Inferring a Status Network of Wikipedia Editors

Andrew Wang

10 August 2017

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Outline

Outline

Motivation

- How is team **structure** related to team **performance**?

Previous Work

- *Cummings and Cross* 2003 — Reply networks of work groups in a Fortune 500 telecommunications firm

Previous Work

- *Cummings and Cross* 2003 — Reply networks of work groups in a Fortune 500 telecommunications firm
- *Brandes et al.* 2009 — Wikipedia edit network: who supports or opposes whose edits

Our Contribution

- How is the **balance of status** in a team related to **performance**?

Outline

Method

- Take a **team** to be the set of editors who work on a particular Wikipedia article

Method

- Take a **team** to be the set of editors who work on a particular Wikipedia article
- Call a team **productive** if it works on a "featured" or "good" article

Method

- Take a **team** to be the set of editors who work on a particular Wikipedia article
- Call a team **productive** if it works on a "featured" or "good" article
- Construct a **status network** for each team

Method

- Take a **team** to be the set of editors who work on a particular Wikipedia article
- Call a team **productive** if it works on a "featured" or "good" article
- Construct a **status network** for each team
- See if these networks are systematically different between productive and regular teams

Method

- Construct a **status network** for each team

Status Networks

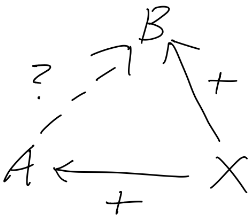

Status Networks

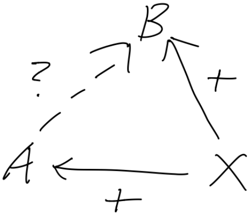

- X has rated B positively, **so B is higher status than average**, so A is more likely to rate B positively than he is to rate a random user positively

Status Networks

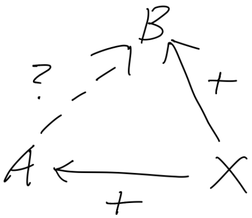

- X has rated B positively, **so B is higher status than average**, so A is more likely to rate B positively than he is to rate a random user positively
  - Generative model

Status Networks

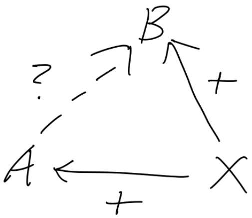

- X has rated B positively, **so B is higher status than average**, so A is more likely to rate B positively than he is to rate a random user positively
  - Generative model
- X has rated A positively, **so A is higher status than average**, so B is less likely to receive a positive rating from A than he is to receive a positive rating from a random user

Status Networks

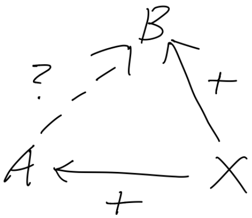

- X has rated B positively, **so B is higher status than average**, so A is more likely to rate B positively than he is to rate a random user positively
  - Generative model
- X has rated A positively, **so A is higher status than average**, so B is less likely to receive a positive rating from A than he is to receive a positive rating from a random user
  - Receptive model

Status Networks

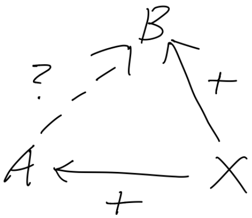

- X has rated B positively, **so B is higher status than average**, so A is more likely to rate B positively than he is to rate a random user positively
  - Generative model
- X has rated A positively, **so A is higher status than average**, so B is less likely to receive a positive rating from A than he is to receive a positive rating from a random user
  - Receptive model
- *(Leskovec et al. 2010)*

