Viral marketing

with Stochastic optimal control of TPP

HUMAN-CENTERED MACHINE LEARNING

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

MAX PLANCK INSTITUTE
FOR SOFTWARE SYSTEMS

Maximizing activity in a social network

Can we **steer users' behavior to maximize activity** in a **social network**?

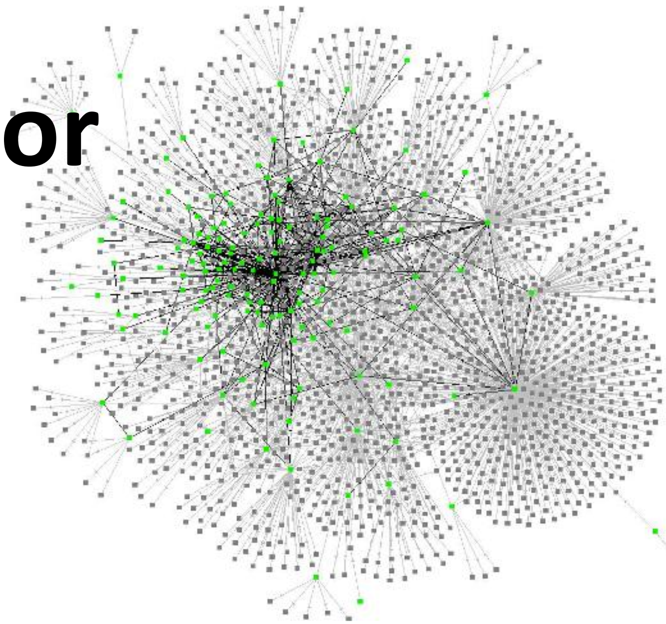

Twitter Stock Tumbles After Drop in User Engagement

7 Ways to Increase Your Social Media Engagement

Endogenous and exogeneous events

Exogenous activity

Users' actions due to drives external to the network

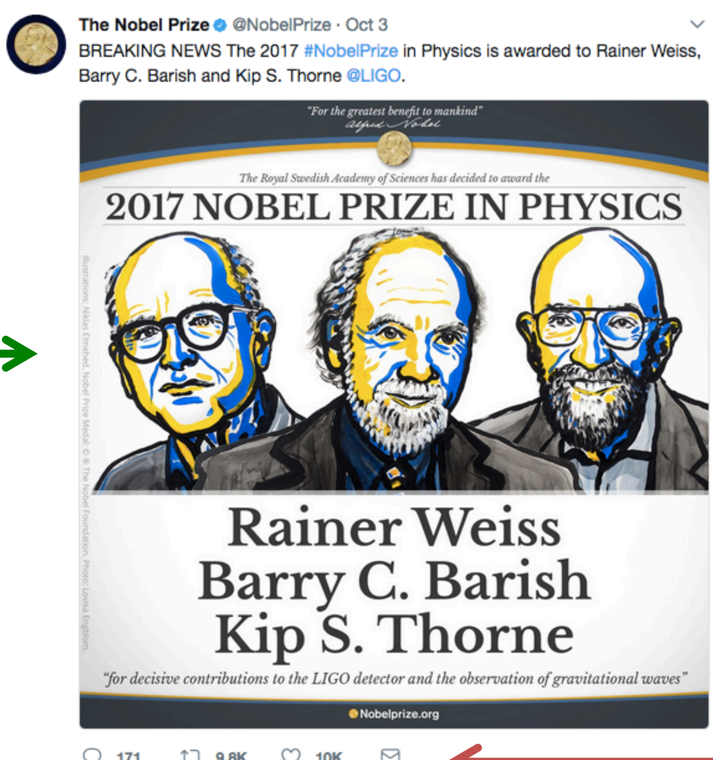

