A Viral Marketing-Based Model For Opinion Dynamics in Online Social Networks

Sijing Tu [1] Stefan Neumann [1]

[1]KTH Royal Institute of Technology

The WebConf 2022

Social Networks Have Been Under Attack

Facebook data: How it was used by Cambridge Analytica

Facebook has said it believes that up to 87 million users' data was improperly shared with Cambridge Analytica.

(BBC)

Trolling of women in politics 'worrying'

Trolling towards female politicians has a "nasty, misogynistic element", Wales' former education minister has said.

Social Networks Have Been Under Attack

Facebook data: How it was used by Cambridge Analytica

Facebook has said it believes that up to 87 million users' data was improperly shared with Cambridge Analytica.

(BBC)

Trolling of women in politics 'worrying'

Trolling towards female politicians has a "nasty, misogynistic element", Wales' former education minister has said.

Are influence campaigns trolling your social media feeds?

New research shows how citizens can know in real time.

(Washington Post)

Social Networks Have Been Under Attack

Facebook data: How it was used by Cambridge Analytica

Facebook has said it believes that up to 87 million users' data was improperly shared with Cambridge Analytica.

(BBC)

Trolling of women in politics 'worrying'

Trolling towards female politicians has a "nasty, misogynistic element", Wales' former education minister has said.

Are influence campaigns trolling your social media feeds?

New research shows how citizens can know in real time.

(Washington Post)

Who's a Bot? Who's Not?

It sometimes seems that automated bots are taking over social media and driving human discourse. But some (real) researchers aren't so sure.

(The New York Times)

Polarizing vs Marketing Content

Emmanuel Macron
Republic on the Move

Marine Le Pen
National Rally

Polarizing vs Marketing Content

Polarizing vs Marketing Content

Polarizing content: like it or hate it

Polarizing vs Marketing Content

Polarizing content: like it or hate it

Polarizing vs Marketing Content

Polarizing content: like it or hate it

Marketing content: like it or ignore it

What Do We Need to Understand?

- How does the content spread through the network?
- When exposed to new content, how do people form their opinions?
- How does spreading content affect the polarization of the network?

What Do We Need to Understand?

- How does the content spread through the network?
- When exposed to new content, how do people form their opinions?
- How does spreading content affect the polarization of the network?

What Do We Need to Understand?

- How does the content spread through the network?
- When exposed to new content, how do people form their opinions?
- How does spreading content affect the polarization of the network?

Existing Models

How Does Information Spread?

- **Independent Cascade Model** (Kempe, Kleinberg, Tardos; 2003)
  - Models how viral content spreads in social networks
  - Each node has one shot to influence its neighbors
  - Once activated, nodes stay activated

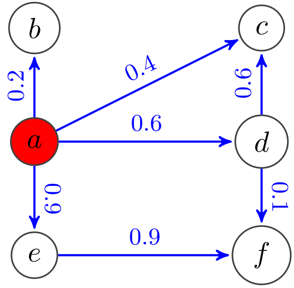

a initially gets activated at time step 0.

How Does Information Spread?

- Independent Cascade Model (Kempe, Kleinberg, Tardos; 2003)
  - Models how viral content spreads in social networks
  - Each node has one shot to influence its neighbors
  - Once activated, nodes stay activated

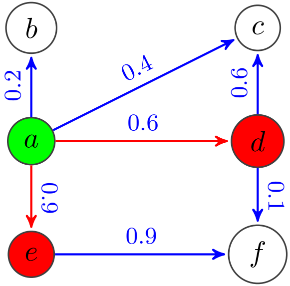

At time step 1, a activates d and e, but fails to activate b and c.

How Does Information Spread?

- Independent Cascade Model (Kempe, Kleinberg, Tardos; 2003)
  - Models how viral content spreads in social networks
  - Each node has one shot to influence its neighbors
  - Once activated, nodes stay activated

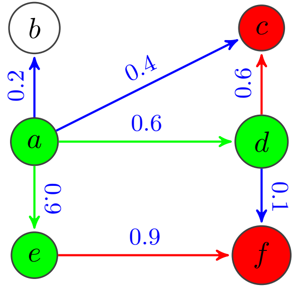

At time step 2, d activates c, e activates f

How Does Information Spread?

