Opinion Dynamics

with Marked Temporal Point Processes (TPPs)

HUMAN-CENTERED MACHINE LEARNING

http://courses.mpi-sws.org/hcml-ws18/

MAX PLANCK INSTITUTE
FOR SOFTWARE SYSTEMS

Opinions in social media

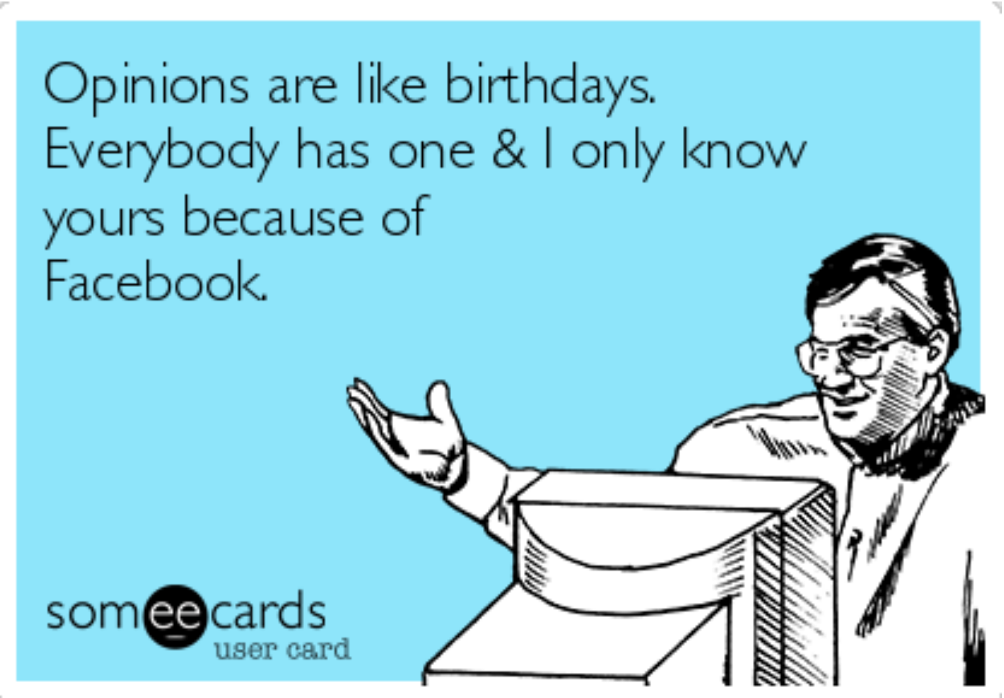

Use social media to sense opinions

CNN
How social media is revolutionizing debates

People's opinion about political discourse

The New York Times
Campaigns Use Social Media to Lure Younger Voters

Investors' sentiment about stocks

Nasdaq
Leveraging Social Media

Business Insider UK
Startups are setting up funds based on what is trending on Twitter

Brand sentiment and reputation

TC
Twitter Unveils A New Set Of Brand-Centric Analytics

The New York Times
Social Media Are Giving a Voice to Taste Buds

What about opinion dynamics?

Complex stochastic processes

(often over a network)

