Social Choice on Social Networks

Amirali Salehi-Abari

University of Toronto, University of Waterloo

September 22, 2016

Social Choice/Group Decision Making

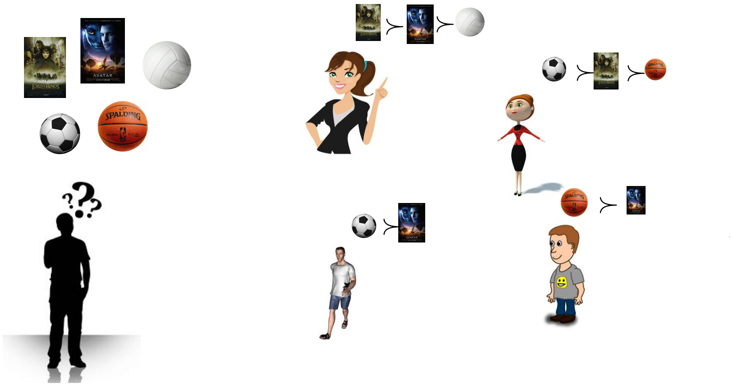

Addresses the problem of choosing a decision, from a set of alternatives/outcomes, for a group of individuals who have their own personal preferences over the set of alternatives.

Group Decision Making

- A group of friends chooses
  - a movie to watch.

Group Decision Making

- A group of friends chooses
  - a movie to watch.
  - a resort at which to vacation.

Group Decision Making

- A group of friends chooses
  - a movie to watch.
  - a resort at which to vacation.
  - a restaurant to dine at.

Group Decision Making

- A group of friends chooses
  - a movie to watch.
  - a resort at which to vacation.
  - a restaurant to dine at.
- Target marketing.

Group Decision Making

- A group of friends chooses
  - a movie to watch.
  - a resort at which to vacation.
  - a restaurant to dine at.
- Target marketing.
- Selecting a policy for an online system or a nation.

Preference Understanding

- Central and necessary for any group recommendation.

Preference Understanding

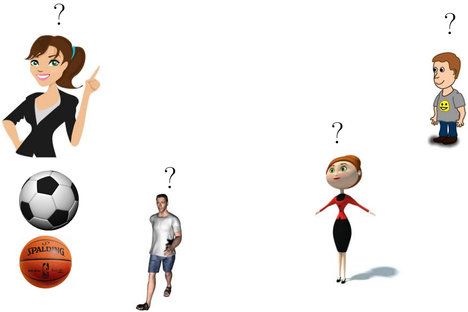

- Central and necessary for any group recommendation.
- Through preference learning or preference elicitation.

Preference Understanding

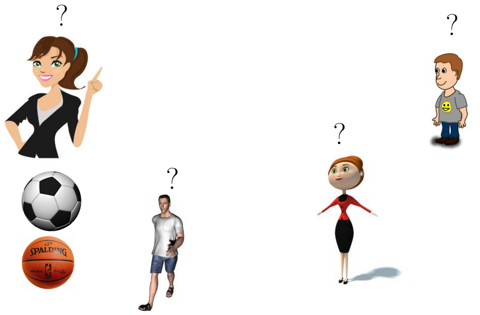

- Central and necessary for any group recommendation.
- Through preference learning or preference elicitation.
- Reducing cognitive and communication burden of users.

Preference Understanding

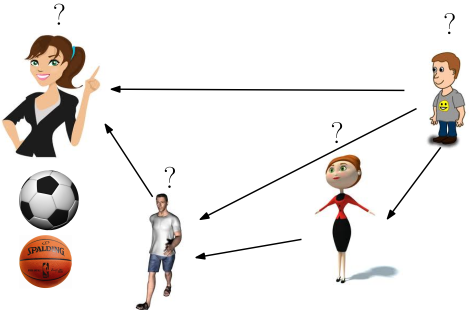

- Central and necessary for any group recommendation.
- Through preference learning or preference elicitation.
- Reducing cognitive and communication burden of users.
- Hypothesis: social networks are useful for preference understanding and consequently group recommendations.

Social Influence or Dynamics

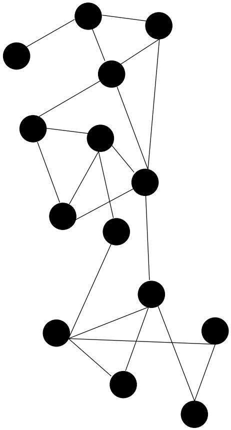

Social Influence or Dynamics

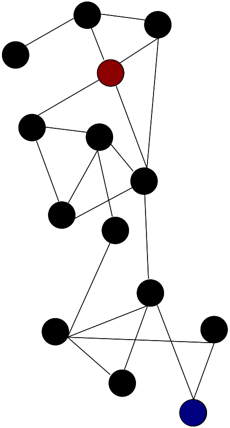

Social Influence or Dynamics

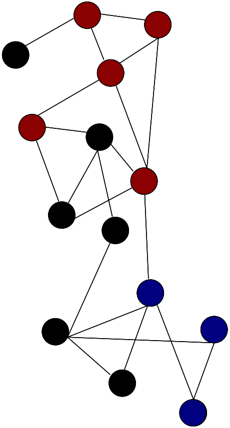

Social Influence or Dynamics

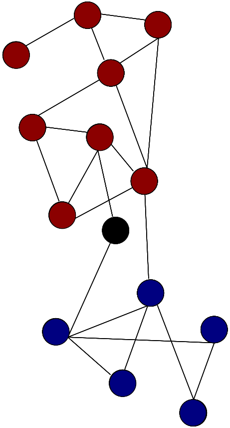

Social Influence or Dynamics

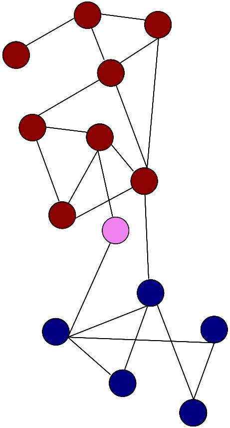

Social Influence or Dynamics

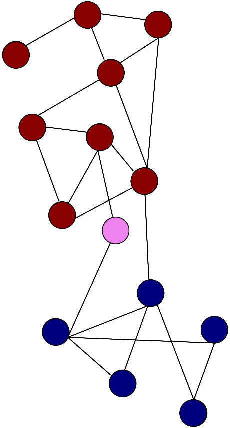

Social Selection or Homophily

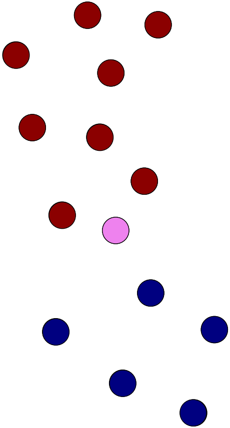

Social Influence or Dynamics

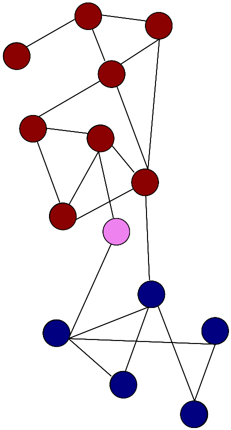

Social Selection or Homophily

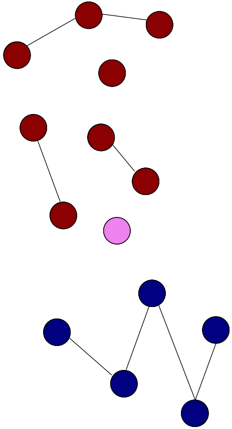

Social Influence or Dynamics

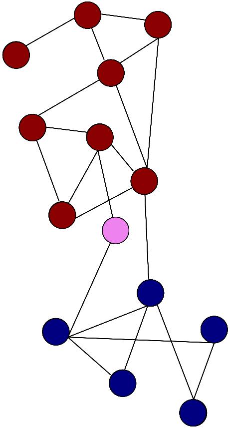

Social Selection or Homophily

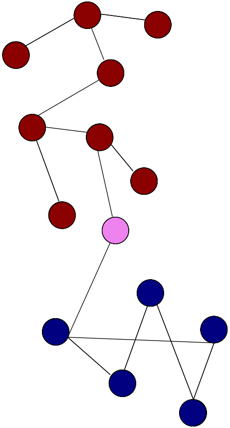

Social Influence or Dynamics

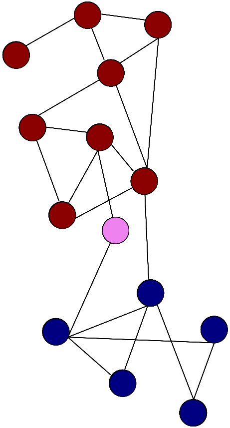

Social Selection or Homophily

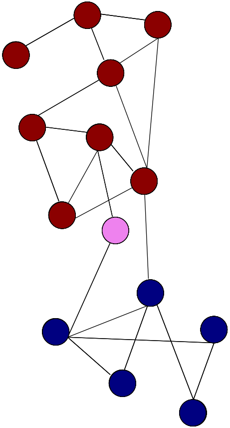

Social Influence or Dynamics

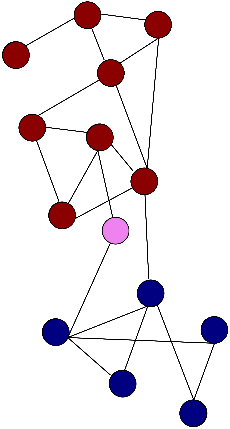

Social Selection or Homophily

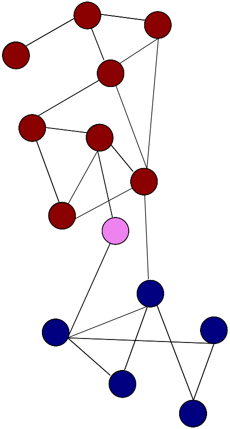

Observation

Dynamics or Selection $\Longrightarrow$ Correlation.

Preference Learning and Decision Making on Social Nets

Research Objective

To advance the understanding and mathematical modelling of preference dynamics and correlations over social networks and to exploit the computational and predictive power of these models for developing efficient algorithms of group decision making and recommendations, with less required user data, and lower cognitive and communication burden.

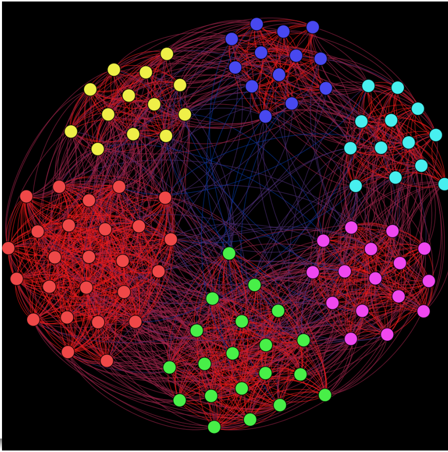

Outline

1. Preference Representation and Group Decision Making
2. Social Choice and Preference Dynamics
3. Sociak Choice and Social Selection in Preferences

Preferences

$\mathcal{A}$: 

- A set of m options or alternatives $\mathcal{A} = \{a_1, \ldots, a_m\}$.

Preferences

$\mathcal{A}$: 

j

- A set of m options or alternatives $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- A set of n agents $\mathcal{N} = \{1, \ldots, n\}$.

Preferences

- A set of m options or alternatives $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- A set of n agents $\mathcal{N} = \{1, \ldots, n\}$.
- For each $j \in \mathcal{N}$, j's *preferences* over $\mathcal{A}$:
  - Utility function $u_j : \mathcal{A} \to \mathbb{R}$.

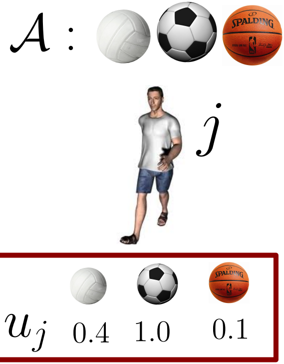

Preferences

- A set of m options or alternatives $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- A set of n agents $\mathcal{N} = \{1, \ldots, n\}$.
- For each $j \in \mathcal{N}$, j's *preferences* over $\mathcal{A}$:
  - Utility function $u_j : \mathcal{A} \to \mathbb{R}$.
  - Strict orderings $\succ_j$.

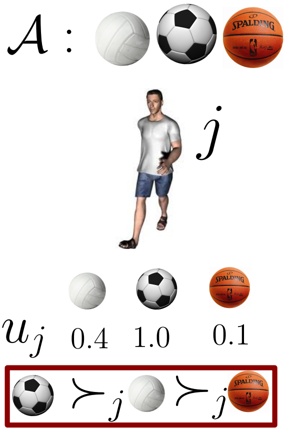

Preferences

- A set of m options or alternatives $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- A set of n agents $\mathcal{N} = \{1, \ldots, n\}$.
- For each $j \in \mathcal{N}$, j's *preferences* over $\mathcal{A}$:
  - Utility function $u_j : \mathcal{A} \rightarrow \mathbb{R}$.
  - Strict orderings $\succ_j$.
  - Ranking $r_j : \mathcal{A} \rightarrow \{1, \cdots, m\}$.

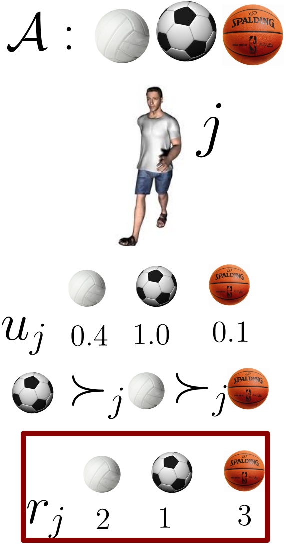

Consensus/Group Decision Making

Social Welfare

- Captures overall societal satisfaction.
- E.g., utilitarian social welfare:

$$sw(a) = \sum_j u_j(a)$$

Group Decision Making

Select $a^* \in \mathcal{A}$ such that maximize social welfare:

$$a^* = \arg\max_{a \in \mathcal{A}} sw(a)$$

Consensus/Group Decision Making

Social Welfare

- Captures overall societal satisfaction.
- E.g., utilitarian social welfare:

$$sw(a) = \sum_{j} u_j(a)$$

Group Decision Making

Select $a^* \in \mathcal{A}$ such that maximize social welfare:

$$a^* = \arg\max_{a \in \mathcal{A}} sw(a)$$

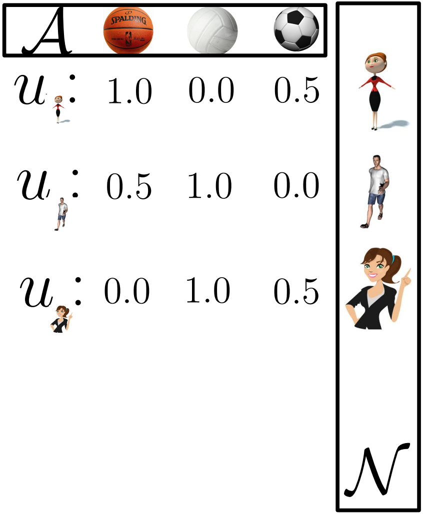

Consensus/Group Decision Making

Social Welfare

- Captures overall societal satisfaction.
- E.g., utilitarian social welfare:

$$sw(a) = \sum_j u_j(a)$$

Group Decision Making

Select $a^* \in \mathcal{A}$ such that maximize social welfare:

$$a^* = \arg\max_{a \in \mathcal{A}} sw(a)$$

$\mathcal{A}$			
$\mathcal{U}$:	1.0	0.0	0.5
$\mathcal{U}$:	0.5	1.0	0.0
$\mathcal{U}$:	0.0	1.0	0.5

sw: 1.5

$\mathcal{N}$

Consensus/Group Decision Making

Social Welfare

- Captures overall societal satisfaction.
- E.g., utilitarian social welfare:

$$sw(a) = \sum_j u_j(a)$$

Group Decision Making

Select $a^* \in \mathcal{A}$ such that maximize social welfare:

