Learning Quadratic Games on Networks

Xiaowen Dong
Department of Engineering Science
University of Oxford
(joint work with Yan Leng, Junfeng Wu, and Alex Pentland)

University of Bath, July 2020

Motivation

- Consider a group of students making choices on educational effort
  - making effort is costly
  - I will benefit from my own effort
  - I will also benefit from my friends' effort

Motivation

- Consider a group of students making choices on educational effort
  - making effort is costly
  - I will benefit from my own effort
  - I will also benefit from my friends' effort

➡ **tend to make effort if friends do**

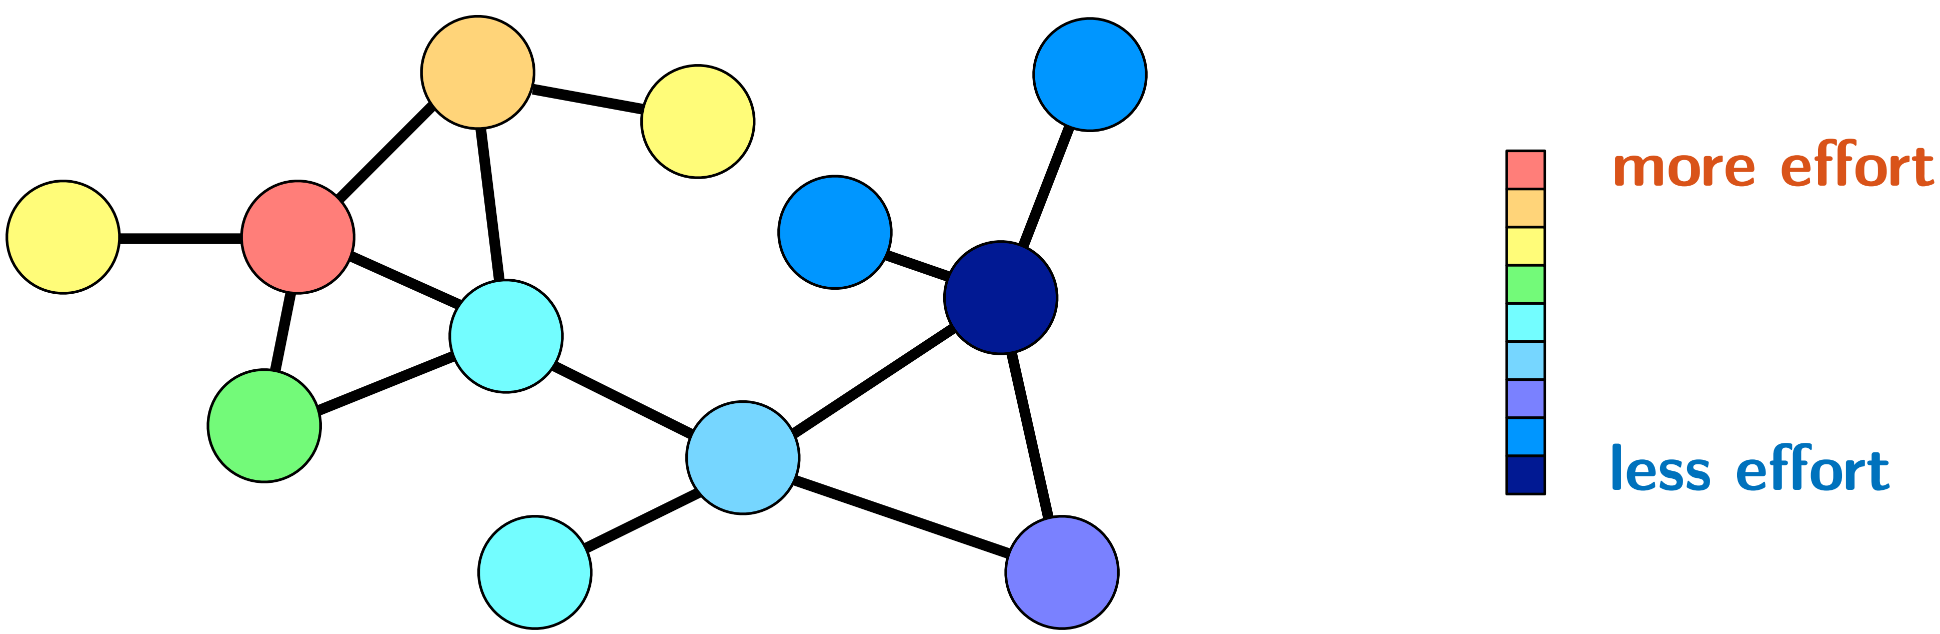

Motivation

- Consider a group of students making choices on buying a book
  - buying a book is costly
  - if a friend of mine will buy, then I will not buy
  - but if none of my friends will buy, then I will buy

Motivation

- Consider a group of students making choices on buying a book
  - buying a book is costly
  - if a friend of mine will buy, then I will not buy
  - but if none of my friends will buy, then I will buy

→ **tend not to make effort if friends do**

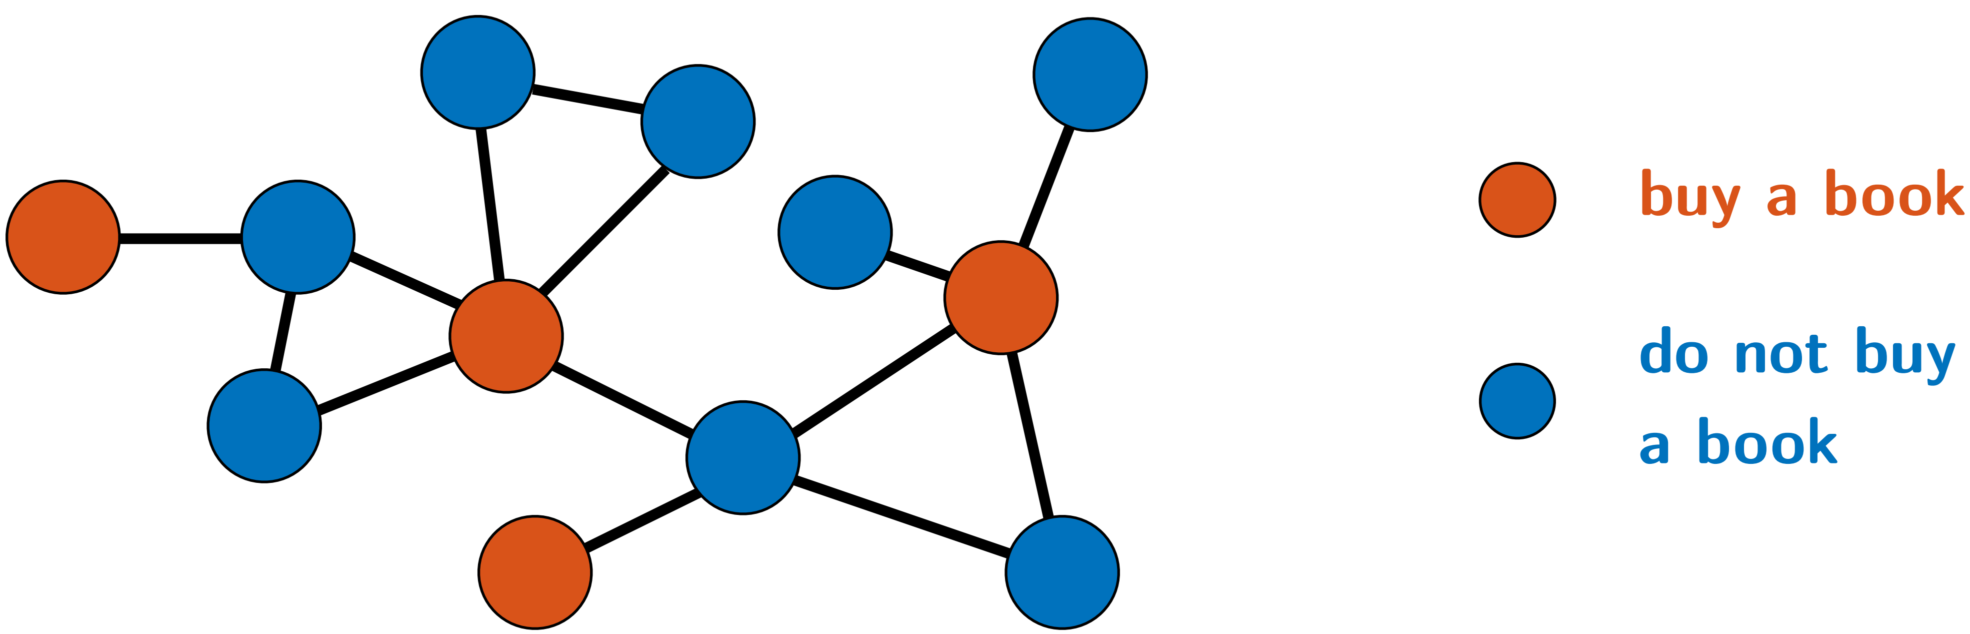

Motivation

- Such strategic interactions can be modelled as games on networks
  - players, actions, payoffs, **interaction network**