Example of opinion dynamics

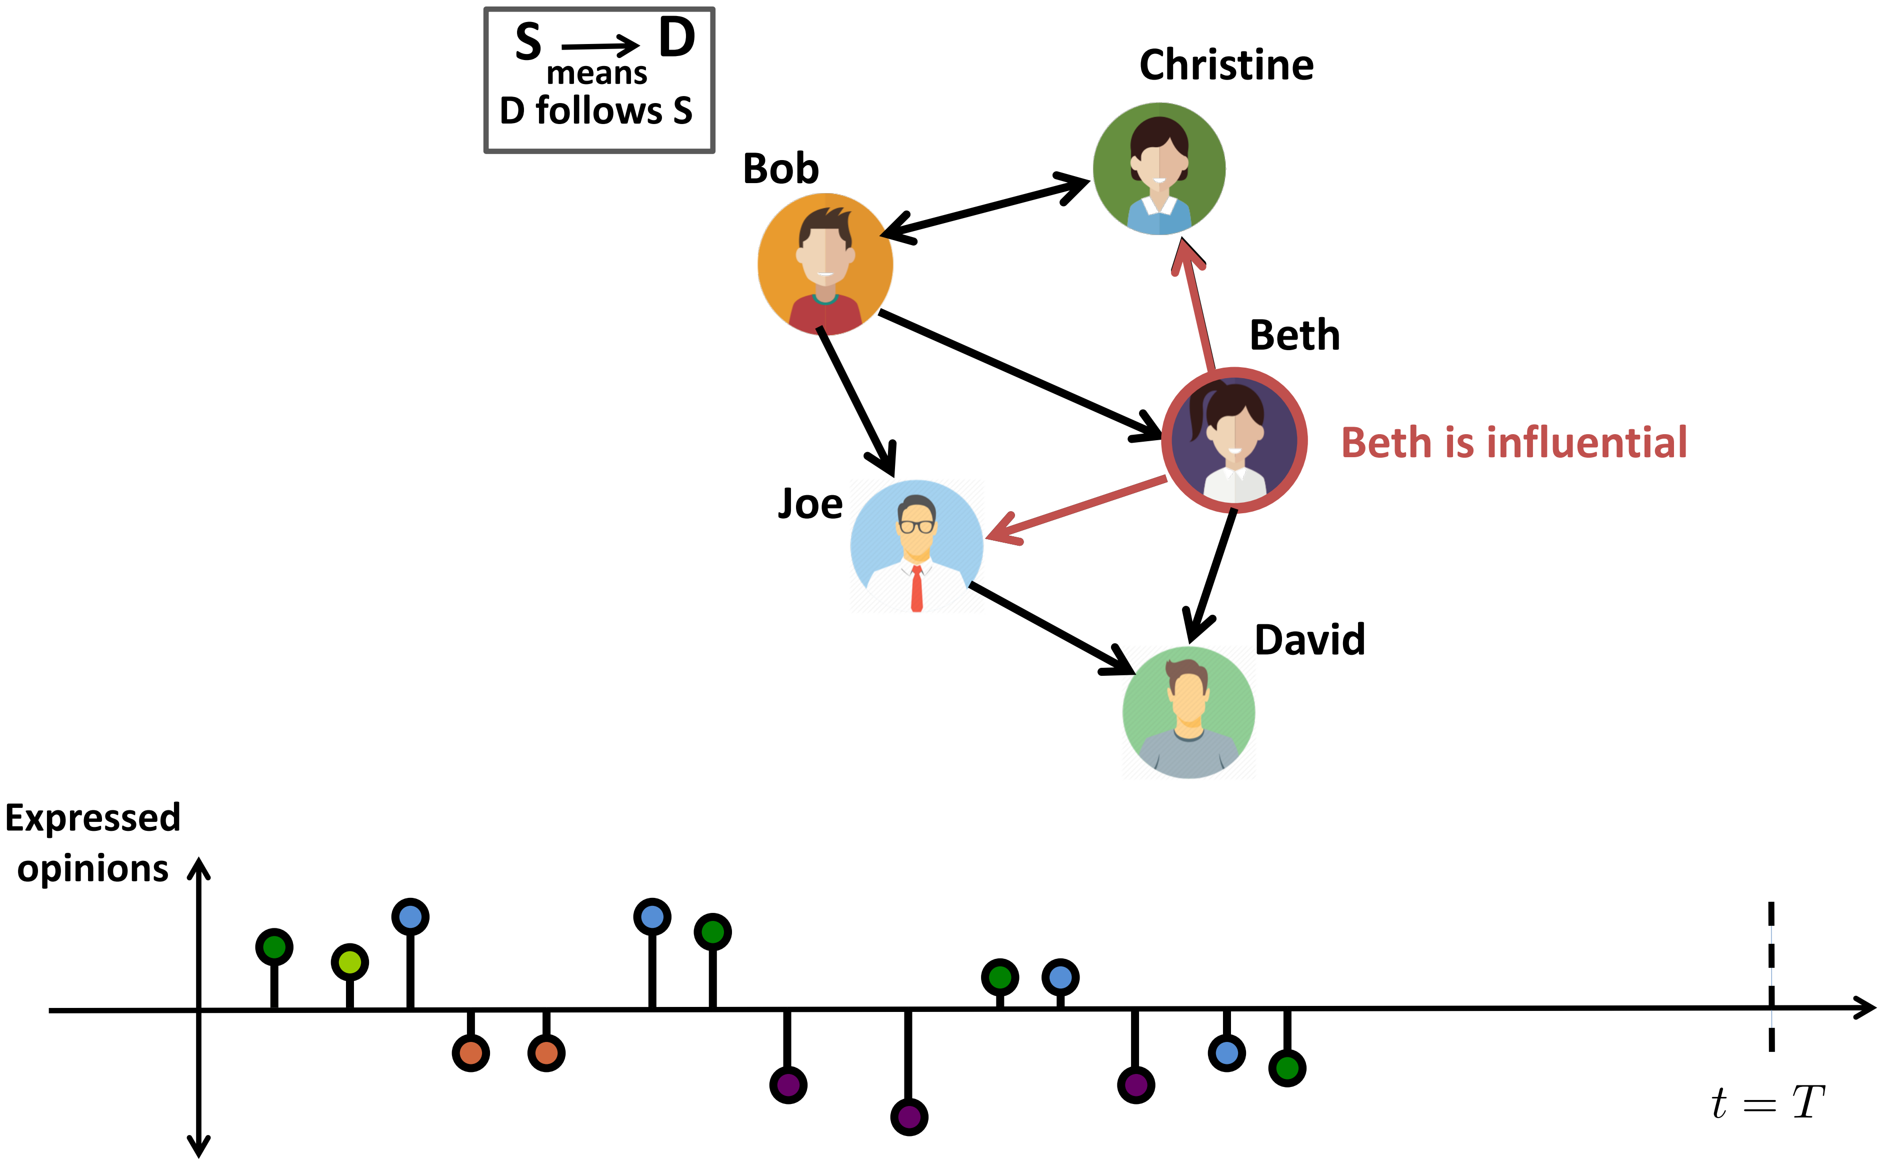

Model of opinion dynamics

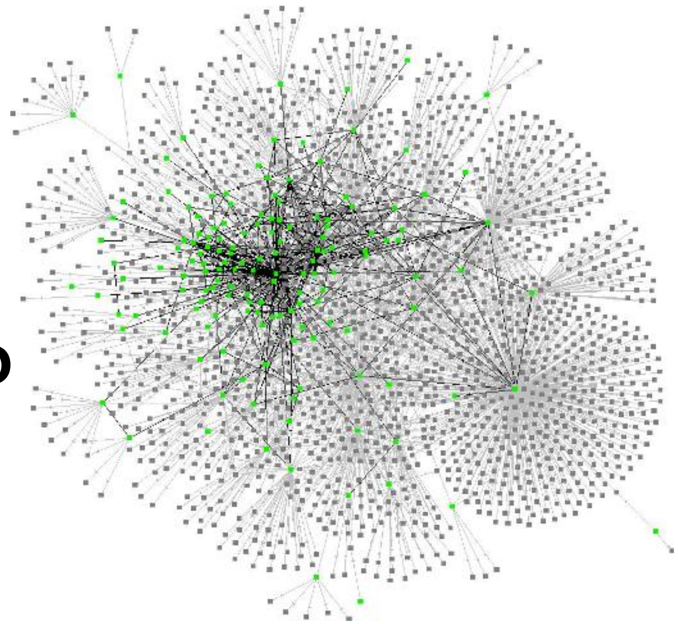

Can we **design a realistic model** that **fits real fine-grained opinion traces**?

Why this goal?

Predict (infer) opinions, even if not expressed!

