
Learning to Infer Structures of Network Games

Emanuele Rossi^{1,2} Federico Monti¹ Yan Leng³ Michael M. Bronstein^{1,4} Xiaowen Dong⁴

Abstract

Strategic interactions between a group of individuals or organisations can be modelled as games played on networks, where a player’s payoff depends not only on their actions but also on those of their neighbours. Inferring the network structure from observed game outcomes (equilibrium actions) is an important problem with numerous potential applications in economics and social sciences. Existing methods mostly require the knowledge of the utility function associated with the game, which is often unrealistic to obtain in real-world scenarios. We adopt a transformer-like architecture which correctly accounts for the symmetries of the problem and learns a mapping from the equilibrium actions to the network structure of the game without explicit knowledge of the utility function. We test our method on three different types of network games using both synthetic and real-world data, and demonstrate its effectiveness in network structure inference and superior performance over existing methods.

1. Introduction

Individuals or organisations cooperate with or compete against each other in a wide range of practical situations. Strategic interactions between individuals are often modelled as games played on networks (Jackson & Zenou, 2014), where an individual’s utility depends not only on their actions (and characteristics) but also on those of their neighbours. In a network game (and games in general), the utility function is of fundamental importance, as rational agents maximise their utilities when making real-life decisions. The current literature on network games has primarily focused on the scenarios where the utility function is predefined, and the structure of the network (represented by a

graph) is known beforehand. However, in practical settings, while it is common to observe the actions of the players, the underlying interaction network often remains hidden or uncertain due to privacy reasons or the dynamic nature of interactions. Furthermore, the utility function is often unobservable. This makes it challenging to exploit network information and utility function for behavioural predictions and network-based interventions (Valente, 2012), e.g., marketing campaigns or information diffusion.

In this paper, we focus on the problem of inferring the structure of the interaction network from observed equilibrium actions of a network game. A few recent studies have tackled similar problems (Irfan & Ortiz, 2011; Honorio & Ortiz, 2015; Ghoshal & Honorio, 2017a;b; Garg & Jaakkola, 2016; 2017; Barik & Honorio, 2019; 2020; Leng et al., 2020); however, several major limitations remain. First, all these methods require the explicit knowledge of the utility function to infer the underlying network structure, which may be impractical to assume and may also change over time. The work of Honorio & Ortiz (2015) considers a more general hypothesis space of games for linear influence game, but they only focus on binary actions and linear payoffs. Second, the methodology in each of these studies has been designed for the specific game under consideration, thereby limiting its scope in handling a wide range of strategic interactions in real-world scenarios.

To address these limitations, we first summarise three common network games studied in the recent literature in a generic form, which is based on both individual and network factors that impact one’s utility and the corresponding equilibrium actions. Despite the different nature of the games, it permits the relation between the equilibrium actions and network structure to be written in a unified manner. This motivates us to propose a data-driven model where we learn the functional mapping between the equilibrium actions and network structure *without* explicit knowledge about the utility function of the game. Our model is based on an encoder-decoder approach where the encoder is a transformer architecture and the decoder can be chosen flexibly, and is trained using pairs of actions and network structures. Once trained, the model can be deployed to infer the network structure from only observed actions. Synthetic and real-world experiments demonstrate the superiority of our method against several state-of-the-art approaches.

*Equal contribution ¹Twitter, London, UK ²Imperial College London, London, UK ³The University of Texas at Austin, Austin, TX, USA ⁴University of Oxford, Oxford, UK. Correspondence to: Emanuele Rossi <emanuele.rossi1909@gmail.com>.

Main contributions. First, we compare three network games using a unified parameterisation, which helps reveal the different nature of these games and interpret the strategic interactions they represent. Second, to our knowledge, our framework is one of the first that is able to infer the network structure behind the games without explicit knowledge about the utility function. This capability is important in real-world scenarios where the nature of the interactions remains hidden or may even evolve over time. Finally, our work contributes to the emerging field of data-driven structural inference by proposing a model based on a novel architecture which adopts a permutation-invariant transformer. Overall, our work contributes to the inference of strategic decision making in various network settings.

2. Related Work

Network games, a class of problems in game theory, have been studied extensively in computer science and economics. The majority of works in the network game literature study the characteristics of games on a known and static graph (Ballester et al., 2006; Bramoullé et al., 2014). While these studies are useful in understanding collective actions and designing interventions (Galeotti et al., 2017), it is increasingly acknowledged that networks are difficult to obtain in practice. Furthermore, the utility function associated with the game is usually unknown as well. We are interested in the *inverse problem* of inferring the network structure based on observed actions. This inverse setting is related to *graph* or *network inference*, a problem that has attracted interest in statistics (Koller & Friedman, 2009; Friedman et al., 2008a), physics (Gomez-Rodriguez et al., 2010; Gomez-Rodriguez et al., 2011), signal processing (Mateos et al., 2019; Dong et al., 2019). Our study differs from these works in accounting for the strategic interactions and the game theoretical framework underlying the observed data.

Deep learning models have been recently proposed for latent graph inference in a number of settings. Kipf et al. (2018) proposes a graph neural network (GNN) model to infer the interactions of an underlying dynamical system from observational data. Differently from this work, their model is trained on predictions of the future state of the system, where there is lack of validation for the learned network interactions. Similarly, in latent graph learning (Cosmo et al., 2020a;b; Wang et al., 2019), the graph is learned jointly with a downstream task, conversely to our scenario where the structure itself is the learning objective. Methods for link prediction (Zhang & Chen, 2018) predict edges in a graph, but they typically require part of the true links to be provided as input (and only predict the missing ones), whereas in our scenario we are interested in inferring networks without observing any link in the test data. Among the few works where the network structure is the learning objective, Yu

et al. (2019) and Zheng et al. (2020) propose to infer a DAG from the observed actions. However, these approaches are limited to predicting acyclic graphs, whereas the graphs we are interested in are often cyclic. The method of Belilovsky et al. (2017) is the mostly related to ours, since they propose a supervised model to infer an undirected graphical model from observed covariates using a series of dilated convolutions. However, their model is not permutation-equivariant w.r.t. the order of nodes and the number of layers depends on the number of nodes. Both issues cause the statistical efficiency of the model to scale poorly with the size of the graph. Different from the above studies, our framework aims to learn game-theoretical relationships in a supervised manner, while leveraging the structural symmetries of this problem.

Finally, there has been a recent stream of literature in learning network games from actions of players (Irfan & Ortiz, 2011; Honorio & Ortiz, 2015; Ghoshal & Honorio, 2017a;b; Garg & Jaakkola, 2016; 2017; Barik & Honorio, 2019; 2020; Leng et al., 2020). Most of these methods focus on either a binary or a finite discrete action space. For continuous actions, Leng et al. (2020) formulate an optimisation problem to learn the structure and marginal benefits of linear quadratic games, while Barik & Honorio (2019) aim at inferring the network structure from an action-conforming graphical game. Our work differs from existing methods in the literature in that it does not assume a particular game-theoretic structure. Instead, we build a transformer-like model that learns a mapping from the equilibrium actions to the network structure of the games without explicit knowledge of the utility functions.

3. Setting

In this section, we start by analysing three commonly studied network games (Sections 3.1 and 3.2). Based on the specific utility function of these games, we establish a generic relationship between the equilibrium actions and network structure (Section 3.3). This eventually motivates the proposed framework that learns network structure without any knowledge of the utility function¹ (Section 4).

3.1. Continuous-action network games

Consider a network of N individuals represented by a weighted and undirected 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 , $w_{ij} = w_{ji} > 0$ if $(i, j) \in \mathcal{E}$ and $w_{ij} = w_{ji} = 0$ otherwise, where w_{ij} is the ij -th entry of the adjacency matrix \mathbf{W} . In this work, we assume the graph

¹The utility function is used only in analysing three classical games in a universal framework as well as generating data in synthetic experiments. It is not used by the proposed method itself.

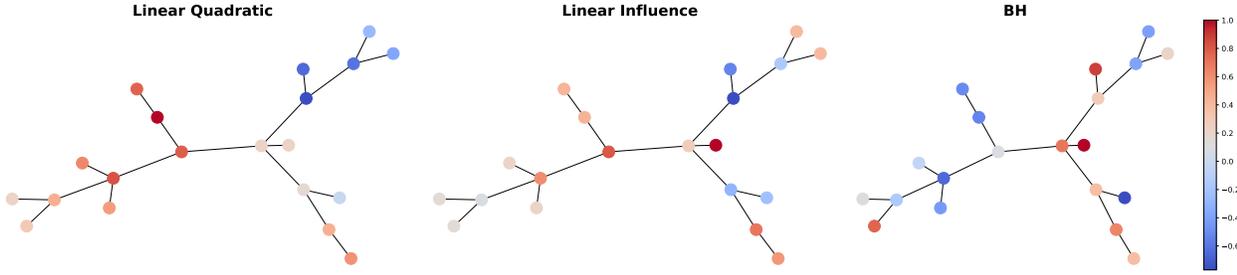


Figure 1. Example actions for different types of games on a Barabási-Albert graph. Actions are normalised to be in $[-1, 1]$ and are displayed as colors on nodes.

\mathcal{G} is connected, and consider the normalised adjacency matrix $\mathbf{A} = \mathbf{D}^{-\frac{1}{2}}\mathbf{W}\mathbf{D}^{-\frac{1}{2}}$ where $\mathbf{D} = \text{diag}(\mathbf{W}\mathbf{1})$, where $\mathbf{1}$ is an N -dimensional vector of ones. In a network game, the payoff u_i of an individual i depends on their action x_i as well as the actions of neighbours $j \in \mathcal{N}_i$. We consider three commonly studied network games: linear quadratic games (Ballester et al., 2006), a variation of the linear influence games developed in Irfan & Ortiz (2011), and the graphical game studied in Barik & Honorio (2019) (which we will refer to as Barik-Honorio or BH graphical game).