Endogenous activity

Users' *responses* to other users' actions in the network

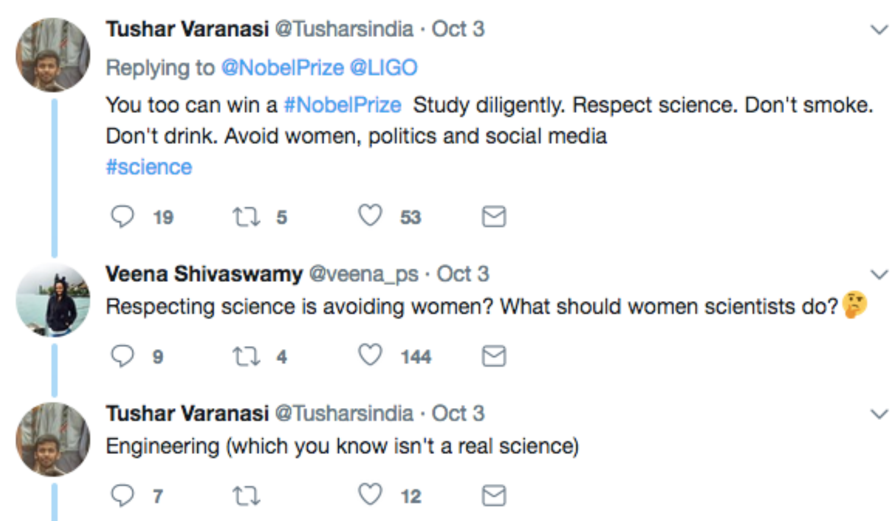

Multidimensional Hawkes process

For each user u, actions as a counting process $N_u(t)$

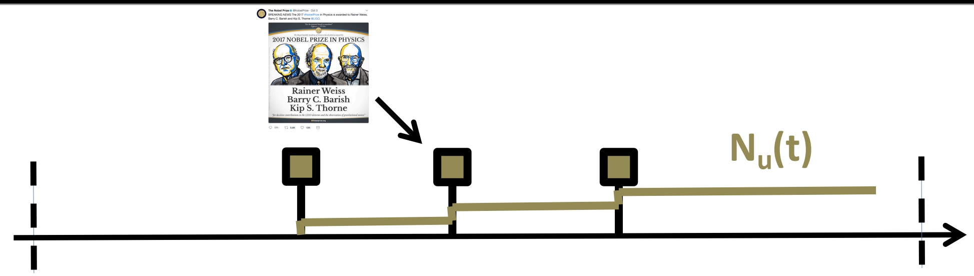

Intensities or rates
(Actions per time unit)

User influence matrix

Non-negative kernel (memory)

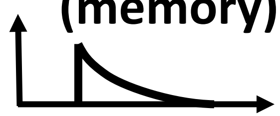

$$\boldsymbol{\lambda}^*(t) = \underbrace{\boldsymbol{\mu}_0}_{\text{Exogenous actions}} + \underbrace{\boldsymbol{A}\int_0^t \kappa(t-s)\,d\boldsymbol{N}(s)}_{\text{Endogenous actions}}$$