Identify opinion leaders

Sergey Bilak

Traditional models of opinion dynamics

There are a lot of
theoretical models of opinion dynamics,
but…

1. **Do not distinguish between latent and expressed opinions**

2. **Opinions are updated sequentially in discrete time**

3. **Difficult to learn from fine-grained data and thus inaccurate predictions**

4. **Focus on steady state, neglecting transient behavior**

Key ideas of Marked TPP model

Latent opinions
vs
expressed opinions

Informational and social influence

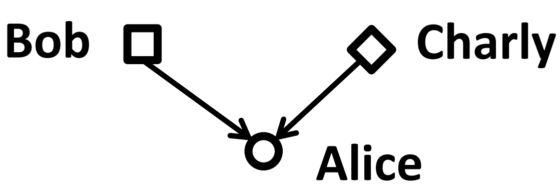

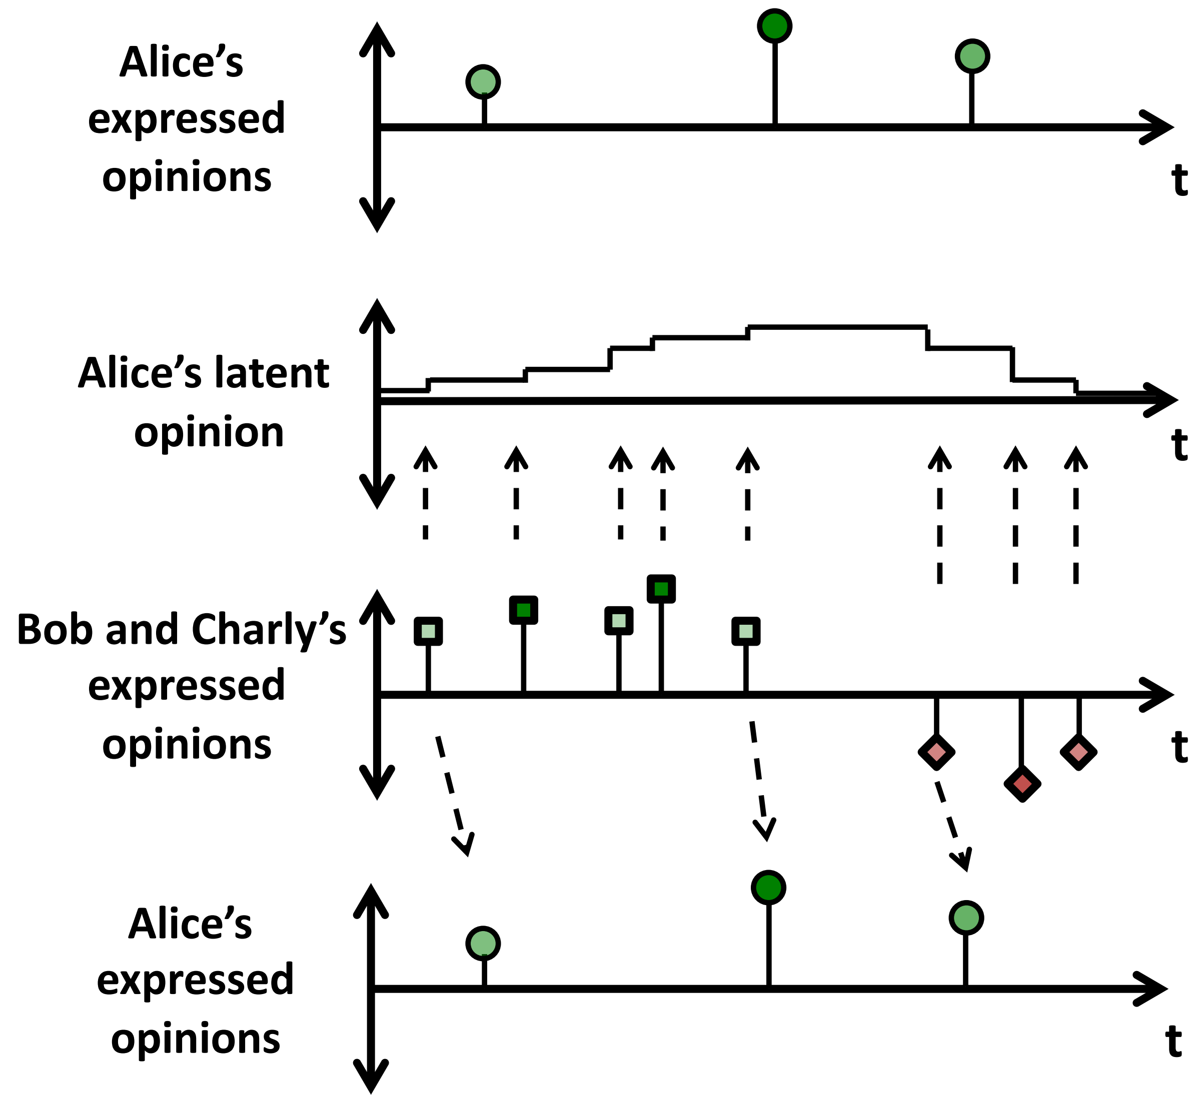

[De et al., NIPS 2016]

Message representation

We represent messages using **marked temporal point processes**:

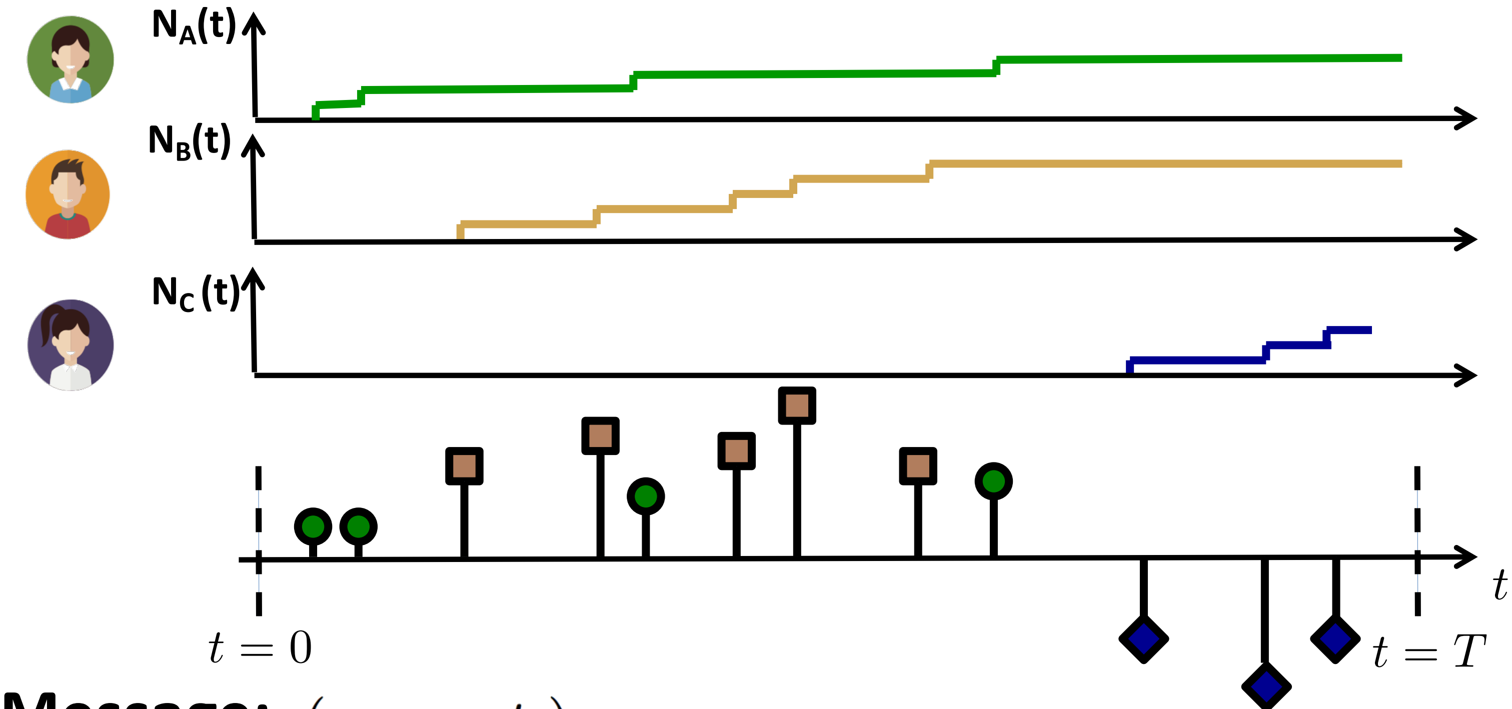

Message: (u_i, m_i, t_i)

User

Sentiment (mark) → **Noisy observation of latent opinion**

Time

[De et al., NIPS 2016]

Message intensity

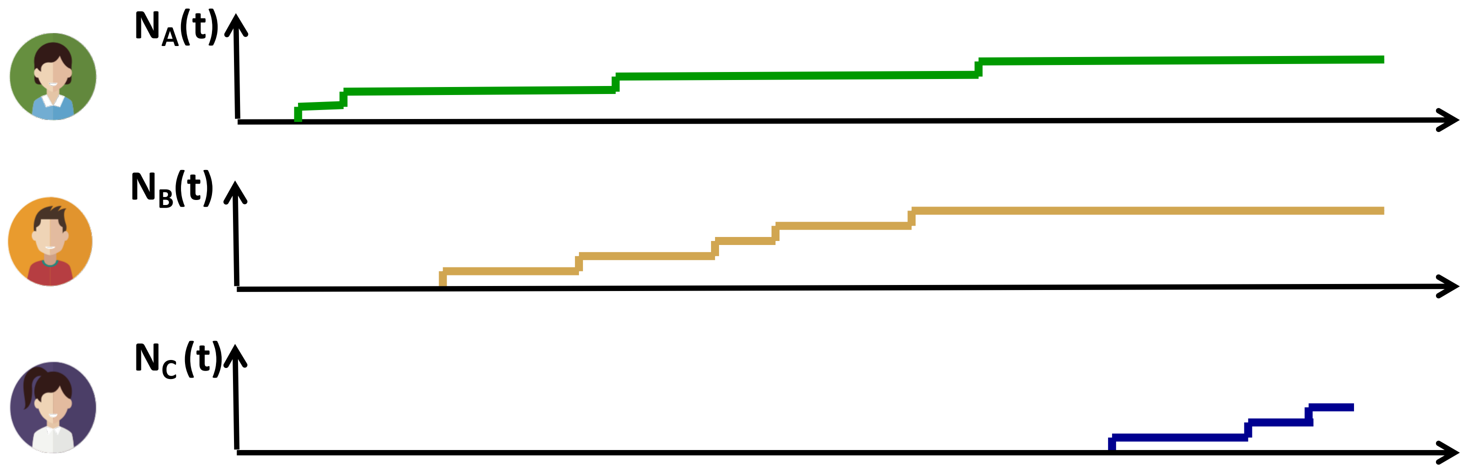

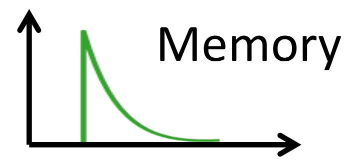

$$\underbrace{\lambda_u^*(t)}_{\text{User's intensity}} = \underbrace{\mu_u}_{\text{Messages on her own initiative}} + \sum_{v\in u\cup\mathcal{N}(u)} b_{vu} \underbrace{\sum_{e_i\in\mathcal{H}_v(t)} \kappa(t-t_i)}_{\text{Previous messages by user v}}$$

Influence from user v on user u

Hawkes process

[De et al., NIPS 2016]