Motivation

- Such strategic interactions can be modelled as games on networks
  - players, actions, payoffs, **interaction network**
  - payoff of an individual depends on **her action** as well as **her neighbours' actions**

Motivation

- Such strategic interactions can be modelled as games on networks
  - players, actions, payoffs, **interaction network**
  - payoff of an individual depends on **her action** as well as **her neighbours' actions**
  - strategic **complements** or strategic **substitutes**

Motivation

- Such strategic interactions can be modelled as games on networks
  - players, actions, payoffs, **interaction network**
  - payoff of an individual depends on **her action** as well as **her neighbours' actions**
  - strategic **complements** or strategic **substitutes**

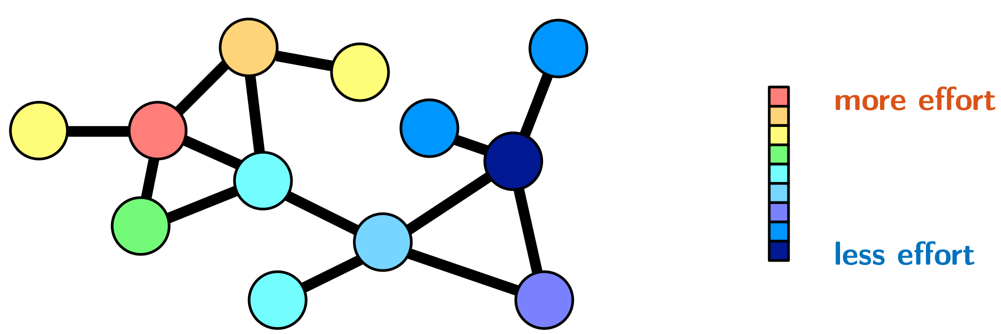

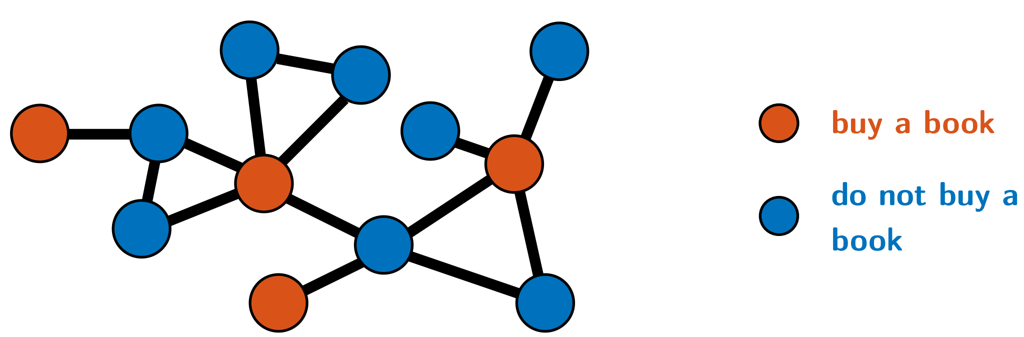

Motivation

- Economics
  - existence of equilibrium or how action/payoff depends on network structure
  - on a **given or predefined** network

Motivation

- Economics
  - existence of equilibrium or how action/payoff depends on network structure
  - on a **given or predefined** network
- Computer science (graphical games)
  - algorithms for computing equilibrium
  - **binary or finite discrete** action space

Motivation

- Economics
  - existence of equilibrium or how action/payoff depends on network structure
  - on a **given or predefined** network
- Computer science (graphical games)
  - algorithms for computing equilibrium
  - **binary or finite discrete** action space
- This work
  - **learning network** given **continuous actions**

Motivation

- Economics
  - existence of equilibrium or how action/payoff depends on network structure
  - on a **given or predefined** network
- Computer science (graphical games)
  - algorithms for computing equilibrium
  - **binary or finite discrete** action space
- This work
  - **learning network** given **continuous actions**

- Many examples
  - observe individual decisions (e.g., adoptions), but not social relationship
  - observe R&D activities of firms, but not collaboration networks
  - observe international policies of countries, but not political alliance

Outline

- Background
  - learning network structure from data
  - network games with linear-quadratic payoffs
- Learning games with linear-quadratic payoffs
  - independent marginal benefits
  - homophilous marginal benefits
- Experimental results
- Discussion

Outline

- Background
  - learning network structure from data
  - network games with linear-quadratic payoffs
- Learning games with linear-quadratic payoffs
  - independent marginal benefits
  - homophilous marginal benefits
- Experimental results
- Discussion

Learning network structure from data

- Different perspectives in the literature
  - statistical: graph captures **data distribution** (e.g., probabilistic graphical model)
  - physics: data correspond to **physical process** on graph (e.g., network cascade)
  - signal processing: graph enforces **signal property** (e.g., smoothness)

Dong et al., "Learning graphs from data: A signal representation perspective," IEEE Signal Processing Magazine, 2019.
Mateos et al., "Connecting the dots: Identifying network structure via graph signal processing," IEEE Signal Processing Magazine, 2019.

Learning network structure from data

- Different perspectives in the literature
  - statistical: graph captures **data distribution** (e.g., probabilistic graphical model)
  - physics: data correspond to **physical process** on graph (e.g., network cascade)
  - signal processing: graph enforces **signal property** (e.g., smoothness)
- **No game-theoretic aspect** of strategic interactions

Dong et al., "Learning graphs from data: A signal representation perspective," IEEE Signal Processing Magazine, 2019.
Mateos et al., "Connecting the dots: Identifying network structure via graph signal processing," IEEE Signal Processing Magazine, 2019.

Network games with linear-quadratic payoffs

- Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with edge weights G_{ij}

Network games with linear-quadratic payoffs

- Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with edge weights G_{ij}
- Payoff of player i: $u_i \left(a_i, \{a_j\}_{j \in \mathcal{N}(i)}, G_{ij}\right)$

Network games with linear-quadratic payoffs

- Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with edge weights G_{ij}
- Payoff of player i: $u_i\left(a_i, \{a_j\}_{j \in \mathcal{N}(i)}, G_{ij}\right)$
- Games with linear-quadratic payoffs

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

Network games with linear-quadratic payoffs

- Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with edge weights G_{ij}
- Payoff of player i: $u_i\left(a_i, \{a_j\}_{j\in\mathcal{N}(i)}, G_{ij}\right)$
- Games with linear-quadratic payoffs

individual action

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

Network games with linear-quadratic payoffs

- Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with edge weights G_{ij}
- Payoff of player i: $u_i\left(a_i, \{a_j\}_{j\in\mathcal{N}(i)}, G_{ij}\right)$
- Games with linear-quadratic payoffs

marginal benefit **individual action**

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

Network games with linear-quadratic payoffs

- Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with edge weights G_{ij}
- Payoff of player i: $u_i\left(a_i, \{a_j\}_{j \in \mathcal{N}(i)}, G_{ij}\right)$
- Games with linear-quadratic payoffs

marginal benefit **individual action** **network factor**

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

Network games with linear-quadratic payoffs

- Consider a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ with edge weights G_{ij}
- Payoff of player i: $u_i\left(a_i, \{a_j\}_{j \in \mathcal{N}(i)}, G_{ij}\right)$
- Games with linear-quadratic payoffs

marginal benefit **individual action** **network factor**

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

individual effect **network effect**

Network games with linear-quadratic payoffs

marginal benefit | individual action | network factor

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

individual effect | network effect

remarks

- continuous actions
- for strategic complements ($\beta > 0$) and substitutes ($\beta < 0$)
- can be used to approximate complex non-linear payoffs
- widely adopted in literature [Jackson15,Bramoullé16]

Jackson and Zezou, "Games on networks," Handbook of Game Theory with Economic Applications, 2015.
Bramoullé and Kranton, "Games played on networks," The Oxford Handbook of the Economics of Networks, 2016.