Linear quadratic games. Linear quadratic games are widely studied in the economics literature (Jackson & Zenou, 2014; Ballester et al., 2006). In this game, a player i chooses their action by maximising the following utility function:

$$\max_{\{x_i\}} u_i = b_i x_i - \frac{1}{2} x_i^2 + \beta \sum_{j \in \mathcal{N}_i} a_{ij} x_i x_j, \quad (1)$$

where b_i represents the marginal benefit of i by taking action x_i and β is the strength of dependencies between actions of neighbours in the network, respectively. Note that this utility function can also be thought of as a second-order approximation to non-linear utility functions of more complex games. The pure-strategy Nash equilibrium (PSNE) of this game is

$$\mathbf{x}^* = (\mathbf{I} - \beta \mathbf{A})^{-1} \mathbf{b}, \quad (2)$$

where \mathbf{x}^* and \mathbf{b} are N -dimensional vectors collecting actions and marginal benefits for all individuals, and \mathbf{I} is the $N \times N$ identity matrix. Under the assumption that $|\beta| < 1$, the matrix inverse is guaranteed to exist as the spectral radius of \mathbf{A} is 1. Furthermore, when $\beta > 0$, the game corresponds to a strategic complement relationship (i.e., intuitively, the incentive of a player to take a higher action is increasing in the number of their neighbours also taking a higher action); when $\beta < 0$, it corresponds to strategic substitute (i.e., intuitively, the incentive of a player to take a higher action is decreasing in the number of their neighbours also taking a higher action).

Linear influence games. Inspired by the threshold model

(Granovetter, 1978), Irfan & Ortiz (2011) proposed the linear influence games, where an individual chooses the action that maximises the following utility function²:

$$\max_{\{x_i\}} u_i = \sum_{j \in \mathcal{N}_i} a_{ij} x_i x_j - b_i x_i, \quad (3)$$

where b_i can be understood as a threshold parameter for i 's level of tolerance for negative effects. Under the assumption that \mathbf{A} is invertible, the PSNE satisfies the following condition:

$$\mathbf{x}^* = \mathbf{A}^{-1} \mathbf{b}. \quad (4)$$

BH graphical game. Finally, in the specific graphical game introduced in Barik & Honorio (2019), an individual maximises the following utility function:

$$\max_{\{x_i\}} u_i = - \left\| x_i - \sum_{j \in \mathcal{V}} a_{ij} x_j \right\|_2. \quad (5)$$

This utility can be used to model games where an individual prefers to conform to the social norm (actions of their neighbours), since their own utility decreases as their action deviates from those of their neighbours. The PSNE for this game satisfies the following condition:

$$\mathbf{x}^* = \mathbf{A} \mathbf{x}^*. \quad (6)$$

This suggests that \mathbf{x}^* is the eigenvector \mathbf{u}_1 of \mathbf{A} which is associated with the largest eigenvalue (which is 1). We consider actions from the set of ϵ -PSNE (Barik & Honorio, 2019), which are obtained by adding noise independently per player. The observed actions \mathbf{x} are: $\mathbf{x} = \mathbf{x}^* + \mathbf{e}$, where \mathbf{e} is Gaussian noise.

Example actions for the three types of games on a Barabási-Albert graph have been illustrated in Fig. 1. By investigating Eq. (2), (4), and (6), we can write the condition for the equilibrium actions \mathbf{x}^* in a generic form³ (see Table 1):

$$\mathbf{x}^* = \mathcal{F}(\mathbf{A}) \mathcal{H}(\mathbf{b}), \quad (7)$$

² The actions are discrete in the originally proposed game. We adapt the game to a continuous setting.

³ We tacitly assume that $|\beta| < 1$ and \mathbf{A} is invertible.

Table 1. Parameterisation of three network games.

	$\mathcal{F}(\mathbf{A})$	$\mathcal{H}(\mathbf{b})$
Linear quadratic	$(\mathbf{I} - \beta\mathbf{A})^{-1}$	\mathbf{b}
Linear influence	\mathbf{A}^{-1}	\mathbf{b}
Barik-Honorio	\mathbf{u}_1	1

where $\mathcal{F}(\mathbf{A})$ is a function of the network structure and $\mathcal{H}(\mathbf{b})$ is a function of additional parameters (if any) associated with the game. That is, $\mathcal{F}(\mathbf{A})$ accounts for the influence from the actions of one’s neighbours in the network. Conversely, $\mathcal{H}(\mathbf{b})$ is only affected by one’s idiosyncratic (individual) characteristics.

3.2. Modeling of Individual Idiosyncratic Characteristics

Under the linear quadratic or linear influence games, the parameter \mathbf{b} captures the marginal benefits or tolerance levels as idiosyncratic characteristics of players in the corresponding games. In the presence of the homophily effect (McPherson et al., 2001), we may assume that this parameter is associated with the network structure. To this end, we propose to model \mathbf{b} as follows:

$$\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \mathbf{L}_\alpha^\dagger), \quad \mathbf{L}_\alpha = (1 - \alpha)\mathbf{I} + \alpha\mathbf{L} \quad (8)$$

where $\mathbf{L} = \mathbf{I} - \mathbf{A}$ is the normalised graph Laplacian matrix, and \dagger represents pseudoinverse in case the matrix is not invertible (this happens when $\alpha = 1$ as \mathbf{L} has a smallest eigenvalue of 0). The parameter $\alpha \in [0, 1]$ controls the relation of the individual idiosyncratic characteristics to the network structure. The two corner cases $\alpha = 0$ (for which we have $\mathbf{L}_0 = \mathbf{I}$) and $\alpha = 1$ (when $\mathbf{L}_1 = \mathbf{L}$) correspond to independent idiosyncratic characteristics and homophilous idiosyncratic characteristics (individuals with similar characteristics tend to be connected), respectively. By varying α from 0 to 1 we can achieve increasing levels of homophily or smoothness (see Section 3.3) of \mathbf{b} on the graph.

3.3. Analysis of Equilibrium Actions

With the conditions for equilibrium actions in Section 3.1 and the modeling of individual idiosyncratic characteristics in Section 3.2, we can analyse explicitly the characteristics of these actions.

Linear quadratic games. Assuming \mathbf{b} of the form (8) and using Eq. (2) and $\mathbf{L} = \mathbf{I} - \mathbf{A}$, we have that the equilibrium actions \mathbf{x}^* follow a multivariate Gaussian distribution:

$$\begin{aligned} \mathbf{x}^* &\sim \mathcal{N}\left(\mathbf{0}, (\mathbf{I} - \beta\mathbf{A})^{-1}(\mathbf{I} - \alpha\mathbf{A})^\dagger(\mathbf{I} - \beta\mathbf{A})^{-1}\right) \\ &= \mathcal{N}\left(\mathbf{0}, \mathbf{U}[(\mathbf{I} - \beta\mathbf{A})^2(\mathbf{I} - \alpha\mathbf{A})]^\dagger\mathbf{U}^\top\right), \end{aligned} \quad (9)$$

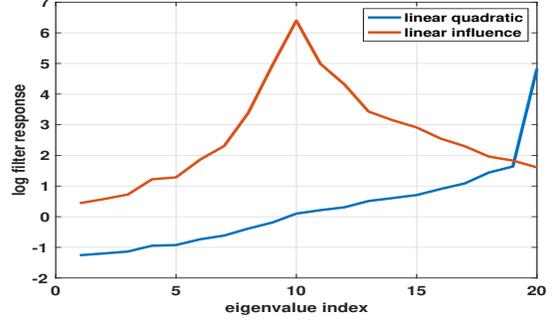


Figure 2. Interpreting actions of linear quadratic/influence games: the equilibrium actions for the two games are dominated by different sets of eigenvectors.

with the eigendecomposition $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$. Eq. (9) illustrates the relationship between the actions \mathbf{x} and the network structure \mathbf{A} and motivates the learning framework proposed in the next section. Furthermore, the covariance in Eq. (9) may be interpreted as a graph filter whose frequency response $\frac{1}{(1-\beta\lambda)^2(1-\alpha\lambda)}$ may shed light on the behaviour of the actions. For $\beta \rightarrow 1$ and $\alpha \rightarrow 1$, the action vector \mathbf{x} tends to behave like the leading eigenvectors of \mathbf{A} (‘low frequency’), which are smooth on the graph⁴. An example of the filter response (with $\beta = 0.8$ and $\alpha = 0.8$) applied on the eigenvalues of an instance of a 20-node Erdős-Rényi graph is shown in Fig. 2 (blue). This shows that in this case the actions of linear-quadratic games are dominated by the leading eigenvector of \mathbf{A} which is smooth.

Linear influence games. Similarly, from Eq. (4) and (8), we have that the equilibrium actions \mathbf{x} follow a multivariate Gaussian distribution:

$$\begin{aligned} \mathbf{x}^* &\sim \mathcal{N}\left(\mathbf{0}, \mathbf{A}^{-1}(\mathbf{I} - \alpha\mathbf{A})^\dagger\mathbf{A}^{-1}\right) \\ &= \mathcal{N}\left(\mathbf{0}, \mathbf{U}[\mathbf{\Lambda}^2(\mathbf{I} - \alpha\mathbf{\Lambda})]^\dagger\mathbf{U}^\top\right). \end{aligned} \quad (10)$$

Interpreting Eq. (10) as a spectral filter of the form $\frac{1}{\lambda^2(1-\alpha\lambda)}$ we can also conclude that an $\alpha \rightarrow 1$ tends to lead to smoother actions on the graph. However, the exact behaviour of actions in this case depends on the magnitude of the eigenvalue of \mathbf{A} closest to 0. Given that the spectrum of \mathbf{A} lies in the range of $[-1, 1]$, the actions are likely to behave like mid-spectrum eigenvectors, which are not necessarily smooth signals on the graph. Similarly, this can be seen from the filter response, under $\alpha = 0.8$, shown in Fig. 2 (red). This shows that in linear influence games the actions are dominated by mid-spectrum eigenvectors which are not necessarily smooth.

BH graphical game. We see from Eq. (6) that the equilibrium actions correspond to the largest eigenvector \mathbf{u}_1 of \mathbf{A} .

⁴By ‘smoothness’ here we mean the Dirichlet energy of the features, i.e., $\text{trace}(\mathbf{X}^\top\mathbf{L}\mathbf{X})$.

Although in this setting the observed actions are ϵ -PSNE, they would still tend to be smooth on the graph.