Sentiment distribution

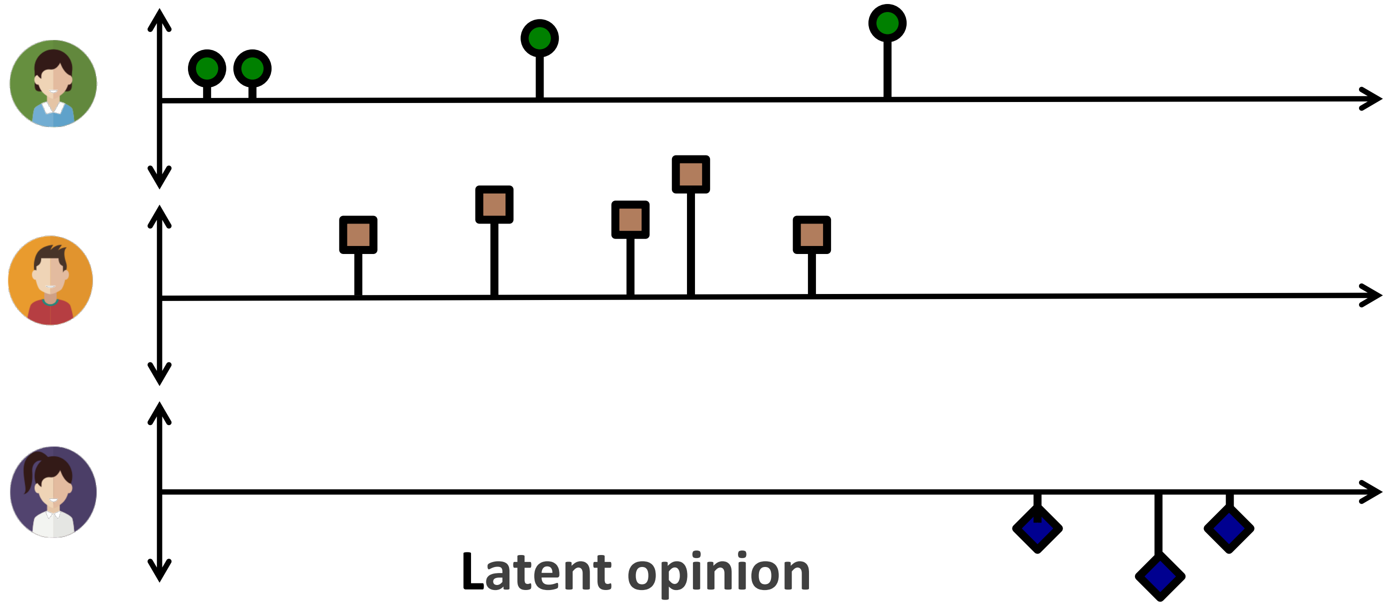

Latent opinion

Sentiment: $m_u \sim p(m|x_u^*(t))$

It depends on the recorded data

Continuous (based on sentiment analysis):

$$p(m|x_u^*(t)) = \mathcal{N}(x_u(t), \sigma_u)$$

Discrete (based on upvotes/downvotes):

$$p(m|x_u^*(t)) = 1/(1 + \exp(-m \cdot x_u(t)))$$

[De et al., NIPS 2016]

Stochastic process for (latent) opinions

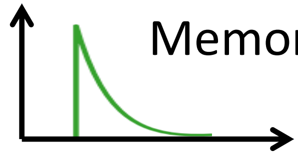

$$\underbrace{x_u^*(t)}_{\text{User's latent opinion}} = \underbrace{\alpha_u}_{\text{User's initial opinion}} + \sum_{v\in\mathcal{N}(u)} \underset{\text{influence from user v on user u}}{a_{vu}} \underbrace{\sum_{e_i\in\mathcal{H}_v(t)} m_i g(t-t_i)}_{\text{Previous sentiment by user v}}$$

[De et al., NIPS 2016]

Stubbornness, conformity, and compromise

The model allows for:

Stubborn users $x_u^*(t) = \alpha_u$

Compromised users

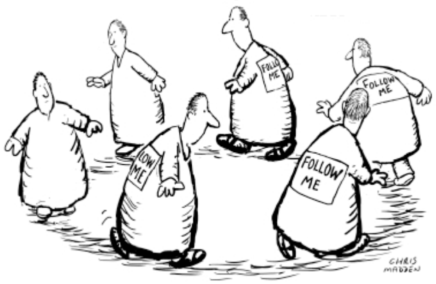

Conforming users $x_u^*(t) = \sum_{v \in \mathcal{N}(u)} a_{vu} \sum_{e_i \in \mathcal{H}_v(t)} m_i g(t - t_i)$