Network games with linear-quadratic payoffs

marginal benefit — individual action — network factor

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

individual effect — network effect

remarks

- continuous actions
- for strategic complements ($\beta > 0$) and substitutes ($\beta < 0$)
- can be used to approximate complex non-linear payoffs
- widely adopted in literature [Jackson15,Bramoullé16]

examples

- education: action is educational effort, utility is achievement
- collaboration: action is joint R&D activities, utility is firm profit
- urban dynamics: action is mobility behaviour, utility is convenience/satisfaction

Jackson and Zezou, "Games on networks," Handbook of Game Theory with Economic Applications, 2015.
Bramoullé and Kranton, "Games played on networks," The Oxford Handbook of the Economics of Networks, 2016.

Network games with linear-quadratic payoffs

- Pure-strategy Nash equilibrium (PSNE)

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

Network games with linear-quadratic payoffs

- Pure-strategy Nash equilibrium (PSNE)

$$u_i = b_i a_i - \frac{1}{2} a_i^2 + \beta a_i \sum_{j \in \mathcal{V}} G_{ij} a_j \quad \Rightarrow \quad \frac{\partial u_i}{\partial a_i} = b_i - a_i + \beta(\mathbf{G}\mathbf{a})_i$$

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

Network games with linear-quadratic payoffs

- Pure-strategy Nash equilibrium (PSNE)

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

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

assumption: spectral radius $\rho(\beta \mathbf{G})$ is smaller than 1

- guarantees matrix inversion
- ensures uniqueness and stability of equilibrium action [Ballester06]

Ballester et al., "Who's who in networks. Wanted: The key player," Econometrica, 2006.

Network games with linear-quadratic payoffs

- Pure-strategy Nash equilibrium (PSNE)

$$u_i = b_i a_i - \frac{1}{2} a_i^2 + \beta a_i \sum_{j \in \mathcal{V}} G_{ij} a_j \quad \Rightarrow \quad \frac{\partial u_i}{\partial a_i} = b_i - a_i + \beta (\mathbf{G}\mathbf{a})_i$$

$$\Rightarrow \quad \mathbf{a} = (\mathbf{I} - \beta \mathbf{G})^{-1} \mathbf{b} = \sum_{p=0}^{+\infty} \beta^p \mathbf{G}^p \ \mathbf{b}$$

assumption: spectral radius $\rho(\beta \mathbf{G})$ is smaller than 1

- guarantees matrix inversion
- ensures uniqueness and stability of equilibrium action [Ballester06]

properties

- equilibrium related to Katz-Bonacich centrality
- payoff interdependency spreads indirectly through network

Ballester et al., "Who's who in networks. Wanted: The key player," Econometrica, 2006.

Network games with linear-quadratic payoffs

- Pure-strategy Nash equilibrium (PSNE)

$$u_i = b_i a_i - \frac{1}{2} a_i^2 + \beta a_i \sum_{j \in \mathcal{V}} G_{ij} a_j \quad \Rightarrow \quad \frac{\partial u_i}{\partial a_i} = b_i - a_i + \beta (\mathbf{G}\mathbf{a})_i$$

$$\Rightarrow \quad \mathbf{a} = (\mathbf{I} - \beta \mathbf{G})^{-1} \mathbf{b} \xlongequal{\mathbf{G} = \boldsymbol{\chi} \boldsymbol{\Lambda} \boldsymbol{\chi}^T} \boldsymbol{\chi} (\mathbf{I} - \beta \boldsymbol{\Lambda})^{-1} \boldsymbol{\chi}^T \mathbf{b}$$

assumption: spectral radius $\rho(\beta \mathbf{G})$ is smaller than 1

- guarantees matrix inversion
- ensures uniqueness and stability of equilibrium action [Ballester06]

properties

- equilibrium related to Katz-Bonacich centrality
- payoff interdependency spreads indirectly through network
- action is filtered version of marginal benefit on graph

Ballester et al., "Who's who in networks. Wanted: The key player," Econometrica, 2006.

Outline

- Background
  - learning network structure from data
  - network games with linear-quadratic payoffs
- Learning games with linear-quadratic payoffs
  - independent marginal benefits
  - homophilous marginal benefits
- Experimental results
- Discussion

Learning with independent marginal benefits

Nash equilibrium $\mathbf{a} = (\mathbf{I} - \beta\mathbf{G})^{-1}\mathbf{b}$ ⇒ $(\mathbf{I} - \beta\mathbf{G})\,\mathbf{a} = \mathbf{b}$

Learning with independent marginal benefits

Nash equilibrium $\mathbf{a} = (\mathbf{I} - \beta\mathbf{G})^{-1}\,\mathbf{b} \Rightarrow (\mathbf{I} - \beta\mathbf{G})\,\mathbf{a} = \mathbf{b}$

consider K games $\mathbf{B} = [\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \cdots, \mathbf{b}^{(k)}] \in \mathbb{R}^{N \times K}$ marginal

$\mathbf{A} = [\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \cdots, \mathbf{a}^{(K)}] \in \mathbb{R}^{N \times K}$ action

Learning with independent marginal benefits

Nash equilibrium $\mathbf{a} = (\mathbf{I} - \beta\mathbf{G})^{-1}\mathbf{b}$ ⇒ $(\mathbf{I} - \beta\mathbf{G})\mathbf{a} = \mathbf{b}$

consider K games $\mathbf{B} = [\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \cdots, \mathbf{b}^{(k)}] \in \mathbb{R}^{N \times K}$ **marginal**

$\mathbf{A} = [\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \cdots, \mathbf{a}^{(K)}] \in \mathbb{R}^{N \times K}$ **action**

joint learning

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

Learning with independent marginal benefits

Nash equilibrium $\mathbf{a} = (\mathbf{I} - \beta\mathbf{G})^{-1}\mathbf{b}$ ➡ $(\mathbf{I} - \beta\mathbf{G})\mathbf{a} = \mathbf{b}$

consider K games $\mathbf{B} = [\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \cdots, \mathbf{b}^{(k)}] \in \mathbb{R}^{N \times K}$ **marginal**

$\mathbf{A} = [\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \cdots, \mathbf{a}^{(K)}] \in \mathbb{R}^{N \times K}$ **action**

joint learning

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

Learning with independent marginal benefits

Nash equilibrium $\mathbf{a} = (\mathbf{I} - \beta \mathbf{G})^{-1} \mathbf{b} \Rightarrow (\mathbf{I} - \beta \mathbf{G}) \mathbf{a} = \mathbf{b}$

consider K games $\mathbf{B} = [\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \cdots, \mathbf{b}^{(k)}] \in \mathbb{R}^{N \times K}$ **marginal**

$\mathbf{A} = [\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \cdots, \mathbf{a}^{(K)}] \in \mathbb{R}^{N \times K}$ **action**

joint learning

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

Learning with independent marginal benefits

Nash equilibrium $\mathbf{a} = (\mathbf{I} - \beta\mathbf{G})^{-1}\,\mathbf{b}$ ➡ $(\mathbf{I} - \beta\mathbf{G})\,\mathbf{a} = \mathbf{b}$

consider K games $\mathbf{B} = [\mathbf{b}^{(1)}, \mathbf{b}^{(2)}, \cdots, \mathbf{b}^{(k)}] \in \mathbb{R}^{N\times K}$ **marginal**