- Independent Cascade Model (Kempe, Kleinberg, Tardos; 2003)
  - Models how viral content spreads in social networks
  - Each node has one shot to influence its neighbors
  - Once activated, nodes stay activated

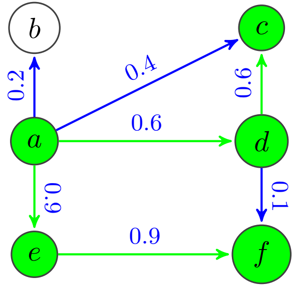

At time step 3, c and f don't have neighbors, the spreading stops here.

How Do People Form Opinions?

- Friedkin–Johnsen model (FJ model) (Friedkin, Johnsen; 1990)
  - Each node i has innate and expressed opinions
  - innate opinion $s_i \in [0, 1]$: fixed, kept private
  - expressed opinion $z_i^t \in [0, 1]$: depends on time t, public

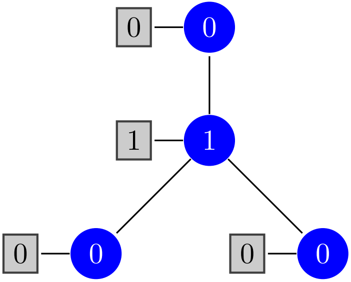

How Do People Form Opinions?

- Friedkin–Johnsen model (FJ model) (Friedkin, Johnsen; 1990)
  - Each node i has innate and expressed opinions
  - innate opinion $s_i \in [0,1]$: fixed, kept private
  - expressed opinion $z_i^t \in [0,1]$: depends on time t, public

Update rule

$$z_i^{t+1} = \frac{s_i + \sum_{j \in N(i)} w_{i,j} z_j^t}{1 + \sum_{j \in N(i)} w_{i,j}}$$

with edge weights $w_{i,j}$ and $N(i) =$ neighbors of i

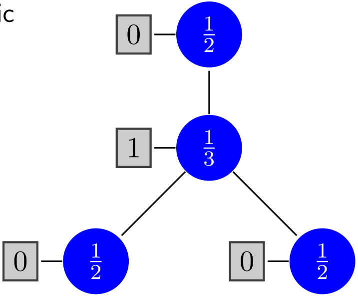

(After several rounds...)

How Do People Form Opinions?

- Friedkin–Johnsen model (FJ model) (Friedkin, Johnsen; 1990)
  - Each node i has innate and expressed opinions
  - innate opinion $s_i \in [0,1]$: fixed, kept private
  - expressed opinion $z_i^t \in [0,1]$: depends on time t, public

Update rule

$$z_i^{t+1} = \frac{s_i + \sum_{j \in N(i)} w_{i,j} z_j^t}{1 + \sum_{j \in N(i)} w_{i,j}}$$

with edge weights $w_{i,j}$ and $N(i)$ = neighbors of i

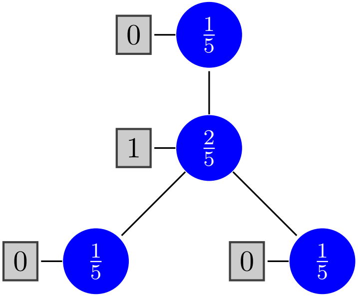

How Do People Form Opinions?

- Friedkin–Johnsen model (FJ model) (Friedkin, Johnsen; 1990)
  - Each node i has innate and expressed opinions
  - innate opinion $s_i \in [0,1]$: fixed, kept private
  - expressed opinion $z_i^t \in [0,1]$: depends on time t, public

Update rule

$$z_i^{t+1} = \frac{s_i + \sum_{j \in N(i)} w_{i,j} z_j^t}{1 + \sum_{j \in N(i)} w_{i,j}}$$

with edge weights $w_{i,j}$ and $N(i)$ = neighbors of i

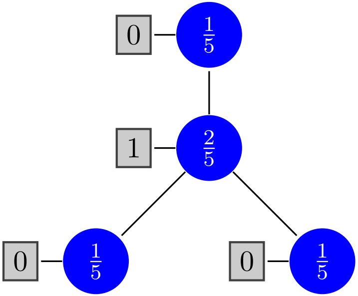

After convergence $\mathbf{z} = (I + L)^{-1}\mathbf{s}$,
where L is the graph Laplacian.