$$a^* = \arg\max_{a \in \mathcal{A}} sw(a)$$

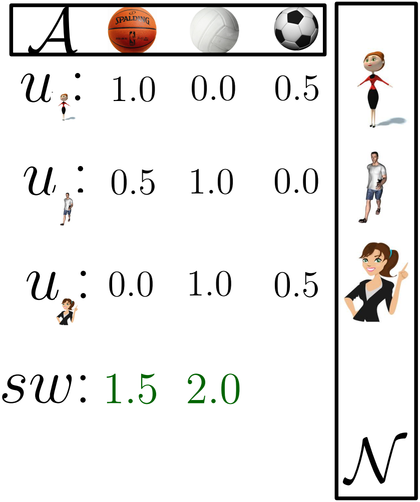

Consensus/Group Decision Making

Social Welfare

- Captures overall societal satisfaction.
- E.g., utilitarian social welfare:

$$sw(a) = \sum_j u_j(a)$$

Group Decision Making

Select $a^* \in \mathcal{A}$ such that maximize social welfare:

$$a^* = \arg\max_{a \in \mathcal{A}} sw(a)$$

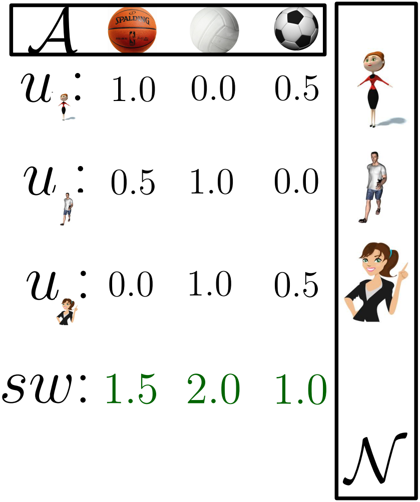

Consensus/Group Decision Making

Social Welfare

- Captures overall societal satisfaction.
- E.g., utilitarian social welfare:

$$sw(a) = \sum_j u_j(a)$$

Group Decision Making

Select $a^* \in \mathcal{A}$ such that maximize social welfare:

$$a^* = \arg\max_{a \in \mathcal{A}} sw(a)$$

$\mathcal{A}$			
u:	1.0	0.0	0.5
u:	0.5	1.0	0.0
u:	0.0	1.0	0.5
sw:	1.5	$\boxed{2.0 \; a^*}$	1.0

$\mathcal{N}$

Standard Group Recommendation with Rankings

Group Recommendation

- Goal: Select option $a^* \in \mathcal{A}$ for group $\mathcal{S} \subseteq \mathcal{N}$.

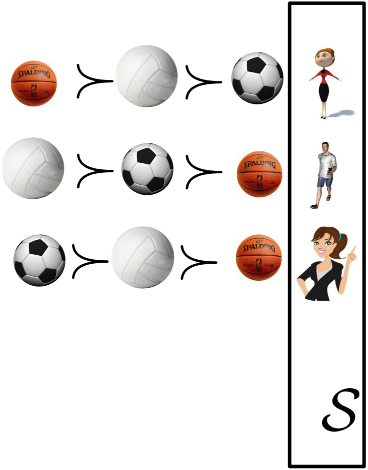

Standard Group Recommendation with Rankings

Group Recommendation

- Goal: Select option $a^* \in \mathcal{A}$ for group $\mathcal{S} \subseteq \mathcal{N}$.
- *Scoring rule* $g(a, r)$:
  - Option a's score in ranking r.
  - E.g., Borda.

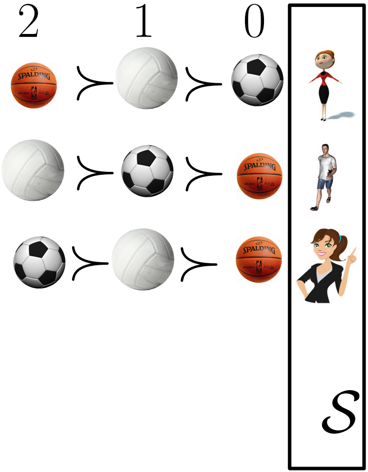

Standard Group Recommendation with Rankings

Group Recommendation

- Goal: Select option $a^* \in \mathcal{A}$ for group $\mathcal{S} \subseteq \mathcal{N}$.
- *Scoring rule* $g(a, r)$:
  - Option a's score in ranking r.
  - E.g., Borda.
- *Social welfare* of $a \in \mathcal{A}$: $sw(a, \mathcal{S}) = \sum_{i \in \mathcal{S}} g(a, r_i)$

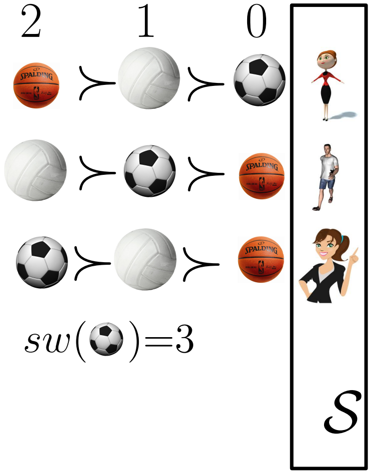

Standard Group Recommendation with Rankings

Group Recommendation

- Goal: Select option $a^* \in \mathcal{A}$ for group $\mathcal{S} \subseteq \mathcal{N}$.
- *Scoring rule* $g(a, r)$:
  - Option a's score in ranking r.
  - E.g., Borda.
- *Social welfare* of $a \in \mathcal{A}$: $sw(a, \mathcal{S}) = \sum_{i \in \mathcal{S}} g(a, r_i)$

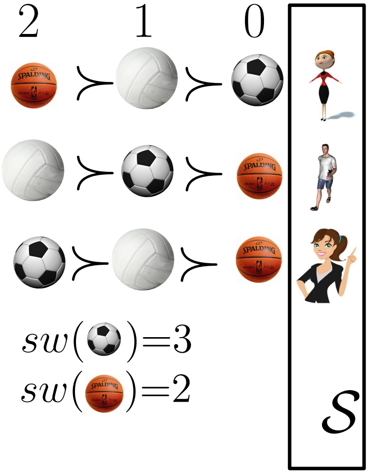

Standard Group Recommendation with Rankings

Group Recommendation

- Goal: Select option $a^* \in \mathcal{A}$ for group $\mathcal{S} \subseteq \mathcal{N}$.
- *Scoring rule* $g(a, r)$:
  - Option a's score in ranking r.
  - E.g., Borda.
- *Social welfare* of $a \in \mathcal{A}$: $sw(a, \mathcal{S}) = \sum_{i \in \mathcal{S}} g(a, r_i)$

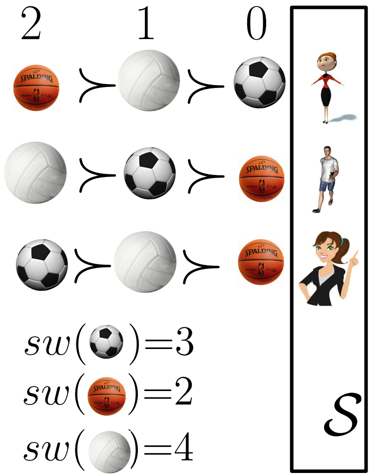

Standard Group Recommendation with Rankings

Group Recommendation

- Goal: Select option $a^* \in \mathcal{A}$ for group $\mathcal{S} \subseteq \mathcal{N}$.
- *Scoring rule* $g(a, r)$:
  - Option a's score in ranking r.
  - E.g., Borda.
- *Social welfare* of $a \in \mathcal{A}$: $sw(a, \mathcal{S}) = \sum_{i \in \mathcal{S}} g(a, r_i)$
- Select a^* so as to maximize *social welfare* $sw(., \mathcal{S})$:
  $$a^* = \arg\max_{a \in \mathcal{A}} sw(a, \mathcal{S})$$

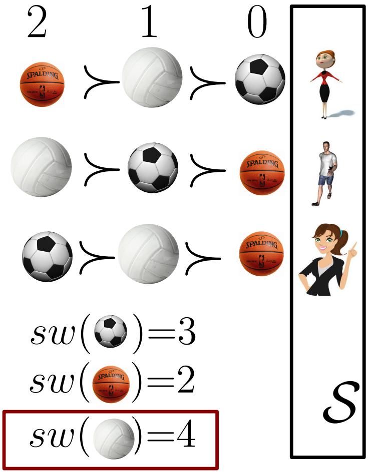

Part I:

Social Choice and Empathetic Preference Dynamics

Empathy–Movie Night

You

Empathy–Movie Night

You

Empathy–Movie Night

You

Empathy–Movie Night

You

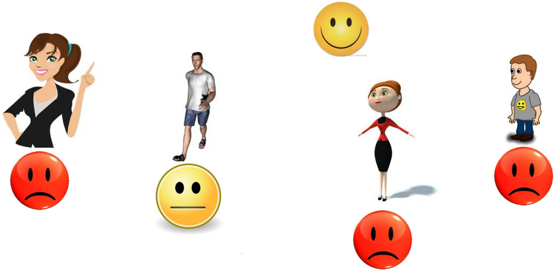

Empathy–Movie Night

You

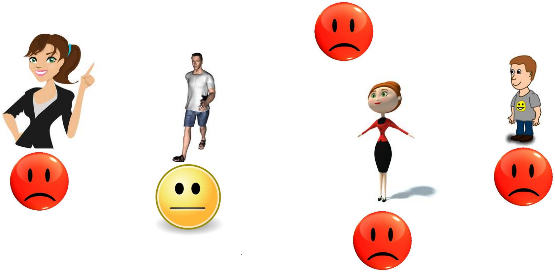

Empathy–Movie Night

You

Empathy–Movie Night

You

Empathy–Movie Night

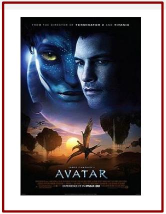

You

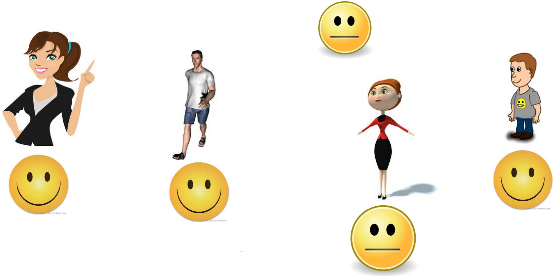

Empathy–Movie Night

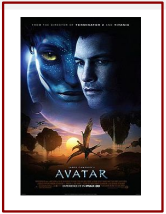

You

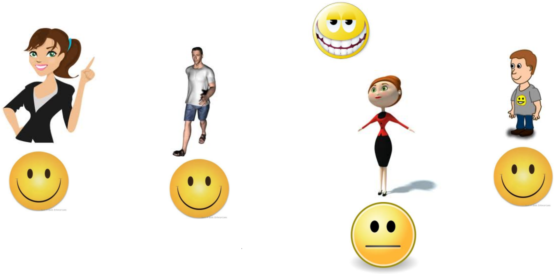

Economic Networks

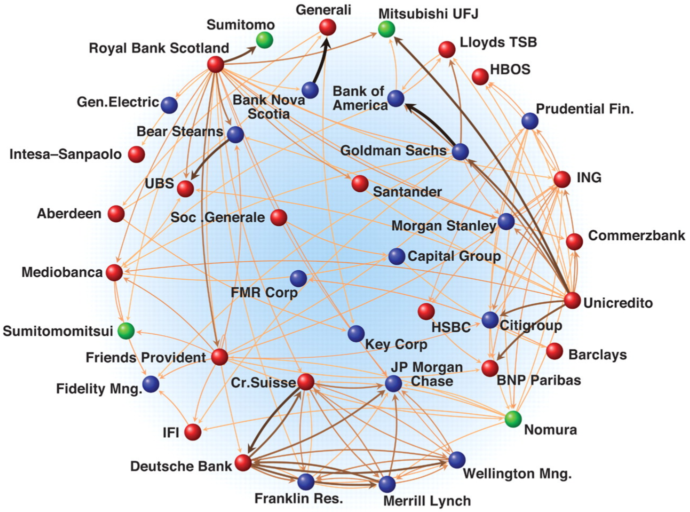

An example of economic networks; image is taken from *Economic Networks: The New Challenges (Frank Schweitzer et. al, 2009).*

Empathetic Models [1]

Empathetic Social Networks

- *Intrinsic preferences* over $\mathcal{A}$: intrinsic utility function $u_j^{\mathcal{I}}(.)$.

Empathetic Models [1]

Empathetic Social Networks

- *Intrinsic preferences* over $\mathcal{A}$: intrinsic utility function $u_j^{\mathcal{I}}(.)$.
- *Complete* preference $u_j(.)$: combining j's intrinsic utility with other individuals' utilities.

Empathetic Models [1]

Empathetic Social Networks

- *Intrinsic preferences* over $\mathcal{A}$: intrinsic utility function $u_j^{\mathcal{I}}(.)$.
- *Complete* preference $u_j(.)$: combining j's intrinsic utility with other individuals' utilities.
- A directed weighted graph $G = (\mathcal{N}, E)$.

Empathetic Models [1]

Empathetic Social Networks

- *Intrinsic preferences* over $\mathcal{A}$: intrinsic utility function $u_j^{\mathcal{I}}(.)$.
- *Complete* preference $u_j(.)$: combining j's intrinsic utility with other individuals' utilities.
- A directed weighted graph $G = (\mathcal{N}, E)$.

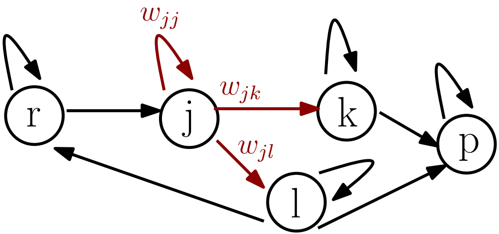

Empathetic Models [1]

Empathetic Social Networks

- *Intrinsic preferences* over $\mathcal{A}$: intrinsic utility function $u_j^{\mathcal{I}}(.)$.
- *Complete* preference $u_j(.)$: combining j's intrinsic utility with other individuals' utilities.
- A directed weighted graph $G = (\mathcal{N}, E)$.

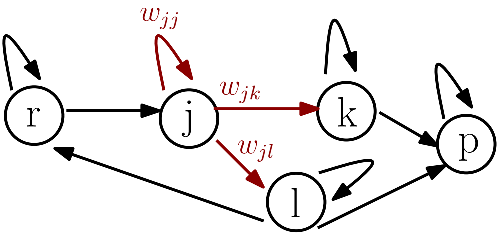

- Weight matrix $\mathbf{W} = [w_{ij}]$.