$\mathbf{A} = [\mathbf{a}^{(1)}, \mathbf{a}^{(2)}, \cdots, \mathbf{a}^{(K)}] \in \mathbb{R}^{N\times K}$ **action**

joint learning

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

Learning with independent marginal benefits

joint learning

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

Learning with independent marginal benefits

joint learning

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

remarks

- quadratic programming jointly convex in **G** and **B**

Learning with independent marginal benefits

joint learning

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

remarks

- quadratic programming jointly convex in **G** and **B**
- spectral radius $\rho(\beta\mathbf{G})$ impact learning performance
  - approaching 0: action independent from graph structure
  - approaching 1: action related to eigenvector centrality

Learning with independent marginal benefits

joint learning

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

remarks

- quadratic programming jointly convex in **G** and **B**
- spectral radius $\rho(\beta\mathbf{G})$ impact learning performance
    - approaching 0: action independent from graph structure
    - approaching 1: action related to eigenvector centrality
- other factors: number of games, noise level, network density

Learning with independent marginal benefits

Algorithm 1 Learning Games with Independent Marginal Benefits

---

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

Learning with homophilous marginal benefits

- Phenomenon of homophily in social networks [McPherson01]

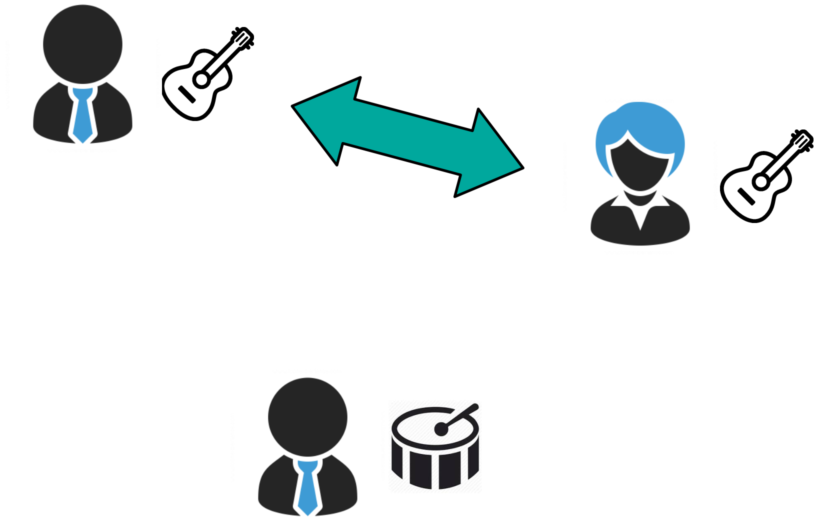

McPherson et al., "Birds of a feather: Homophily in social networks," Annual Review of Sociology, 2001.

Learning with homophilous marginal benefits

- Phenomenon of homophily in social networks [McPherson01]
- Given homophily marginal benefits are **smooth** functions on graph

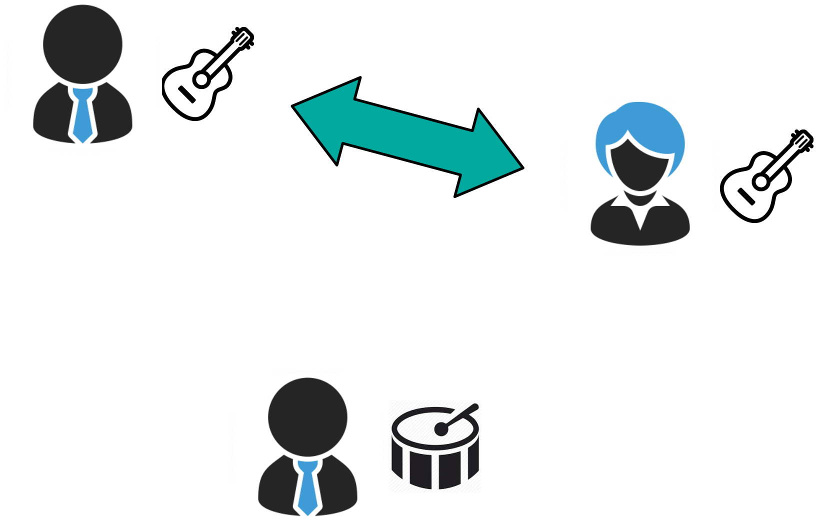

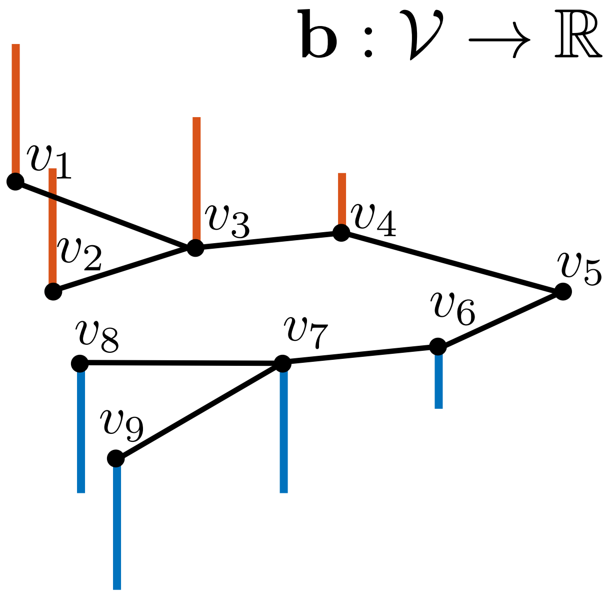

McPherson et al., "Birds of a feather: Homophily in social networks," Annual Review of Sociology, 2001.

Learning with homophilous marginal benefits

- Phenomenon of homophily in social networks [McPherson01]
- Given homophily marginal benefits are **smooth** functions on graph

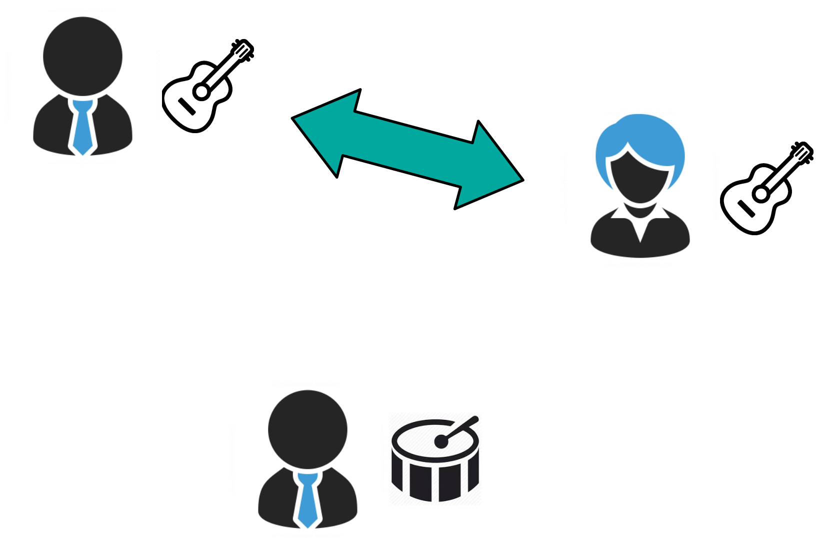

$$\mathbf{b} : \mathcal{V} \to \mathbb{R}$$

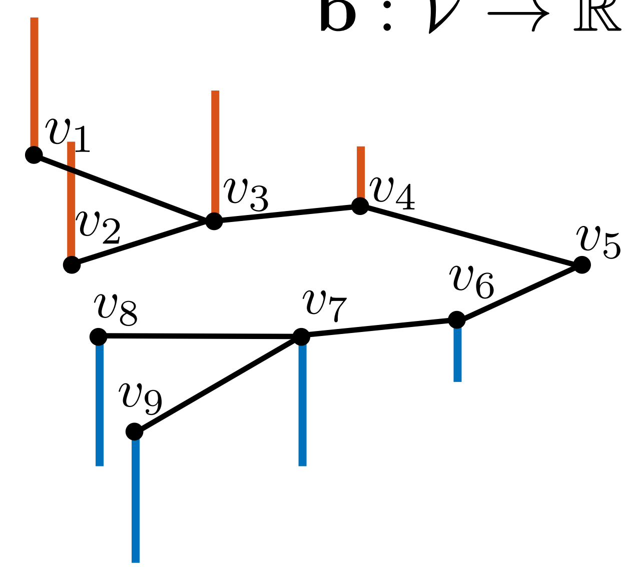

$$\mathbf{b}^T \mathbf{L} \mathbf{b} = \frac{1}{2} \sum_{i,j=1}^{N} G_{ij} (b_i - b_j)^2$$

measure of “smoothness” [Zhou04]

McPherson et al., “Birds of a feather: Homophily in social networks,” Annual Review of Sociology, 2001.
Zhou and Schölkopf, “A regularization framework for learning from graph data,” ICML Workshop, 2004.

