Opinion Diffusion in Similarity-Driven Networks

Edoardo Baccini, Zoé Christoff and Rineke Verbrugge
e.baccini@rug.nl

Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence
University of Groningen

PhDs in Logic XIII, Turin

university of groningen

Introduction

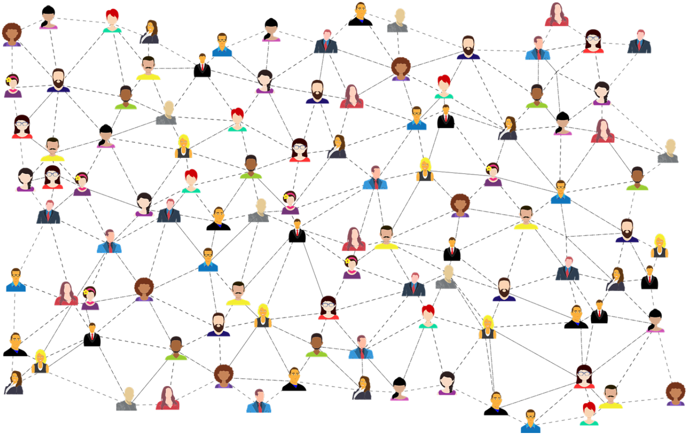

Introduction: Social Network Dynamics

Social influence: an agent's opinions can be influenced by their friends or peers (Granovetter 1978; Easley and Kleinberg 2010).

Similarity-driven network change: agents tend to connect to other agents with whom they share similar opinions, and to disconnect from those who do not (McPherson, Smith-Lovin, and Cook 2001).

Social network logics for:

- diffusion on networks e.g. (Baltag et al. 2018; Christoff and Naumov 2019)
- network change e.g. (Smets and Velázquez-Quesada 2020)
- both happening at different times (Smets and Velázquez-Quesada 2019).

Introduction

Threshold models of diffusion (Easley and Kleinberg 2010)

- set of agents initially accepting an issue
- uniform *influenceability threshold*

Adoption Rule

Agents accept an issue when a proportion of their influencers who accept the issue meets the threshold.

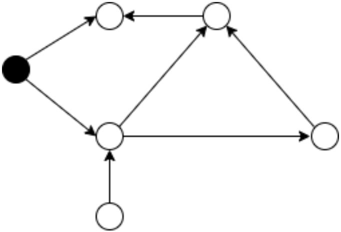

Introduction

Threshold models of diffusion (Easley and Kleinberg 2010)

- set of agents initially accepting an issue
- uniform *influenceability threshold*

Adoption Rule

Agents accept an issue when a proportion of their influencers who accept the issue meets the threshold.

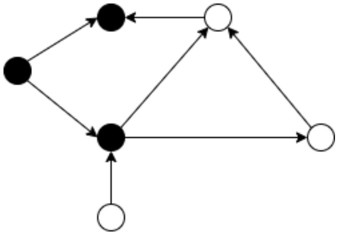

Introduction

Threshold models of diffusion (Easley and Kleinberg 2010)

- set of agents initially accepting an issue
- uniform *influenceability threshold*

Adoption Rule

Agents accept an issue when a proportion of their influencers who accept the issue meets the threshold.

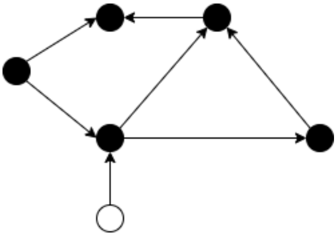

Introduction

Threshold-based network change (McPherson, Smith-Lovin, and Cook 2001)

- set of issues;
- uniform *similarity threshold*.

> **Connection Rule**
>
> Two agents connect if and only if the set of issue they agree on meets the similarity threshold.

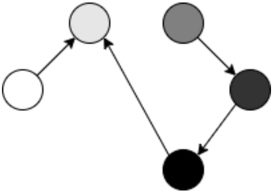

Introduction

Threshold-based network change (McPherson, Smith-Lovin, and Cook 2001)

- set of issues;
- uniform *similarity threshold*.