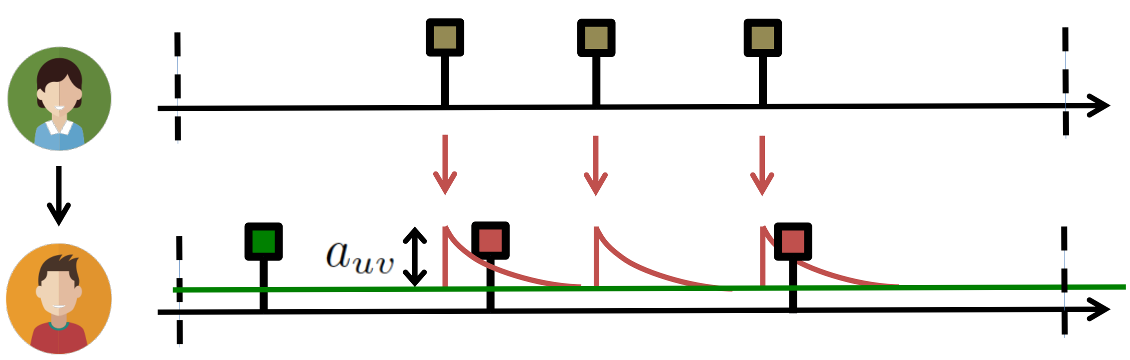

Steering endogenous actions

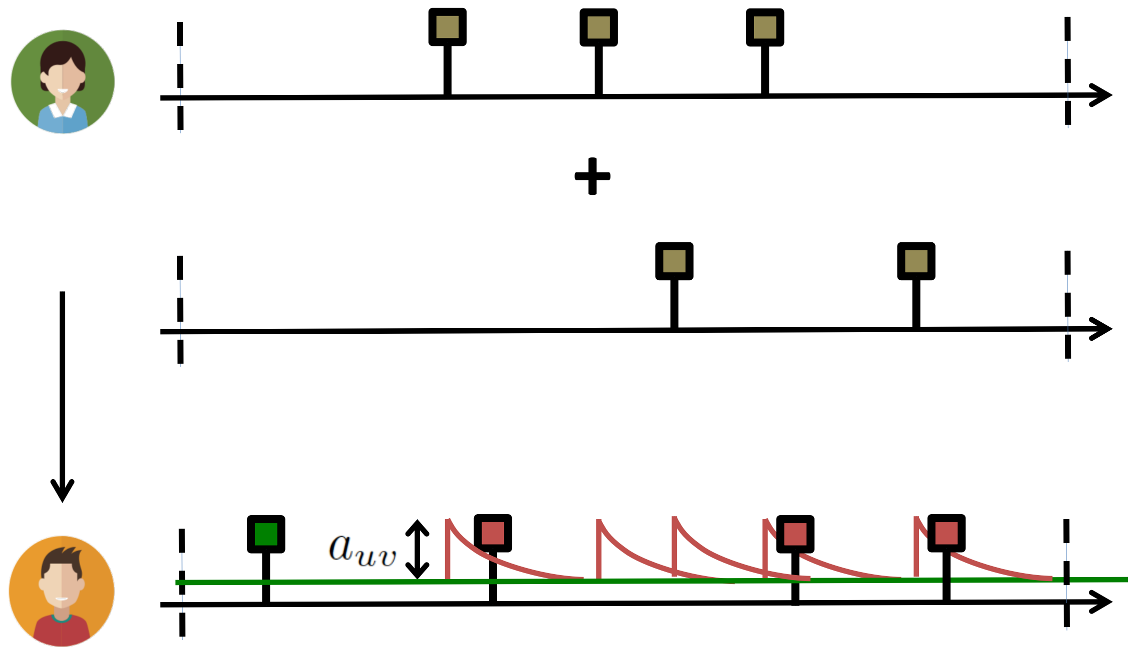

Organic actions

Directly incentivized actions

$$\boldsymbol{\lambda}^*(t) = \boldsymbol{\mu}_0 + \boldsymbol{A} \int_0^t \kappa(t-s)\, d\boldsymbol{N}(s) + \boldsymbol{A} \int_0^t \kappa(t-s)\, d\boldsymbol{M}(s)$$

Intensities of directly incentivized actions $\mathbb{E}[d\boldsymbol{M}(t)|\mathcal{H}(t)] = \boldsymbol{u}(t)\,dt$

Directly incentivized actions

[Zarezade et al., 2018]

Cost to go & Bellman's principle of optimality

Optimization problem

$$\begin{cases} \underset{\boldsymbol{u}(t_0,t_f]}{\text{minimize}} \quad \mathbb{E}_{(\boldsymbol{N},\boldsymbol{M})(t_0,t_f]}\left[\phi(\boldsymbol{\lambda}(t_f)) + \int_{t_0}^{t_f} \overbrace{\ell(\boldsymbol{\lambda}(t),\boldsymbol{u}(t))}^{\textbf{Loss}}\,dt\right] \\ \text{subject to} \quad u_i(t) \geq 0,\ \forall t \in (t_0,t_f],\ i=1,\ldots,n \end{cases}$$

Dynamics defined by Jump SDEs

$$d\boldsymbol{\lambda}(t) = [w\boldsymbol{\mu}_0 - w\boldsymbol{\lambda}(t)]\,dt + \boldsymbol{A}\,d\boldsymbol{N}(t) + \boldsymbol{A}\,d\boldsymbol{M}(t)$$

To solve the problem, we first define the corresponding **optimal cost-to-go:**

$$J(\boldsymbol{\lambda}(t),t) = \min_{\boldsymbol{u}(t,t_f]} \mathbb{E}_{(\boldsymbol{N},\boldsymbol{M})(t,t_f]}\left[\phi(\boldsymbol{\lambda}(t_f)) + \int_t^{t_f} \ell(\boldsymbol{\lambda}(s),\boldsymbol{u}(s))\,ds\right]$$

The cost-to-go, evaluated at t_0, recovers the optimization problem!