Learning with homophilous marginal benefits

joint learning

$$\underset{\mathbf{G},\mathbf{B}}{\text{minimize}} \;\; h(\mathbf{G},\mathbf{B}) = ||(\mathbf{I}-\beta\mathbf{G})\mathbf{A}-\mathbf{B}||_F^2 + \theta_1||\mathbf{G}||_F^2 + \theta_2 \text{tr}(\mathbf{B}^T\mathbf{L}\mathbf{B}),$$

$$\text{subject to} \;\; G_{ij} = G_{ji},\; G_{ij} \geq 0,\; G_{ii} = 0 \;\text{ for } \forall i,j \in \mathcal{V},$$

$$||\mathbf{G}||_1 = N,$$

$$\mathbf{L} = \text{diag}(\sum_{j\in\mathcal{V}} G_{ij}) - \mathbf{G}$$

Learning with homophilous marginal benefits

joint learning

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

remarks

- not jointly convex in **G** and **B**
- convex in subproblems of solving for one while fixing other

Learning with homophilous marginal benefits

Algorithm 2 Learning Games with Homophilous Marginal Benefits

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

Outline

- Background
  - learning network structure from data
  - network games with linear-quadratic payoffs
- Learning games with linear-quadratic payoffs
  - independent marginal benefits
  - homophilous marginal benefits
- Experimental results
- Discussion

Experiments on synthetic data

- Random graphs with 20 nodes
  - Erdős-Rényi (ER): edges created independently with certain probability
  - Watts-Strogatz (WS): regular graph followed by random rewiring
  - Barabási-Albert (BA): graph generated using preferential attachment

Experiments on synthetic data

- Random graphs with 20 nodes
  - Erdős-Rényi (ER): edges created independently with certain probability
  - Watts-Strogatz (WS): regular graph followed by random rewiring
  - Barabási-Albert (BA): graph generated using preferential attachment
- Compute β so that $\rho(\beta \mathbf{G}) \in (0, 1)$
- Initialise marginal benefits for 50 games
  - homophilous: $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \mathbf{L}^{\dagger} + \frac{1}{10}\mathbf{I})$
- Generate equilibrium actions: $\mathbf{a} = (\mathbf{I} - \beta \mathbf{G})^{-1}\,\mathbf{b}$

Setting

- Evaluate on area under the curve (AUC)
- Baselines
  - sample correlation as edge weights
  - graph learned by regularised graphical Lasso [Lake10]

$$\begin{aligned} \underset{\boldsymbol{\Theta},\ \sigma^2}{\text{maximize}} \quad & \log\ \det\boldsymbol{\Theta} - \mathrm{tr}(\frac{1}{M}\mathbf{X}\mathbf{X}^T\boldsymbol{\Theta}) - \rho||\boldsymbol{\Theta}||_1, \\ \text{subject to} \quad & \boldsymbol{\Theta} = \mathbf{L} + \frac{1}{\sigma^2}\mathbf{I},\ \mathbf{L} \in \mathcal{L}, \end{aligned}$$

Lake and Tenenbaum, "Discovering structure by learning sparse graph," Annual Cognitive Science Conference, 2010.