Connection Rule

Two agents connect if and only if the set of issue they agree on meets the similarity threshold.

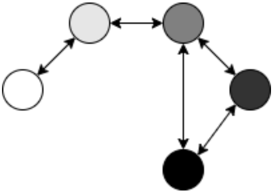

Overview

1. Logic for synchronous change
2. Stabilization with a single binary issue
3. Summary and further work

Overview

Logic for Synchronous Change

Definition

A model $\mathcal{M}$ is a tuple $\langle \mathcal{A}, \mathcal{S}, \mathcal{F}, \mathcal{V}, \omega, \tau \rangle$ such that:

- $\mathcal{A}$ is a non-empty finite set of agents;
- $\mathcal{S} \subseteq \mathcal{A} \times \mathcal{A}$ is a social influence relation between agents;
- $\mathcal{F}$ is a non-empty finite set of issues;
- $\mathcal{V} : \mathcal{A} \longrightarrow \mathcal{P}(\mathcal{F})$ is a valuation function, assigning to each agent a set of issues they accept;
- $\omega, \tau \in \mathbb{Q}$, s.t. $0 \leq \omega \leq 1$ and $0 \leq \tau \leq 1$, interpreted, respectively, as *similarity threshold* and *influenceability threshold*.

Model Update: Opinion Diffusion and Network Change

Opinion Update

An agent will accept an issue if and only if either:

- The influenceability threshold τ is 0.
- The agent has no influencer and already accepts the issue.
- The proportion of its influencers that accept the issue is at least τ.

Model Update: Opinion Diffusion and Network Change

Opinion Update

An agent will accept an issue if and only if either:

- The influenceability threshold τ is 0.
- The agent has no influencer and already accepts the issue.
- The proportion of its influencers that accept the issue is at least τ.

Network Update

Agents will be connected if and only if the proportion of issues they agree on is at least ω.

Model Update

Example

Model update from $\mathcal{M}$ to $\mathcal{M}_{\odot}$:

- $\mathcal{F} = \{f\}$;
- similarity threshold $\omega = 1$;
- influenceability threshold $\tau = \frac{1}{2}$.

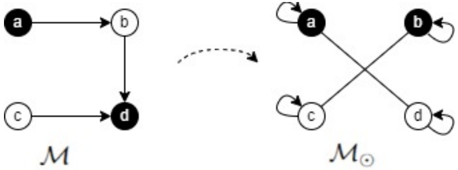

Note: the updated model is reflexive and symmetric.

Model Update

Given a model $\mathcal{M} = \langle \mathcal{A}, \mathcal{S}, \mathcal{F}, \mathcal{V}, \omega, \tau \rangle$, the updated model $\mathcal{M}_{\odot} = \langle \mathcal{A}, \mathcal{S}', \mathcal{F}, \mathcal{V}', \omega, \tau \rangle$ is such that for any $a, b \in \mathcal{A}$ and any $f \in \mathcal{F}$:

$$(a,b) \in \mathcal{S}' \text{ iff } \frac{|(\mathcal{V}(a) \cap \mathcal{V}(b)) \cup (\overline{\mathcal{V}(a)} \cap \overline{\mathcal{V}(b)})|}{|\mathcal{F}|} \geq \omega$$

$$f \in \mathcal{V}'(a) \text{ iff } \left\{ \begin{array}{ll} f \in \mathcal{V}(a), & \text{if } N(a) = \emptyset \\ f \in \mathcal{F}, & \text{if } \tau = 0 \\ \frac{|N_f(a)|}{|N(a)|} \geq \tau, & \text{otherwise} \end{array} \right\}$$

where $N_f(a) := \{b \in A : (b,a) \in \mathcal{S} \text{ and } f \in \mathcal{V}(b)\}$ and $N(a) := \{b \in A : (b,a) \in \mathcal{S}\}$.