Limitations of Existing Models

- Existing models focus on either information spreading or opinion formation
- To understand how viral content influences user opinions, we need both

$\Rightarrow$ Can we get a model in which users update their opinions when exposed to viral content?

Our Contribution:

The Spread–Acknowledge Model

The Spread–Acknowledge Model

- We introduce the Spread–Acknowledge Model
- The model allows to quantify how much viral contents impacts user opinions

→ Combines the independent cascade and FJ models

- Nodes spread content similar to the independent cascade model
- When nodes are exposed to the content, they might change their innate opinion

→ The content and the opinion updates trickle through the network

- The model proceeds in rounds
- In each round, the nodes first update their opinions and then share content

The Spread–Acknowledge Model

- We introduce the Spread–Acknowledge Model
- The model allows to quantify how much viral contents impacts user opinions

→ Combines the independent cascade and FJ models

- Nodes spread content similar to the independent cascade model
- When nodes are exposed to the content, they might change their innate opinion

→ The content and the opinion updates trickle through the network

- The model proceeds in rounds
- In each round, the nodes first update their opinions and then share content

The Spread–Acknowledge Model

- We introduce the Spread–Acknowledge Model
- The model allows to quantify how much viral contents impacts user opinions

→ Combines the independent cascade and FJ models

- Nodes spread content similar to the independent cascade model
- When nodes are exposed to the content, they might change their innate opinion

→ The content and the opinion updates trickle through the network

- The model proceeds in rounds
- In each round, the nodes first update their opinions and then share content

The Spread–Acknowledge Model

- Initially, all nodes have innate and expressed opinion 0

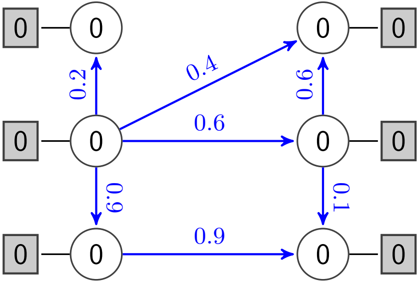

The node is not influenced.

The Spread–Acknowledge Model

- After being exposed to the viral content,
  the red node changes its innate opinion and spreads information to its neighbors

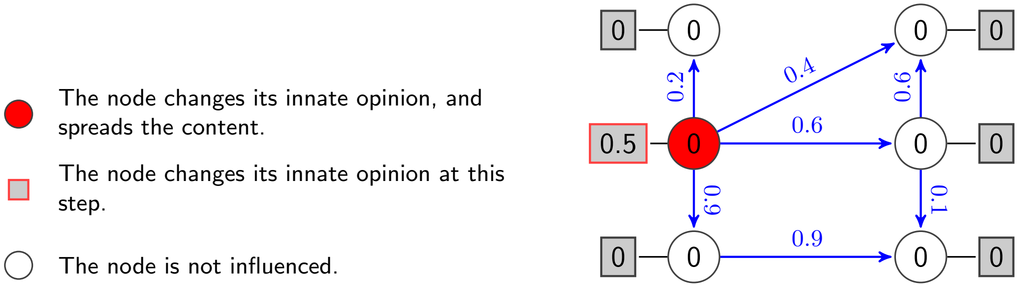

The Spread–Acknowledge Model

- The expressed opinions get updated as a consequence of the changed innate opinions

- The node changes its innate opinion, and spreads the content.
- The node changes its innate opinion at this step.
- The node is not influenced.

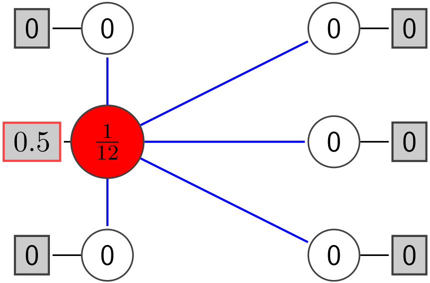

The Spread–Acknowledge Model

- The green node shares the content
  The yellow node changes its innate opinion (but does *not* share it)
  The red node changes its innate opinion and shares the content in the next round

The node only changes its innate opinion.

The node already spreads information.

The node changes its innate opinion, and spreads the content.

The influencing edge.