Learning interaction network

independent marginal benefits

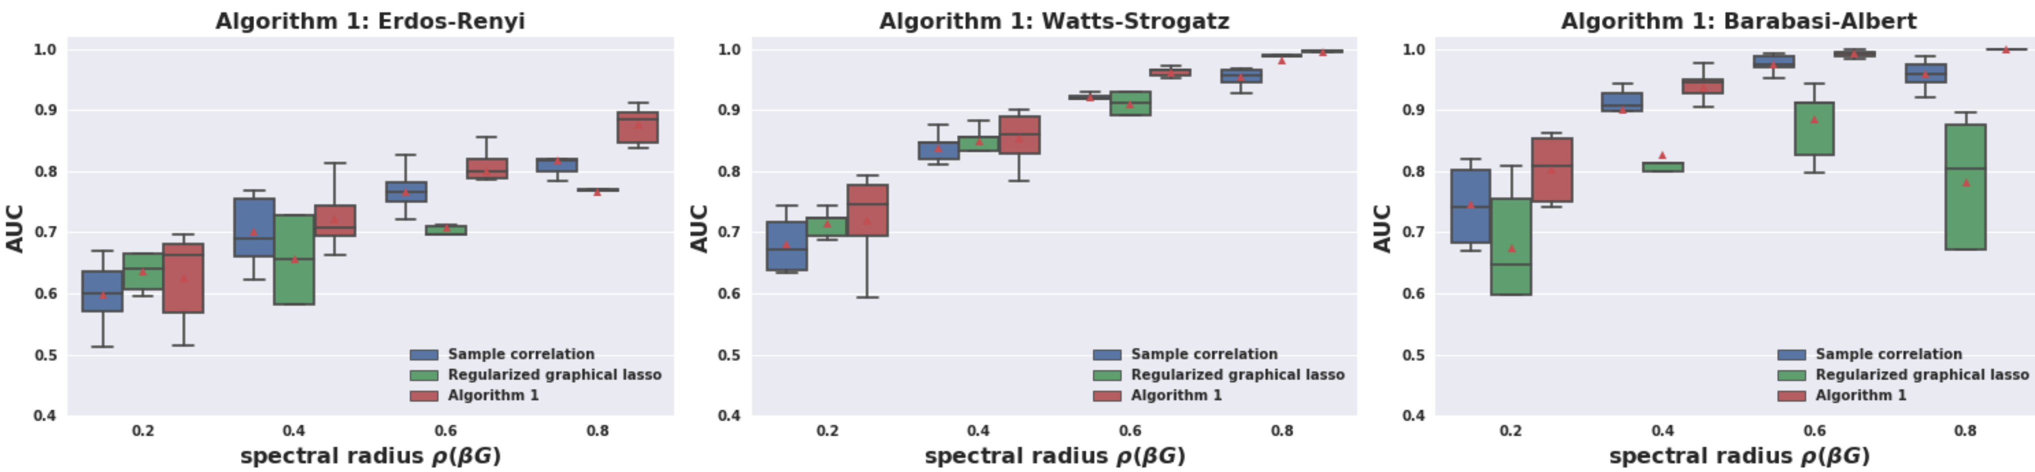

homophilous marginal benefits

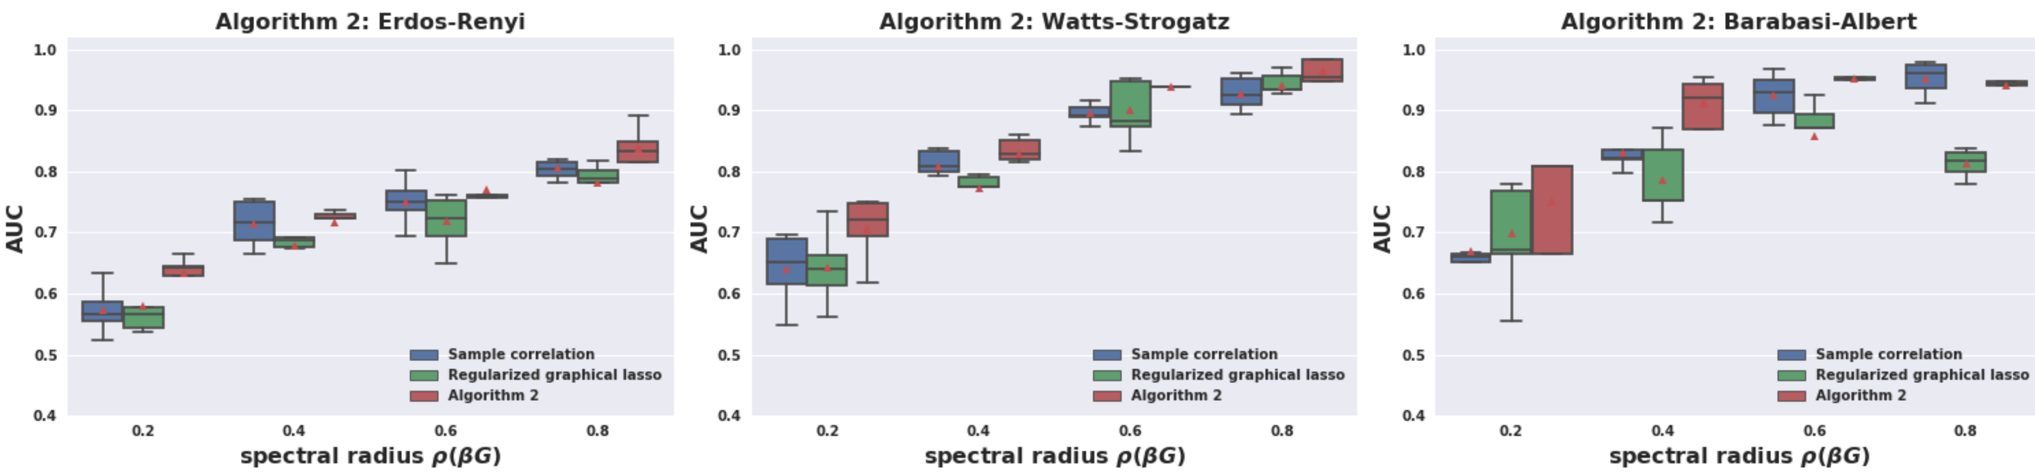

Performance vs. # games & noise

- Homophilous marginal benefits with $\rho(\beta\mathbf{G}) = 0.6$

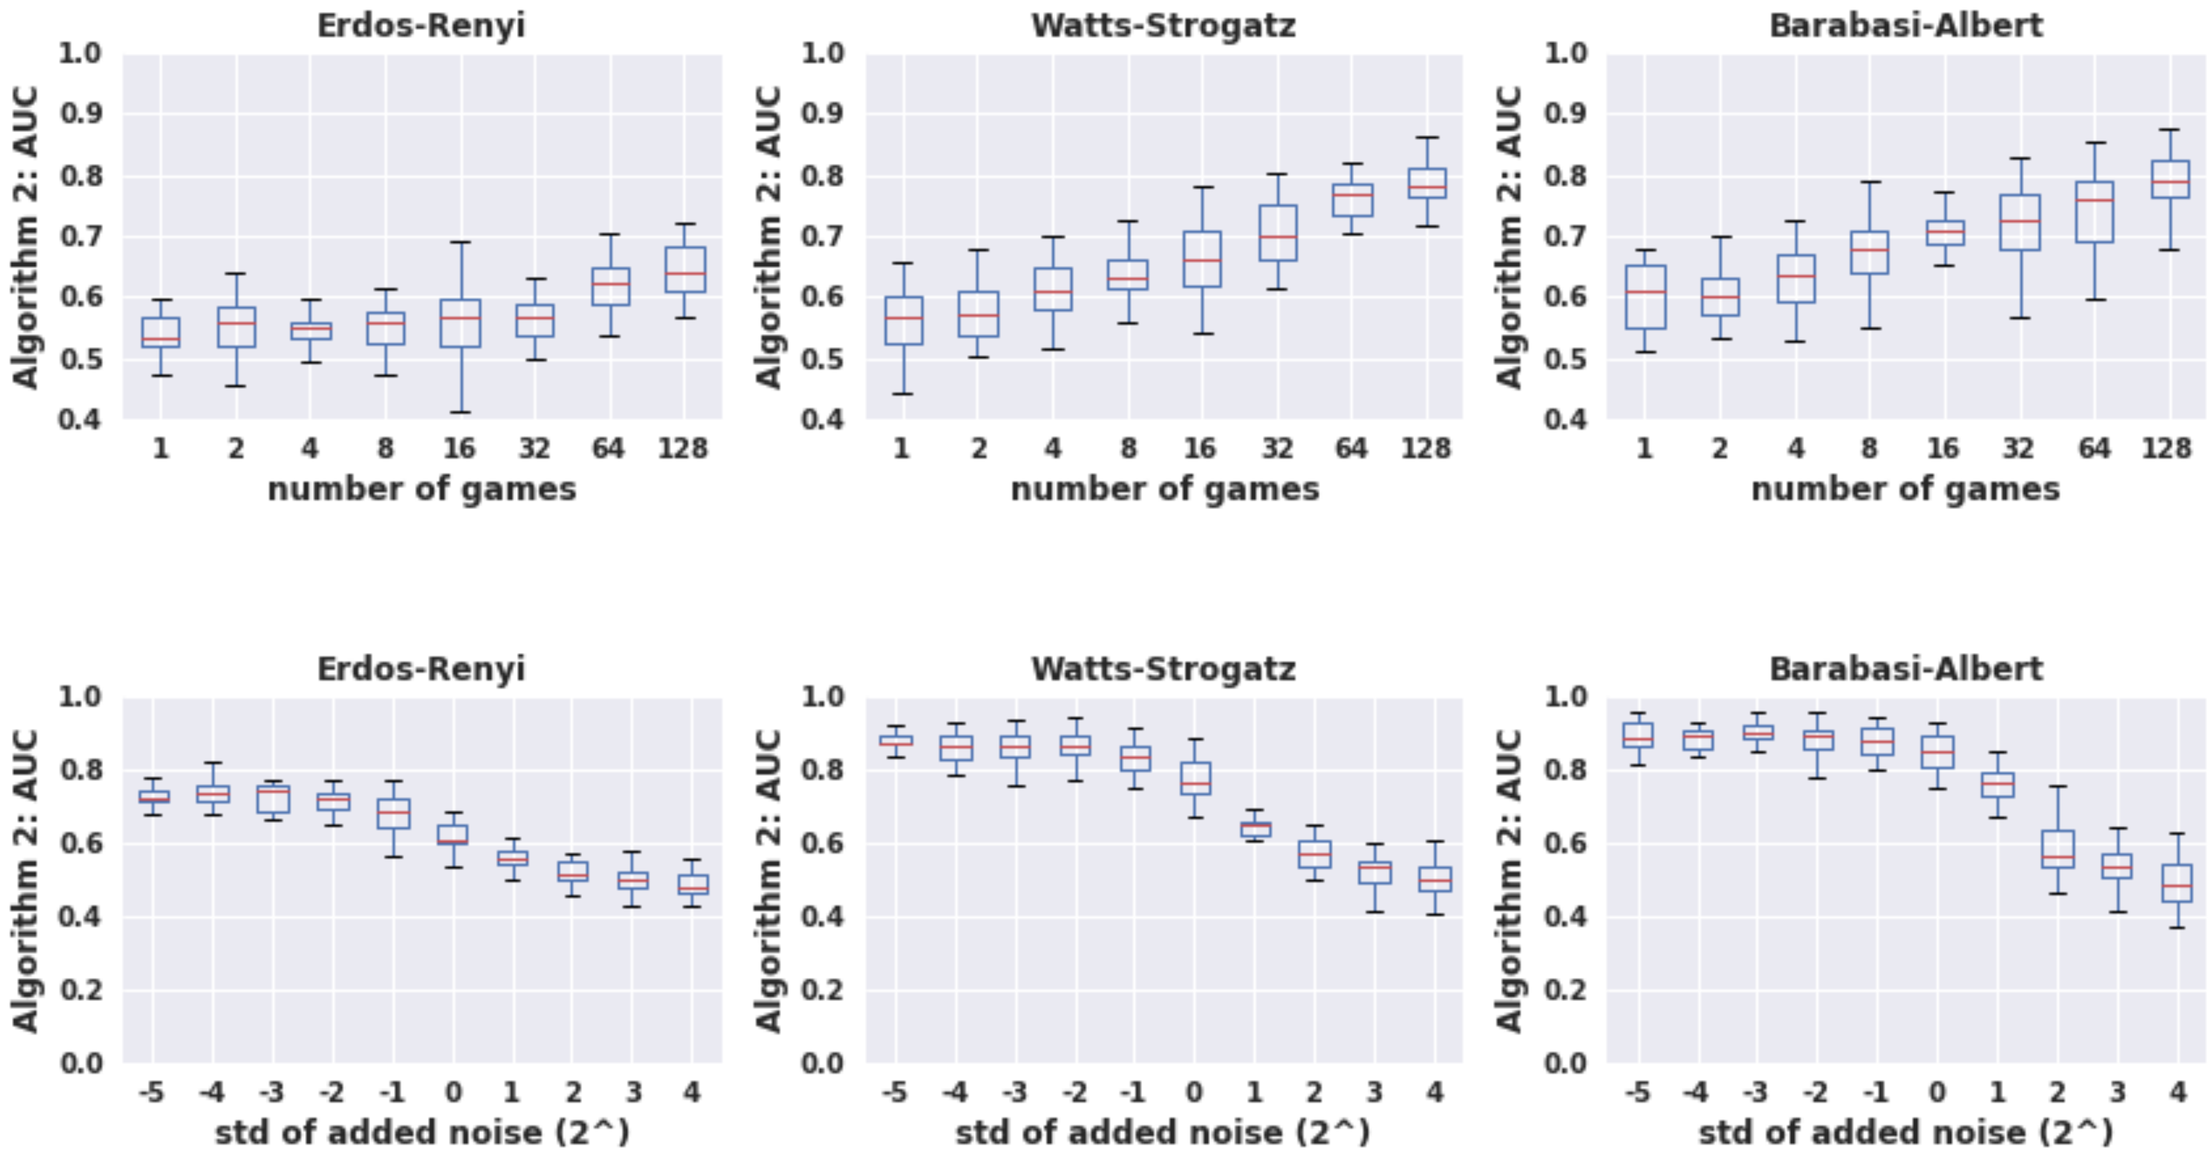

Performance vs. network structure

- Homophilous marginal benefits with $\rho(\beta \mathbf{G}) = 0.6$
- Parameters in random graph models