Example: positive opinions *win*

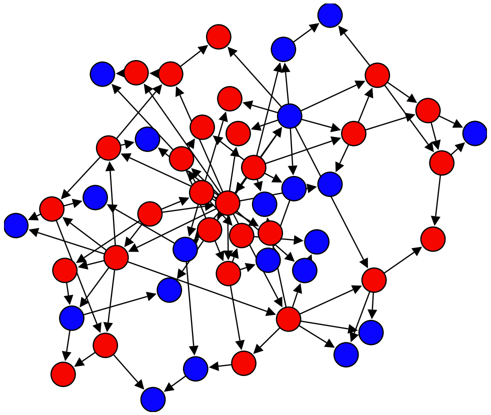

t = 0

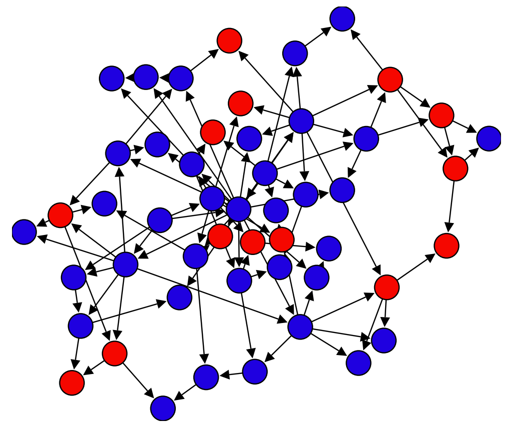

t=T

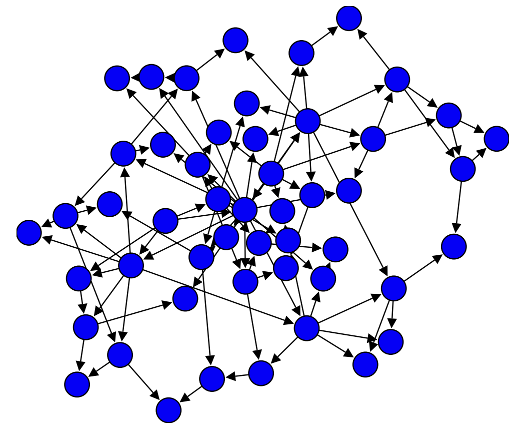

t=∞

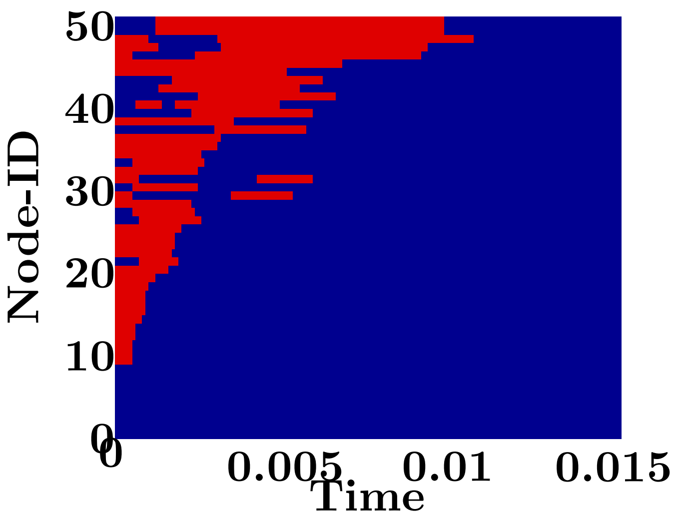

Opinions per node

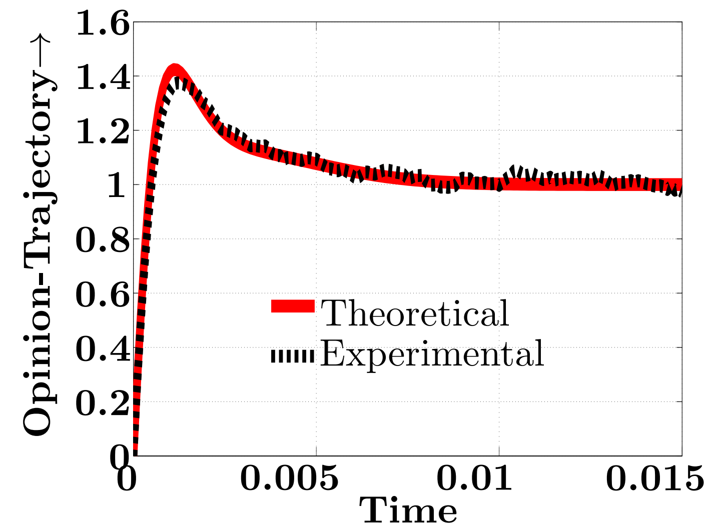

Average opinion

Example: negative opinions *win*

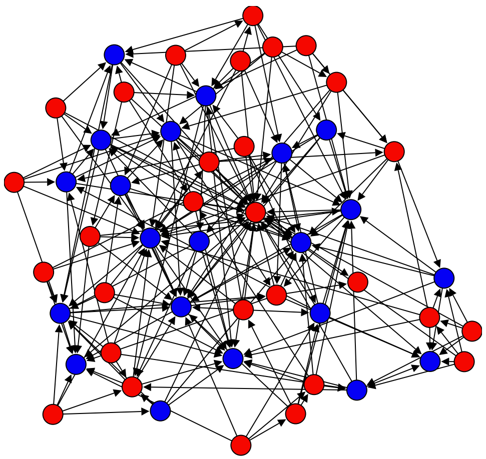

t = 0

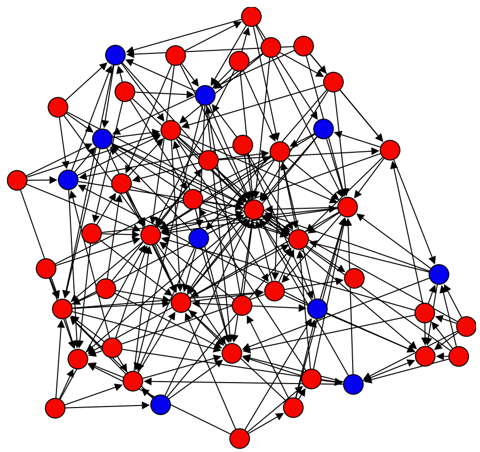

t=T

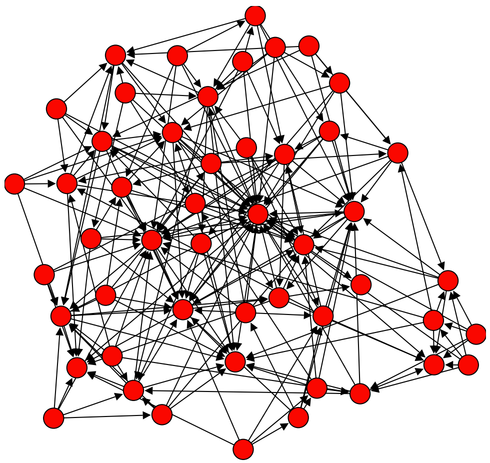

t=∞

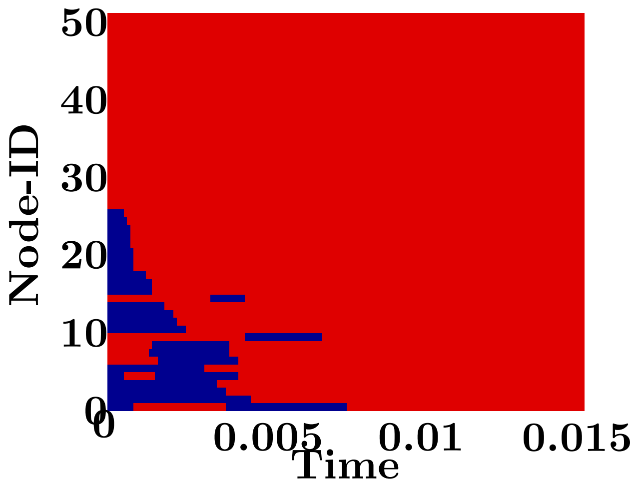

Opinions per node

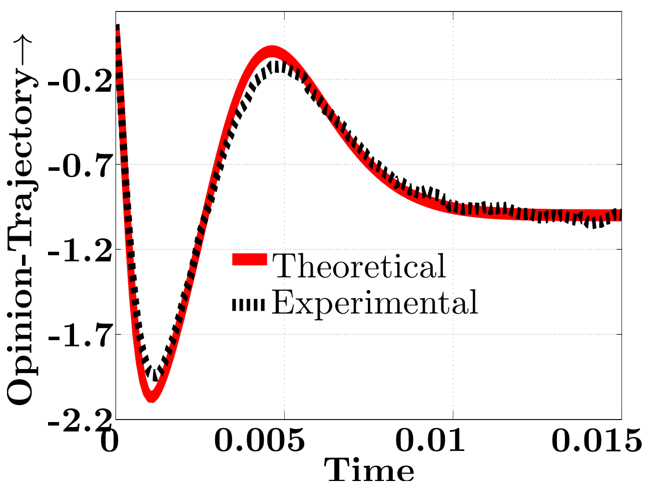

Average opinion

Example: opinions get polarized

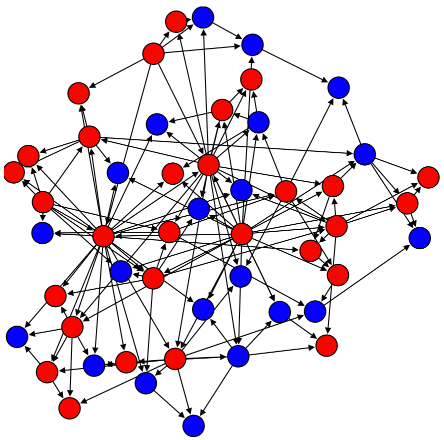

t = 0

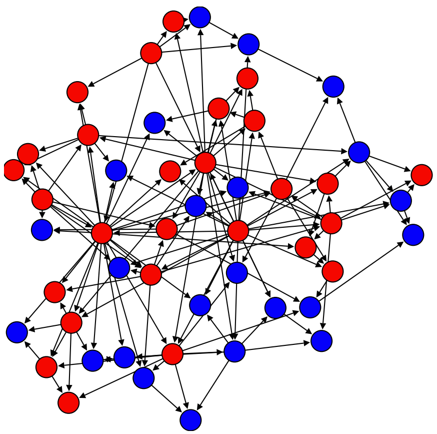

t=T

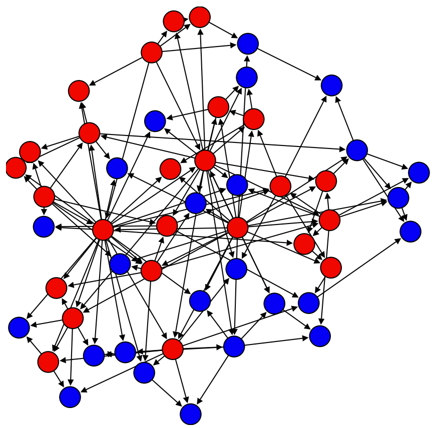

t=∞

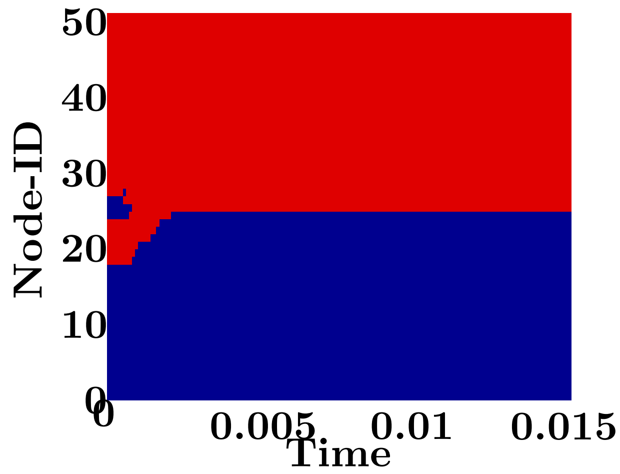

Opinions per node

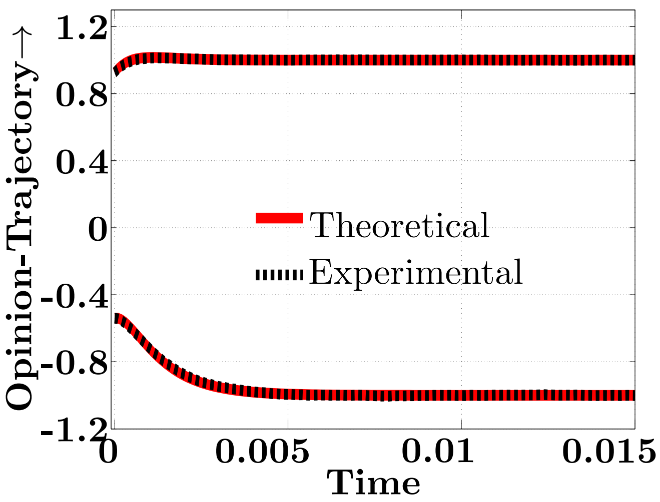

Average opinion

Model inference from opinion data

Events likelihood

$$\underbrace{\sum_{e_i \in \mathcal{H}(T)} \log p(m_i | x^*_{u_i}(t_i))}_{\text{Message sentiments (marks)}} + \underbrace{\sum_{e_i \in \mathcal{H}(T)} \log \lambda^*_{u_i}(t_i) - \sum_{u \in \mathcal{V}} \int_0^T \lambda^*_u(\tau)\, d\tau}_{\text{Message times}}$$

> **Theorem.** The **maximum likelihood** problem is **convex in the model parameters.**

Markov property $\Rightarrow$ **Sums and integrals in linear time!**

[De et al., NIPS 2016]

Opinion model as Jump SDEs