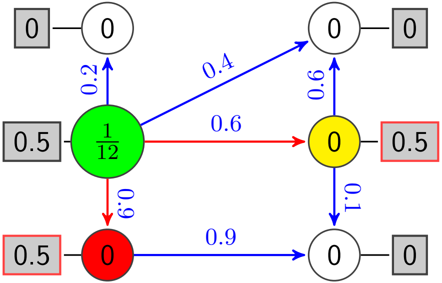

The Spread–Acknowledge Model

- As the innate opinions get updated, the expressed opinions are also updated

- The node only changes its innate opinion.
- The node already spreads information.
- The node changes its innate opinion, and spreads the content.
- The influencing edge.

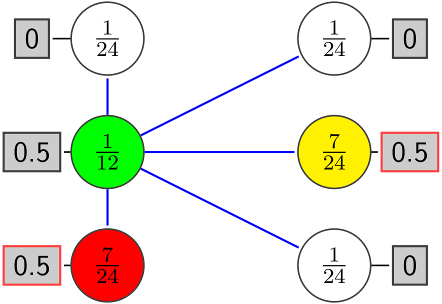

The Spread–Acknowledge Model

- The bottom-left node influences the bottom-right node
  The bottom-right node changes its innate opinion, but it does not spread the content

The node only changes its innate opinion.

The node already spreads information.

The node changes its innate opinion, and spreads the content.

The influencing edge.

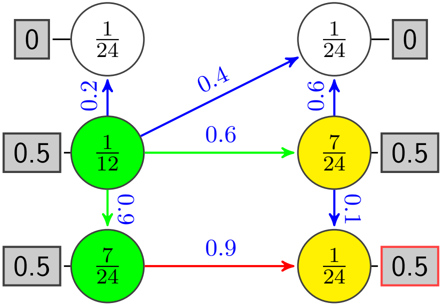

The Spread–Acknowledge Model

- When no more contents get shared, the expressed opinions are updated as in the FJ model

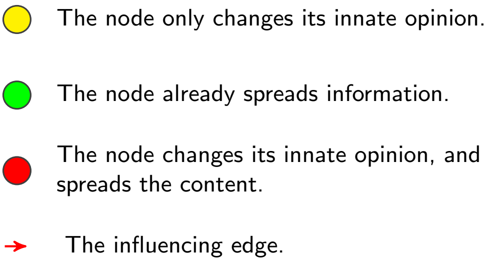

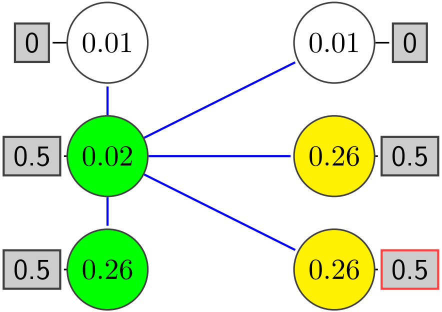

Updating Innate Opinions

When nodes acknowledge the viral content they update their innate opinions:

- Let $\epsilon > 0$ be a parameter
- Marketing content
  - Acknowledging users increase their opinion
  - New opinion $s'_v = s_v + \epsilon$
- Polarizing content
  - Proponents like it more, opponents are repelled
  - Let $\tau > 0$ be a parameter
  - New opinion $s'_v = \begin{cases} s_v + \epsilon, & \text{if } s_v \geq \tau \\ s_v - \epsilon, & \text{if } s_v < \tau \end{cases}$

$\Rightarrow$ In the experiments, we will study how these different types of contents influence polarization, disagreement, ...

Updating Innate Opinions

When nodes acknowledge the viral content they update their innate opinions:

- Let $\epsilon > 0$ be a parameter
- Marketing content
  - Acknowledging users increase their opinion
  - New opinion $s'_v = s_v + \epsilon$
- Polarizing content
  - Proponents like it more, opponents are repelled
  - Let $\tau > 0$ be a parameter
  - New opinion $s'_v = \begin{cases} s_v + \epsilon, & \text{if } s_v \geq \tau \\ s_v - \epsilon, & \text{if } s_v < \tau \end{cases}$

$\Rightarrow$ In the experiments, we will study how these different types of contents influence polarization, disagreement, ...