In summary, our analysis in this section shows that the smoothness of the equilibrium actions on the graph depends on β and α in the linear quadratic game. The actions in linear influence game are likely to be nonsmooth, whereas those in the BH graphical game are likely to be smooth. We empirically validate this analysis in Section 5.1. Regardless of the smoothness, the relationship between the equilibrium actions and the network structure demonstrated in this section motivates us to propose a learning framework in the next section to infer the network structure from the observed actions.

4. Proposed Approach

Motivated by the analysis in Section 3.3, we propose a model that learns a direct mapping from the observed actions to the network structure. Such a model is agnostic to the utility function of the game and avoids strong assumptions on it (as long as it conforms to the broad class of games whose equilibrium actions can be parameterised by Eq. (7)). Specifically, we consider the scenario where social network and decision data exist for a small population. The objective is to learn the mapping from decisions to the network structure on this small population such that it can be used to infer a large-scale unobserved network. For example, for cost-effective data collection, government, public agencies, and researchers can collect social network data on a small population (by asking individuals to nominate their friends) and then use the proposed method to learn the network interactions for a larger population.

In our setting, we assume to have a training set \mathcal{D} of action-graph pairs $(\mathbf{X}^{(\ell)}, \mathbf{A}^{(\ell)})$ coming from games with the same (but unknown) utility function. For each ℓ , the model takes as input an $N \times K$ matrix $\mathbf{X}^{(\ell)}$ containing the actions of N players over K independent games as columns, and outputs a predicted $N \times N$ adjacency matrix $\hat{\mathbf{A}}^{(\ell)} = g_{\Theta}(\mathbf{X}^{(\ell)})$ for the network game (we drop ℓ from now on for simplicity). Once trained, our model can then infer the network structure corresponding to previously unseen actions, as long as they are generated following a similar utility function. Moreover, this framework makes our model generalisable to learning networks with the number of nodes different from what was observed during training.

The model parameters are learned such that they minimise the cross-entropy between the binary ground truth adjacency matrix \mathbf{A} and the predicted continuous one $\hat{\mathbf{A}}$. Binary cross-entropy is a standard loss function for link prediction tasks using graph neural networks (Hamilton et al., 2017), of which our problem is an instance. We also experimented with weighting the loss based on the proportion of edges

in the graph, but it did not change the performance of the model.

Our model follows an encoder-decoder architecture $g_{\Theta}(\mathbf{X}) = \text{dec}(\text{enc}(\mathbf{X}))$, which is a standard solution for link-prediction problems in graph neural networks (Kipf & Welling, 2016) since the resulting number of parameters of the model is independent on the size of the graph, allowing for more statistical efficiency, and for the same model to work on graphs of different sizes. The *encoder* outputs an $N \times K \times F$ tensor $\mathbf{Z} = \text{enc}(\mathbf{X})$, mapping each node i in game k to an F -dimensional latent embedding \mathbf{z}_{ik} . The *decoder* outputs the predicted $N \times N$ adjacency matrix $\hat{\mathbf{A}} = \text{dec}(\mathbf{Z})$, where the score \hat{a}_{ij} for each edge (i, j) is computed using the $K \times F$ -dimensional embeddings $\mathbf{Z}_i, \mathbf{Z}_j$ of the respective nodes.

Given that the ordering of nodes in the graph is arbitrary, the model has to be a *permutation-equivariant* (Bronstein et al., 2021) function over the set of nodes: $g_{\Theta}(\mathbf{P}_1\mathbf{X}) = \mathbf{P}_1g_{\Theta}(\mathbf{X})\mathbf{P}_1^{\top}$, where \mathbf{P}_1 is an $N \times N$ permutation matrix interpreted as reordering of the nodes of the graph. Additionally, in the case where there is no correspondence between game k in one graph and game k in another one (e.g., actions corresponding to user rating where however users in different graphs have rated different items), the model should also be *permutation-invariant* over the set of games, i.e., $g_{\Theta}(\mathbf{X}\mathbf{P}_2) = g_{\Theta}(\mathbf{X})$, where \mathbf{P}_2 is another $K \times K$ permutation matrix on the games. Overall, this combined symmetry condition can be written as $g_{\Theta}(\mathbf{P}_1\mathbf{X}\mathbf{P}_2) = \mathbf{P}_1g_{\Theta}(\mathbf{X})\mathbf{P}_1^{\top}$.

Encoder The input to our model is a variable-length set (of actions), which excludes using a multi-layer perceptron (since it would not be able to handle the variable length) and sequence models such as LSTMs (Hochreiter & Schmidhuber, 1997) or GRUs (Cho et al., 2014) (since they treat the input as an ordered sequence, while it is an un-ordered set). Moreover, since the ground truth graph between the players is not known a priori, but it is however important for the model to exchange information between players (as the values of agents' actions are meaningless if not compared to the ones of the others), the encoder needs to perform message passing on the fully connected graph of players. Therefore, we propose a Transformer-like (Vaswani et al., 2017) encoder (which we refer to as Network Game Transformer or *NuGgeT*). *NuGgeT* processes a *set* of games happening on a given network and outputs for each node a *set* of K different game-specific embeddings (one for each game that happens on the same network). In this architecture, the value of action x_{ik} of player i in game k is never mixed with the actions of other players in a different game k' , if not for the computation of an aggregated attention score α_{ij} that captures the overall similarity between players i and j . Intuitively, in settings where we are given the outcomes of multiple games played on multiple graphs, there is no

correspondence between such games. Thus, the value of an action does not bring any information about the role of a node in a graph and the only useful information that needs to be exchanged across games is the similarity of the nodes.

Specifically, for each node $i = 1, \dots, N$ and game $k = 1, \dots, K$, the embedding \mathbf{z}_{ik} is computed as follows:

$$\begin{aligned} \mathbf{y}_{ik} &= \text{ReLU}(x_{ik}\mathbf{w} + \mathbf{b}) \\ \alpha_{ij}^{(h)} &= \text{softmax}_j \left(\square_{k=1}^K \mathbf{y}_{ik}^\top \mathbf{W}_Q^{(h)} \mathbf{W}_K^{(h)} \mathbf{y}_{jk} \right) \\ \mathbf{z}_{ik} &= \phi \left(\mathbf{y}_{ik}, \sum_{j=1}^N \alpha_{ij}^{(1)} \mathbf{y}_{jk}, \dots, \sum_{j=1}^N \alpha_{ij}^{(H)} \mathbf{y}_{jk} \right) \end{aligned} \quad (11)$$

where \square denotes a general permutation-invariant aggregation operator (e.g. max, mean $\frac{1}{K} \sum_{k=1}^K$ or sum $\sum_{k=1}^K$), ψ is some learnable function, $\mathbf{w} \in \mathbb{R}^F$, $\mathbf{b} \in \mathbb{R}^F$, $\mathbf{W}_K^{(h)} \in \mathbb{R}^{F \times F'}$, and $\mathbf{W}_Q^{(h)} \in \mathbb{R}^{F' \times F}$ are learnable parameters, and $h = 1, \dots, H$ denotes the attention heads. For a generic node i in game k , NuGgeT first expands i 's action x_{ik} into a vector \mathbf{y}_{ik} of F features, then computes H attention scores (one for each head) for each pair of nodes (i, j) via multiplicative attention on the expanded actions $\mathbf{y}_{.k}$ ⁵ (aggregating the unnormalised scores across games), and finally refines vector \mathbf{y}_{ik} via a function ψ that processes the H aggregated representation of the neighbours obtained from the attention mechanism.

Decoder In the case of directed graphs, there is no specific requirement for the decoder. If the graph is undirected, and therefore the adjacency matrix is symmetric, the decoder should be symmetric w.r.t. pairs of nodes (i.e., (i, j) and (j, i) treated the same way). Additionally, we impose invariance w.r.t. ordering of the games. *NuGgeT*'s decoder computes the predicted adjacency \hat{a}_{ij} by aggregating the game-specific embeddings computed in (11) for a pair of nodes (i, j) using a general permutation-invariant aggregation operator \square , which is then passed through a learnable function ψ :

$$\hat{a}_{ij} \psi \left(\square_{k=1}^K \mathbf{z}_{ik} \odot \mathbf{z}_{jk} \right). \quad (12)$$

Here \odot denotes element-wise product, whose use ensures symmetry w.r.t. node pairs (since $\mathbf{z}_{ik} \odot \mathbf{z}_{jk} = \mathbf{z}_{jk} \odot \mathbf{z}_{ik}$), while the permutation invariant operator over k ensures invariance to the ordering of games. We empirically observe this approach to work better compared to simpler

⁵For a given player i , using the expanded actions $\mathbf{y}_{.k}$ rather than input scalars $x_{.k}$ allows the attention mechanism to produce attention scores which are not necessarily linearly dependent on the value of neighbours' action x_{jk} , thus producing richer attention scores (see Fig. 12 in Appendix).

permutation-invariant functions such as the dot product of the concatenation of the embeddings from multiple games. In the SM we prove that *NuGgeT* satisfies the symmetry conditions outlined above.

5. Experiments

Our implementation of *NuGgeT* uses the sum $\sum_{k=1}^K$ as the permutation-invariant functions \square and \odot , and two different 2-layer MLPs for ϕ and ψ . We use the Adam optimiser (Kingma & Ba, 2015) with a learning rate of 0.001, a batch size of 100 and a patience of 50 epochs. We did not perform any particular hyperparameter tuning for our method, since we found it to be quite robust to the choice of hyperparameters and perform well with standard choices (see Appendix). In all our experiments, we report the mean and the standard error of the mean over the test graphs. Note that we ignore diagonal elements of the adjacency matrix both for training and evaluation. Since they are always zero, the model could easily memorise them, influencing the metrics. We use an AWS p3.16xlarge machine with 8 GPUs. While the training of our model takes between 5 and 10 minutes on a single GPU, the whole set of experiments conducted in the paper necessitate roughly 4 days of GPU time.