Syntax $\mathcal{L}_\odot$ and Semantics

Definition (Syntax $\mathcal{L}_\odot$)

Fix a set of issues $\mathcal{F}$ and a set of agents $\mathcal{A}$. The syntax $\mathcal{L}_\odot$ is the following:

$$\phi := \; S_{ab} \mid f_a \mid \neg\phi \mid \phi \wedge \phi \mid \; \odot\, \phi$$

where $f \in \mathcal{F}$ and $a, b \in \mathcal{A}$.

Definition (Semantic clauses for $\mathcal{L}_\odot$)

The truth of a formula ϕ in $\mathcal{M}$ is inductively defined as follows:

$\mathcal{M} \models f_a$ if and only if $f \in \mathcal{V}(a)$

$\mathcal{M} \models S_{ab}$ if and only if $(a, b) \in \mathcal{S}$

$\mathcal{M} \models \neg\phi$ if and only if $\mathcal{M} \not\models \phi$

$\mathcal{M} \models \phi \wedge \psi$ if and only if $\mathcal{M} \models \phi$ and $\mathcal{M} \models \psi$

$\mathcal{M} \models \odot\phi$ if and only if $\mathcal{M}_\odot \models \phi$, where $\mathcal{M}_\odot$ is the updated model.

Conformity pressure and Similarity pressure

Conformity pressure

The *conformity pressure* of an agent can be captured with the formula:

$$f_a^\tau := (f_a \wedge \neg \bigvee_{b \in \mathcal{A}} S_{ba}) \vee f^\tau_{N(a)}$$

$$f^\tau_{N(a)} := \bigvee_{\{G \subseteq N \subseteq \mathcal{A} : \frac{|G|}{|N|} \geq \tau\}} (\bigwedge_{b \in N} S_{ba} \wedge \bigwedge_{b \notin N} \neg S_{ba} \wedge ((\bigwedge_{b \in G} f_b \wedge \neg f_b) \vee (\bigvee_{b \in N} S_{ba} \wedge \bigwedge_{b \in G} f_b)))$$

Conformity pressure and Similarity pressure

Conformity pressure

The *conformity pressure* of an agent can be captured with the formula:

$$f_a^\tau := (f_a \wedge \neg \bigvee_{b\in\mathcal{A}} S_{ba}) \vee f^\tau_{N(a)}$$

$$f^\tau_{N(a)} := \bigvee_{\{G\subseteq N\subseteq\mathcal{A}:\frac{|G|}{|N|}\geq\tau\}} (\bigwedge_{b\in N} S_{ba} \wedge \bigwedge_{b\notin N} \neg S_{ba} \wedge ((\bigwedge_{b\in G} f_b \wedge \neg f_b) \vee (\bigvee_{b\in N} S_{ba} \wedge \bigwedge_{b\in G} f_b)))$$

Similarity pressure

The *similarity pressure* of one agent with respect to another agent is expressed by the formula:

$$sim^\omega_{ab} := \bigvee_{\{E\subseteq\mathcal{F}:\frac{|E|}{|\mathcal{F}|}\geq\omega\}} \bigwedge_{f\in E} (f_a \leftrightarrow f_b)$$

Reduction Axioms

$\odot S_{ab} \leftrightarrow sim^{\omega}_{ab}$	$\odot$-network
$\odot f_a \leftrightarrow f^{\tau}_a$	$\odot$-opinions
$\odot(\phi \wedge \psi) \leftrightarrow \odot\phi \wedge \odot\psi$	$\odot$-distributivity
$\odot\neg\phi \leftrightarrow \neg\odot\phi$	$\odot$-neg

Soundness and Completeness of $L^{\omega\tau}$

Fix ω, τ. The logic $L^{\omega\tau}$ is any complete set of axioms and derivation rules of propositional logic, together with the above reduction axioms, a necessitation rule for $\odot$, and substitution of equivalents.