Updating Innate Opinions

When nodes acknowledge the viral content they update their innate opinions:

- Let $\epsilon > 0$ be a parameter
- Marketing content
  - Acknowledging users increase their opinion
  - New opinion $s'_v = s_v + \epsilon$
- Polarizing content
  - Proponents like it more, opponents are repelled
  - Let $\tau > 0$ be a parameter
  - New opinion $s'_v = \begin{cases} s_v + \epsilon, & \text{if } s_v \geq \tau \\ s_v - \epsilon, & \text{if } s_v < \tau \end{cases}$

$\Rightarrow$ In the experiments, we will study how these different types of contents influence polarization, disagreement, . . .

Simulating the Model

Simulating the Model

How can we efficiently simulate our model?

- Previously:
  - Friedkin–Johnsen model has nice closed form solution $z = (I + L)^{-1}s$
  - Independent cascade model can be simulated using reverse reachable sets

- Now:
  - Opinion formation and information spreading are interleaved
  - Not clear how to use existing techniques for our model
  - $\rightarrow$ How can we efficiently simulate our model?

Simulating the Model

$\boxed{\text{Update opinions}}$ $z_i^1 = \frac{s_i^0 + \sum_{j \in N(i)} w_{i,j} z_j^0}{1 + \sum_{j \in N(i)} w_{i,j}}$

$\downarrow$

Round 1: $\boxed{\text{Information spread; Nudge innate opinions}}$ $s_i^1 = f_0(s^0)$

Simulating the Model

Update opinions $\quad z_i^1 = \frac{s_i^0 + \sum_{j \in N(i)} w_{i,j} z_j^0}{1 + \sum_{j \in N(i)} w_{i,j}}$

↓

Information spread; Nudge innate opinions $\quad s_i^1 = f_0(s^0)$

↓

Update opinions $\quad z_i^2 = \frac{s_i^1 + \sum_{j \in N(i)} w_{i,j} z_j^1}{1 + \sum_{j \in N(i)} w_{i,j}}$

↓

Round 2: Information spread; Nudge innate opinions $\quad s_i^2 = f_1(s^1)$

Simulating the Model

Update opinions $\quad z_i^1 = \frac{s_i^0 + \sum_{j \in N(i)} w_{i,j} z_j^0}{1 + \sum_{j \in N(i)} w_{i,j}}$

↓

Information spread; Nudge innate opinions $\quad s_i^1 = f_0(s^0)$

↓

Update opinions $\quad z_i^2 = \frac{s_i^1 + \sum_{j \in N(i)} w_{i,j} z_j^1}{1 + \sum_{j \in N(i)} w_{i,j}}$

↓

And so on...

Information spread; Nudge innate opinions $\quad s_i^2 = f_1(s^1)$

↓

...

The Two-Stage Model

- In our model the information spreading and the opinion updates are interleaved
- Consider the simpler Two-Stage Model:
  - Stage 1: Run the independent cascade model

  $\rightarrow$ this reveals which users change their innate opinions
  - Stage 2: Run the FJ model with the updated innate opinions

$\Rightarrow$ can be simulated using known results

- **Lemma:** The Spread–Acknowledge Model and the Two-Stage Model generate the same distributions over the updated innate and expressed opinions.

$\Rightarrow$ we can efficiently simulate our model by simulating the Two-Stage model

The Two-Stage Model

- In our model the information spreading and the opinion updates are interleaved
- Consider the simpler Two-Stage Model:
  - Stage 1: Run the independent cascade model

  $\rightarrow$ this reveals which users change their innate opinions
  - Stage 2: Run the FJ model with the updated innate opinions

$\Rightarrow$ can be simulated using known results

- **Lemma:** The Spread–Acknowledge Model and the Two-Stage Model generate the same distributions over the updated innate and expressed opinions.

$\Rightarrow$ we can efficiently simulate our model by simulating the Two-Stage model

Further Results

Further Results

- We give algorithms for the efficient simulation of our model
- We prove their correctness and error guarantees

- We give approximation algorithms for maximizing the sum of opinions

→ Builds upon framework of reverse reachable sets (Borgs et al.; 2014)

- We give heuristics for maximizing disagreement, polarization, . . .

- Our model can easily be generalized in many ways
  (e.g., using linear threshold model, different ways to nudge opinions, . . .)

Further Results

- We give algorithms for the efficient simulation of our model
- We prove their correctness and error guarantees

- We give approximation algorithms for maximizing the sum of opinions

→ Builds upon framework of reverse reachable sets (Borgs et al.; 2014)

- We give heuristics for maximizing disagreement, polarization, ...