5.1. Synthetic Data

Data Generation We follow the setup in Leng et al. (2020) for generating the synthetic graphs using three different random graph models: Erdős-Rényi (ER), Watts-Strogatz (WS), and Barabási-Albert (BA). More details and exact parameters for the synthetic data are provided in the Appendix. All the graphs have $N = 20$ vertices in our experiments. For each type of graphs above, we simulate equilibrium actions for linear quadratic, linear influence, and BH graphical games using their respective utility function. For linear quadratic games, once the graphs are constructed, we compute $\beta > 0$ such that the spectral radius $\rho(\beta A)$ varies between 0 and 1. For linear influence games, when the adjacency matrix is not invertible we take its pseudoinverse. For BH games, we set $\text{std}(\mathbf{e}) = 1$ and ensure the resulting actions are ϵ -PSNE with $\epsilon = 0.2$. The generated actions together with the ground truth network structure are used to train the model. We use 850 graphs in the training set, 50 graphs for validation and 100 graphs for testing and verify that there is no overlap between them. It is important to notice that while α , β and ϵ are used to generate the synthetic data, they are not used or known by the model.

Baselines We compare with the following general baselines: *Correlation*, *Anticorrelation*, *Graphical Lasso* (Friedman et al., 2008b) and *DeepGraph* (Belilovsky et al., 2017). We also compare to game-specific baselines: *LinQuadOpt*

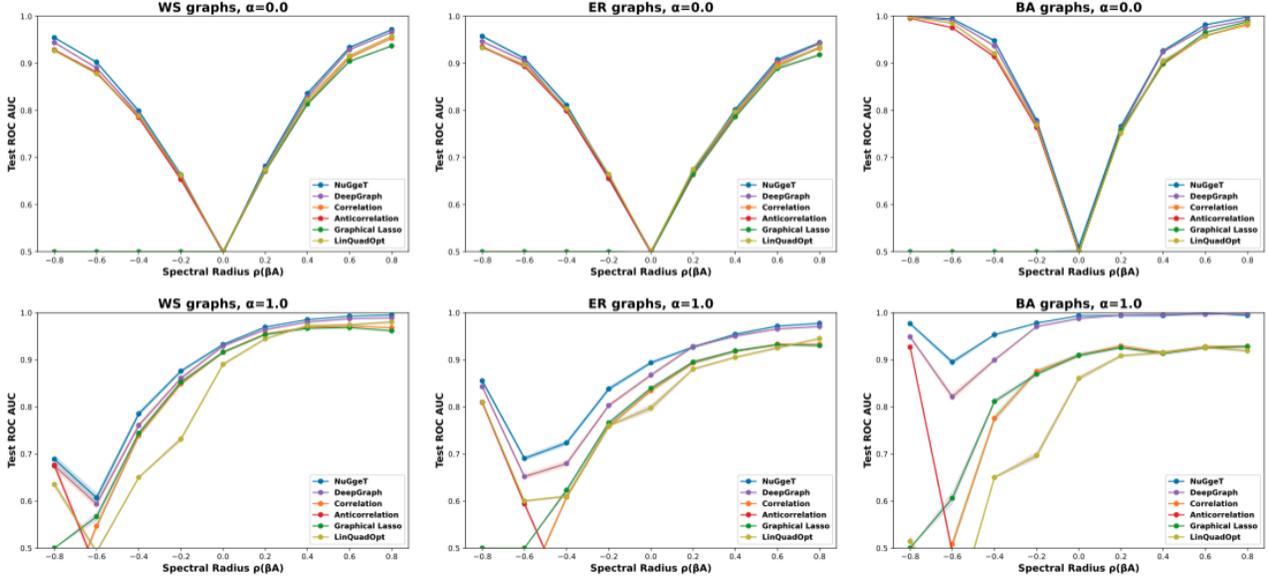


Figure 3. Results for linear quadratic games with varying α and spectral radius $\rho(\beta A)$.

(independent) (Leng et al., 2020) and *LinQuadOpt* (homophilous) (Leng et al., 2020) for linear quadratic games, and *BlockRegression* (Barik & Honorio, 2019) for BH Graphical Games. A more detailed description of the baselines and how they are tuned is provided in the Appendix.

Results Results for Linear Quadratic Games are reported in Fig. 3. Columns are different types of graphs (ER, WS, BA), rows are different values of α that controls smoothness of the marginal benefits ($\alpha = 0$ and $\alpha = 1$, results on $\alpha = 0.5$ are presented in Fig. 6 in Appendix), while the x-axis represents the spectral radius $\rho(\beta A)$. *NuGgeT* is on par or superior to other methods in all scenarios, with *DeepGraph* being the runner-up competitor. We observe larger performance gap with $\alpha = 1$, i.e., when the distribution of marginal benefits largely depends on graph structure. Interestingly, it can be observed that in this case *NuGgeT* and *DeepGraph* perform well in both cases of strategic complements ($\rho(\beta A) > 0$, neighbours take similar actions) and strategic substitutes ($\rho(\beta A) < 0$, neighbours take opposite actions), whereas other baselines only perform well in one of the two cases. This is due to the former two methods learning directly a mapping between actions and graph structure that may correspond to different characteristics of the equilibrium actions.

For Linear Influence Games, we report the results in Fig. 4. Each plot corresponds to a different type of graph, and the x-axis represents the benefit smoothness α . As expected, all methods improve their performance as α grows, since the actions generally become smoother over the graph (Section 3.3). Again, *NuGgeT* outperforms the baselines in all scenarios. Interestingly, the performance on WS and ER

graphs seems to be much lower than for BA graphs. This can be understood empirically by analysing the eigenvalues of the normalised adjacency matrix A for different graphs. For WS and ER graphs, the smallest absolute (non-zero) eigenvalue of A is on average much smaller than for BA graphs (Fig. 7 in Appendix). As explained in section 3.3, eigenvalues with very small absolute values will result in actions behaving like mid-spectrum eigenvectors (which are not necessarily smooth). We verify this further in Fig. 8 in Appendix, which shows the spectral coefficients of the actions for all combinations of graphs and games. Linear Influence Games on ER and WS graphs are indeed the only scenarios where we empirically observe large graph Fourier coefficients for mid-spectrum eigenvectors. On BA graphs we do not observe this behaviour; mid-spectrum eigenvectors are most often not represented at all by the actions (the BA graphs we generate are trees that have zero eigenvalues associated with mid-spectrum eigenvectors which are discarded when taking the pseudoinverse), and actions tend to be smoother in this case. The results in Fig. 4 suggest that when actions are less smooth (weaker association of actions with graph structure) all methods tend to perform less well.

Results for BH Graphical Games are reported in Fig. 5. Since there are no parameters controlling the game, we only have 3 configurations corresponding to different graph types. Again, *NuGgeT* outperforms other methods in all configurations, with the gap being largest on BA graphs. We also perform ablation studies on the number of games, number of training graphs and size of the graphs. The results are presented and discussed in Appendix.

In many real-world scenarios, the observed actions will not be exactly at equilibrium but close to it, or trying to re-

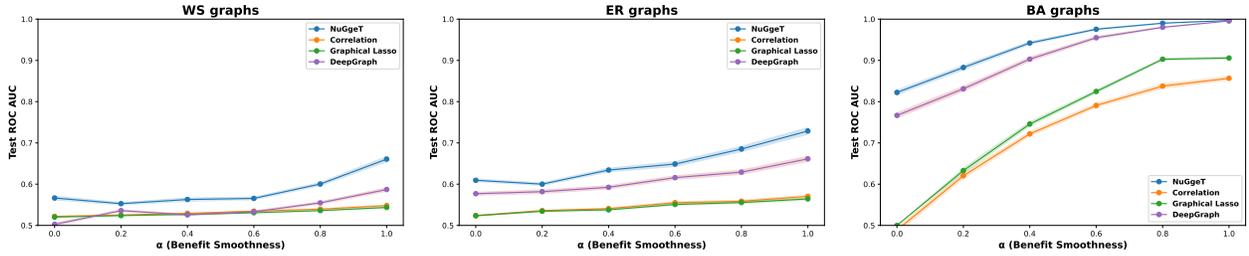


Figure 4. Results on linear influence games when varying the smoothness α of the marginal benefits.

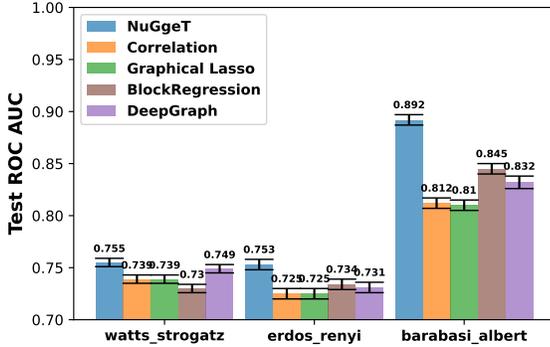


Figure 5. Results for BH graphical games.

Table 2. Test accuracy on Barabási-Albert graphs for NuGgeT given different levels of Gaussian noise added to player actions.

Noise Std	Linear Quadratic	Linear Influence	BH
no noise	99.87 \pm 0.02	99.87 \pm 0.02	99.87 \pm 0.02
0.10	99.33 \pm 0.10	99.33 \pm 0.10	99.33 \pm 0.10
0.20	95.50 \pm 0.38	95.50 \pm 0.38	95.50 \pm 0.38

converge to the equilibrium after some perturbation. We therefore investigate how the performance of our model degrades with “noisy” samples where the sampled actions are only near equilibrium. Table 2 shows the results for BA graphs and signals with unit norm. We can see performance remains satisfactory given reasonable amount of noise.

To conclude, our method outperforms the baselines in all three types of games, most significantly in the scenarios where the actions are not necessarily smooth on the network and our method was able to infer such relationship from the observed data. While learning the mapping from actions to network, the model will learn not only the utility function of the games represented in the data, but also the typical structure of the networks. If all the networks in the training data are Barabási-Albert graphs, the prediction of the model will likely also be a Barabási-Albert graph. On the other hand, optimisation methods such as LinQuadOpt are not able to automatically bias their prediction to a particular class of graphs.

5.2. Real-World Data

Players in real-world scenarios often act according to strategic interactions. It has been shown in the sociology and game theory literature that people have an incentive to conform to social norms (Young, 2009; Montanari & Saberi, 2010) or be influenced by their social network neighbours (Leng et al., 2020). That is, individual utilities are higher if their behaviours are similar to that of their neighbours in the social network. Such a mechanism will lead to strategic complement relationships. Following this assumption, we validate our model on two datasets, comparing with the baseline methods introduced in section 5.1.