Local Empathetic Model

- $u_j(a) = \sum_k w_{jk} u_k^{\mathcal{I}}(a)$

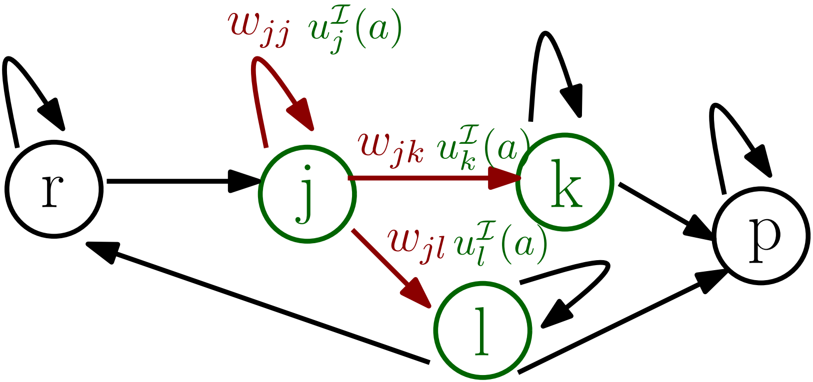

Local Empathetic Model

- $u_j(a) = \sum_k w_{jk} u_k^{\mathcal{I}}(a)$

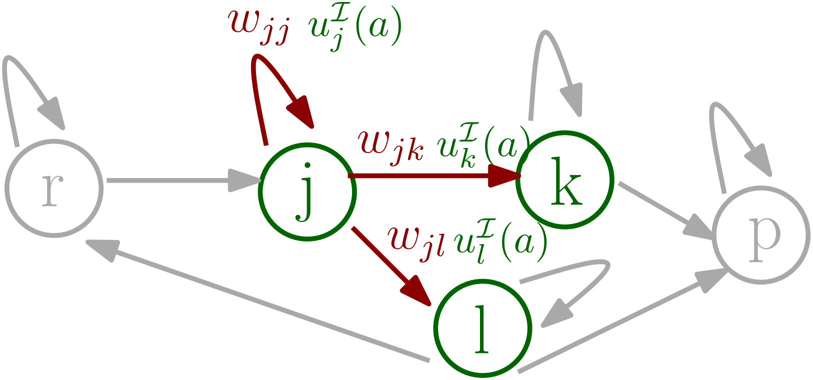

Global Empathetic Model

- $u_j(a) = w_{jj} u_j^{\mathcal{I}}(a) + \sum_{k \neq j} w_{jk} u_k(a)$

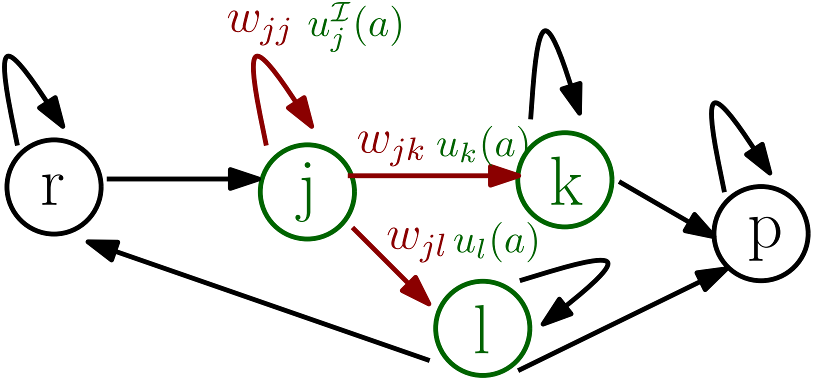

Global Empathetic Model

- $u_j(a) = w_{jj} u_j^{\mathcal{I}}(a) + \sum_{k \neq j} w_{jk} u_k(a)$

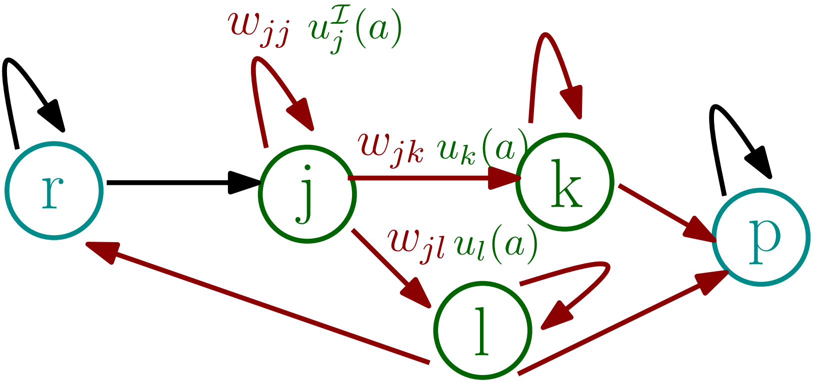

Global Empathetic Model

- $u_j(a) = w_{jj} u_j^{\mathcal{I}}(a) + \sum_{k \neq j} w_{jk} u_k(a)$

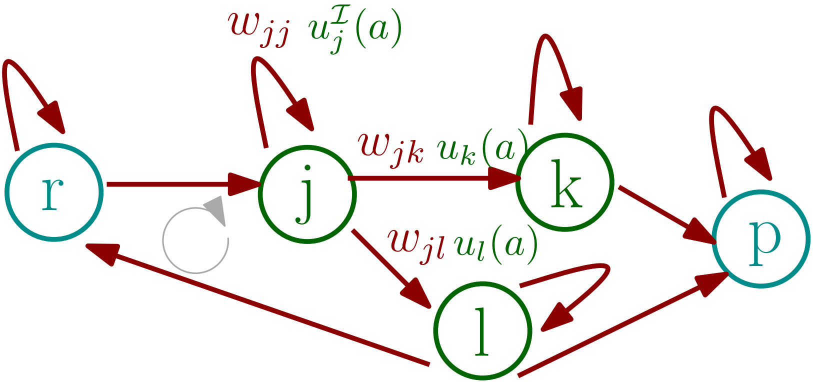

Fixed-Point Utilities

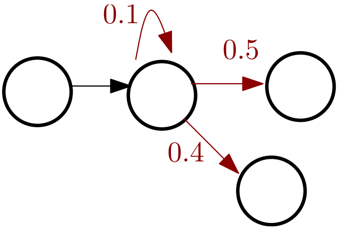

(Mild) Assumptions

- **Non-negativity**: weights are not negative ($w_{ij} \geq 0$).
- **Normalization**: outgoing weights sum up to 1 ($\sum_k w_{jk} = 1$).
- **Positive self-loop**: self-loops greater than 0 ($w_{jj} > 0$).

Fixed-Point Utility Thm.

Assuming non-negativity, normalization, and positive self-loop, there is a fixed-point solution $\mathbf{u}(a) = (\mathbf{I} - \mathbf{W} + \mathbf{D})^{-1}\mathbf{D}\mathbf{u}^{\mathcal{I}}(a)$.

Agents with Empathetic Utilities.

Agents with Empathetic Utilities.

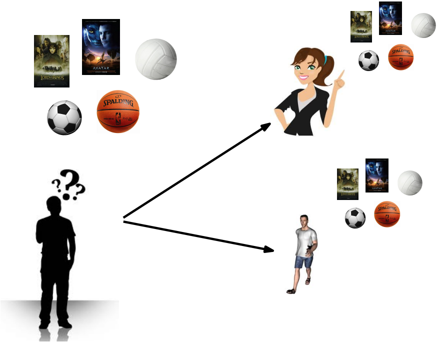

Agents with Empathetic Utilities.

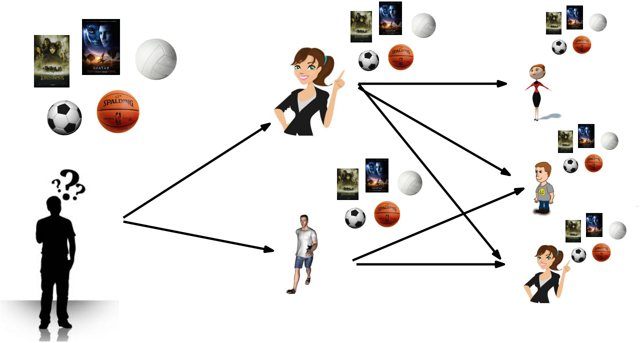

What do empathetic models offer?

- Simplifying preference aggregation.
    - Reporting intrinsic preferences + relationship weights.
- Reducing informational, and cognitive burden.
- Decreasing communications between agents.

Recall: Group Decision Making

- Select $a^* \in \mathcal{A}$ such that maximize social welfare $sw(.)$
- Non-empathetic: $sw(a) = \sum_j u_j^{\mathcal{I}}(a)$

Social Welfare as Weighted Intrinsic Utilities

Recall: Group Decision Making

- Select $a^* \in \mathcal{A}$ such that maximize social welfare $sw(.)$
- Non-empathetic: $sw(a) = \sum_j u_j^{\mathcal{I}}(a)$

Social Welfare as Weighted Intrinsic Utilities

- The social welfare in both models can be written as

$$sw(a) = \sum_j \omega_j u_j^{\mathcal{I}}(a).$$

Recall: Group Decision Making

- Select $a^* \in \mathcal{A}$ such that maximize social welfare $sw(.)$
- Non-empathetic: $sw(a) = \sum_j u_j^{\mathcal{I}}(a)$

Social Welfare as Weighted Intrinsic Utilities

- The social welfare in both models can be written as

$$sw(a) = \sum_j \omega_j u_j^{\mathcal{I}}(a).$$

- In local model, ω_j is the sum of the j's incoming edge weights:

$$\omega_j = \sum_i w_{ij}$$

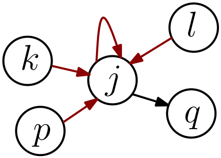

Recall: Group Decision Making

- Select $a^* \in \mathcal{A}$ such that maximize social welfare $sw(.)$
- Non-empathetic: $sw(a) = \sum_j u_j^{\mathcal{I}}(a)$

Social Welfare as Weighted Intrinsic Utilities

- The social welfare in both models can be written as

$$sw(a) = \sum_j \omega_j u_j^{\mathcal{I}}(a).$$

- In local model, ω_j is the sum of the j's incoming edge weights:

$$\omega_j = \sum_i w_{ij}$$

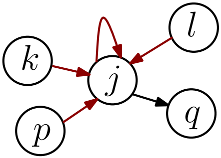

- Global model: $(\omega_1, \ldots, \omega_n) = \mathbf{e}^\top(\mathbf{I} - \mathbf{W} + \mathbf{D})^{-1}\mathbf{D}$.

Different Models, Different Weights, Different Decisions.

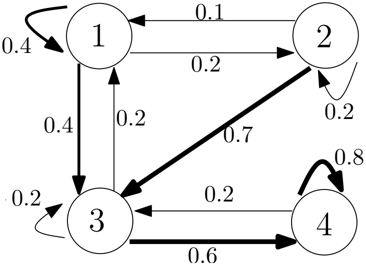

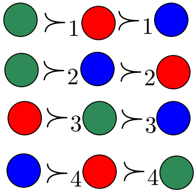

Different Models, Different Weights, Different Decisions.

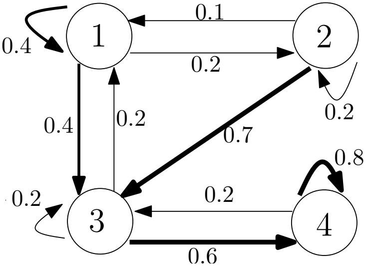

Plurality 1 0 0

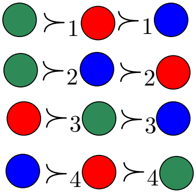

Different Models, Different Weights, Different Decisions.

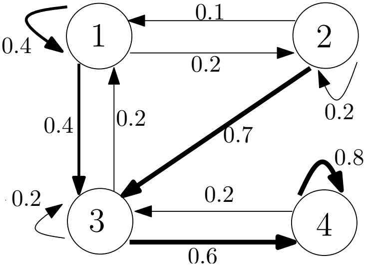

Plurality 1 0 0

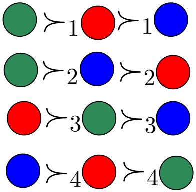

non-empathetic

local empathetic

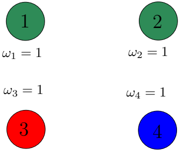

Different Models, Different Weights, Different Decisions.

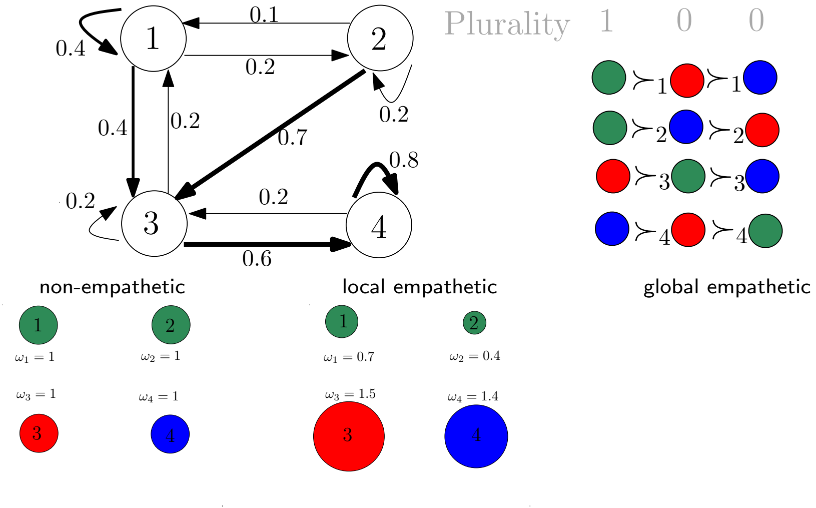

Different Models, Different Weights, Different Decisions.

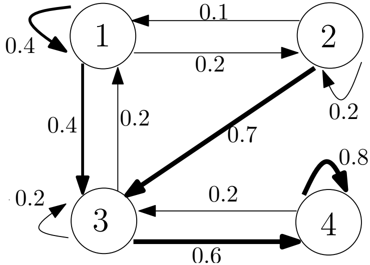

Plurality 1 0 0

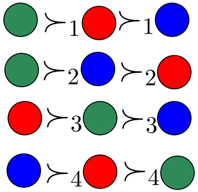

non-empathetic

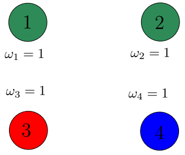

local empathetic

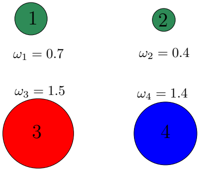

global empathetic

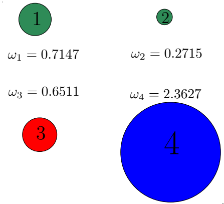

Computing Consensus Winner

- If $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ are pre-computed, then the consensus winner a^* is computed by $a^* = \arg\max_{a \in A} \sum_{n} \omega_j u_j^{\mathcal{I}}(a)$ in $O(nm)$ time.

Computing Consensus Winner

- If $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ are pre-computed, then the consensus winner a^* is computed by $a^* = \arg\max_{a \in A} \sum_n \omega_j u_j^{\mathcal{I}}(a)$ in $O(nm)$ time.
- Local model: $\boldsymbol{\omega}$ in $O(|E|)$.

Computing Consensus Winner

- If $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ are pre-computed, then the consensus winner a^* is computed by $a^* = \arg\max_{a \in A} \sum_{n} \omega_j u_j^{\mathcal{I}}(a)$ in $O(nm)$ time.
- Local model: $\boldsymbol{\omega}$ in $O(|E|)$.
  - Computation of a^* in $O(|E| + nm)$.

Computing Consensus Winner

- If $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ are pre-computed, then the consensus winner a^* is computed by $a^* = \arg\max_{a \in A} \sum_{n} \omega_j u_j^{\mathcal{I}}(a)$ in $O(nm)$ time.
- Local model: $\boldsymbol{\omega}$ in $O(|E|)$.
  - Computation of a^* in $O(|E| + nm)$.
- Global model: computation of $\boldsymbol{\omega}$ requires $O(n^3)$.

Computing Consensus Winner

- If $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ are pre-computed, then the consensus winner a^* is computed by $a^* = \arg\max_{a \in A} \sum_n \omega_j u_j^{\mathcal{I}}(a)$ in $O(nm)$ time.
- Local model: $\boldsymbol{\omega}$ in $O(|E|)$.
  - Computation of a^* in $O(|E| + nm)$.
- Global model: computation of $\boldsymbol{\omega}$ requires $O(n^3)$.
  - Computation of a^* in $O(n^3 + nm)$.

Computing Consensus Winner

- If $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ are pre-computed, then the consensus winner a^* is computed by $a^* = \arg\max_{a \in A} \sum_{n} \omega_j u_j^{\mathcal{I}}(a)$ in $O(nm)$ time.
- Local model: $\boldsymbol{\omega}$ in $O(|E|)$.
  - Computation of a^* in $O(|E| + nm)$.
- Global model: computation of $\boldsymbol{\omega}$ requires $O(n^3)$.
  - Computation of a^* in $O(n^3 + nm)$.
- $O(n^3)$ algorithm is not scalable for large networks.