Theorem

Let $\omega, \tau \in [0, 1]$ *be given For any* $\phi \in \mathcal{L}_{\odot}$,

$$\models_{\mathcal{C}^{\omega\tau}} \phi \text{ iff } \vdash_{L^{\omega\tau}} \phi.$$

where $\mathcal{C}_{\omega,\tau}$ *is the class of models with thresholds* ω, τ.

Overview

Stabilization

Objective: investigate the limit properties of the sequences of models obtained by iterative updates of synchronous change.

Question 1: Given an initial model $\mathcal{M}$, does $\mathcal{M}$ stabilize?

Question 2: Given a model $\mathcal{M}$ that does not stabilize, what is its limit behavior?

university of groningen

Stabilization with a single binary issue

Definition

A model $\mathcal{M}$ is stable if and only if $\mathcal{M}_{\odot} = \mathcal{M}$. $\mathcal{M}$ stabilizes if and only if there is an $n \in \mathbb{N}$ such that $\mathcal{M}_{\odot^n} = \mathcal{M}_{\odot^{n+1}}$.

Example (Stabilization sequence of the initial model $\mathcal{M}$)

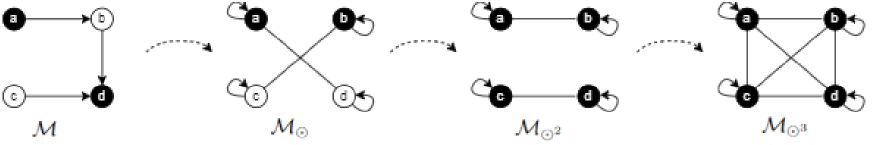

Oscillation

We look at stabilization as a special kind of oscillation (van Benthem 2015).

Example ($\mathcal{M}_\odot$ oscillates with length 2)

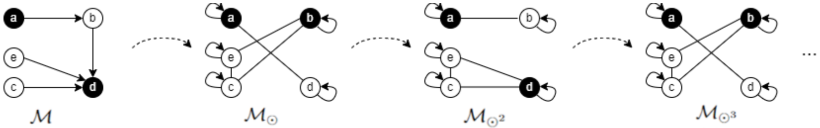

Note: stabilization is an oscillation of length 1.

Partitioning the set of agents

The central concept for characterizing stability is that of an f-partition.

Definition

(f-partition) Given any model $\mathcal{M} = \langle \mathcal{A}, \mathcal{S}, \mathcal{F}, \mathcal{V}, \omega, \tau \rangle$ and any $f \in \mathcal{F}$, we call f-partition the partition of $\mathcal{A}$ in the four following sets:

- $\mathcal{K}_f := \{a \in \mathcal{A} : \mathcal{M} \models f_a \wedge f_a^\tau\}$ ("f-keepers");
- $\mathcal{D}_f := \{a \in \mathcal{A} : \mathcal{M} \models f_a \wedge \neg f_a^\tau\}$ ("f-droppers");
- $\mathcal{K}_{\neg f} := \{a \in \mathcal{A} : \mathcal{M} \models \neg f_a \wedge \neg f_a^\tau\}$ ("$\neg f$-keepers");
- $\mathcal{D}_{\neg f} := \{a \in \mathcal{A} : \mathcal{M} \models \neg f_a \wedge f_a^\tau\}$ ("$\neg f$-droppers").

The relative sizes of these four sets determine whether a model stabilizes or not.

They also determine the length of a model's future oscillation.

Steps

We characterize stabilization with the following steps:

1. models with $\tau = 0$ or $\omega = 0$ stabilize;
2. if at least two sets of the set $\mathcal{K}_f, \mathcal{D}_f, \mathcal{D}_{\neg f}, \mathcal{K}_{\neg f}$ of the f-partition are empty, then the model stabilizes.
3. if at most one of the sets in the f-partition is empty, a model stabilizes if some specific conditions on the relative sizes of the sets $\mathcal{K}_f, \mathcal{D}_f, \mathcal{D}_{\neg f}, \mathcal{K}_{\neg f}$ do not hold.