Indian Villages The *Indian Villages* dataset⁶ (Banerjee et al., 2013) is a well-known dataset in the economics literature that contains data from a survey of social networks in 75 villages in rural southern Karnataka, a state in India. Each village constitutes a social network graph, where nodes are households and edges are self-reported friendships. Following the setup in Leng et al. (2020), we consider as actions the number of rooms, number of beds and other decisions families have to make related to their household. The reasoning is that if neighbours adopt a specific facility, villagers tend to gain higher payoff by doing the same, i.e., complying with social norms. We only consider the 48 villages for which we have the ground truth actions and network. 40 are used for training, 3 for validation and 5 for testing. Categorical actions are one-hot encoded, while numerical actions are treated as continuous features. The resulting dataset has graphs with 10 actions and a number of nodes ranging between 77 and 356. It can be seen from Table 3 that *NuGgeT* outperforms all other methods by at least 5.01%. *DeepGraph* fails to learn altogether on this dataset.

Yelp Ratings The *Yelp Ratings* dataset⁷ consists of rating of users to business, as well as the social connectivity between user. Similarly to the previous case, on deciding to review a local business, people may have an incentive

⁶The Indian Villages dataset can be accessed at <https://doi.org/10.7910/DVN/U3BIHX>.

⁷The Yelp dataset can be accessed at <https://www.yelp.com/dataset>.

to conform to the social norms they perceive, which are formed by the ratings from their neighbours. Yelpers tend to gain higher payoff with similar ratings due to social conformity, i.e., strategic complements in a game-theoretic context. From the raw data we extract 5000 sub-graphs representing communities, where the actions are the average rating of users to 22 categories of businesses. The task is to reconstruct the social connectivity of the users given their actions (ratings). 4250 graphs are used for training, 250 for validation and 500 for testing. More details of the dataset construction are provided in the appendix. On this dataset, NuGgeT outperforms all other baselines by at least 2.79%. Overall, the results on real-world data show the efficacy of NuGgeT in cases where the game utility is not explicitly known.

Table 3. Test ROC AUC for Indian Villages and Yelp Ratings data.

Model	Indian Villages	Yelp Ratings
<i>Correlation</i>	0.5816±0.0135	0.6222±0.0043
<i>Anticorrelation</i>	0.4184±0.0135	0.3778±0.0043
<i>Graphical Lasso</i>	0.5823±0.0152	0.6523±0.0038
<i>Baraki and Honorio</i>	0.5715±0.0164	0.6786±0.0032
<i>LinQuadOpt (indep.)</i>	0.5557±0.0108	0.6796±0.0033
<i>LinQuadOpt (homop.)</i>	0.5789±0.0174	0.6310±0.0036
<i>DeepGraph</i>	0.4965±0.0143	0.6776±0.0039
<i>NuGgeT</i>	0.6324±0.0167	0.7057±0.0035

6. Conclusion and Future Work

In this work, we propose a novel framework to infer the network structure behind the games from their equilibrium actions. Unlike existing methods, we achieve so by learning a mapping from the actions to the network structure without knowing the utility function of the game. This is especially beneficial in real-world scenarios where the nature of strategic interactions between players of the game remains hidden or may evolve over time.

Limitations and Future Work. The current work only deals with static games and networks of small scale. Moreover, we do not deal with repeated games, i.e., where players have to take multiple actions sequentially. A promising future direction is therefore to extend it to dynamic games and networks, where both the utility function and the structure of the network may change over time, as well as dealing with larger graphs and repeated games. Moreover, we currently deal with test graphs that have a similar structure to the graphs the model has been trained on, e.g., we use the Barabási-Albert graphs during both training and testing. An interesting future work would be to work on the generalisation capability of this approach, i.e., the ability to be trained on one type of graph (e.g., Erdős-Rényi) but generalise to different types as well (e.g., Barabási-Albert or Watts-Strogatz). In addition, our method cannot guarantee the uniqueness of the learned network; indeed, the

main focus of the study is to propose a first and efficient data-driven learning framework without assuming the utility function and prior knowledge about the network structure. We leave the identification of the network structure, which is a challenging problem in itself, for future work.

Acknowledgements

Y.L. is supported by the NSF grant IIS-2153468. X.D. acknowledges support from the Oxford-Man Institute of Quantitative Finance and the EPSRC (EP/T023333/1).

References

- Ballester, C., Calvó-Armengol, A., and Zenou, Y. Who’s who in networks. wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. The Diffusion of Microfinance, 2013. URL <https://doi.org/10.7910/DVN/U3BIHX>.
- Barik, A. and Honorio, J. Provable computational and statistical guarantees for efficient learning of continuous-action graphical games. *arXiv preprint arXiv:1911.04225*, 2019.
- Barik, A. and Honorio, J. Provable sample complexity guarantees for learning of continuous-action graphical games with nonparametric utilities. *CoRR*, abs/2004.01022, 2020. URL <https://arxiv.org/abs/2004.01022>.
- Belilovsky, E., Kastner, K., Varoquaux, G., and Blaschko, M. B. Learning to discover sparse graphical models. In *International Conference on Machine Learning*, pp. 440–448. PMLR, 2017.
- Bramoullé, Y., Kranton, R., and D’amours, M. Strategic interaction and networks. *American Economic Review*, 104(3):898–930, 2014.
- Bronstein, M. M., Bruna, J., Cohen, T., and Velicković, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv:2104.13478*, 2021.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1179. URL <https://aclanthology.org/D14-1179>.

- Cosmo, L., Kazi, A., Ahmadi, S., Navab, N., and Bronstein, M. M. Latent-graph learning for disease prediction. In Martel, A. L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M. A., Zhou, S. K., Racoceanu, D., and Joskowicz, L. (eds.), *Medical Image Computing and Computer Assisted Intervention - MICCAI 2020 - 23rd International Conference, Lima, Peru, October 4-8, 2020, Proceedings, Part II*, volume 12262 of *Lecture Notes in Computer Science*, pp. 643–653. Springer, 2020a. doi: 10.1007/978-3-030-59713-9_62. URL https://doi.org/10.1007/978-3-030-59713-9_62.
- Cosmo, L., Kazi, A., Ahmadi, S., Navab, N., and Bronstein, M. M. Latent patient network learning for automatic diagnosis. *CoRR*, abs/2003.13620, 2020b. URL <https://arxiv.org/abs/2003.13620>.
- Dhillon, I. S., Guan, Y., and Kulis, B. J. Weighted graph cuts without eigenvectors: A multilevel approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 29(0), nov 2007.
- 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.
- Friedman, J., Hastie, T., and Tibshirani, R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008a.
- Friedman, J., Hastie, T., and Tibshirani, R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, July 2008b. ISSN 1465-4644, 1468-4357. doi: 10.1093/biostatistics/kxm045. URL <http://biostatistics.oxfordjournals.org/content/9/3/432>.
- 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. Learning graphical games from behavioral data: Sufficient and necessary conditions. *arXiv preprint arXiv:1703.01218*, 2017a.
- Ghoshal, A. and Honorio, J. Learning sparse polymatrix games in polynomial time and sample complexity. *arXiv preprint arXiv:1706.05648*, 2017b.
- Gomez-Rodriguez, M., Leskovec, J., and Krause, A. Inferring networks of diffusion and influence. In *Proc. of the 16th ACM SIGKDD Inter. Conf. on Knowledge Discovery and Data Mining*, pp. 1019–1028, Washington, DC, USA, 2010.
- Gomez-Rodriguez, M., Balduzzi, D., and Schölkopf, B. Uncovering the temporal dynamics of diffusion networks. In *Proc. of the 28th Inter. Conf. on Machine Learning*, pp. 561–568, Bellevue, Washington, USA, 2011.
- Granovetter, M. Threshold models of collective behavior. *American journal of sociology*, 83(6):1420–1443, 1978.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7e9-Paper.pdf>.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- 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: 1157–1210, 2015.
- Irfan, M. T. and Ortiz, L. E. A game-theoretic approach to influence in networks. 2011.
- Jackson, M. O. and Zenou, Y. Games on networks. *Handbook of Game Theory, vol. 4, Peyton Young and Shmuel Zamir, eds.*, 2014.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. In *International Conference on Machine Learning*, pp. 2688–2697. PMLR, 2018.
- Kipf, T. N. and Welling, M. Variational graph auto-encoders. *NIPS Workshop on Bayesian Deep Learning*, 2016.
- Koller, D. and Friedman, N. *Probabilistic graphical models: Principles and techniques*. MIT Press, 2009.
- Leng, Y., Dong, X., Wu, J., and Pentland, A. Learning quadratic games on networks. In *International conference on Machine Learning*, 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.
- Montanari, A. and Saberi, A. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.
- Valente, T. W. Network interventions. *Science*, 337(6090):49–53, 2012.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- Wang, Y., Sun, Y., Liu, Z., Sarma, S. E., Bronstein, M. M., and Solomon, J. M. Dynamic graph cnn for learning on point clouds. *ACM Trans. Graph.*, 38(5), October 2019. ISSN 0730-0301. doi: 10.1145/3326362. URL <https://doi.org/10.1145/3326362>.
- Young, H. P. Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American economic review*, 99(5):1899–1924, 2009.
- Yu, Y., Chen, J., Gao, T., and Yu, M. Dag-gnn: Dag structure learning with graph neural networks. *arXiv preprint arXiv:1904.10098*, 2019.
- Zhang, M. and Chen, Y. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, pp. 5165–5175, 2018.
- Zheng, X., Dan, C., Aragam, B., Ravikumar, P., and Xing, E. Learning sparse nonparametric dags. In *International Conference on Artificial Intelligence and Statistics*, pp. 3414–3425. PMLR, 2020.

A. Proof of NuGgeT symmetries

We would like to show that the *NuGgeT* model g is *permutation-invariant* over the set of games and *permutation-equivariant* over the set of nodes.