Computing Consensus Winner

- If $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$ are pre-computed, then the consensus winner a^* is computed by $a^* = \arg\max_{a \in A} \sum_n \omega_j u_j^{\mathcal{I}}(a)$ in $O(nm)$ time.
- Local model: $\boldsymbol{\omega}$ in $O(|E|)$.
  - Computation of a^* in $O(|E| + nm)$.
- Global model: computation of $\boldsymbol{\omega}$ requires $O(n^3)$.
  - Computation of a^* in $O(n^3 + nm)$.
- $O(n^3)$ algorithm is not scalable for large networks.
- Propose an iterative algorithm: each iteration in $O(nm)$.

Iterative Methods

The Jacobi Method

$$u_j^{(t+1)}(a) = w_{jj} u_j^{\mathcal{I}}(a) + \sum_{k \neq j} w_{jk} u_k^{(t)}(a)$$

- Natural interpretation: each individual repeatedly
  1. Observes her friends' revealed utilities.
  2. Updates her own utility.

Iterative Candidate Elimination Algorithm (ICE)

- Exploited Jacobi iterative method to estimate utilities.
- Approx. bounds on estimated utilities and social welfare.
- Select the consensus winner based on approximations.

Iterated Candidate Elimination (ICE)

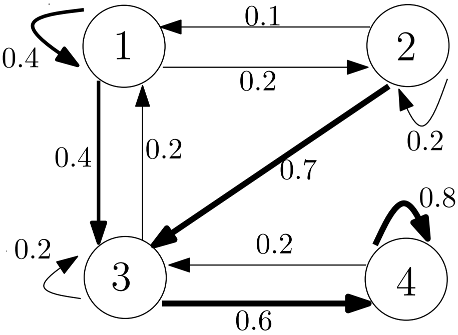

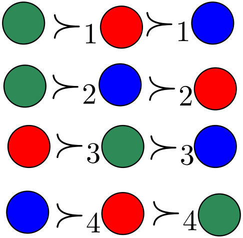

Iterated Candidate Elimination (ICE)

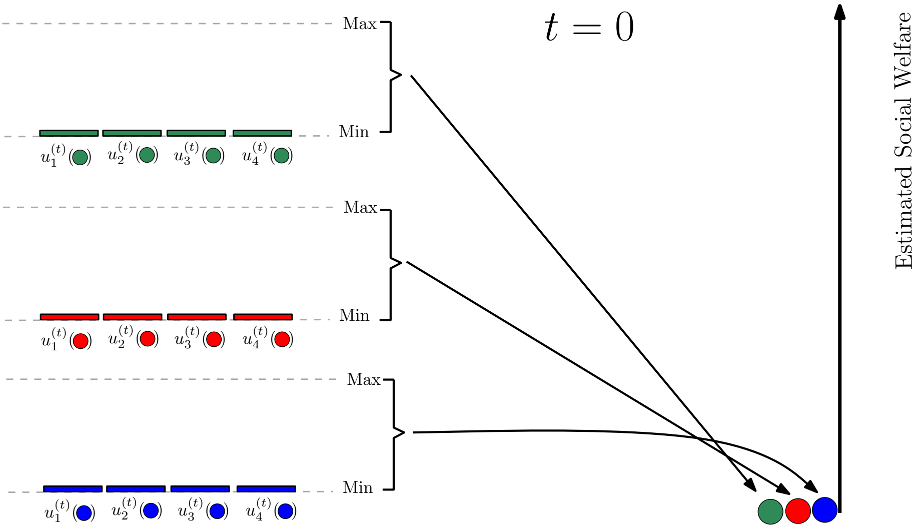

Iterated Candidate Elimination (ICE)

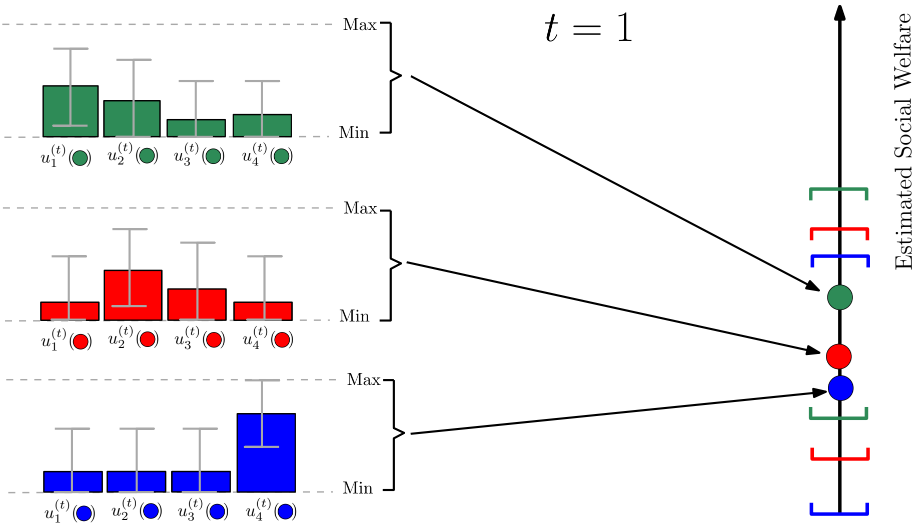

Iterated Candidate Elimination (ICE)

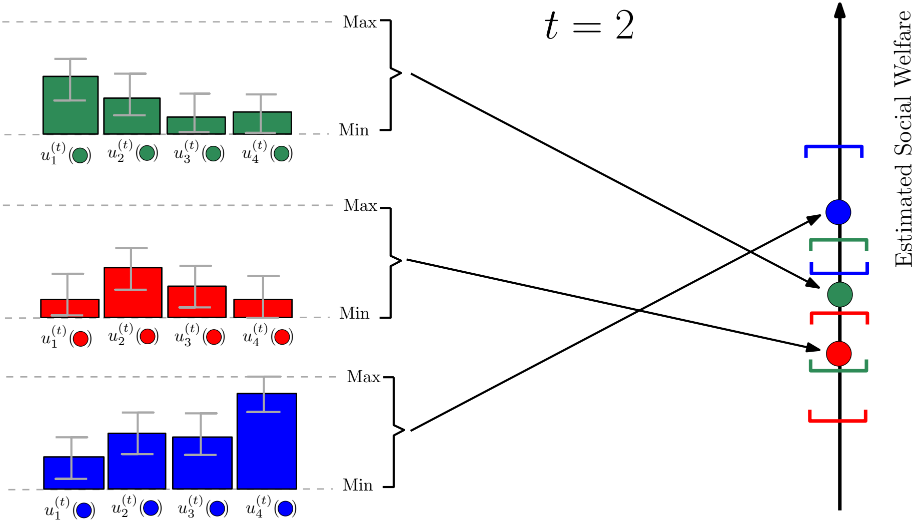

Iterated Candidate Elimination (ICE)

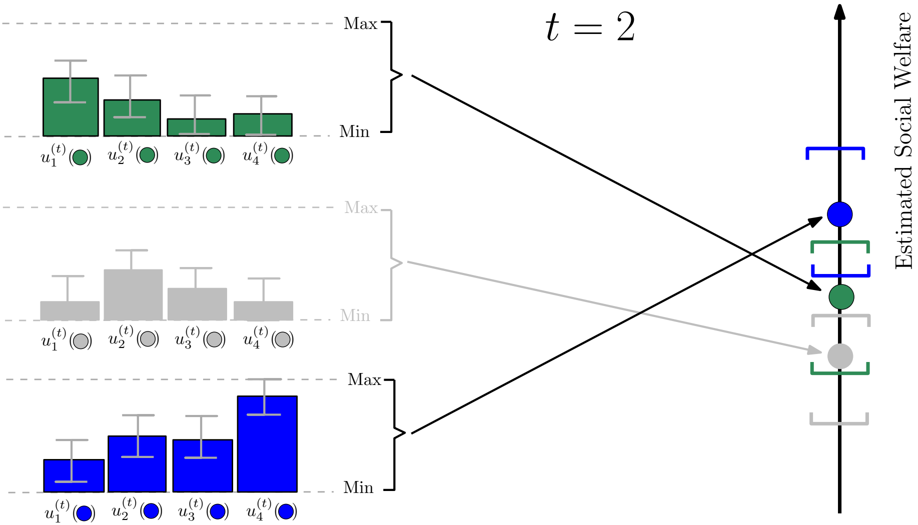

Iterated Candidate Elimination (ICE)

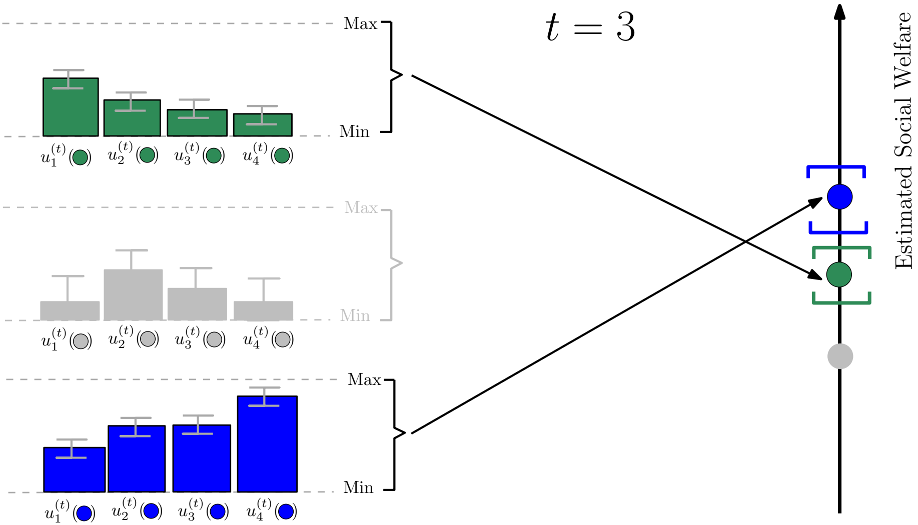

Iterated Candidate Elimination (ICE)

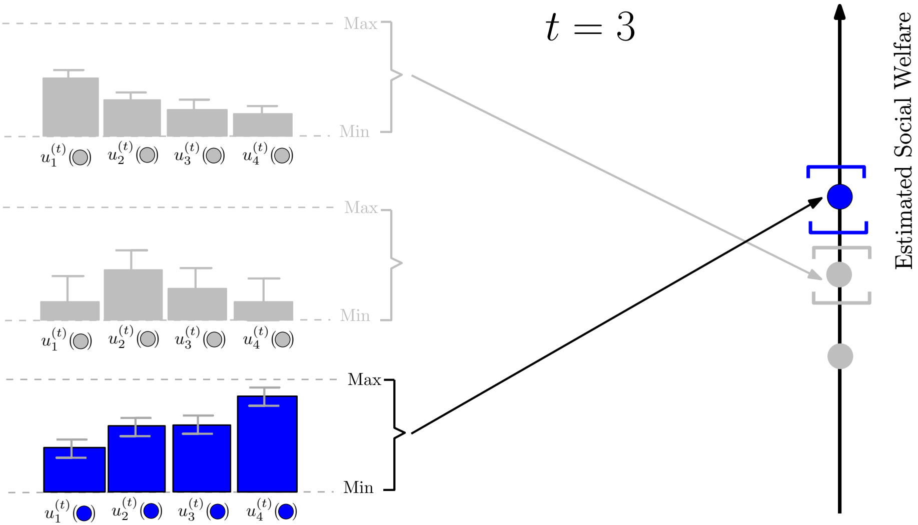

Empirical Experiments

Objectives

1. Demonstrate the importance of "empathetic" modelling.
2. Analyse the performance of our ICE algorithm.

Experimental Setup

- Generated a network by *preferential attachment* model.
- Sampled a node intrinsic preferences either from the Irish voting dataset or the impartial culture model.
- Used the Borda or plurality rule to attain intrinsic utilities.

Importance of Modelling Empathy.

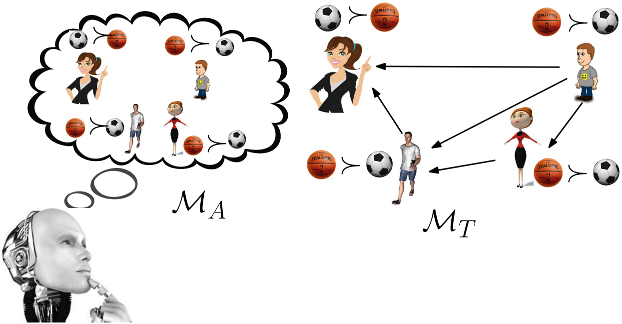

- True model $\mathcal{M}_T$ vs. Assumed model $\mathcal{M}_A$.

Importance of Modelling Empathy.

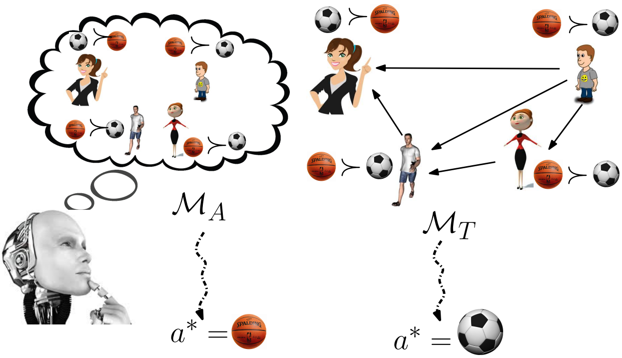

- True model $\mathcal{M}_T$ vs. Assumed model $\mathcal{M}_A$.

Importance of Modelling Empathy.

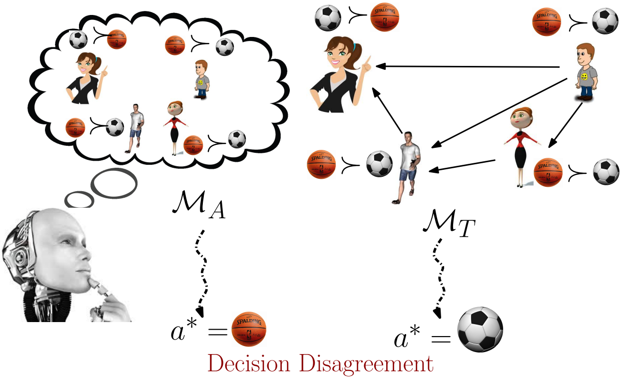

- True model $\mathcal{M}_T$ vs. Assumed model $\mathcal{M}_A$.
- % Decision Disgareement (DD) = % that true and assumed models disagree in optimal decisions.

Importance of Modelling Empathy.

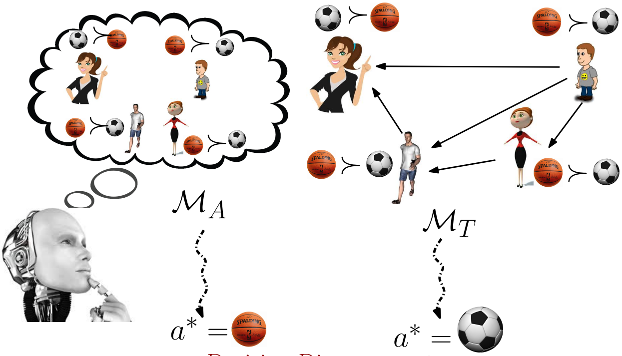

Decision Disagreement

- True model $\mathcal{M}_T$ vs. Assumed model $\mathcal{M}_A$.
- % Decision Disgareement (DD) = % that true and assumed models disagree in optimal decisions.
- Relative Social Welfare Loss (RSWL) = % of loss in social welfare due to model disagreement.