Step 1: $\tau = 0$ or $\omega = 0$

Any model that has either $\tau = 0$ or $\omega = 0$ stabilizes.

If $\tau = 0$, after one step every agent will accept the target issue, thus becoming similar to every other agent.

If $\omega = 0$, after one step every agent will be connected to any other agent, thus having the same conformity pressure as every other agent.

Example

$\mathcal{M}$ with $\tau = 0$, $\omega = 1$.

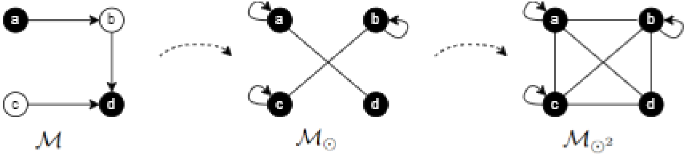

Step 2: At least two sets of the f-partition are empty

If at least two sets of the partition are empty, then the model stabilizes.

Example

$\mathcal{M}$ with $\tau, \omega > 0$ in which there are only f-keepers and $\neg f$-keepers.

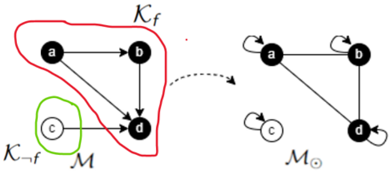

Step 3: At most one set in the f-partition is empty.

After an update:

- Models are composed of two disjoint components.
- Within each component every agent is influenced by every agent.
- Within each component, every agent has the same conformity pressure as every other agent.

.

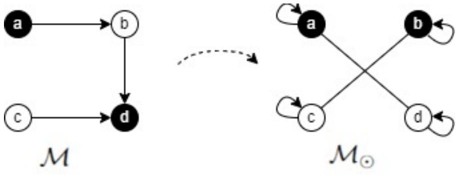

Step 3: At most one set in the f-partition is empty

If at any point of a sequence of update the two components agree on the opinion to adopt further, a stable state will be reached.

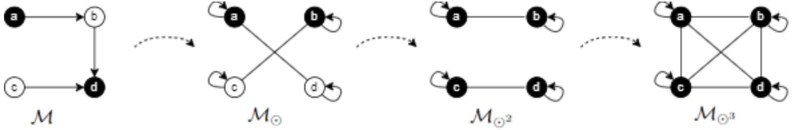

Step 3: At most one set in the f-partition is empty.

For a model not to stabilize, the two generated components must disagree after every update.

Question: Which are all the possible ways in which the two components generated by an update can keep disagreeing with each other?

Equivalently: Which are all the possible ways in which $\mathcal{M}$ can oscillate with length $l > 1$?

Step 3: Oscillations with $l > 1$

All the possible ways in which a model $\mathcal{M}$ can oscillate with length $l > 1$ are:

(1) $\frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_f|} < \tau,\ \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{D}_{\neg f}|+|\mathcal{K}_{\neg f}|} \geq \tau,\ \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{K}_f|+|\mathcal{D}_{\neg f}|} \geq \tau,\ \frac{|\mathcal{K}_{\neg f}|}{|\mathcal{K}_{\neg f}|+|\mathcal{D}_f|} < \tau$

(2) $\frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_f|} \geq \tau, \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{D}_{\neg f}|+|\mathcal{K}_{\neg f}|} < \tau, \frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_{\neg f}|} \geq \tau,\ \frac{|\mathcal{D}_f|}{|\mathcal{D}_f|+|\mathcal{K}_{\neg \{}|} < \tau$