[Zarezade et al., 2018]

Cost to go & Bellman's principle of optimality

Loss

Op

This is a stochastic optimal control problem for jump SDEs
(we know how to solve this!)

Jump SDEs

To solve the problem, we first define the corresponding **optimal cost-to-go:**

$$J(\boldsymbol{\lambda}(t), t) = \min_{\boldsymbol{u}(t,t_f]} \mathbb{E}_{(\boldsymbol{N},\boldsymbol{M})(t,t_f]} \left[\phi(\boldsymbol{\lambda}(t_f)) + \int_t^{t_f} \ell(\boldsymbol{\lambda}(s), \boldsymbol{u}(s))\, ds \right]$$

The cost-to-go, evaluated at t_0, recovers the optimization problem!

[Zarezade et al., 2018]

Hamilton-Jacobi-Bellman (HJB) equation

Lemma. The **optimal cost-to-go** satisfies **Bellman's Principle of Optimality**

$$J(\boldsymbol{\lambda}(t), t) = \min_{\boldsymbol{u}(t,t+dt]} \left\{ \mathbb{E}_{(\boldsymbol{N},\boldsymbol{M})(t,t+dt]} \left[J(\boldsymbol{\lambda}(t+dt), t+dt) \right] + \ell(\boldsymbol{\lambda}(t), \boldsymbol{u}(t))\, dt \right\}$$

$$dJ(\boldsymbol{\lambda}(t), t) = J(\boldsymbol{\lambda}(t+dt), t+dt) - J(\boldsymbol{\lambda}(t), t)$$

$$0 = \min_{\boldsymbol{u}(t,t+dt]} \left\{ \mathbb{E}_{(\boldsymbol{N},\boldsymbol{M})(t,t+dt]} \left[dJ(\boldsymbol{\lambda}(t), t) \right] + \ell(\boldsymbol{\lambda}(t), \boldsymbol{u}(t))\, dt \right\}$$

$$d\boldsymbol{\lambda}(t) = [w\boldsymbol{\mu}_0 - w\boldsymbol{\lambda}(t)]\, dt + \boldsymbol{A}\, d\boldsymbol{N}(t) + \boldsymbol{A}\, d\boldsymbol{M}(t)$$

Hamilton-Jacobi-Bellman (HJB) equation

Partial differential equation in J (with respect to λ and t)

[Zarezade et al., 2018]

Solving the HJB equation

Consider a quadratic loss

$$\ell(\boldsymbol{\lambda}(t), \boldsymbol{u}(t)) = \underbrace{-\frac{1}{2}\boldsymbol{\lambda}^T(t)\,\boldsymbol{Q}\,\boldsymbol{\lambda}(t)}_{\text{Rewards organic actions}} + \underbrace{\frac{1}{2}\boldsymbol{u}^T(t)\,\boldsymbol{S}\,\boldsymbol{u}(t)}_{\text{Penalizes directly incentivizes actions}}$$

We propose $J(\boldsymbol{\lambda}(t), t)$ and then show that the optimal intensity is:

$$\boldsymbol{u}^*(t) = -\boldsymbol{S}^{-1}\left[\boldsymbol{A}^T\boldsymbol{g}(t) + \boldsymbol{A}^T\boldsymbol{H}(t)\boldsymbol{\lambda}(t) + \frac{1}{2}\,\mathrm{diag}(\boldsymbol{A}^T\boldsymbol{H}(t)\boldsymbol{A})\right]$$

Computed offline once!