Starting with the first one, we need to show that $g(\mathbf{X}\mathbf{P}_2) = g(\mathbf{X})$, where \mathbf{P}_2 is a $K \times K$ permutation matrix on the games. For what concerns the encoder, we only need to prove that $\alpha_{ij}^{(h)}$ is invariant w.r.t. permutation of the games. Let p_k be the index of the non-zero entry of the k -th column of \mathbf{P}_2 , we have:

$$\begin{aligned}\alpha_{ij}^{(h)} &= \prod_{k=1}^K \mathbf{y}_{ik}^\top \mathbf{W}_Q^{(h)} \mathbf{W}_K^{(h)} \mathbf{y}_{jk} \\ &= \prod_{k=1}^K \mathbf{y}_{ip_k}^\top \mathbf{W}_Q^{(h)} \mathbf{W}_K^{(h)} \mathbf{y}_{jp_k},\end{aligned}$$

since $\prod_{k=1}^K$ is chosen to be a permutation invariant operator. This shows that $\alpha_{ij}^{(h)}$ does not depend on the particular permutation p chosen, i.e., it is invariant w.r.t. permutation of the games. Similarly, the same holds for the decoder as $\prod_{k=1}^K$ is a permutation invariant operator and therefore $\prod_{k=1}^K \mathbf{z}_{ik} \odot \mathbf{z}_{jk} = \prod_{k=1}^K \mathbf{z}_{ip_k} \odot \mathbf{z}_{jp_k}$ for any permutation p .

Regarding the second part of the proof, we need to show that $g(\mathbf{P}_1\mathbf{X}) = \mathbf{P}_1g(\mathbf{X})\mathbf{P}_1^\top$, where \mathbf{P}_1 is an $N \times N$ permutation matrix interpreted as reordering of the nodes of the graph. Letting p_k be the index of the non-zero entry of the k -th row of \mathbf{P}_1 . Note that:

$$\begin{aligned}\text{enc}(\mathbf{P}_1\mathbf{X})_{ik} &= \psi\left(\mathbf{y}_{p_ik}, \sum_{j=1}^N \alpha_{p_ip_j}^{(1)} \mathbf{y}_{p_jk}, \dots, \sum_{j=1}^N \alpha_{p_ip_j}^{(H)} \mathbf{y}_{p_jk}\right) \\ &= \psi\left(\mathbf{y}_{p_ik}, \sum_{j=1}^N \alpha_{p_ip_j}^{(1)} \mathbf{y}_{jk}, \dots, \sum_{j=1}^N \alpha_{p_ip_j}^{(H)} \mathbf{y}_{jk}\right) \\ &= \mathbf{z}_{p_ik}\end{aligned}$$

where the second equality stems from the fact that summation is a permutation invariant operator (i.e., $\sum_{j=1}^N x_i = \sum_{j=1}^N x_{p_i}$ for any permutation p). We then have:

$$\begin{aligned}\left(\mathbf{P}_1 g(\mathbf{X}) \mathbf{P}_1^\top\right)_{ij} &= \hat{a}_{p_i p_j} \\ &= \psi\left(\prod_{k=1}^K \mathbf{z}_{p_ik} \odot \mathbf{z}_{p_jk}\right) \\ &= \psi\left(\prod_{k=1}^K \text{enc}(\mathbf{P}_1\mathbf{X})_{ik} \odot \text{enc}(\mathbf{P}_1\mathbf{X})_{jk}\right) \\ &= g(\mathbf{P}_1\mathbf{X})_{ij}.\end{aligned}$$

Since this holds for all indices ij , it follows that $g(\mathbf{P}_1\mathbf{X}) = \mathbf{P}_1g(\mathbf{X})\mathbf{P}_1^\top$.

B. Experiments

B.1. Synthetic Data Generation

We generate the synthetic data from three different graph models: Erdős-Rényi (ER), Barabási-Albert (BA) and Watts-Strogatz (WS). In ER graphs, an edge is present with a probability of $p = 0.2$, independently from all other possible edges.

In WS graphs, we set the exact degree of the nodes to be $k = \log_2(N)$, with a probability of $p = 0.2$ for the random rewiring process. Finally, in BA graphs, nodes are added one at a time and each new node has $m = 1$ edges which are preferentially attached to existing nodes with already high degree (this results in a tree graph).

B.2. Yelp Dataset Generation

Starting from the raw data at <https://www.yelp.com/dataset>, we create a dataset by performing the following steps:

1. We compute the rating for every user to every businesses category, by averaging the ratings a user has given to all business of each category
2. We weight each edge in the original user-user graph with the fraction of common categories the two users rated at least once
3. We cluster the above weighted graph using Graclus (Dhillon et al., 2007) with the objective of minimising the normalised cut.
4. We discard all the clusters with less than 10 nodes, or with very sparse ratings (less than 25% of the categories with a rating for at least 25% of the users). The result of this step is $\sim 27k$ graphs of different users.
5. We rank the clusters extracted above by their density of ratings and keep the 5000 most dense clusters. Each cluster constitutes a graph associated with node attributes, which can be used to train and test our model.

B.3. Baselines

In both the synthetic and real-world data experiments we compare with the following baselines:

Correlation The Pearson correlation coefficient between the actions of two nodes. This works particularly well when the actions are homophilous over the graph, i.e., nodes connected in the graph tend to take similar actions. We implement this baseline ourselves.

Anticorrelation The negative of the Pearson correlation coefficient between the actions of two nodes. This works well in case of strategic substitutes, i.e., when nodes connected in the graph tend to take different actions. We implement this baseline ourselves.

Graphical Lasso (Friedman et al., 2008b) Computes a sparse penalised estimation of the inverse of the covariance matrix. We use SKGGM⁸ QuicGraphicalLasso with empirical covariance initialisation for this.

LinQuadOpt Algorithm presented in (Leng et al., 2020) which assumes the form of the game to be linear quadratic. It has two version, one where the benefits are assumed to be independent (*LinQuadOpt (Independent)*), and another for homophilous benefits (*LinQuadOpt (Homophilous)*). We implement the algorithm ourselves (no public code was provided).

BlockRegression Algorithm presented in (Barik & Honorio, 2019), which has been designed specifically for BH graphical games. We re-implement this baseline ourselves following the paper (no public code was provided).

DeepGraph Algorithm presented in (Belilovsky et al., 2017) which recovers the graph from the covariance matrix using a series of dilated convolutions. We re-implement *DeepGraph* in PyTorch ourselves taking as inspiration the public TensorFlow implementation of the authors.

All the baselines are tuned on the validation set. For both *Graphical Lasso* and *BlockRegression* we tune the regularisation parameter in the range $10^k : k \in [-5, 5]$ with an interval of one, whereas for *LinQuadOpt* both regularisation parameters are tuned in $10^k : k \in [-6, 1]$, also with an interval of one. We train *DeepGraph* using the same hyperparameters used for NuGget (Adam optimiser, learning rate of 0.001, batch size of 100 and a patience of 50), with $\lceil \log_2(N_{\max}) \rceil$ convolutional layers (N_{\max} corresponds here to the maximum number of nodes of any graph in the dataset) and dilation coefficient equal to $d_k = 2^{k-1}$ for layer k as specified in (Belilovsky et al., 2017).

⁸<https://github.com/skggm/skggm>

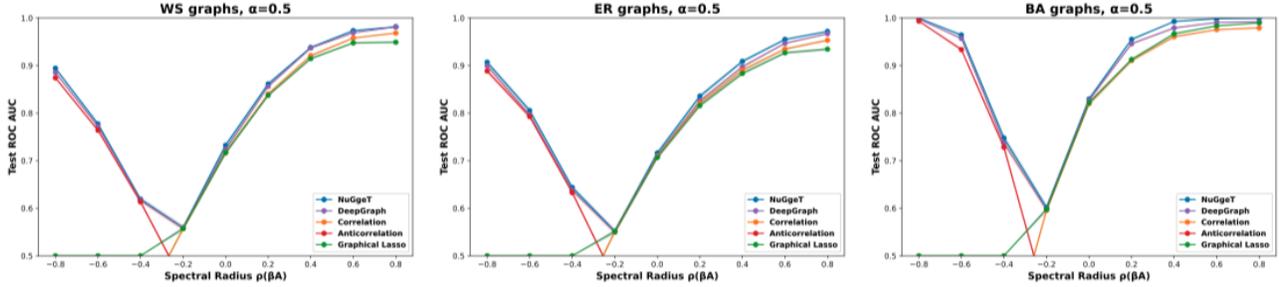


Figure 6. Results for linear quadratic games with $\alpha = 0.5$ and variable spectral radius $\rho(\beta A)$.

B.4. NuGgeT Hyperparameters

We do not perform any extensive hyperparameter tuning for the NuGgeT model, but instead use the same standard choice of hyperparameters (reported in Table 4) for all experiments.

Hyperparameter Name	Value
F	10
F'	10
H (num of heads)	10
ψ 's num layers	2
ψ 's hidden dim	100

Table 4. Hyperparameters used for *NuGgeT* in all experiments.

B.5. Additional Results

Additional results on the linear quadratic and BH games are presented in Fig. 6 and Fig. 5, respectively.

C. Spectral Analysis of the Games

Fig. 7 shows that the smallest absolute (non-zero) eigenvalues of the normalised adjacency matrix \mathbf{A} is on average much smaller for ER and WS graphs than for BA graphs. This results in the corresponding eigenvectors having a large influence on the actions for the Linear Influence games, where the Nash equilibrium actions satisfy $\mathbf{x}^* = \mathbf{A}^{-1}\mathbf{b}$. This is confirmed by Fig. 8, which shows the graph Fourier coefficients for different eigenvalue indexes: the actions of linear influence games on ER and WS graphs are dominated by mid-range eigenvectors.