Status Networks

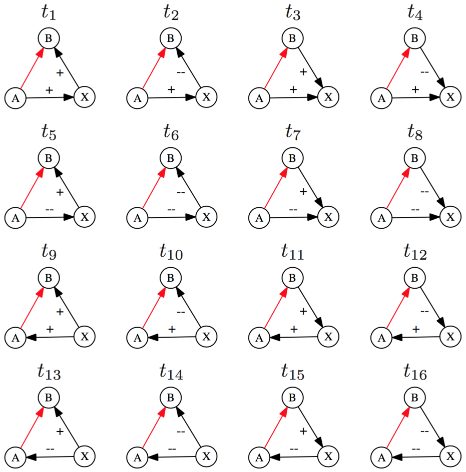

Ingredients for a Status Network

- Nodes represent people

Ingredients for a Status Network

- Nodes represent people
- Signed directed edges represent status ratings

Ingredients for a Status Network

- Nodes represent people
- Signed directed edges represent status ratings
- Triangles should follow generative and receptive predictions

How Do I Get One?

- RfA voting network *(11,381 nodes, 189,001 edges)*

How Do I Get One?

- RfA voting network *(11,381 nodes, 189,001 edges)*
- Coordination network *(48,079 nodes, 210,248 edges)*

How Do I Get One?

- Coordination network *(48,079 nodes, 210,248 edges)*

Coordination

- *Danescu-Niculescu-Mizil et al.* 2012: People subconsciously mirror each other's grammar words in conversations (linguistic coordination)

Coordination

- *Danescu-Niculescu-Mizil et al.* 2012: People subconsciously mirror each other's grammar words in conversations (linguistic coordination)
- Mirroring increases when we talk to people of high status

Coordination

- *Danescu-Niculescu-Mizil et al.* 2012: People subconsciously mirror each other's grammar words in conversations (linguistic coordination)
- Mirroring increases when we talk to people of high status
- $C(a, b) = P(E_{u_a \to u_b} | E_{u_b}) - P(E_{u_a \to u_b})$

Coordination

- *Danescu-Niculescu-Mizil et al.* 2012: People subconsciously mirror each other's grammar words in conversations (linguistic coordination)
- Mirroring increases when we talk to people of high status
- $C(a, b) = P(E_{u_a \to u_b} | E_{u_b}) - P(E_{u_a \to u_b})$
- Can we take $C(a, b)$ as the weight of the edge from a to b?

Extending the Model

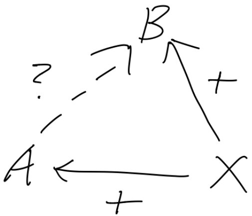

- X has rated B positively, **so B is higher status than average**, so A is more likely to rate B positively than he is to rate a random user positively

Extending the Model

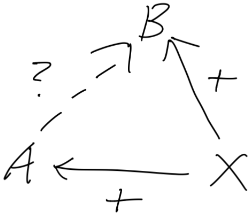

- X has rated B *higher than X rates people on average*, **so B is higher status than average**, so A should rate B *higher than A rates people on average*

Extending the Model

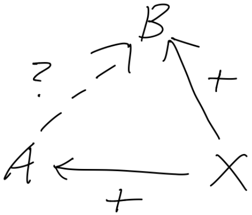

- X has rated B *higher than X rates people on average*, **so B is higher status than average**, so A should rate B *higher than A rates people on average*
- A has received a positive rating from X, **so A is higher status than average**, so B is less likely to receive a positive rating from A than he is to receive a positive rating from a random user

Extending the Model

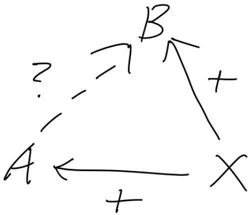

- X has rated B *higher than X rates people on average*, **so B is higher status than average**, so A should rate B *higher than A rates people on average*
- A has received a *higher score from X than A receives from people on average*, **so A is higher status than average**, so B should receive a *lower score from A than B receives from people on average*

Outline

RfA Voting Network

- Generative model: 12/16 predictions correct
- Receptive model: 12/16 predictions correct

Coordination Network

- Generative model: 12/16 predictions correct
- Receptive model: 14/16 predictions correct