  - ER: each node pair connected with probability p
  - WS: k-regular graph with rewiring probability p
  - BA: m nodes added at each graph generation step

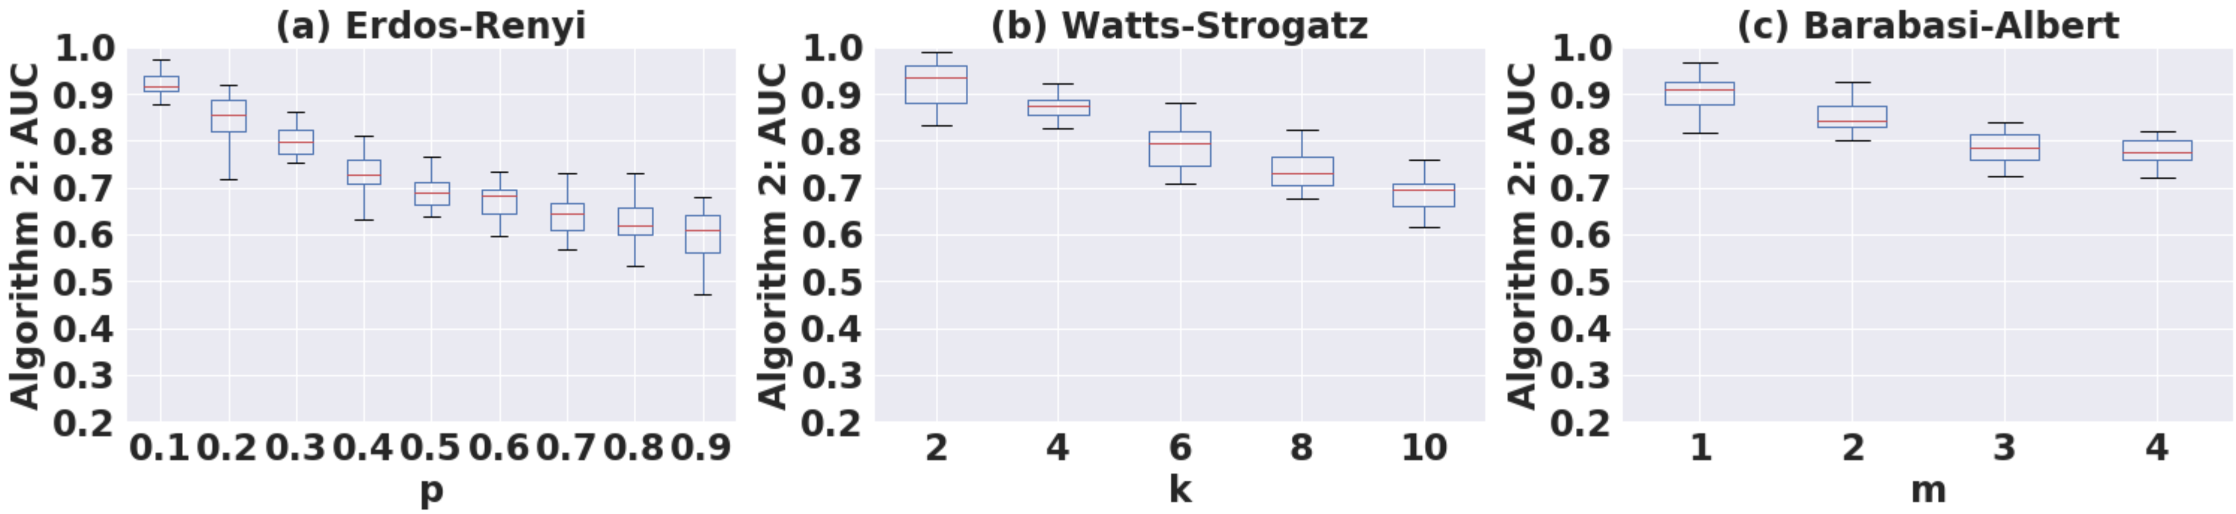

Performance vs. strength of homophily

- Homophilous marginal benefits with $\rho(\beta \mathbf{G}) = 0.6$
- Marginal benefits B as linear combinations of 1st-5th (strong homophily), 6th-10th (medium), 11th-15th (weak) eigenvectors of graph Laplacian