Proposition. The tuple ($\mathbf{x}^*(t)$, $\boldsymbol{\lambda}^*(t)$, $\mathbf{N}(t)$) is a **Markov process**, whose dynamics are defined by the following **marked jumped stochastic differential equations** (SDEs)

Network!

Latent opinions

Informational influence

Expressed opinions

$$d\mathbf{x}^*(t) = \omega(\boldsymbol{\alpha} - \mathbf{x}^*(t))dt + \mathbf{A}(\mathbf{m}(t) \odot d\boldsymbol{N}(t))$$

$$d\boldsymbol{\lambda}^*(t) = \nu(\boldsymbol{\mu} - \boldsymbol{\lambda}^*(t))dt + \mathbf{B}\, d\boldsymbol{N}(t)$$

Message intensities

Temporal influence

Network!

[De et al., NIPS 2016]

Proof sketch of the proposition (I)

Let's do it for one-dimensional Hawkes:

$$\lambda^*(t) = \lambda_0(t) + \alpha \int_0^t g(t-s)dN(s) \qquad g(t) = e^{-wt}\mathbb{I}(t \geq 0)$$

$$d\lambda^*(t) = \lambda_0'(t)dt + \alpha\, d\left(\int_0^t g(t-s)dN(s)\right)$$

$$\int_0^{t+dt} g(t+dt-s)\,dN(s) - \int_0^t g(t-s)\,dN(s)$$

$$= \int_0^{t+dt} \left(g(t-s) + g'(t-s)dt\right) dN(s) - \int_0^t g(t-s)\,dN(s)$$

$$= \int_t^{t+dt} g(t-s)\,dN(s) + dt \int_0^{t+dt} g'(t-s)\,dN(s)$$

Proof sketch of the proposition (II)

$$= \int_t^{t+dt} g(t-s)\, dN(s) + dt \int_0^{t+dt} g'(t-s)\, dN(s)$$

$$= g(0)dN(t) - w\, dt \int_0^{t+dt} g(t-s)\, dN(s)$$

$$= dN(t) - w\, dt \int_0^{t} g(t-s)\, dN(s)$$

$$= dN(t) + \frac{w}{\alpha}[\lambda_0(t) - \lambda^*(t)]dt.$$

$$d\lambda^*(t) = [\lambda_0'(t) + w\lambda_0(t) - w\lambda^*(t)]\, dt + \alpha dN(t)$$

$$d\boldsymbol{\lambda}^*(t) = \nu(\boldsymbol{\mu} - \boldsymbol{\lambda}^*(t))dt + \mathbf{B}\, d\boldsymbol{N}(t)$$