The Effect of m

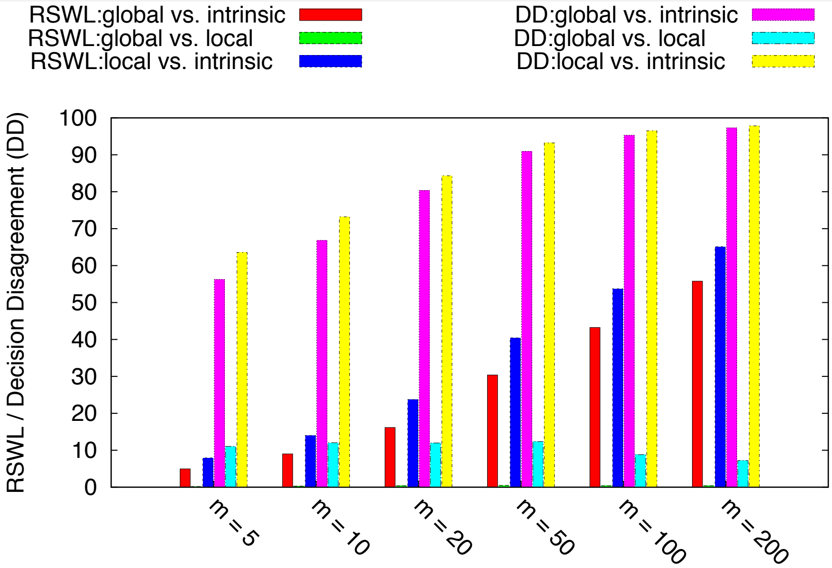

The Effect of m

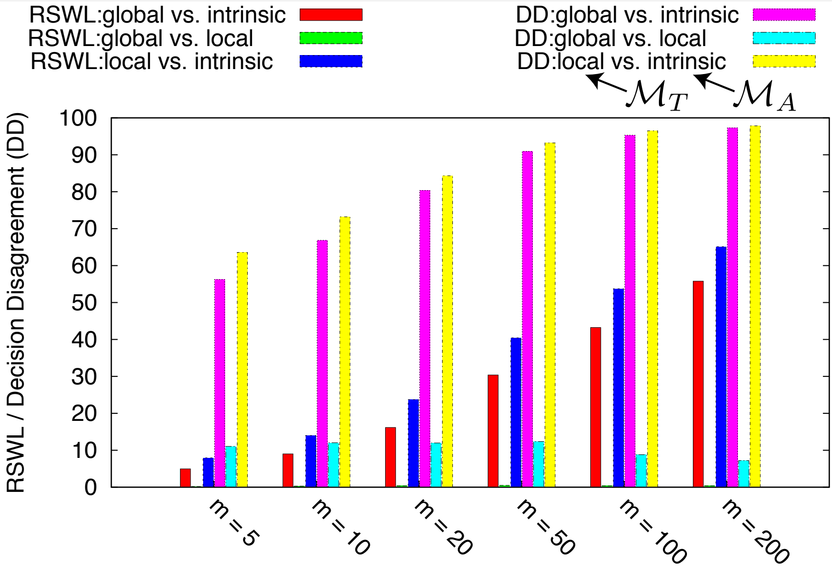

The Effect of m

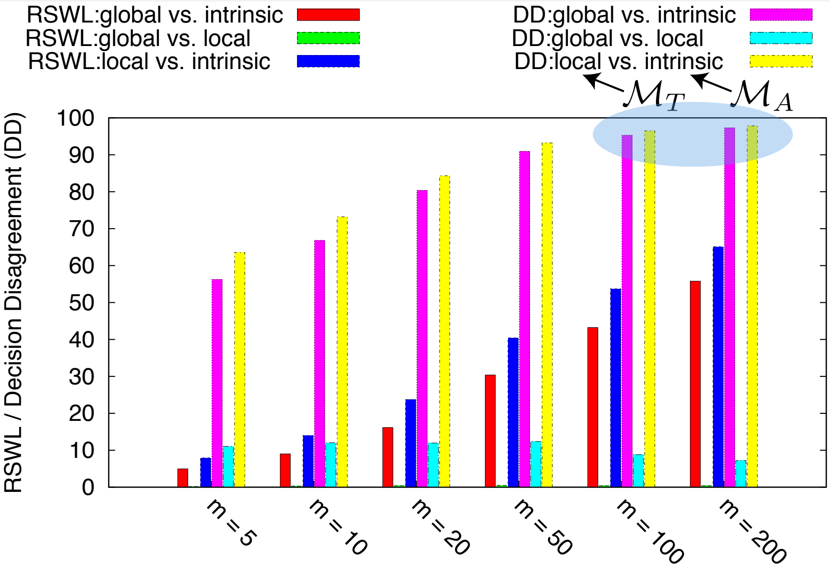

The Effect of m

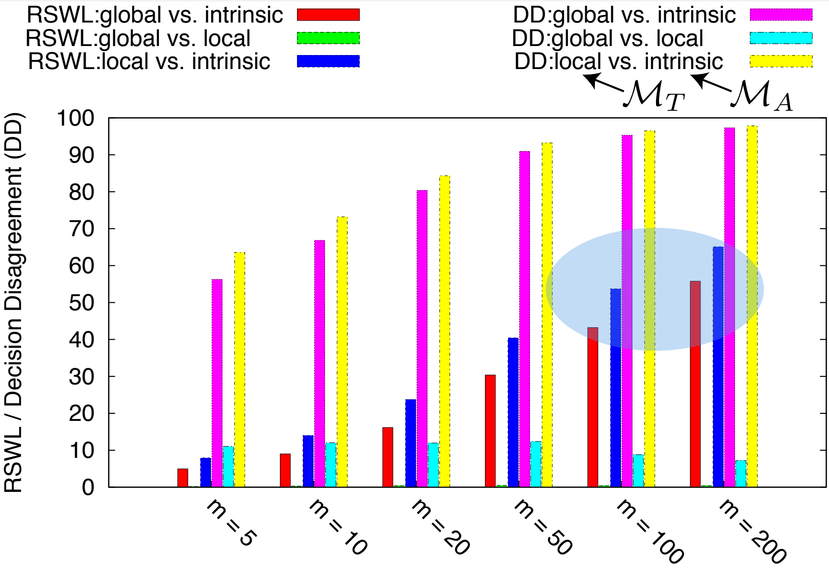

A Sample Run of the ICE Algorithm

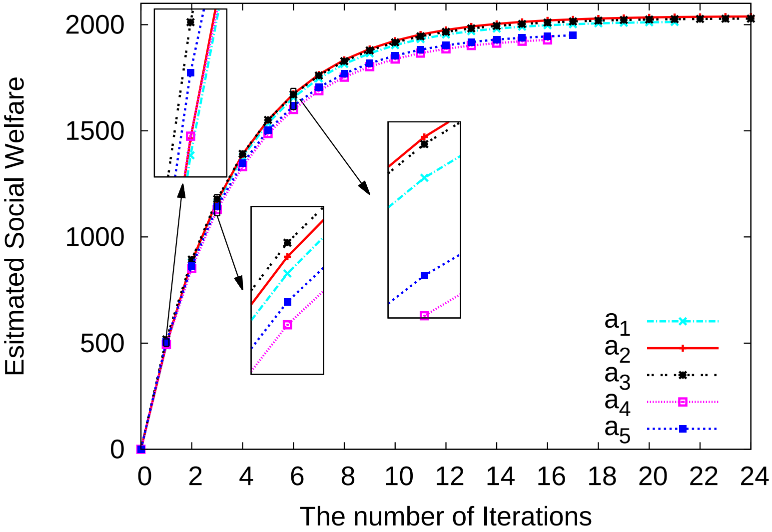

ICE Performance

- Objective: test the speed of convergence of ICE.
- Metric: average number of iterations.

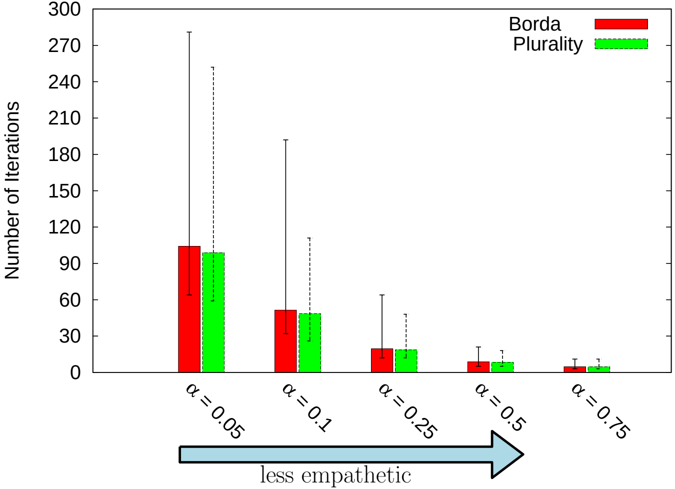

Part II:

Social Choice and Social Selection in Preferences

Ranking Networks [2]

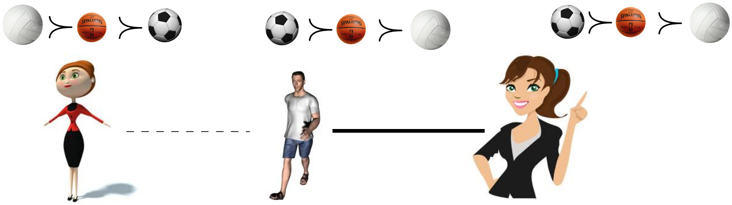

- How to capture the correlations of
  - ranking preferences over social networks?

Ranking Networks [2]

- How to capture the correlations of
  - ranking preferences over social networks?
  - "relevance ranking" or "label ranking" over information networks (e.g., www)?

Ranking Networks [2]

- How to capture the correlations of
  - ranking preferences over social networks?
  - "relevance ranking" or "label ranking" over information networks (e.g., www)?
- How to capture the correlations of *ranked data* over *networks*?

Ranking Networks [2]

- How to capture the correlations of
  - ranking preferences over social networks?
  - "relevance ranking" or "label ranking" over information networks (e.g., www)?
- How to capture the correlations of *ranked data* over *networks*?

Our Proposal: Ranking Networks

Each node possesses a ranking; the more similar two nodes' rankings are, the more likely they will be connected.

Ranking Networks

Generative Model

- m alternatives: $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- n nodes: $\mathcal{N} = \{1, \ldots, n\}$.
- r_i: i's ranking over $\mathcal{A}$.

Ranking Networks

Generative Model

- m alternatives: $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- n nodes: $\mathcal{N} = \{1, \ldots, n\}$.
- r_i: i's ranking over $\mathcal{A}$.
- $r_i \sim \rho(r)$: ranking distribution.

Ranking Networks

Generative Model

- m alternatives: $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- n nodes: $\mathcal{N} = \{1, \ldots, n\}$.
- r_i: i's ranking over $\mathcal{A}$.
- $r_i \sim \rho(r)$: ranking distribution.
- $d(r, r')$ measures their similarity.
  - E.g., *Kendall's* τ distance.

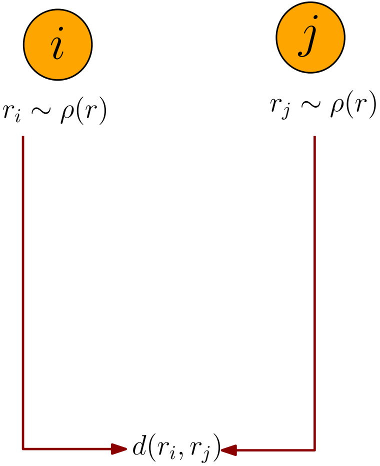

Ranking Networks

Generative Model

- m alternatives: $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- n nodes: $\mathcal{N} = \{1, \ldots, n\}$.
- r_i: i's ranking over $\mathcal{A}$.
- $r_i \sim \rho(r)$: ranking distribution.
- $d(r, r^{'})$ measures their similarity.
  - E.g., *Kendall's* τ distance.
- *Connection probability function.*
  - E.g., $c(d) = (1 + \frac{d}{\beta})^{-\alpha}$.

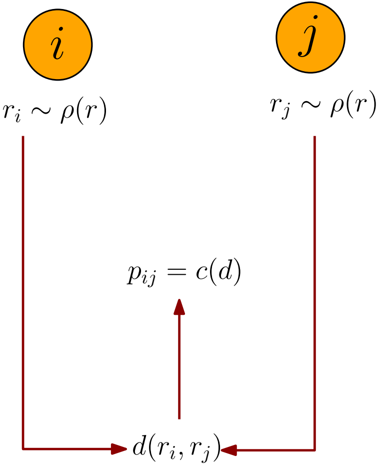

Ranking Networks

Generative Model

- m alternatives: $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- n nodes: $\mathcal{N} = \{1, \ldots, n\}$.
- r_i: i's ranking over $\mathcal{A}$.
- $r_i \sim \rho(r)$: ranking distribution.
- $d(r, r^{'})$ measures their similarity.
  - E.g., *Kendall's* τ distance.
- *Connection probability function.*
  - E.g., $c(d) = (1 + \frac{d}{\beta})^{-\alpha}$.
- Edge probability $p_{ij} = c(d(r_i, r_j))$.

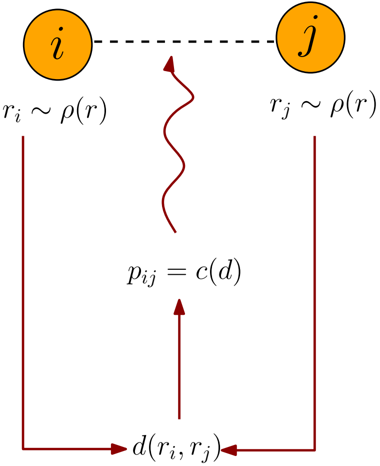

Ranking Networks

Generative Model

- m alternatives: $\mathcal{A} = \{a_1, \ldots, a_m\}$.
- n nodes: $\mathcal{N} = \{1, \ldots, n\}$.
- r_i: i's ranking over $\mathcal{A}$.
- $r_i \sim \rho(r)$: ranking distribution.
- $d(r, r^{'})$ measures their similarity.
  - E.g., *Kendall's* τ distance.
- *Connection probability function.*
  - E.g., $c(d) = (1 + \frac{d}{\beta})^{-\alpha}$.
- Edge probability $p_{ij} = c(d(r_i, r_j))$.
- $e_{ij} \sim Bernoulli(p_{ij})$

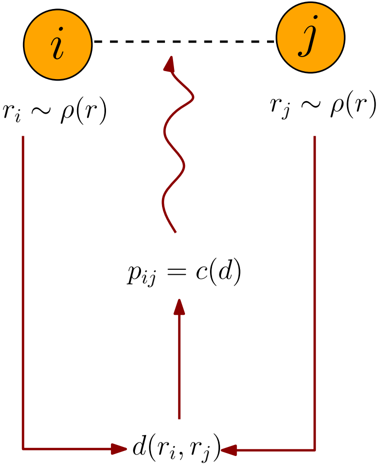

Ranking Networks

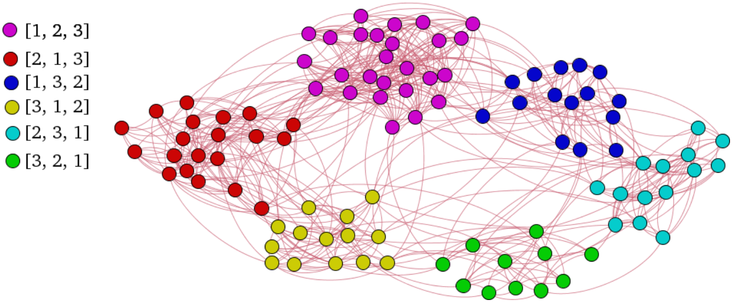

Ranking Networks

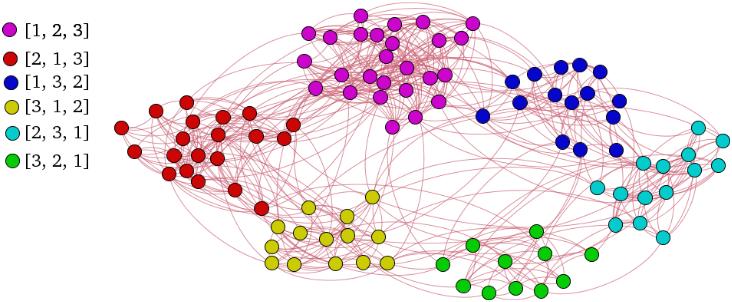

Applications

- Ranking networks help predict the missing preferences.

