we consider the overall trading activities of 235 countries on 96 export products and 96 import products in year 2008 as our observed actions³. 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 2002⁴. In constructing the groundtruth, we consider the edge weight between each pair of nations as the logarithmic of the total amount of trades (import and export) between the two nations.

The utilities from the demand and supply of a nation depends on that of their neighboring nations. These neighboring nations are the ones with which a particular nation traded in 2002. On the demand side, the more demand a nation has, the less utility this nation would gain from trading with a high-demand nation. On the supply side, the more supply a nation has, the less utility this nation would obtain by trading with a nation also with high supply. Therefore, we expect a strategic substitute relationship between the nations.

²The 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.

³Data 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).

⁴The 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.




Figure 8: Clustering of Swiss cantons based on the political network learned by Algorithm 1 (left) and Algorithm 2 (right).

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%, respectively⁵. The larger performance gain in this case is due to the fact that the regularized graphical Lasso approach in [28] is 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.

6.3 Political preference

The third example 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 2012⁶. 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 between these cantons reflected by their votes in the national referendums⁷.

Unlike the previous examples, it is more difficult to define a groundtruth network in this case. Instead, we apply spectral clustering [44] 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. 8(a) and Fig. 8(b), respectively. 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 Aargauis, 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

Extensive research has focused on the understanding of decision-making of individuals in a given network environment. However, there is a lack of effort in addressing the problem in the opposite direction, i.e., inferring the underlying interaction graph, especially when one is not readily available, given strategic decisions of players in a network game. In this paper, we have proposed two novel learning frameworks for a joint inference of graph structure and individual marginal benefits for a

⁵The 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.

⁶The voting statistics were obtained via <http://www.swissvotes.ch>.

⁷We 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.

broad class of network games, i.e., games with linear-quadratic payoffs. Testing our algorithms in both synthetic and real world settings, we show that they achieve superior performance compared to the baseline techniques. Furthermore, we study systematically several factors that affect the learning performance of the proposed algorithms. 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 practical scenarios. For instance, the learned graph, which captures the strategic interactions between the players, may be used for detecting communities formed by the players [45], which 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. 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 utilities of all players, which can be done by adjusting, according to the network topology, the marginal benefits via incentivization [13]. 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 instance, it would be essential to study graph inference given partial or incomplete observations of the players' actions, especially in the case where it is costly to observe the actions of all the network players. It would also be interesting to 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 may need to be adapted accordingly for network games of different payoff functions. We leave these studies as future work.

Acknowledgement

The authors would like to thank Georgios Stathopoulos and Dorina Thanou for helpful discussions about the optimization problems in the paper.

References

- [1] Lars Backstrom, Paolo Boldi, Marco Rosa, Johan Ugander, and Sebastiano Vigna. Four degrees of separation. In *Proceedings of the 4th Annual ACM Web Science Conference*, pages 33–42, 2012.
- [2] Nicholas A. Christakis and James H. Fowler. *Connected: The surprising power of our social networks and how they shape our lives*. Little, Brown, 2009.
- [3] Sinan Aral, Lev Muchnik, and Arun Sundararajan. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51):21544–21549, 2009.
- [4] Yan Leng, Xiaowen Dong, Esteban Moro, and Alex Pentland. The rippling effect of social influence via phone communication network. In *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks*, S. Lehmann and Y.-Y. Ahn (Eds.), pages 323–333, 2018.
- [5] A. Bandura and D. C. McClelland. Social learning theory. 1977.
- [6] Xiaowen Dong, Yoshihiko Suhara, Burcin Bozkaya, Vivek K. Singh, Bruno Lepri, and Alex Pentland. Social bridges in urban purchase behavior. *ACM Transactions on Intelligent Systems and Technology, Special Issue on Urban Intelligence*, 9(3):33:1–33:29, 2018.
- [7] Matthew O. Jackson and Yves Zenou. Games on networks. *Handbook of Game Theory, Vol. 4*, Peyton Young and Shmuel Zamir, eds., 2014.
- [8] Yann Bramoullé and Rachel Kranton. Games played on networks. *The Oxford Handbook of the Economics of Networks*, Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds., 2016.
- [9] Sanjeev Goyal and José Luis Moraga-González. R&D networks. *Rand Journal of Economics*, 32(4):686–707, 2001.
- [10] Yann Bramoullé and Rachel Kranton. Public goods in networks. *Journal of Economic Theory*, 135(1):478–494, 2007.
- [11] Coralio Ballester, Antoni Calvó-Armengol, and Yves Zenou. Who's who in networks. Wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.