D. Ablation Studies

In Fig. 9, we investigate how the performance changes with the number of games available. The smaller the number of games, the less information to reconstruct the graph from. We use $\alpha = 1$ for both linear quadratic and linear influence games, and a spectral radius of 0.6 for linear quadratic games. In line with our expectations, all methods generally improve as more games are available. We analyse the effect of larger graph sizes (Fig. 10). The more nodes, the more edge combinations exist and the harder the task becomes, which explains the decrease in performance of both methods as the number of nodes increases. Interestingly, the magnitude of the drop depends heavily on the combination of game and graph types, but *NugGeT* seems to be more robust than *DeepGraph*. We also investigate the effect of the number on training graphs on the model performance (Fig. 11). *NugGeT* requires less training graphs compared to *DeepGraph* to obtain a similar performance. Fig. 9 and 10 show respectively the performance of various methods when varying the number of games and the number of nodes respectively. We use $\alpha = 1$ for both linear quadratic and linear influence games, and $\beta = 0.6$ (strategic complements) for linear quadratic games. In line with our expectations, all methods generally improve as more games are available. On the other hand, the more nodes, the more edge combinations exist and the harder the task becomes, which explains the decrease in performance of all methods as the number of nodes increases. Interestingly, the magnitude of the

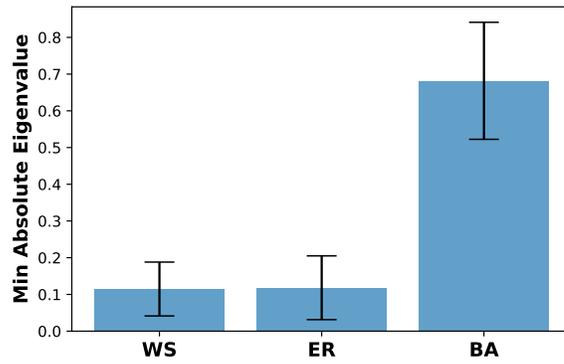


Figure 7. Mean and standard deviation of the minimum absolute non-zero eigenvalue of the adjacency matrix for different types of graph models. Statistics are computed over 1000 graphs with 20 nodes.

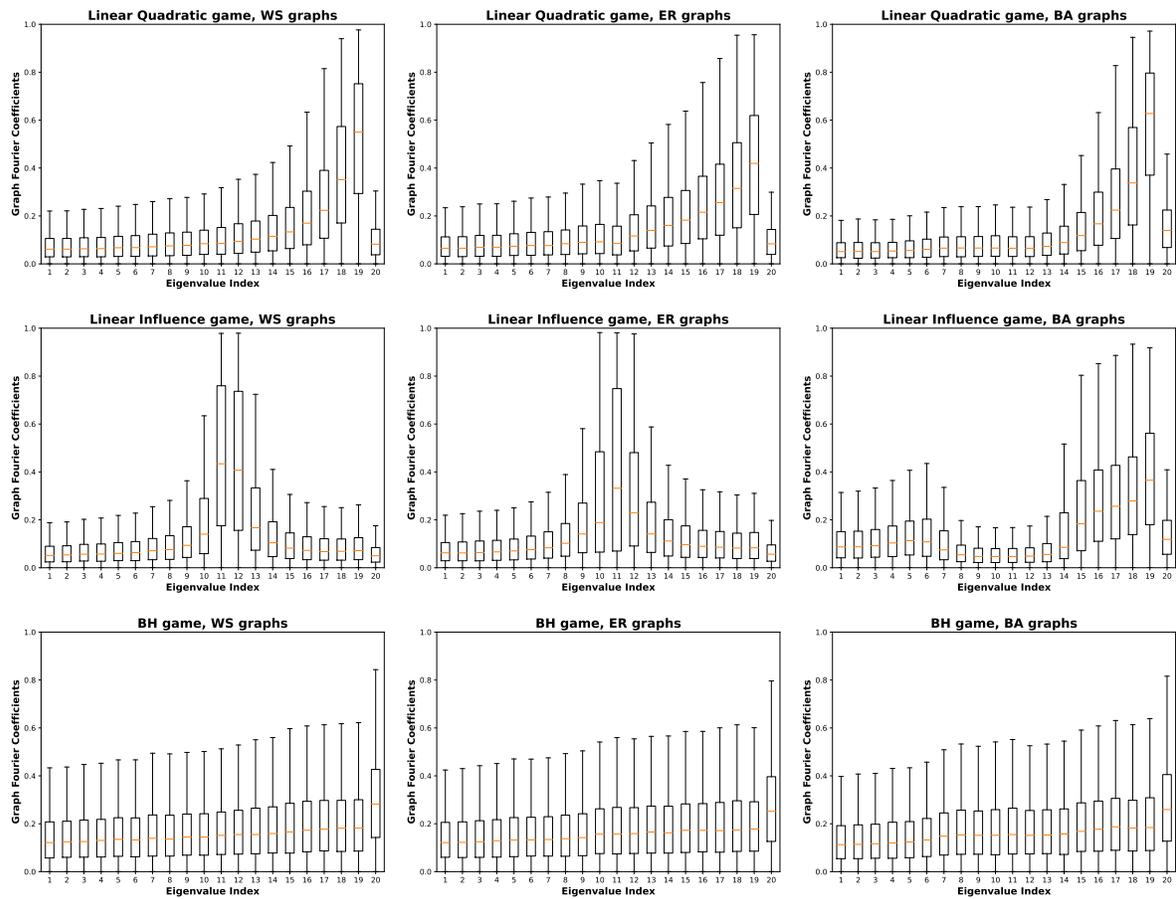


Figure 8. Magnitude of coefficients of normalised action vectors for different games and graphs. These are obtained by taking the Graph Fourier Transform of the actions, i.e., taking the inner product between the action and each of the adjacency matrix eigenvectors.

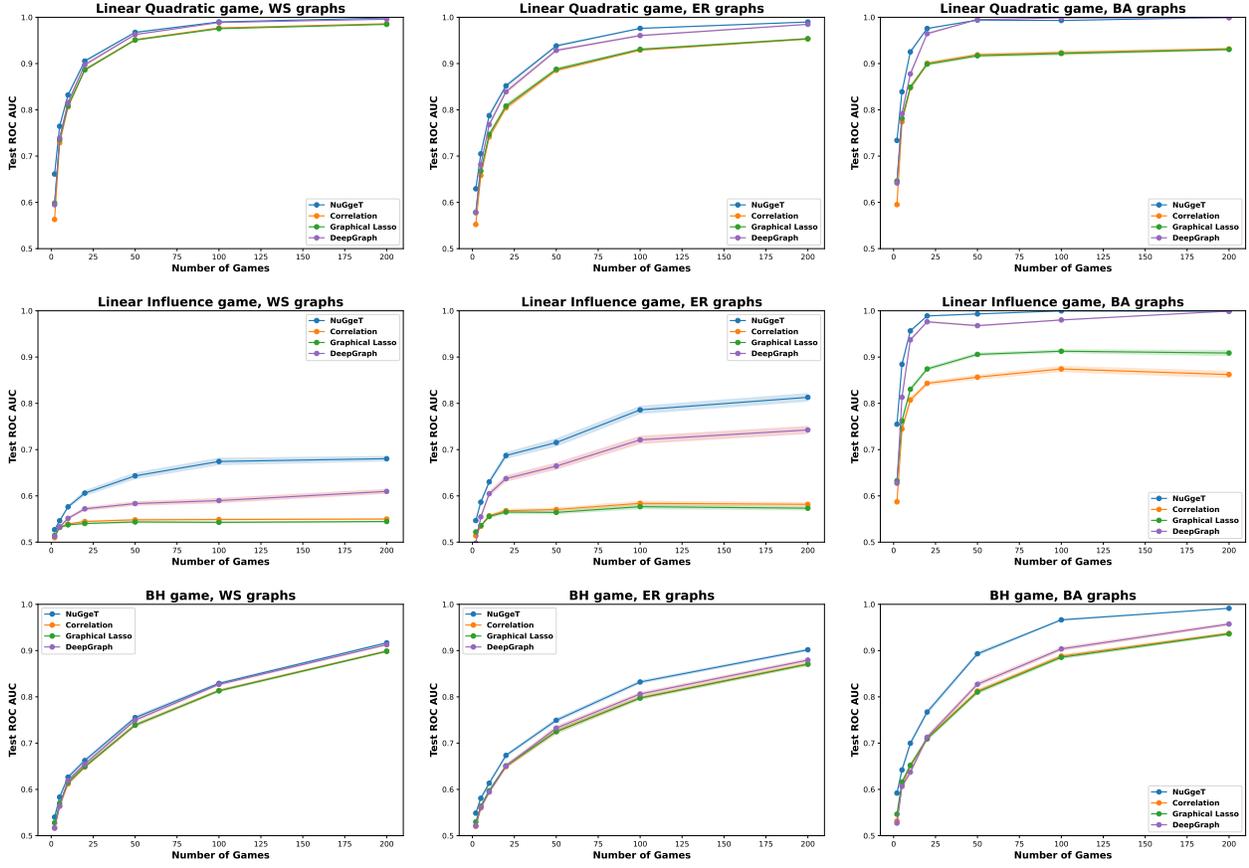


Figure 9. Results with varying number of games.

drop depends heavily on the combination of game and graph types, but *NuGgeT* seems to be more robust than *DeepGraph*.

E. *NuGgeT*'s Attention Mechanism

We investigate the difference between a multiplicative attention mechanism operating directly on scalar actions (i.e., with unnormalised attention coefficients of the form: $\tilde{\alpha}_{ij} = x_i x_j a$ with a a learnable parameter), and one operating on transformed actions (i.e., $\tilde{\alpha}_{ij} = \mathbf{y}_i^\top \mathbf{W}_Q \mathbf{W}_K \mathbf{y}_j$, with $\mathbf{y}_i = \text{ReLU}(x_i \mathbf{w} + \mathbf{b})$). In Fig. 12 we plot the non-normalised attention values of the two different mechanisms for varying values of the input actions x_i, x_j . We use a positive and a negative value of a for showing the capability of the simple scalar-based attention mechanism to identify correlation and anti-correlation in the input actions x and random coefficients drawn from a normal distribution for $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{w}, \mathbf{b}$ for the attention mechanism used in *NuGgeT*. While the scalar attention always shows a linear trend for a fixed x_i (or x_j), the multiplicative layer operating on vectors is able to produce richer non-linear attention patterns with maximum / minimum values for the interior points of $[-1, 1]$.

You can have as much text here as you want. The main body must be at most 8 pages long. For the final version, one more page can be added. If you want, you can use an appendix like this one, even using the one-column format.