Ranking Networks

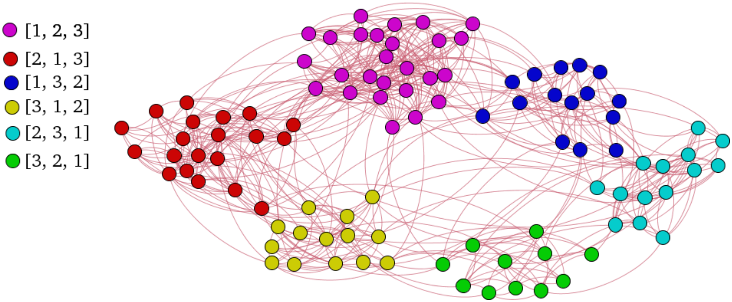

Applications

- Ranking networks help predict the missing preferences.
  - Efficient recommender and advertising mechanisms.

Ranking Networks

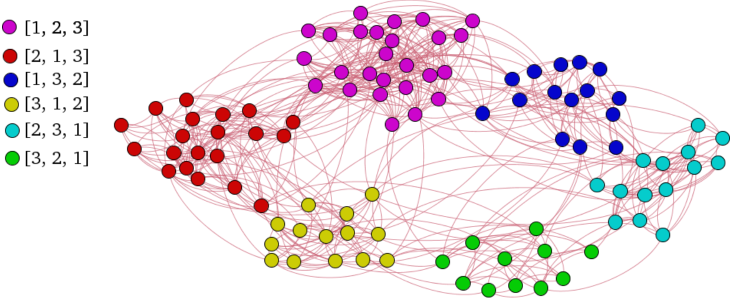

Applications

- Ranking networks help predict the missing preferences.
  - Efficient recommender and advertising mechanisms.
- Application in information retrieval and topic modeling.

Ranking Networks

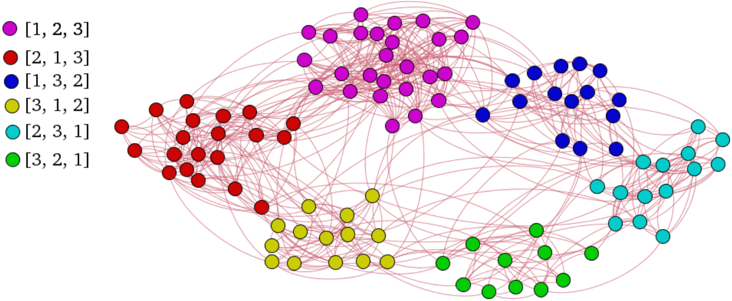

Applications

- Ranking networks help predict the missing preferences.
  - Efficient recommender and advertising mechanisms.
- Application in information retrieval and topic modeling.
- Useful for *link prediction* and *link mining*.

Ranking Networks

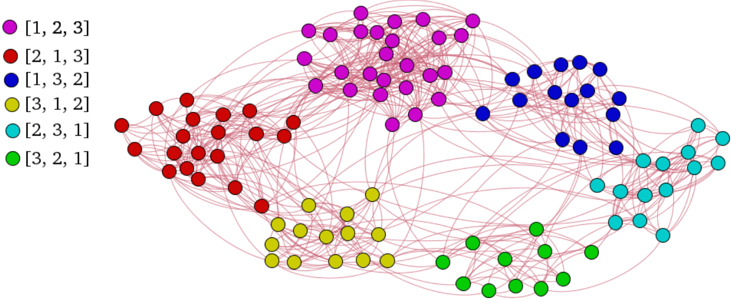

Applications

- Ranking networks help predict the missing preferences.
  - Efficient recommender and advertising mechanisms.
- Application in information retrieval and topic modeling.
- Useful for *link prediction* and *link mining*.

Overview of Group Rec. with Partial Preferences [3]

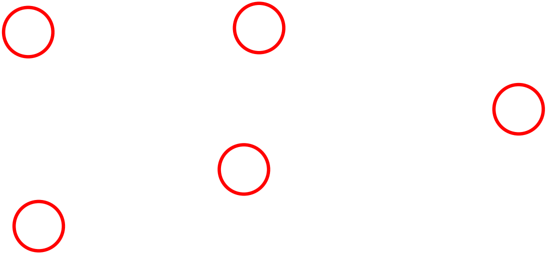

- Make a decision for a group of individuals.

Overview of Group Rec. with Partial Preferences [3]

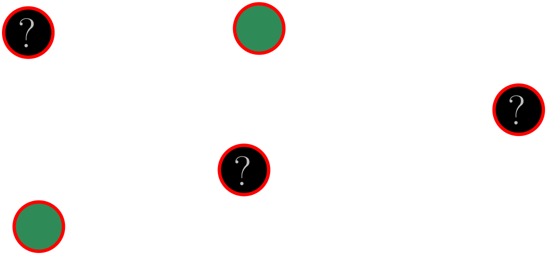

- Make a decision for a group of individuals.
- Some preferences are **missing**.

Overview of Group Rec. with Partial Preferences [3]

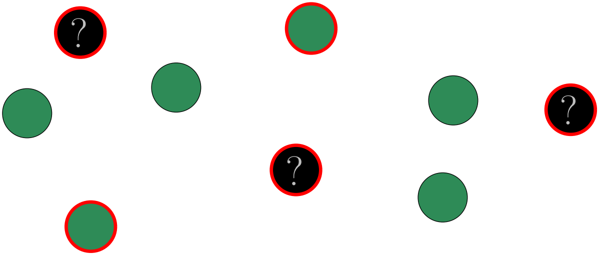

- Make a decision for a group of individuals.
- Some preferences are **missing**.
- Infer missing preferences given
  - Observed individual preferences (inside+outside of group).

Overview of Group Rec. with Partial Preferences [3]

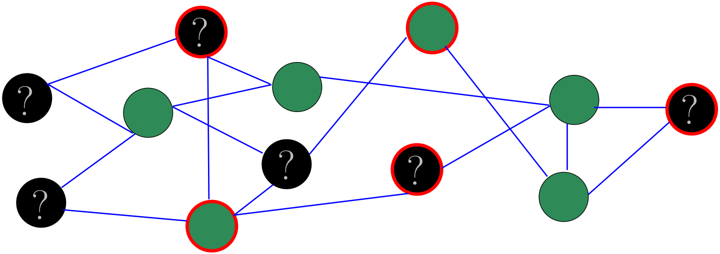

- Make a decision for a group of individuals.
- Some preferences are **missing**.
- Infer missing preferences given
  - Observed individual preferences (inside+outside of group).
  - Observed social network structure.

Overview of Group Rec. with Partial Preferences [3]

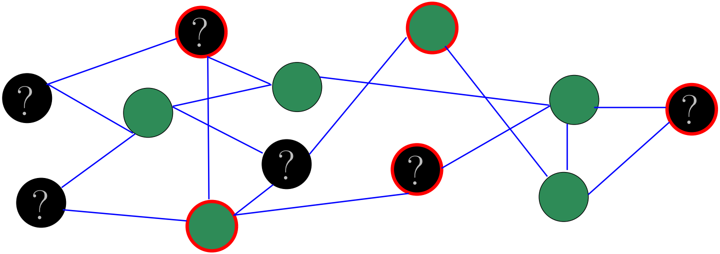

- Make a decision for a group of individuals.
- Some preferences are **missing**.
- Infer missing preferences given
  - Observed individual preferences (inside+outside of group).
  - Observed social network structure.
- Make a decision based on inferred preferences.

Capturing Preference Correlations by POSN

Preference-Oriented Social Network (POSN) Model

- A special instance of ranking networks.
- Exploited for effective group recommendation.

Joint and Conditional Distributions under POSN

- Joint distribution over rankings $R = [r_i]$ and network $\mathbf{E} = [e_{ij}]$:

$$\mathbb{P}(R, \mathbf{E}|\boldsymbol{\theta}) = \prod_i \rho(r_i|\boldsymbol{\eta}) \prod_{j<i} \mathbb{P}(e_{ij}|r_i, r_j, \boldsymbol{\lambda})$$

Graphical Representation

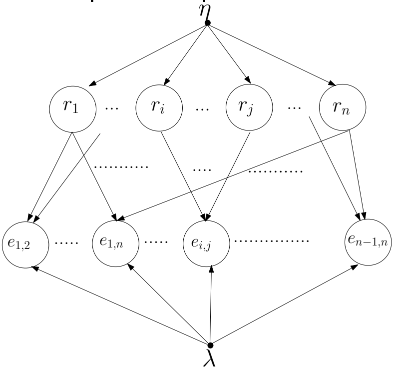

Joint and Conditional Distributions under POSN

- Joint distribution over rankings $R = [r_i]$ and network $\mathbf{E} = [e_{ij}]$:

$$\mathbb{P}(R, \mathbf{E}|\boldsymbol{\theta}) = \prod_i \rho(r_i|\boldsymbol{\eta}) \prod_{j<i} \mathbb{P}(e_{ij}|r_i, r_j, \boldsymbol{\lambda})$$

- R^O: Observed rankings.
- $\mathbf{E}$: Observed social network.
- R^U: Unobserved rankings.
- Conditional distribution:

$$\mathbb{P}(R^U|R^O, \mathbf{E}, \boldsymbol{\eta}, \boldsymbol{\lambda}).$$

Graphical Representation

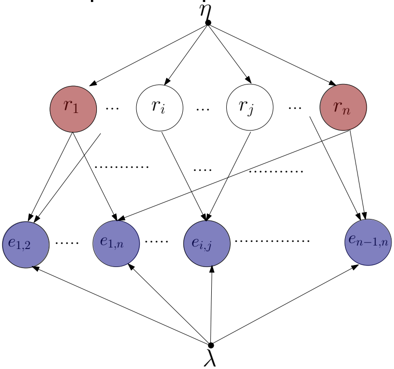

Social Welfare and Missing Preferences

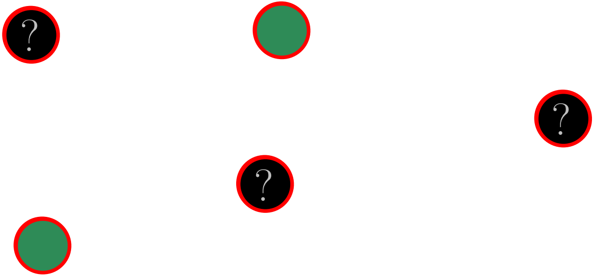

Social Welfare and Missing Preferences

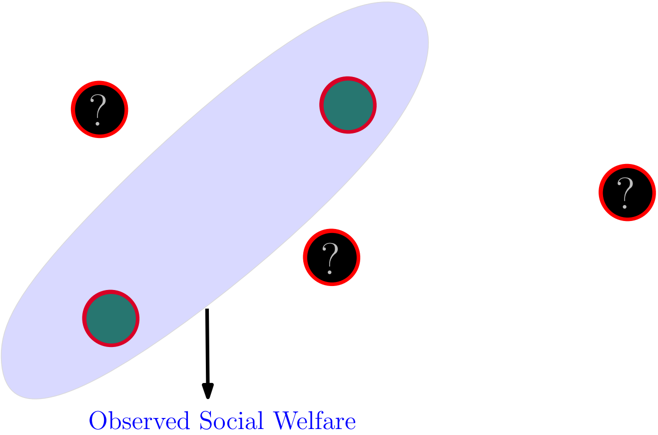

Social Welfare and Missing Preferences

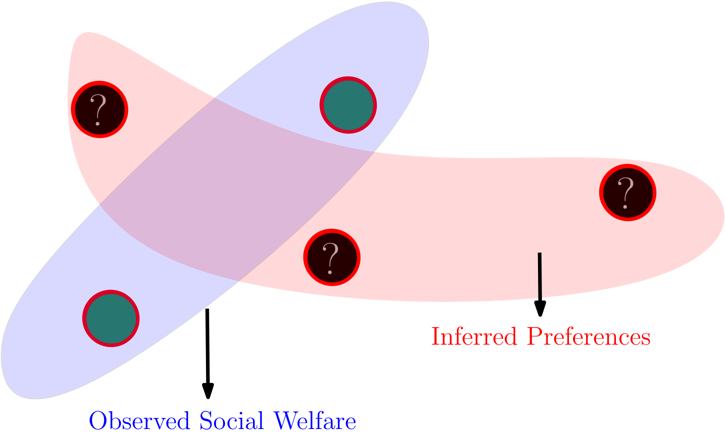

Social Welfare and Missing Preferences

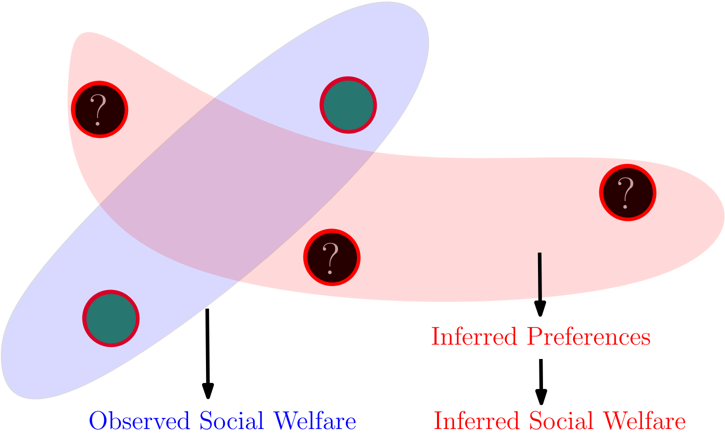

Social Welfare and Missing Preferences

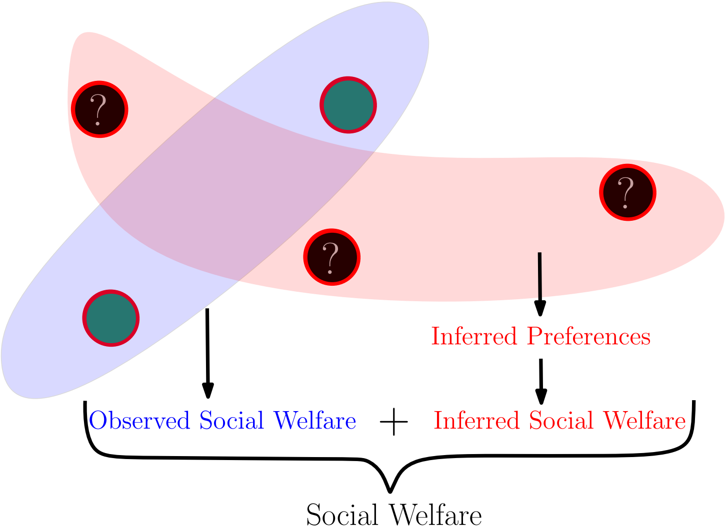

Inference and Sampling

Inference

- Must compute $\mathbb{P}(R^U|\mathbf{E}, R^O, \boldsymbol{\theta})$ for our inference methods.

Inference and Sampling

Inference

- Must compute $\mathbb{P}(R^U|\mathbf{E}, R^O, \boldsymbol{\theta})$ for our inference methods.
- Exact computation is computationally expensive: $O((m!)^{|U|})$.

Inference and Sampling

Inference

- Must compute $\mathbb{P}(R^U | \mathbf{E}, R^O, \boldsymbol{\theta})$ for our inference methods.
- Exact computation is computationally expensive: $O((m!)^{|U|})$.
- Developed sampling methods for approximation.