- [12] Yann Bramoullé, Rachel Kranton, and Martin D’amours. Strategic interaction and networks. *American Economic Review*, 104(3):898–930, 2014.
- [13] Andrea Galeotti, Benjamin Golub, and Sanjeev Goyal. Targeting interventions in networks. *arXiv:1710.06026*, 2017.
- [14] D. Koller and N. Friedman. *Probabilistic graphical models: Principles and techniques*. MIT Press, 2009.
- [15] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- [16] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *Proc. of the 16th ACM SIGKDD Inter. Conf. on Knowledge Discovery and Data Mining*, pages 1019–1028, Washington, DC, USA, 2010.
- [17] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the temporal dynamics of diffusion networks. In *Proc. of the 28th Inter. Conf. on Machine Learning*, pages 561–568, Bellevue, Washington, USA, 2011.
- [18] X. Dong, D. Thanou, M. Rabbat, and P. Frossard. Learning graphs from data: A signal representation perspective. *arXiv preprint arXiv:1806.00848*, 2018.
- [19] G. Mateos, S. Segarra, A. G. Marques, and A. Ribeiro. Connecting the dots: Identifying network structure via graph signal processing. *arXiv preprint arXiv:1810.13066*, 2018.
- [20] Michael Kearns, Michael Littman, and Satinder Singh. Graphical models for game theory. In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, 2001.
- [21] Mohammad Irfan and Luis Ortiz. On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach. *Artificial Intelligence*, 215:79–119, 2014.
- [22] Jean Honorio and Luis E Ortiz. Learning the structure and parameters of large-population graphical games from behavioral data. *Journal of Machine Learning Research*, 16:1157–1210, 2015.
- [23] Asish Ghoshal and Jean Honorio. Learning graphical games from behavioral data: Sufficient and necessary conditions. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 1532–1540, 2017.
- [24] Asish Ghoshal and Jean Honorio. Learning sparse polymatrix games in polynomial time and sample complexity. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, pages 1486–1494, 2018.
- [25] Vikas Garg and Tommi Jaakkola. Learning tree structured potential games. In *Advances in Neural Information Processing Systems*, pages 1552–1560, 2016.
- [26] Vikas Garg and Tommi Jaakkola. Local aggregative games. In *Advances in Neural Information Processing Systems*, pages 5341–5351, 2017.
- [27] Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Networks, shocks, and systemic risk. Working Paper 20931, National Bureau of Economic Research, February 2015.
- [28] Brenden Lake and Joshua Tenenbaum. Discovering structure by learning sparse graph. In *Proceedings of the Annual Cognitive Science Conference*, 2010.
- [29] M. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69:118–121, 1974.
- [30] H.-T. Wai, A. Scaglione, and A. Leshem. Active sensing of social networks. *IEEE Transactions on Signal and Information Processing over Networks*, 2(3):406–419, Sep 2016.
- [31] C. Hu, L. Cheng, J. Sepulcre, K. A. Johnson, G. E. Fakhri, Y. M. Lu, and Q. Li. A spectral graph regression model for learning brain connectivity of alzheimer’s disease. *PLoS ONE*, 10(5):e0128136, May 2015.
- [32] X. Dong, D. Thanou, P. Frossard, and P. Vandergheynst. Learning laplacian matrix in smooth graph signal representations. *IEEE Transactions on Signal Processing*, 64(23):6160–6173, 2016.
- [33] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67, Part 2:301–320, 2005.
- [34] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [35] M. S. Andersen, J. Dahl, and L. Vandenberghe. CVXOPT: A Python package for convex optimization. pages 301–320, version 1.2.0. Available at cvxopt.org, 2018.
- [36] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [37] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001.
- [38] Matthew O. Jackson. *Social and economic networks*. Princeton University Press, 2010.

- [39] F. R. K. Chung. *Spectral graph theory*. American Mathematical Society, 1997.
- [40] D. P. Bertsekas. *Nonlinear programming*. Athena Scientific, 1995.
- [41] Abhijit Banerjee, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. The diffusion of microfinance. *Science*, 341(6144):1236498, 2013.
- [42] H Peyton Young. Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American economic review*, 99(5):1899–1924, 2009.
- [43] Andrea Montanari and Amin Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.
- [44] U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, Dec 2007.
- [45] S. Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.