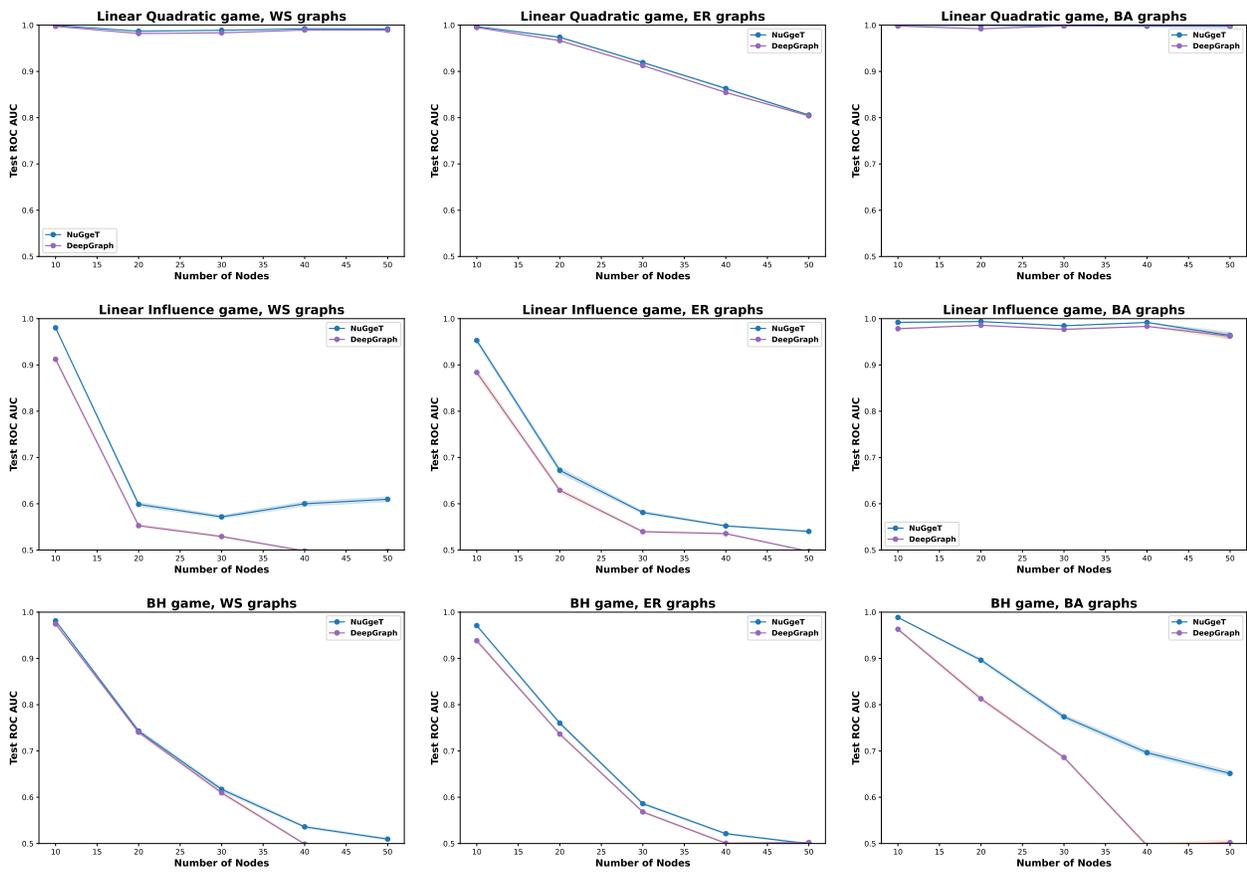


Figure 10. Results with varying number of nodes.

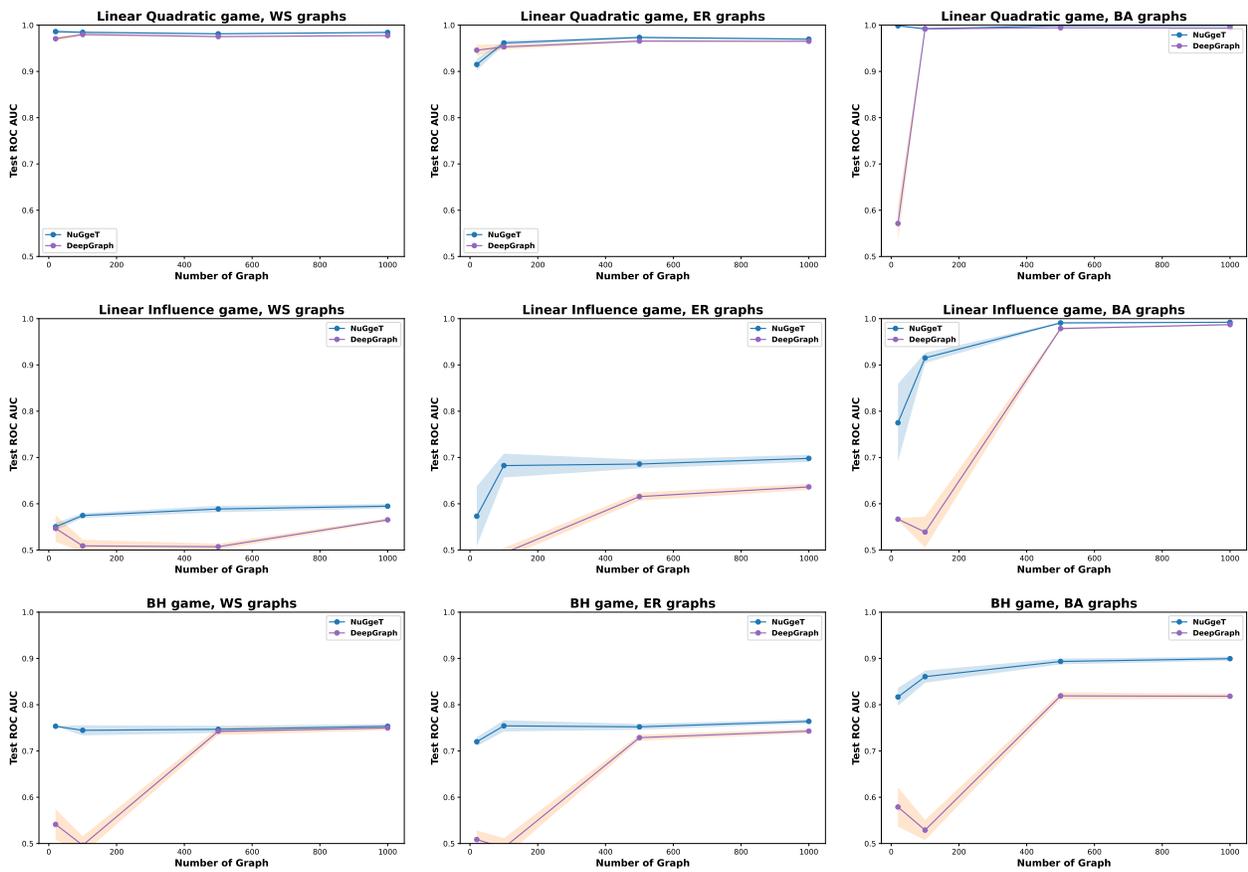


Figure 11. Results with varying number of training graphs.

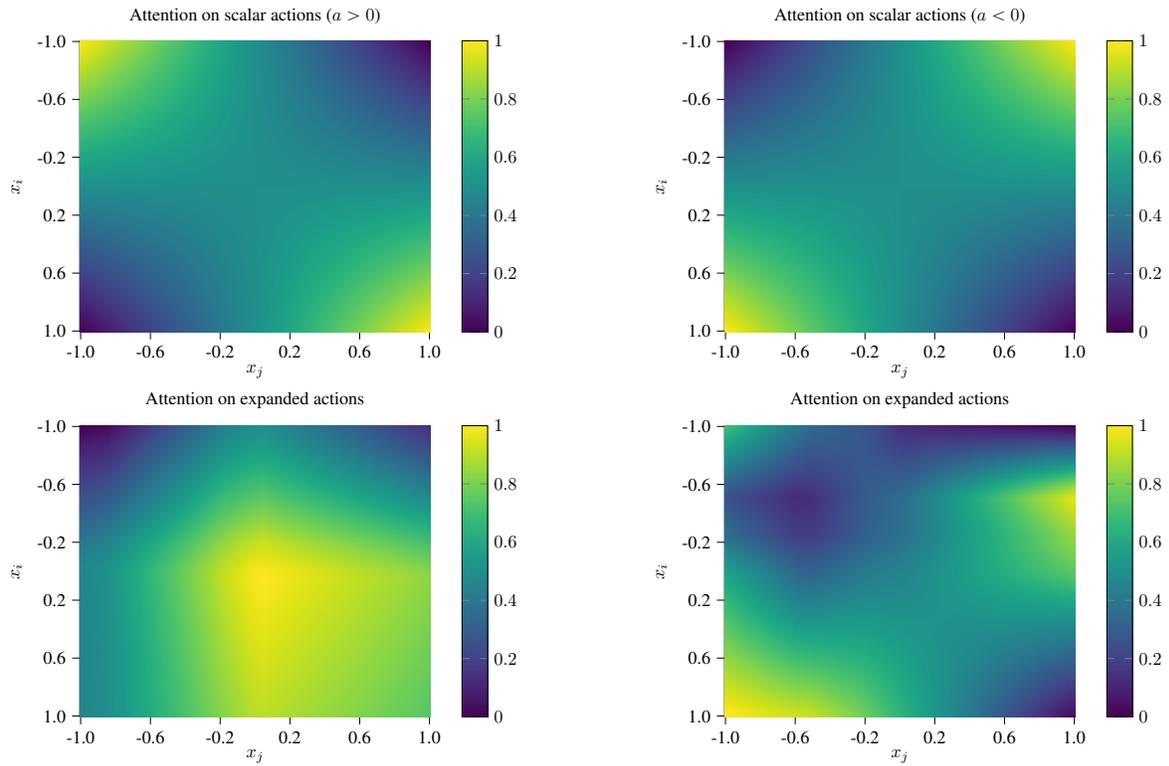


Figure 12. Top, attention values for a multiplicative attention layer directly operating on the input scalars $x_i \in [-1, 1]$, $x_j \in [-1, 1]$ with positive (left) and negative (right) coefficient. Bottom, attention values for a multiplicative attention layer with random coefficients processing the expanded actions y_i, y_j of x_i, x_j . The result is scaled to $[0, 1]$ for visualisation purposes.