Mistakes

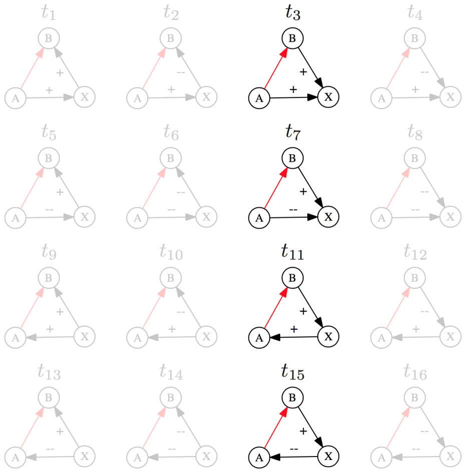

Coordination Network (No Admins)

- Generative model: 11/16 predictions correct
- Receptive model: 14/16 predictions correct

Coordination Network (No Admins)

- Generative model: 11/16 predictions correct
- Receptive model: 14/16 predictions correct
- There's an implicit status hierarchy among users!

Official vs Unofficial Status

- Next steps: look at both networks together to understand the interplay between *de jure* and *de facto* status

Combined Model

- Should evaluate generative and receptive model together

Combined Model

- Should evaluate generative and receptive model together
- Each triangle in the network falls under two triangle types (one for each model), giving 64 cases

Combined Model

Coordination

- 34 cases follow the generative and receptive predictions
- 11 cases follow just the generative prediction
- 15 cases follow just the generative prediction
- 4 cases follow neither prediction

Combined Model

Coordination — Random Baseline

- 16 cases follow the generative and receptive predictions
- 14 cases follow just the generative prediction
- 17 cases follow just the generative prediction
- 17 cases follow neither prediction

Combined Model

Voting

- 40 cases follow the generative and receptive predictions
- 9 cases follow just the generative prediction
- 8 cases follow just the generative prediction
- 7 cases follow neither prediction

Outline

Homophily

- How is a's coordination level toward b influenced by a's neighbors in the network?

Homophily

- How is a's coordination level toward b influenced by a's neighbors in the network?
- If more and more of a's neighbors coordinate to b, does a's deviation from baseline increase more than a non-neighbor's?

Homophily

- How is a's coordination level toward b influenced by a's neighbors in the network?
- If more and more of a's neighbors coordinate to b, does a's deviation from baseline increase more than a non-neighbor's?
- $f(a, b) = d(a) - \frac{1}{|N|} \sum_{n \in N} d(n)$ where N is the set of people n for whom $C(n, b)$ is defined and
  $d(x) = C(x, b) - \text{generative baseline}(x)$

Homophily

- How is a's coordination level toward b influenced by a's neighbors in the network?
- If more and more of a's neighbors coordinate to b, does a's deviation from baseline increase more than a non-neighbor's?
- $f(a, b) = d(a) - \frac{1}{|N|} \sum_{n \in N} d(n)$ where N is the set of people n for whom $C(n, b)$ is defined and $d(x) = C(x, b) - \text{generative baseline}(x)$
- Does f tend to increase as the proportion of a's neighbors who support b increases?

Homophily — Coordination

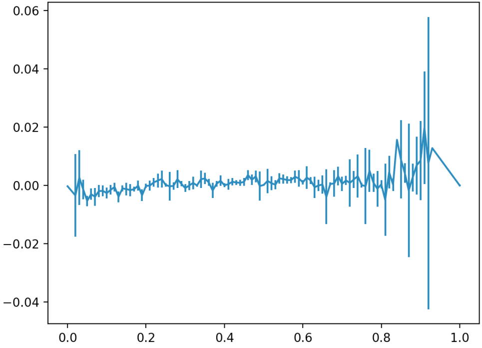

- Regression line: $y = 0.0074x - 0.0023$, $r = 0.54$, slope $\neq 0$ with $p < 1.4 * 10^{-8}$