(3) $\frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_f|} < \tau, \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{D}_{\neg f}|+|\mathcal{K}_{\neg f}|} \geq \tau, \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{K}_f|+|\mathcal{D}_{\neg f}|} < \tau,\ \frac{|\mathcal{K}_{\neg f}|}{|\mathcal{D}_f|+|\mathcal{K}_{\neg f}|} \geq \tau, \frac{|\mathcal{D}_f|}{|\mathcal{K}_f|+|\mathcal{D}_f|} \geq \tau, \frac{|\mathcal{K}_{\neg f}|}{|\mathcal{D}_{\neg f}|+|\mathcal{K}_{\neg f}|} < \tau,\ \frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_{\neg f}|} \geq \tau,\ \frac{|\mathcal{D}_f|}{|\mathcal{D}_f|+|\mathcal{K}_{\neg f}|} < \tau$

(4) $\frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_f|} \geq \tau,\ \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{D}_{\neg f}|+|\mathcal{K}_{\neg f}|} < \tau,\ \frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_{\neg f}|} < \tau,\ \frac{|\mathcal{D}_f|}{|\mathcal{D}_f|+|\mathcal{K}_{\neg f}|} \geq \tau, \frac{|\mathcal{D}_f|}{|\mathcal{K}_f|+|\mathcal{D}_f|} < \tau, \frac{|\mathcal{K}_{\neg f}|}{|\mathcal{D}_{\neg f}|+|\mathcal{K}_{\neg f}|} \geq \tau,\ \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{K}_f|+|\mathcal{D}_{\neg f}|} \geq \tau, \frac{|\mathcal{K}_{\neg f}|}{|\mathcal{D}_f|+|\mathcal{K}_{\neg f}|} < \tau$

Condition (1) and (2) characterize oscillation of length 2; conditions (3) and (4) characterize oscillation of length 4.

university of groningen

Example: Condition (1)

Condition (1): $\frac{|\mathcal{K}_f|}{|\mathcal{K}_f|+|\mathcal{D}_f|}<\tau, \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{D}_{\neg f}|+|\mathcal{K}_{\neg f}|}\geq\tau, \frac{|\mathcal{D}_{\neg f}|}{|\mathcal{K}_f|+|\mathcal{D}_{\neg f}|}\geq\tau, \frac{|\mathcal{K}_{\neg f}|}{|\mathcal{K}_{\neg f}|+|\mathcal{D}_f|}<\tau$

Example

Starting from a model $\mathcal{M}$, representation of the oscillatory behavior determined by condition (1).

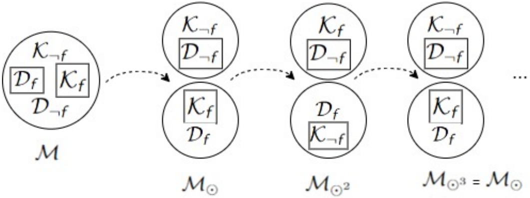

university of groningen

Characterizing stabilization

Theorem

Let $\mathcal{M} = \langle \mathcal{A}, \mathcal{S}, \mathcal{F}, \mathcal{V}, \omega, \tau \rangle$ and $\mathcal{M} \in \mathcal{C}_1$, $\mathcal{M}$ stabilizes if and only if one of the following holds:

1. $\omega = 0$ *or* $\tau = 0$;
2. *at least two of the sets* $\mathcal{K}_f, \mathcal{D}_f, \mathcal{D}_{\neg f}, \mathcal{K}_{\neg f}$ *are empty;*
3. *at most one of the sets* $\mathcal{K}_f, \mathcal{D}_f, \mathcal{D}_{\neg f}, \mathcal{K}_{\neg f}$ *is empty and none of the conditions* (1) *to* (4) *above holds.*

university of groningen

Overview

university of groningen

Summary

- A sound and complete axiom system $L^{\omega\tau}$ for *synchronous* opinion diffusion in similarity-driven networks.
- Characterization of stabilization for the diffusion of a single binary issue in similarity-driven networks.

Further Work

- Characterize stabilization for models with multiple issues.
- Investigate the interdefinability of synchronic and diachronic changes.
- Analyse more refined network-change update: agents connect to similar agents that are *close enough*.
- Refine opinions: acceptance/rejection VS evidence-based approach.

Contact: e.baccini@rug.nl