Closed form solution to a first order ODE

Solution to a matrix Riccati differential equation

[Zarezade et al., 2018]

The Cheshire algorithm

Intuition

Steering actions means sampling action user & times from u*(t)

More in detail

Since the intensity function u*(t) is stochastic, we sample from it using:

- → **Superposition principle**
- → **Standard thinning**

> **It only requires sampling $1^T N(t_f)$ from inhomog. Poisson!**

Easy to implement

Algorithm 1: CHESHIRE: it returns user i and time τ for

Initialization:
Compute $\boldsymbol{H}(t)$ and $\mathbf{g}(t)$;
$\mathbf{u}(t) \leftarrow -\boldsymbol{S}^{-1}\left[\boldsymbol{A}^T(\boldsymbol{g}(t) + \boldsymbol{H}(t)\boldsymbol{\mu}_0) + \frac{1}{2}\,\mathrm{diag}(\boldsymbol{A}^T\boldsymbol{H}(t)\boldsymbol{A})\right]$;
General subroutine:
$(i, \tau) \leftarrow \mathit{Sample}(\mathbf{u}(t))$;
$(j, s) \leftarrow \mathit{NextAction}()$;
while $s < \tau$ **do**
    $\boldsymbol{\lambda}_N(t) \leftarrow \boldsymbol{A}\mathbf{e}_j\kappa(t - s)$;
    $\mathbf{u}_N(t) \leftarrow -\boldsymbol{S}^{-1}\boldsymbol{A}^T\boldsymbol{H}(t)\boldsymbol{\lambda}_N(t)$;
    $(k, r) \leftarrow \mathit{Sample}(\mathbf{u}_N(t))$;
    if $r < \tau$ **then**
        $\tau \leftarrow r$;
        $i \leftarrow k$;
    $\mathbf{u}(t) \leftarrow \mathbf{u}(t) + \mathbf{u}_N(t)$;
    $(j, s) \leftarrow \mathit{NextAction}()$;
$\boldsymbol{\lambda}_M(t) \leftarrow \boldsymbol{A}\mathbf{e}_i\kappa(t - \tau)$;
$\mathbf{u}_M(t) \leftarrow -\boldsymbol{S}^{-1}\boldsymbol{A}^T\boldsymbol{H}(t)\boldsymbol{\lambda}_M(t)$;
$\mathbf{u}(t) \leftarrow \mathbf{u}(t) + \mathbf{u}_M(t)$;
return (i, τ)

[Zarezade et al., 2018]

Experiments on real data

Five **Twitter** datasets (users) where actions are tweets and retweets

1. Fit model parameters

Network inference!

$$d\boldsymbol{\lambda}(t) = [w\boldsymbol{\mu}_0 - w\boldsymbol{\lambda}(t)]\, dt + \boldsymbol{A}\, d\boldsymbol{N}(t)$$

exogeneous rate ($\boldsymbol{\mu}_0$) **influence matrix** ($\boldsymbol{A}$)

2. Simulate steering endogenous actions

$$d\boldsymbol{\lambda}(t) = [w\boldsymbol{\mu}_0 - w\boldsymbol{\lambda}(t)]\, dt + \boldsymbol{A}\, d\boldsymbol{N}(t) + \boldsymbol{A}\, d\boldsymbol{M}(t)$$

directly incentivized tweets chosen by each method ($\boldsymbol{M}(t)$)

[Zarezade et al., 2018]

Evaluation metrics & baselines

Evaluation metrics

- $\bar{N}(t) = \sum_{u \in \mathcal{V}} \mathbb{E}[N_u(t)]$
  - → **Average number of not directly incentivized tweets**
- $\bar{t}_{30K}$
  - → **Average time to reach 30,000 not directly incentivized tweets**

Baselines

- **MSC** [Farajtabar et al., NIPS '16]
- **OPL** [Farajtabar et al., NIPS '14]
- **PRK** (Pagerank)
- **DEG** (Out-degree)

[Zarezade et al., 2018]