- Our model can easily be generalized in many ways
(e.g., using linear threshold model, different ways to nudge opinions, ...)

Further Results

- We give algorithms for the efficient simulation of our model
- We prove their correctness and error guarantees

- We give approximation algorithms for maximizing the sum of opinions

$\rightarrow$ Builds upon framework of reverse reachable sets (Borgs et al.; 2014)

- We give heuristics for maximizing disagreement, polarization, . . .

- Our model can easily be generalized in many ways
  (e.g., using linear threshold model, different ways to nudge opinions, . . .)

Experiments

Impact of Different Content Types on Polarization

- Polarization is the variance of the opinions

$\rightarrow$ $\mathcal{P} = \sum_{i \in V} (z_i - \bar{z})^2$ where $\bar{z} = \frac{1}{|V|} \sum_i z_i$ is the average opinion

- What happens when the 0.5% most influential nodes start sharing a content?
- Marketing contents and polarizing contents have very different impacts!

Impact of Different Content Types on Polarization

- Polarization is the variance of the opinions

$\rightarrow \mathcal{P} = \sum_{i \in V} (z_i - \bar{z})^2$ where $\bar{z} = \frac{1}{|V|} \sum_i z_i$ is the average opinion

- What happens when the 0.5% most influential nodes start sharing a content?
- Marketing contents and polarizing contents have very different impacts!

Impact of Different Content Types on Polarization

- Polarization is the variance of the opinions

$\rightarrow$ $\mathcal{P} = \sum_{i \in V} (z_i - \bar{z})^2$ where $\bar{z} = \frac{1}{|V|} \sum_i z_i$ is the average opinion

- What happens when the 0.5% most influential nodes start sharing a content?
- Marketing contents and polarizing contents have very different impacts!

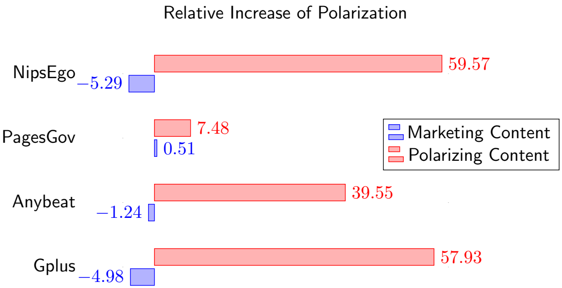

Impact of Different Content Types on Disagreement

- Disagreement measures how much interacting users disagree

$\to$ $\mathcal{D} = \sum_{u,v \in V} (z_u - z_v)^2$

- What happens when the 0.5% most influential nodes start sharing a content?
- As before: very different impacts!

Impact of Different Content Types on Disagreement

- Disagreement measures how much interacting users disagree

$\rightarrow \mathcal{D} = \sum_{u,v \in V} (z_u - z_v)^2$

- What happens when the 0.5% most influential nodes start sharing a content?
- As before: very different impacts!

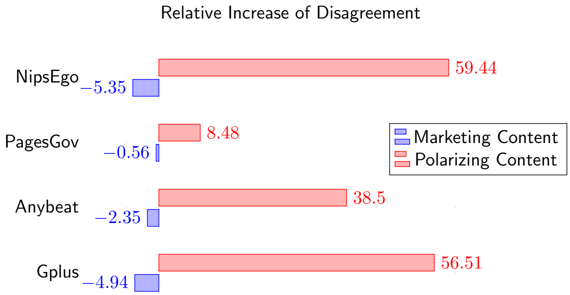

Summary and Future Work

Summary and Future Work

Our contributions:

- Novel model that allows to quantify how viral content influences user opinions
- Algorithms to efficiently simulate it
- Experiments show big difference between controversial marketing contents

Interesting directions for future work:

- Validating the model in practice
- Estimating the model parameters
- Can we get approximation algorithms for optimizing polarization, disagreement?

Summary and Future Work

Our contributions:

- Novel model that allows to quantify how viral content influences user opinions
- Algorithms to efficiently simulate it
- Experiments show big difference between controversial marketing contents

Interesting directions for future work:

- Validating the model in practice
- Estimating the model parameters
- Can we get approximation algorithms for optimizing polarization, disagreement?