Opinion forecasting

For forecasting, we compute **conditional averages:**

History up to t_0

$$\mathbb{E}_{\mathcal{H}_t \setminus \mathcal{H}_{t_0}}[\boldsymbol{x}^*(t) | \mathcal{H}_{t_0}]$$

Sources of Randomness

$p(m|x_u^*(t))$ $\boldsymbol{\lambda}^*(t)$

$$d\mathbf{x}^*(t) = \omega(\boldsymbol{\alpha} - \mathbf{x}^*(t))dt + \mathbf{A}(\mathbf{m}(t) \odot d\mathbf{N}(t))$$

If $p(m|x_u^*(t)) = \mathcal{N}(x_u(t), \sigma_u)$

- $\boldsymbol{\lambda}^*(t) = \boldsymbol{\mu}$ **analytical solution (Th. 2)**
- $\boldsymbol{\lambda}^*(t) : b_{vu} = 0 \quad v \neq u$ **numerical solution (Th. 4)**

Otherwise: **Sampling based solution**

Opinion forecasting

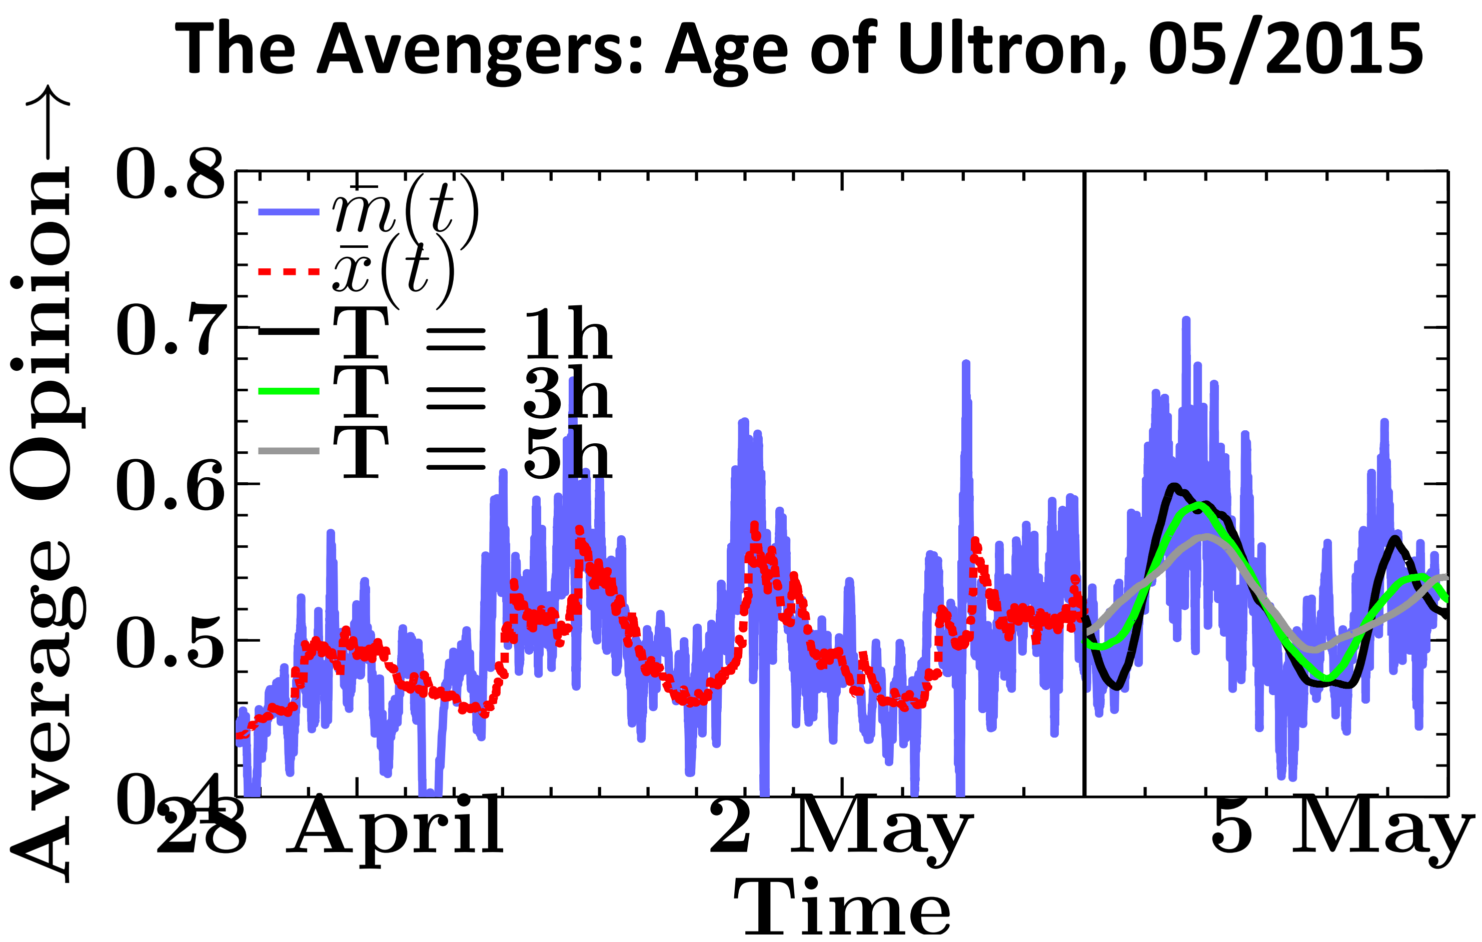

The forecasted opinion becomes less accurate as T increases, as one may expect.

[De et al., NIPS 2016]