Performance vs. time

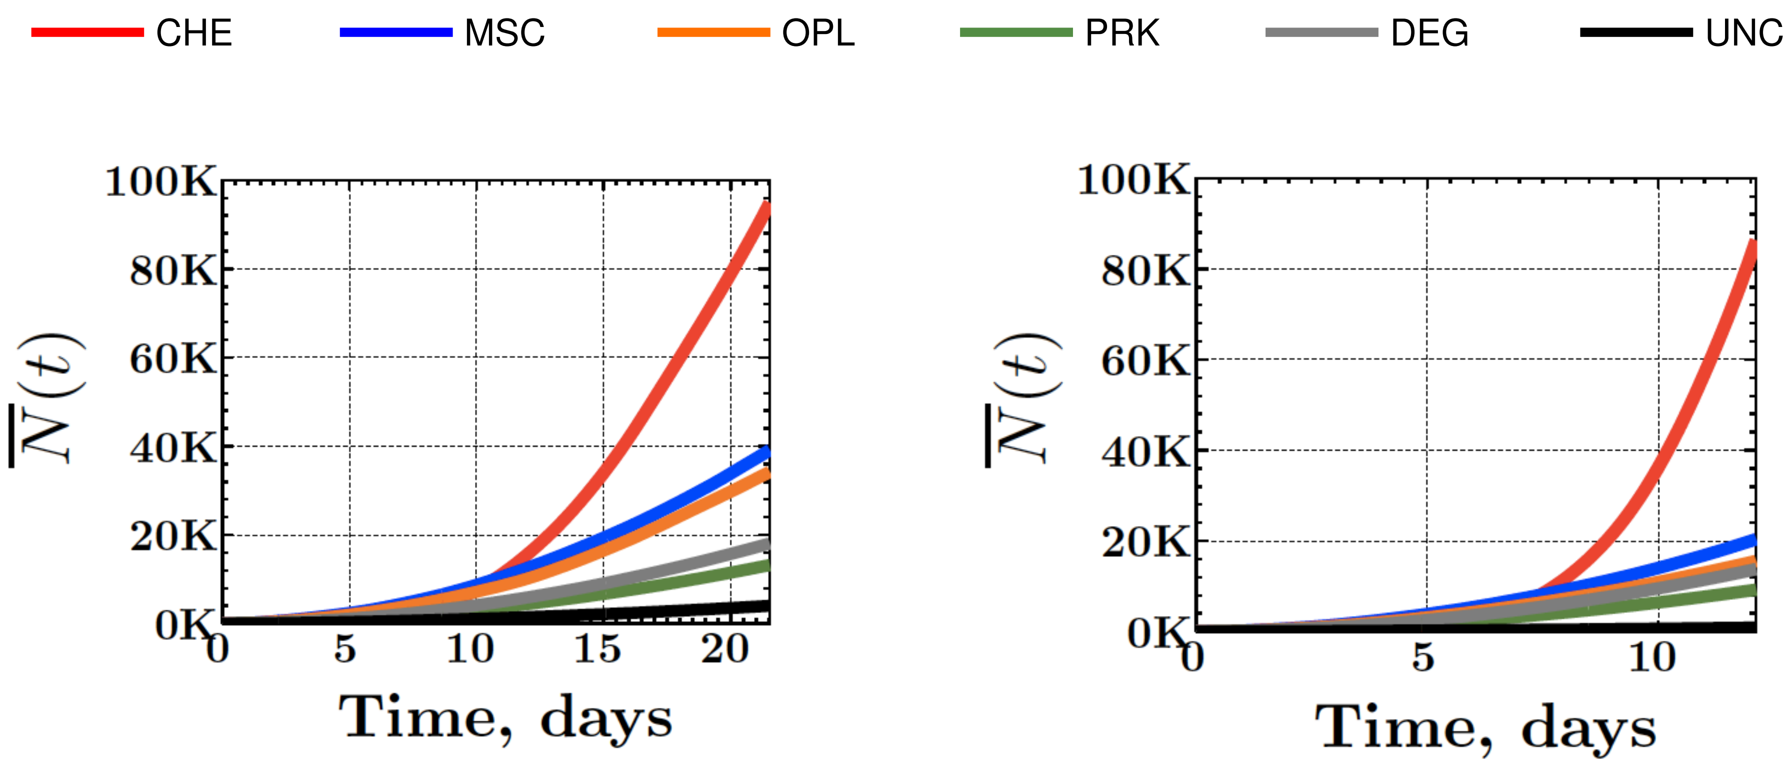

Sports, $M(t_f) \approx 5k$ **Series, $M(t_f) \approx 5k$**

Cheshire (in red) triggers 100%-400% more posts than the second best performer.