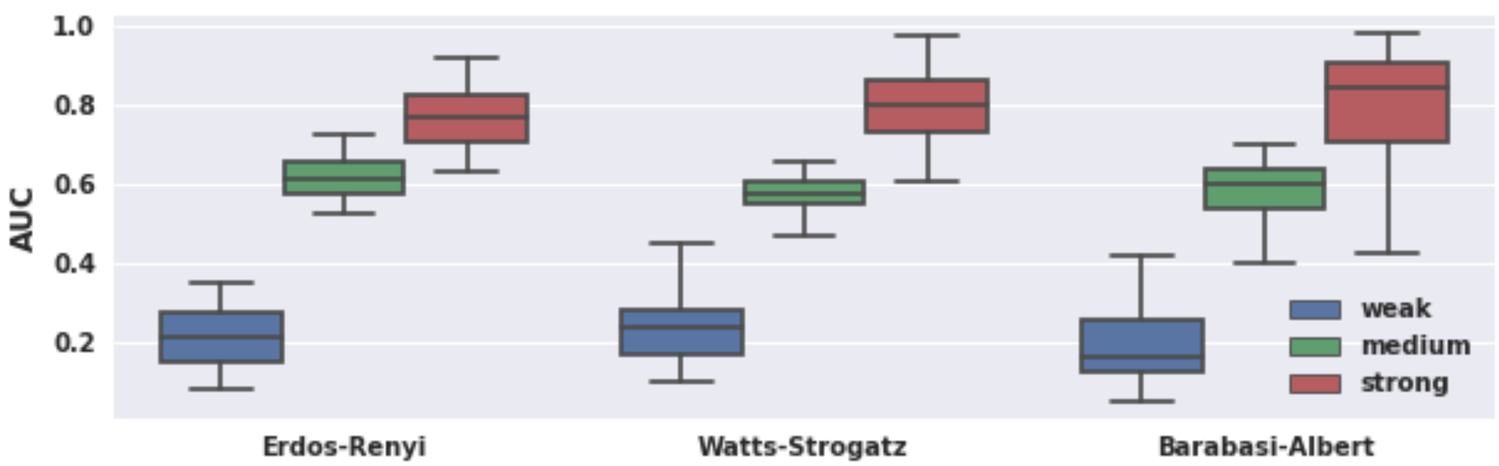

Learning marginal benefits

- Compare groundtruth and learned marginal benefits using coefficients of determination (R^2)

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

Experiments on real world data

- Learning social network
  - players: 182 households in a village in rural India [Banerjee13]
  - actions in 31 games: number of facilities adopted by each household
  - strategic complements: conformity to social norms (decisions made by neighbours)
  - compare with groundtruth: self-reported friendship

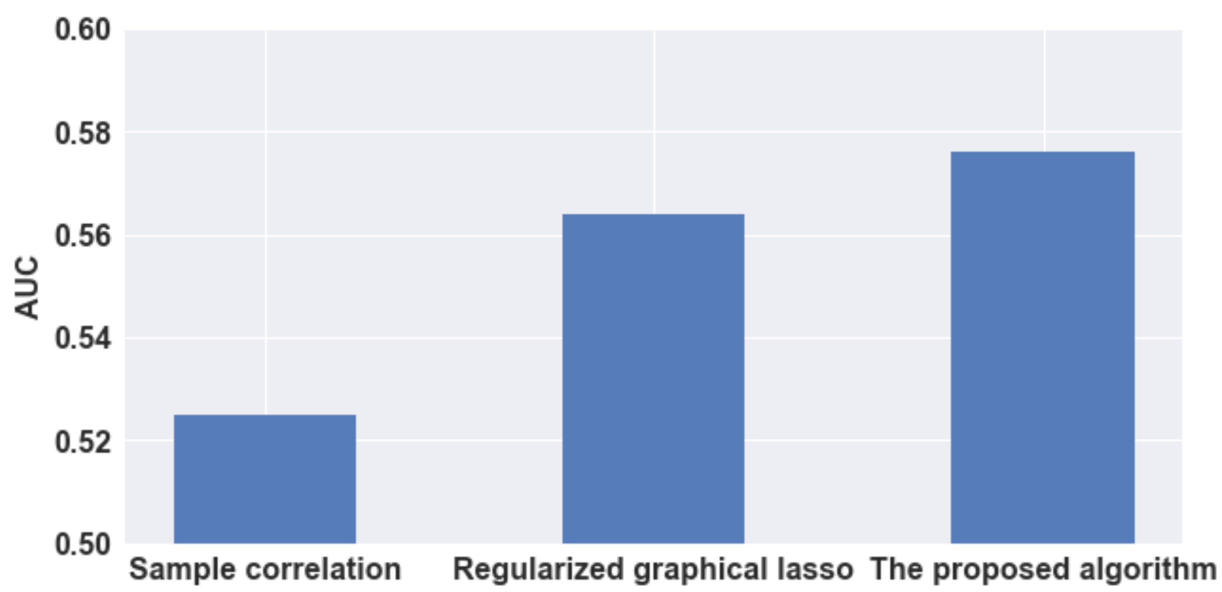

Banerjee et al., "The diffusion of microfinance," Science, 2013.

Experiments on real world data

- Learning trade network
  - players: 235 countries
  - actions in 192 games: import/export of 96 products of countries in 2008
  - strategic substitutes: complementary demand/supply leads to nonsmooth signals on trade network
  - compare with groundtruth: trade network of countries in 2002

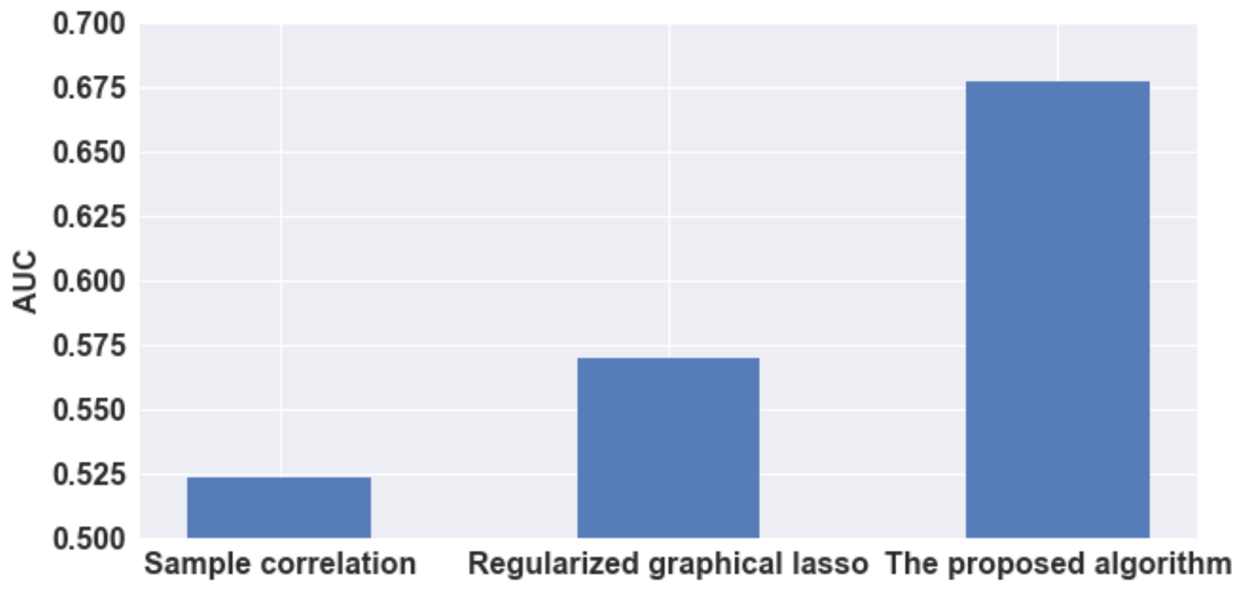

Outline

- Background
  - learning network structure from data
  - network games with linear-quadratic payoffs
- Learning games with linear-quadratic payoffs
  - independent marginal benefits
  - homophilous marginal benefits
- Experimental results
- Discussion

Discussion

- Applications in practical scenarios
  - detect communities of players (for stratification)
  - compute centrality measures (for efficient targeting strategies)
  - design intervention mechanisms to achieve planning objective
    - maximise total utilities of players (via adjusting marginal benefits) [Galeotti17]
    - reduce inequality between players (via adjusting interaction network)

Galeotti et al., "Targeting interventions in networks,", arXiv, 2017.

Discussion

- Applications in practical scenarios
  - detect communities of players (for stratification)
  - compute centrality measures (for efficient targeting strategies)
  - design intervention mechanisms to achieve planning objective
    - maximise total utilities of players (via adjusting marginal benefits) [Galeotti17]
    - reduce inequality between players (via adjusting interaction network)
- Open issues & future directions
  - determination of β (strength of strategic interaction)
  - probabilistic interpretation of learning framework
  - theoretical understanding, e.g., recovery guarantee
  - more general payoff functions
  - real-world intervention

Galeotti et al., "Targeting interventions in networks,", arXiv, 2017.