Homophily — Voting

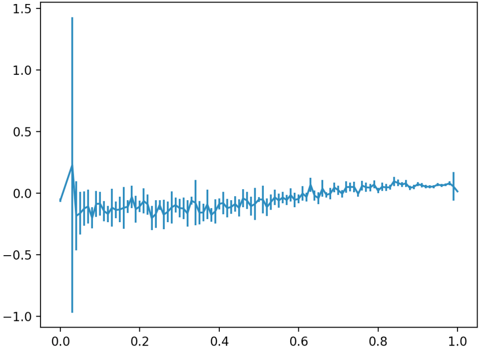

- Regression line: $y = 0.23x - 0.16$, $r = 0.77$, slope $\neq 0$ with $p < 1.4 * 10^{-20}$

Team Performance — Coordination

- In how many admin-nonadmin pairs does the admin coordinate more to the nonadmin than vice versa?

Team Performance — Coordination

- In how many admin-nonadmin pairs does the admin coordinate more to the nonadmin than vice versa?
- Productive teams: **43.46%** of pairs

Team Performance — Coordination

- In how many admin-nonadmin pairs does the admin coordinate more to the nonadmin than vice versa?
- Productive teams: **43.46%** of pairs
- Non-productive teams: **42.55%** of pairs

Team Performance — Coordination

- In how many admin-nonadmin pairs does the admin coordinate more to the nonadmin than vice versa?
- Productive teams: **43.46%** of pairs
- Non-productive teams: **42.55%** of pairs
- Controversial article teams: **44.55%** of pairs

Team Performance — Coordination

- In how many admin-nonadmin pairs does the admin coordinate more to the nonadmin than vice versa?
- Productive teams: **43.46%** of pairs
- Non-productive teams: **42.55%** of pairs
- Controversial article teams: **44.55%** of pairs
- (Differences are significant with $p \ll 0.001$)

Team Performance — Coordination

Coordination Network

- Generative model: in 10/16 triangle types, the “surprise” of the predicted edge is significantly different between productive and non-productive teams

Team Performance — Coordination

Coordination Network

- Generative model: in 10/16 triangle types, the "surprise" of the predicted edge is significantly different between productive and non-productive teams
  - **Surprise**: (Signed) number of standard deviations by which the observed weight of the edge differs from the expected weight

Team Performance — Coordination

Coordination Network

- Generative model: in 10/16 triangle types, the "surprise" of the predicted edge is significantly different between productive and non-productive teams
  - **Surprise**: (Signed) number of standard deviations by which the observed weight of the edge differs from the expected weight
- Receptive model: in 7/16 triangle types, the "surprise" of the predicted edge is significantly different between productive and non-productive teams

Team Performance — Voting

Voting Network

- Generative model: in 11/16 triangle types, the "surprise" of the predicted edge is significantly different between productive and non-productive teams
- Receptive model: in 9/16 triangle types, the "surprise" of the predicted edge is significantly different between productive and non-productive teams

Team Performance — Voting

- Treat the voting network as a **social balance network**

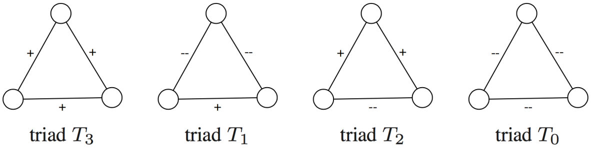

Team Performance — Voting

- We find that the proportion of each triad type is closer to random in productive teams compared to non-productive teams

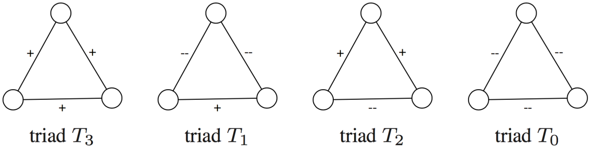

Next Steps

- Compare coordination and voting networks, relate to team performance

Next Steps

- Compare coordination and voting networks, relate to team performance
- Track the progress of a team over time

Thanks for a great summer!

Thoughts and suggestions?