[Zarezade et al., 2018]

Performance vs. # of incentivized tweets

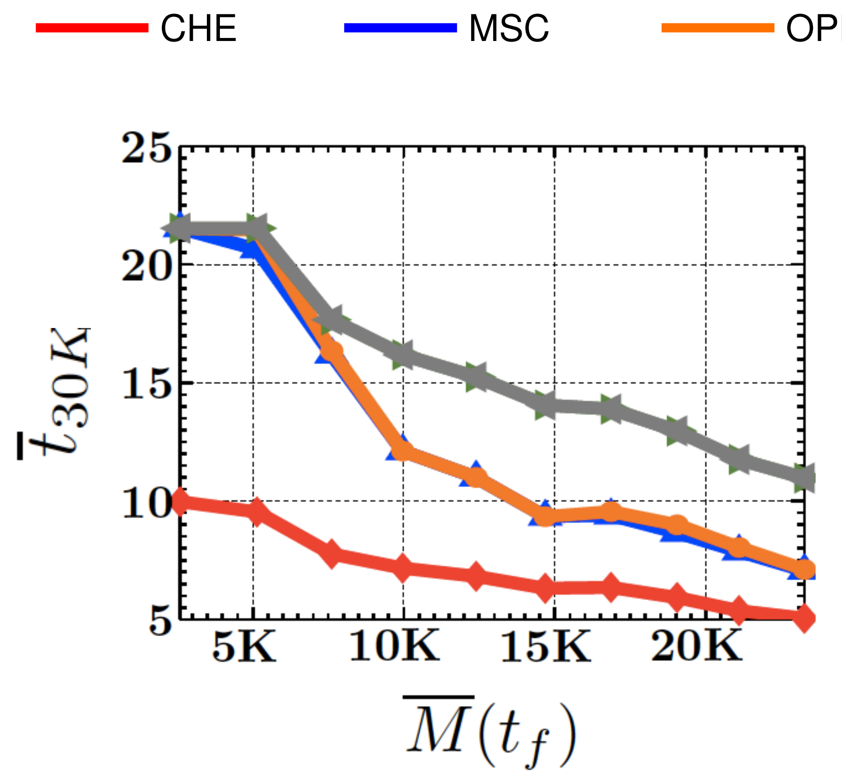

Sports, $M(t_f) \approx$ 5k

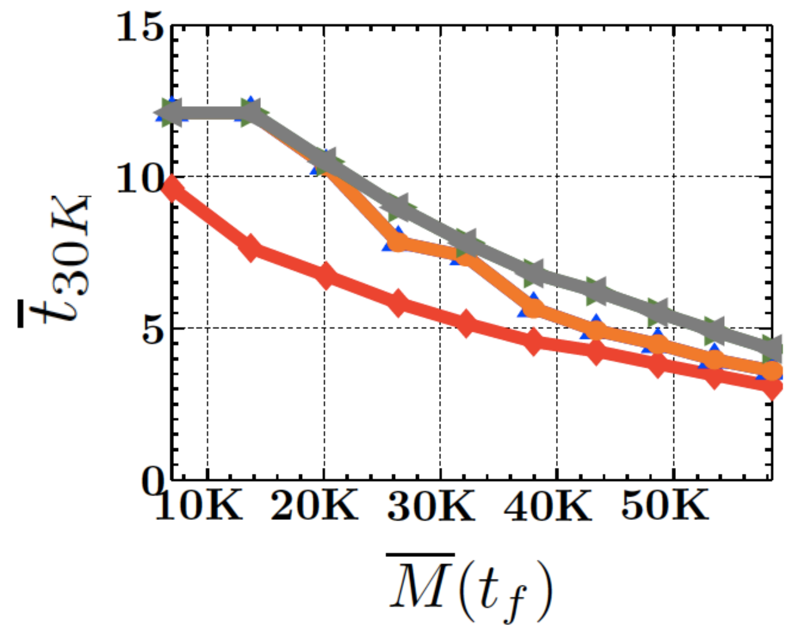

Series, $M(t_f) \approx$ 5k

Cheshire (in red) reaches 30K tweets 20-50% faster than the second best performer

[Zarezade et al., 2018]

Why Cheshire?

"the Cheshire Cat has the ability to appear and disappear in any location"

Alice's Adventures in Wonderland, Lewis Carroll