Inference and Sampling

Inference

- Must compute $\mathbb{P}(R^U|\mathbf{E}, R^O, \boldsymbol{\theta})$ for our inference methods.
- Exact computation is computationally expensive: $O((m!)^{|U|})$.
- Developed sampling methods for approximation.

High-level Idea

- Sample L preference profiles $R^{(1)}, \ldots, R^{(L)}$ from the posterior.
  - Each preference profile consists of all unobserved individuals' preference rankings.

Inference and Sampling

Inference

- Must compute $\mathbb{P}(R^U|\mathbf{E}, R^O, \boldsymbol{\theta})$ for our inference methods.
- Exact computation is computationally expensive: $O((m!)^{|U|})$.
- Developed sampling methods for approximation.

High-level Idea

- Sample L preference profiles $R^{(1)}, \ldots, R^{(L)}$ from the posterior.
  - Each preference profile consists of all unobserved individuals' preference rankings.
- $\mathbb{P}(R^U = R'|\mathbf{A}, R^O, \boldsymbol{\theta}) \approx$ the fraction of number of times that preference profile R' is sampled.

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

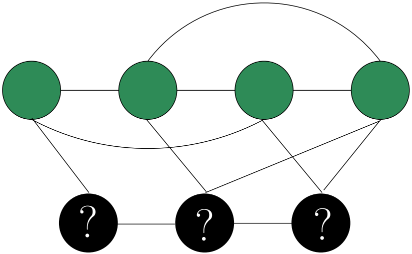

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

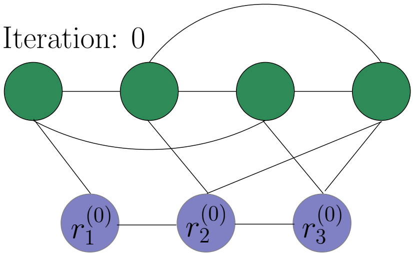

Iterative Metropolis-within-Gibbs Sampling Algorithm

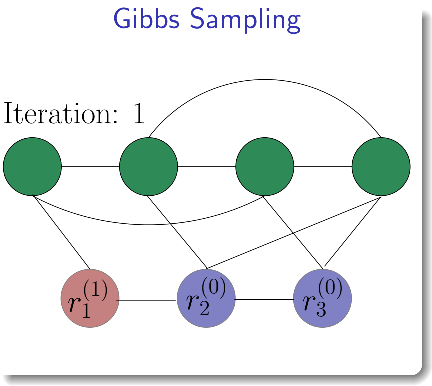

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

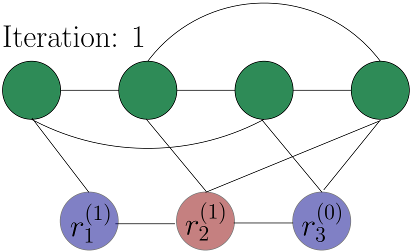

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

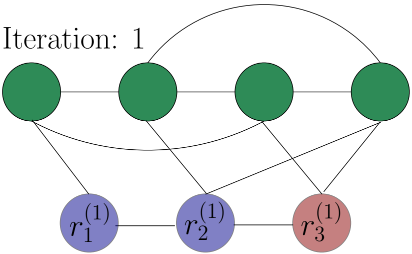

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

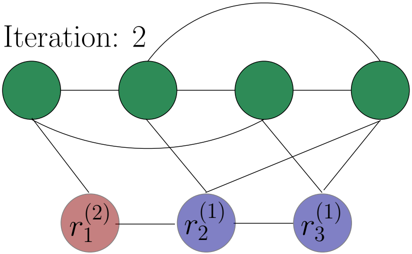

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

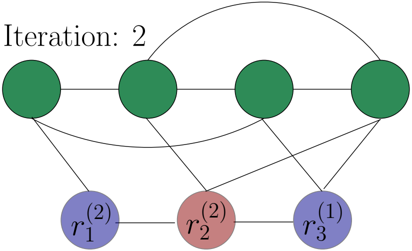

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

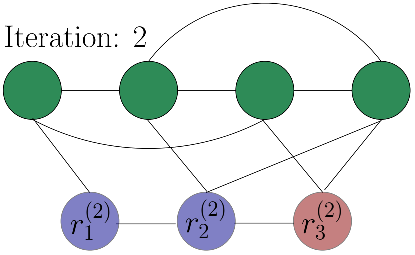

Metropolis Sampling

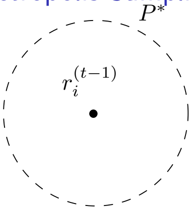

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

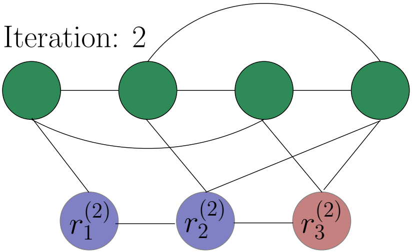

Metropolis Sampling

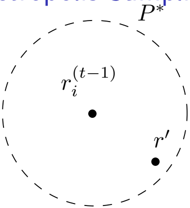

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

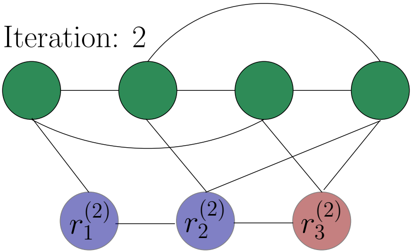

Metropolis Sampling

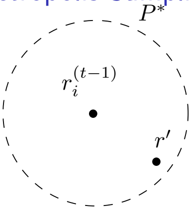

- $q = \dfrac{\mathbb{P}(r')}{\mathbb{P}\left(r_i^{(t-1)}\right)}$

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

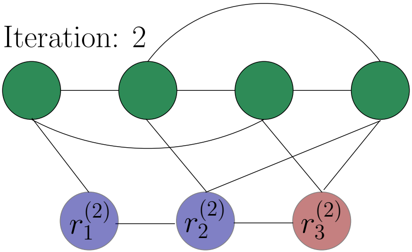

Metropolis Sampling

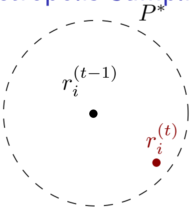

- $q = \dfrac{\mathbb{P}(r')}{\mathbb{P}\left(r_i^{(t-1)}\right)}$
- $q \geq 1 \, : \; r_i^{(t)} \leftarrow r'$

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

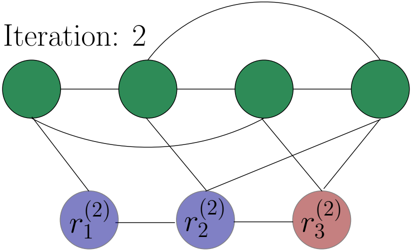

Metropolis Sampling

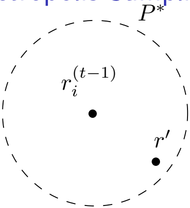

- $q = \frac{\mathbb{P}(r')}{\mathbb{P}\left(r_i^{(t-1)}\right)}$
- $q \geq 1: \; r_i^{(t)} \leftarrow r'$
- $q < 1:$

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

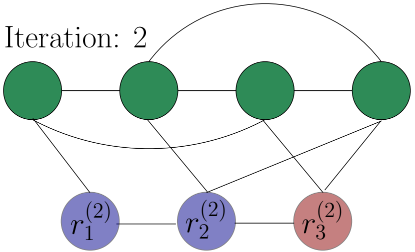

Metropolis Sampling

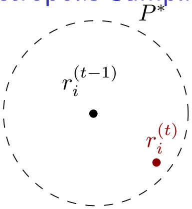

- $q = \frac{\mathbb{P}(r')}{\mathbb{P}\left(r_i^{(t-1)}\right)}$
- $q \geq 1 : \; r_i^{(t)} \leftarrow r'$
- $q < 1$: with prob. q, set $r_i^{(t)} \leftarrow r'$

Iterative Metropolis-within-Gibbs Sampling Algorithm

Gibbs Sampling

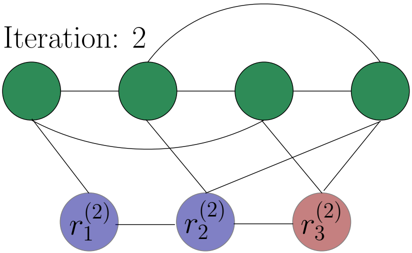

Metropolis Sampling

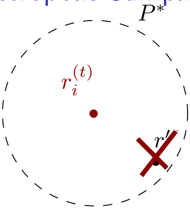

- $q = \dfrac{\mathbb{P}(r')}{\mathbb{P}\left(r_i^{(t-1)}\right)}$
- $q \geq 1 : \; r_i^{(t)} \leftarrow r'$
- $q < 1$: with prob. q, set $r_i^{(t)} \leftarrow r'$, otherwise $r_i^{(t)} \leftarrow r_i^{(t-1)}$

Empirical Analysis

Experiment Goal

To assess the effectiveness of our group recommendation and inference algorithms using various data sets.

Outline

- Experimental Setup:
  - Preference Observability.
  - Group Selection Strategies.
  - Performance Metrics.
  - Benchmarks.
- Results: Irish Data, Flixster Data.

Experimental Setup

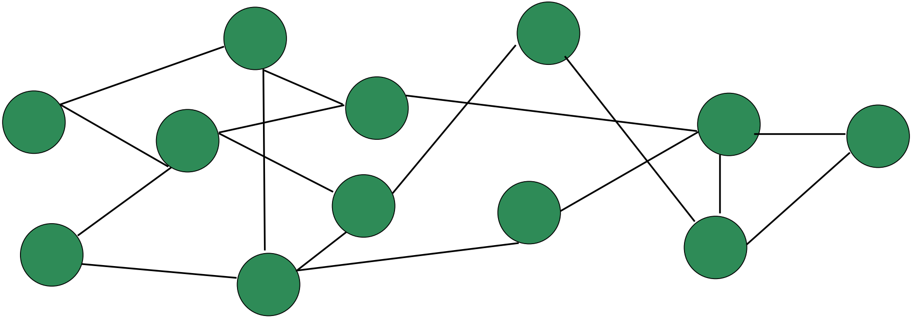

Preference Observability

- Varying the degree to which preferences are observed with $\psi \in [0, 1]$, the probability that a node's ranking is observed.

Experimental Setup

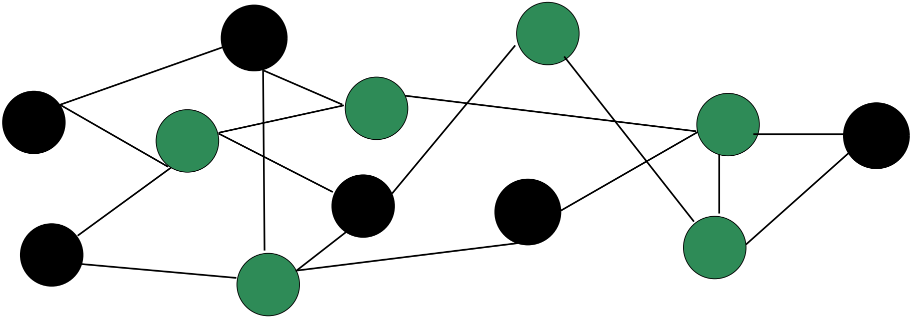

Preference Observability

- Varying the degree to which preferences are observed with $\psi \in [0, 1]$, the probability that a node's ranking is observed.

Experimental Setup

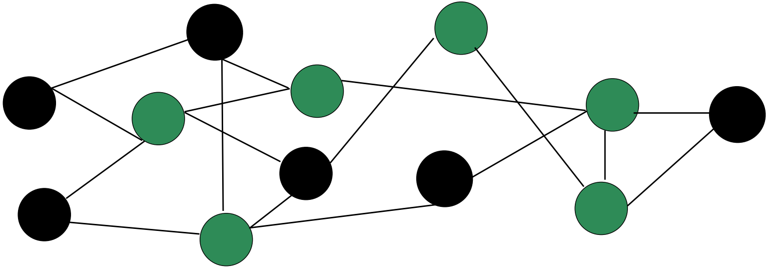

Group Selection Strategies

- Select group $\mathcal{S}$ (with n_s members) by one of

Experimental Setup

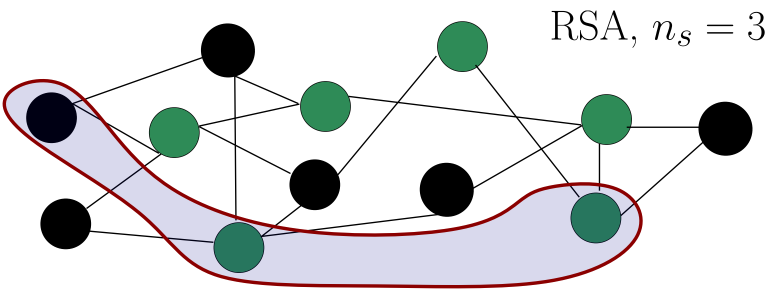

Group Selection Strategies

- Select group $\mathcal{S}$ (with n_s members) by one of
  - *RSA (Random Selection from All).*

Experimental Setup

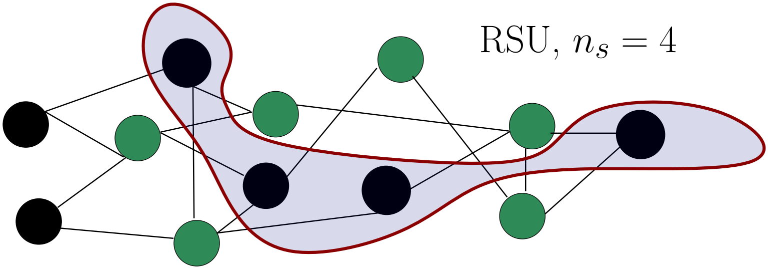

Group Selection Strategies

- Select group $\mathcal{S}$ (with n_s members) by one of
  - *RSA (Random Selection from All).*
  - *RSU (Random Selection from Unobserved).*

Experimental Setup

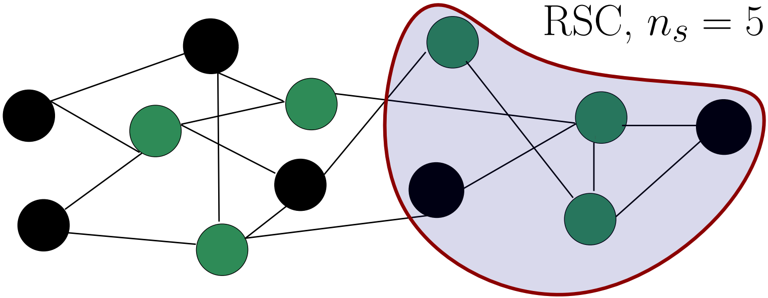

Group Selection Strategies

- Select group $\mathcal{S}$ (with n_s members) by one of
  - *RSA (Random Selection from All).*
  - *RSU (Random Selection from Unobserved).*
  - *RSC (Random Selection from Community).*

Experimental Setup

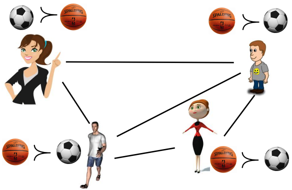

Performance Metrics: Measuring Recommendation Quality

- o^*, o^*_{inf}: optimal under true preferences, inference method.
- $sw(\cdot)$: social welfare with true preferences.
- *Relative social welfare loss (RSWL)*: $[sw(o^*) - sw(o^*_{inf})]/sw(o^*)$

Experimental Setup

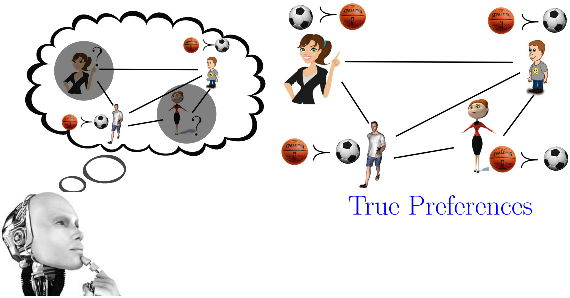

Performance Metrics: Measuring Recommendation Quality

- o^*, o^*_{inf}: optimal under true preferences, inference method.
- $sw(\cdot)$: social welfare with true preferences.
- *Relative social welfare loss (RSWL)*: $[sw(o^*) - sw(o^*_{inf})]/sw(o^*)$

Experimental Setup

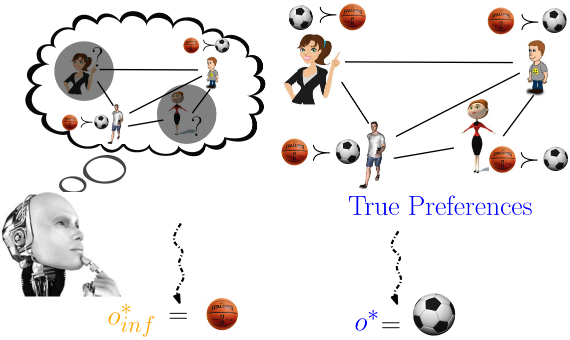

Performance Metrics: Measuring Recommendation Quality

- o^*, o^*_{inf}: optimal under true preferences, inference method.
- $sw(\cdot)$: social welfare with true preferences.
- *Relative social welfare loss (RSWL)*: $[sw(o^*) - sw(o^*_{inf})]/sw(o^*)$

Experimental Setup

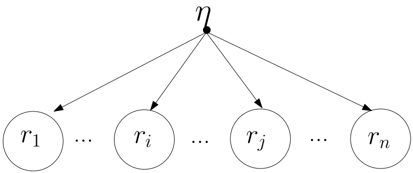

Group Decision Making Benchmarks.

- Considered other ways of dealing with missing preferences:

Experimental Setup

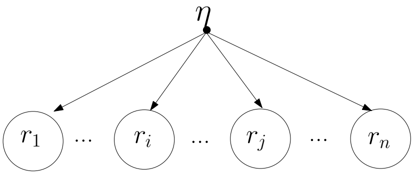

Group Decision Making Benchmarks.

- Considered other ways of dealing with missing preferences:
- ϕ-Mallows inference (PM): all unobserved preferences are independently drawn from a ϕ-Mallows model.

Experimental Setup

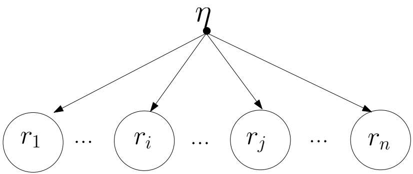

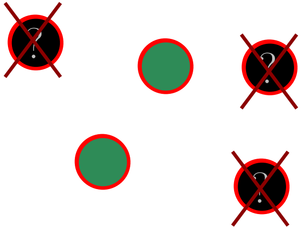

Group Decision Making Benchmarks.

- Considered other ways of dealing with missing preferences:
- ϕ-Mallows inference (PM): all unobserved preferences are independently drawn from a ϕ-Mallows model.
- Discard Unobserved (DU): ignore unobserved preferences and make a decision using only observed preferences.

Irish Data

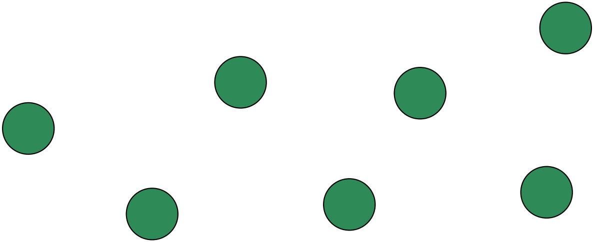

Irish Data

Irish Data

- Real-world preferences, Dublin West, 2002 Irish Election.

Irish Data

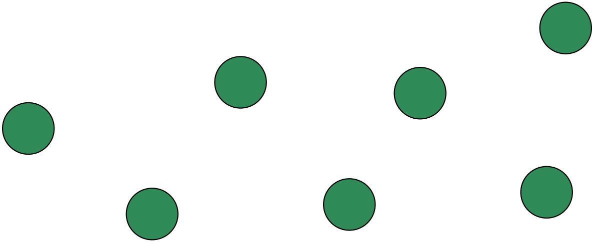

Irish Data

Irish Data

- Real-world preferences, Dublin West, 2002 Irish Election.
- No social network data in Irish Data.

Irish Data

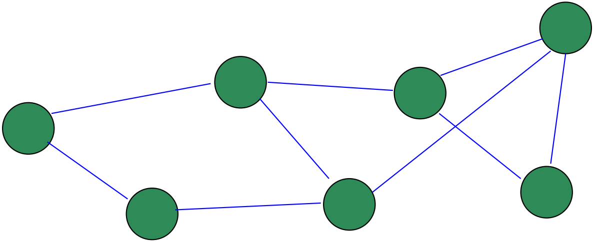

POSN

Irish Data

- Real-world preferences, Dublin West, 2002 Irish Election.
- No social network data in Irish Data.
- For each experiment, generate 10 partially observed POSNs

Irish Data

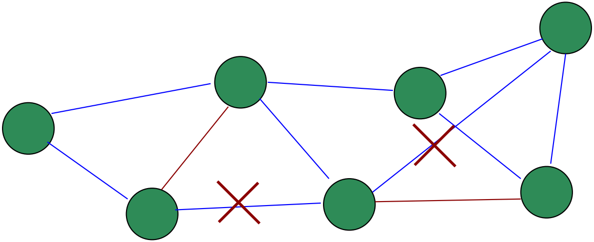

POSN + Noise

Irish Data

- Real-world preferences, Dublin West, 2002 Irish Election.
- No social network data in Irish Data.
- For each experiment, generate 10 partially observed POSNs with additional noise: randomly delete or add an edge with probability ε.

Irish Data

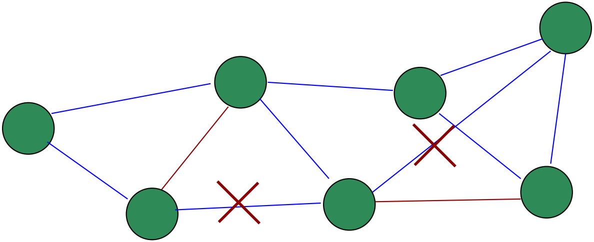

POSN + Noise

Irish Data

- Real-world preferences, Dublin West, 2002 Irish Election.
- No social network data in Irish Data.
- For each experiment, generate 10 partially observed POSNs with additional noise: randomly delete or add an edge with probability ε.
  - Reflecting scenarios when learned model parameters provide a less-than-ideal fit to the underlying data.

Highlight of Results: Irish Data

RSA

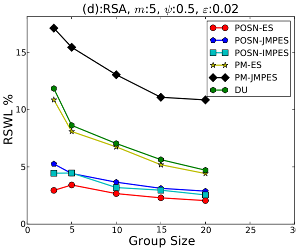

RSU

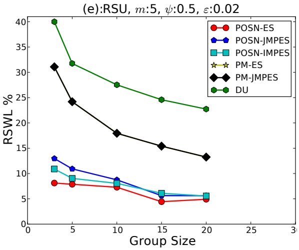

RSC

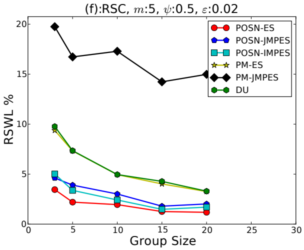

Avg. RSWL for various n_s, group selection methods, but fixed $m = 5$, $\psi = 0.5$, $\varepsilon = 0.02$.

Highlight of Results: Irish Data

RSA

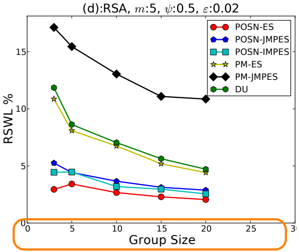

RSU

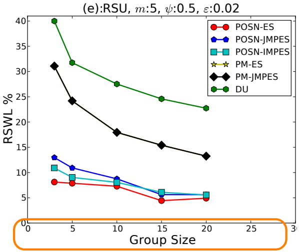

RSC

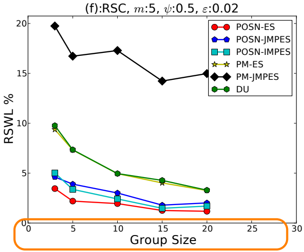

Avg. RSWL for various n_s, group selection methods, but fixed $m = 5$, $\psi = 0.5$, $\varepsilon = 0.02$.

Highlight of Results: Irish Data

RSA RSU RSC

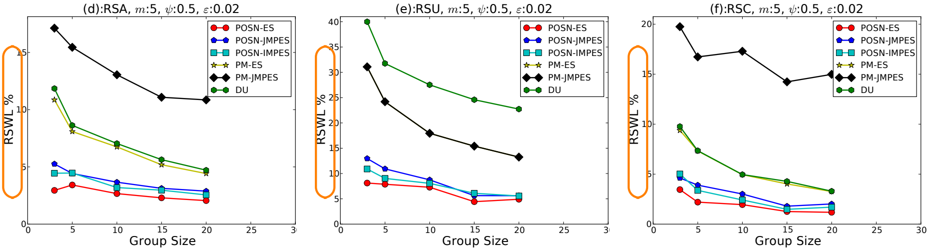

Avg. RSWL for various n_s, group selection methods, but fixed $m = 5$, $\psi = 0.5$, $\varepsilon = 0.02$.

Highlight of Results: Irish Data

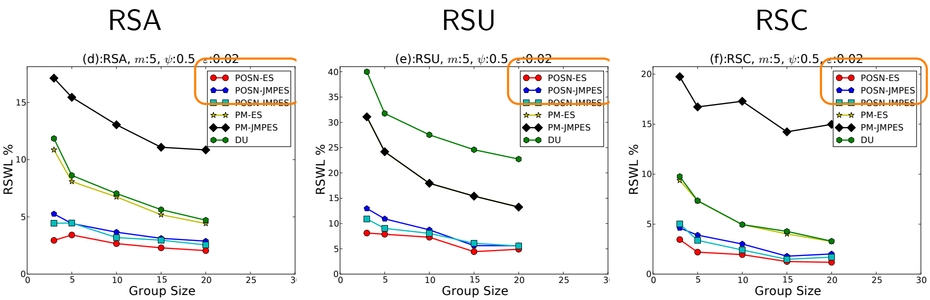

Avg. RSWL for various n_s, group selection methods, but fixed $m = 5$, $\psi = 0.5$, $\varepsilon = 0.02$.

Highlight of Results: Irish Data

RSA

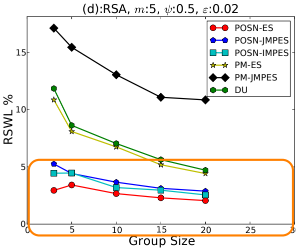

RSU

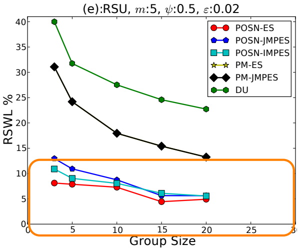

RSC

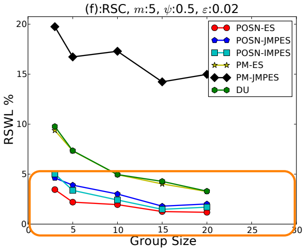

Avg. RSWL for various n_s, group selection methods, but fixed $m = 5$, $\psi = 0.5$, $\varepsilon = 0.02$.

Flixster

Flixster Data and Experimental Setup

- Flixster data (Jamali & Ester): social network, movie ratings.

Flixster

Flixster Data and Experimental Setup

- Flixster data (Jamali & Ester): social network, movie ratings.
- Crawled movie genres from IMDB.

Flixster

Flixster Data and Experimental Setup

- Flixster data (Jamali & Ester): social network, movie ratings.
- Crawled movie genres from IMDB.
- No preference ranking and sparsity of movie ratings:

Flixster

Flixster Data and Experimental Setup

- Flixster data (Jamali & Ester): social network, movie ratings.
- Crawled movie genres from IMDB.
- No preference ranking and sparsity of movie ratings:
    - Aggregated ratings to genre scores, then into preference ranking over movie genres.

Flixster

Flixster Data and Experimental Setup

- Flixster data (Jamali & Ester): social network, movie ratings.
- Crawled movie genres from IMDB.
- No preference ranking and sparsity of movie ratings:
    - Aggregated ratings to genre scores, then into preference ranking over movie genres.
    - 4 diverse genres: Comedy, Drama, Kids/Family, Mystery.

Flixster

Flixster Data and Experimental Setup

- Flixster data (Jamali & Ester): social network, movie ratings.
- Crawled movie genres from IMDB.
- No preference ranking and sparsity of movie ratings:
    - Aggregated ratings to genre scores, then into preference ranking over movie genres.
    - 4 diverse genres: Comedy, Drama, Kids/Family, Mystery.
- Tested our methods on 10 partially observed network, $\psi = 0.5$

Highlight of Results: Flixster Data

Results Summary

- POSN outperforms the ϕ-Mallows benchmarks.
- POSN outperforms DU for small groups but performs comparably for large groups (due to Homophily)

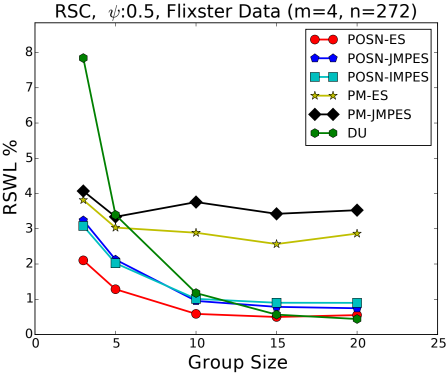

Highlight of Results: Flixster Data

Results Summary

- POSN outperforms the ϕ-Mallows benchmarks.
- POSN outperforms DU for small groups but performs comparably for large groups (due to Homophily)

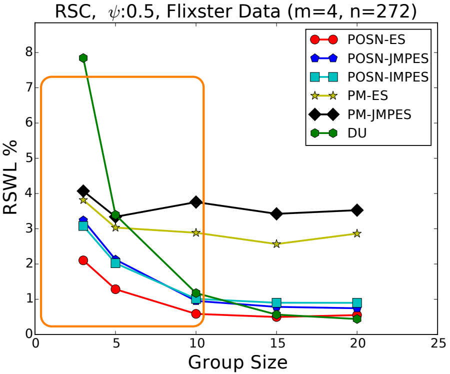

Highlight of Results: Flixster Data

Results Summary

- POSN outperforms the ϕ-Mallows benchmarks.
- POSN outperforms DU for small groups but performs comparably for large groups (due to Homophily)

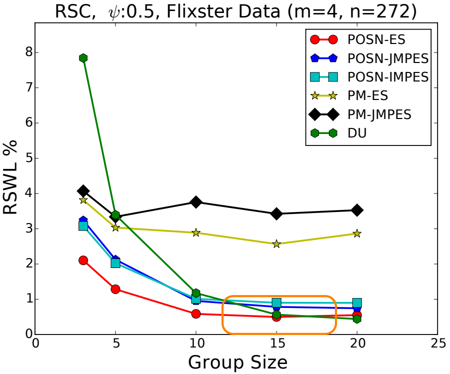

Amirali Salehi-Abari and Craig Boutilier.
Empathetic social choice on social networks.
In *Proceedings of The 13th International Conference on Autonomous Agents and Multiagent Systems*, pages 693–700, 2014.

Amirali Salehi-Abari and Craig Boutilier.
Ranking networks.
In *NIPS-13 Workshop on Frontiers of Network Analysis: Methods, Models, and Applications*, 2013.

Amirali Salehi-Abari and Craig Boutilier.
Preference-oriented social networks: Group recommendation and inference.
In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys-15)*, 2015.