The popularity and engagement of online videos

Marian-Andrei Rizoiu

ComputationalMedia @ANU: http://cm.cecs.anu.edu.au

Popularity is complex

Daily Views

Popularity is complex but predictable

Daily
Views

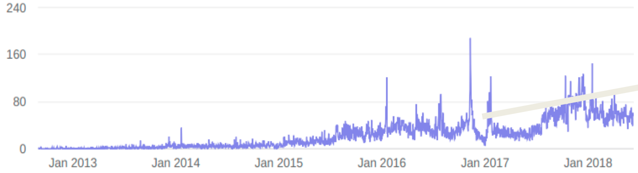

Daily
Shares

Asynchronous multiple sources help

Asynchronous multiple sources help

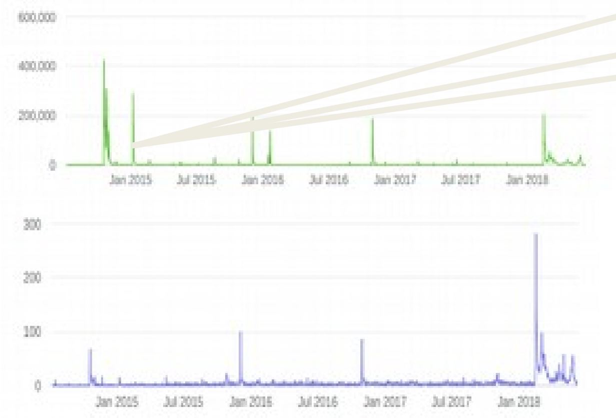

~15k Followers

~7k Followers

~7k Followers

Open questions:

- How can we design a new model for multiple asynchronous streams?
- What about latent/uncaptured sources?
- Is influence of user promoting content important?

Popularity (Current Landscape)

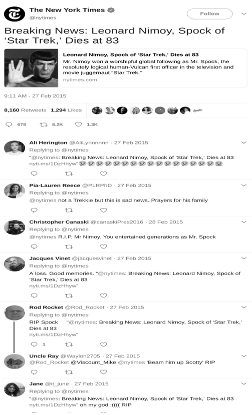

Individual user actions
Point Processes

Aggregated daily views
Daily Shares

Popularity (Current Landscape)

Individual user actions
Point Processes

?

Aggregated daily views
Daily Shares

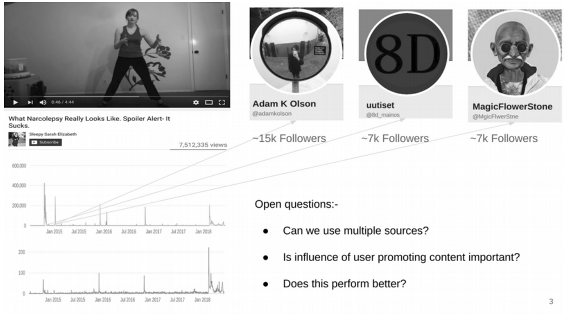

Popularity with Asynchronous Streams

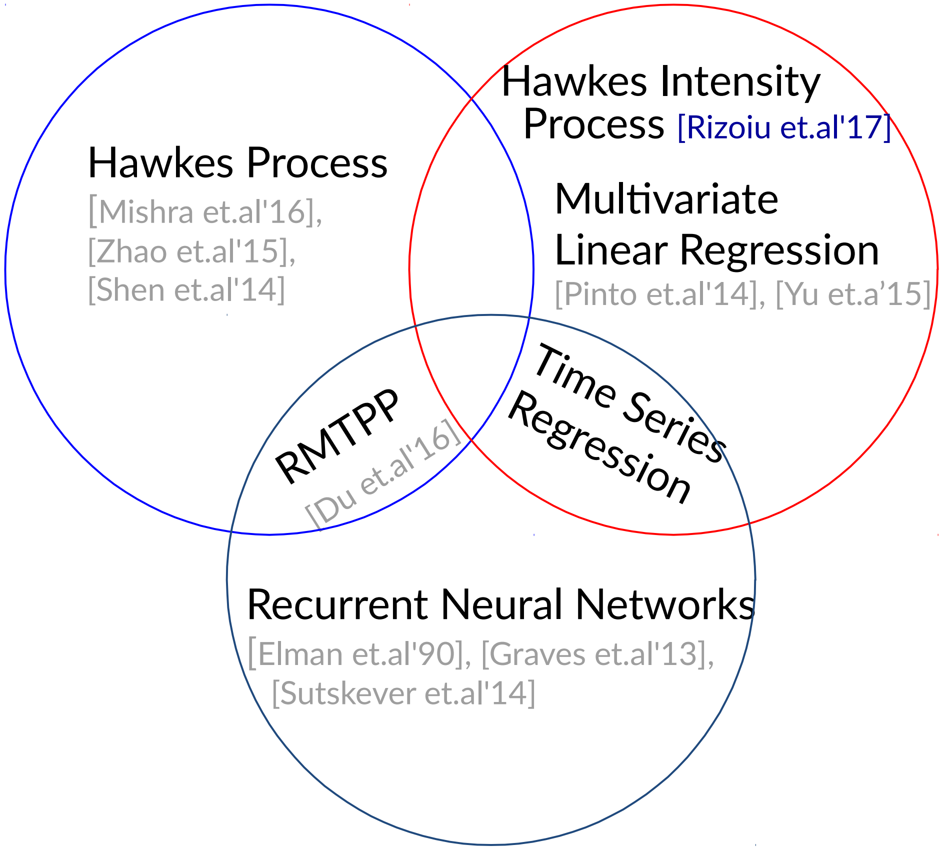

Popularity with Asynchronous Streams

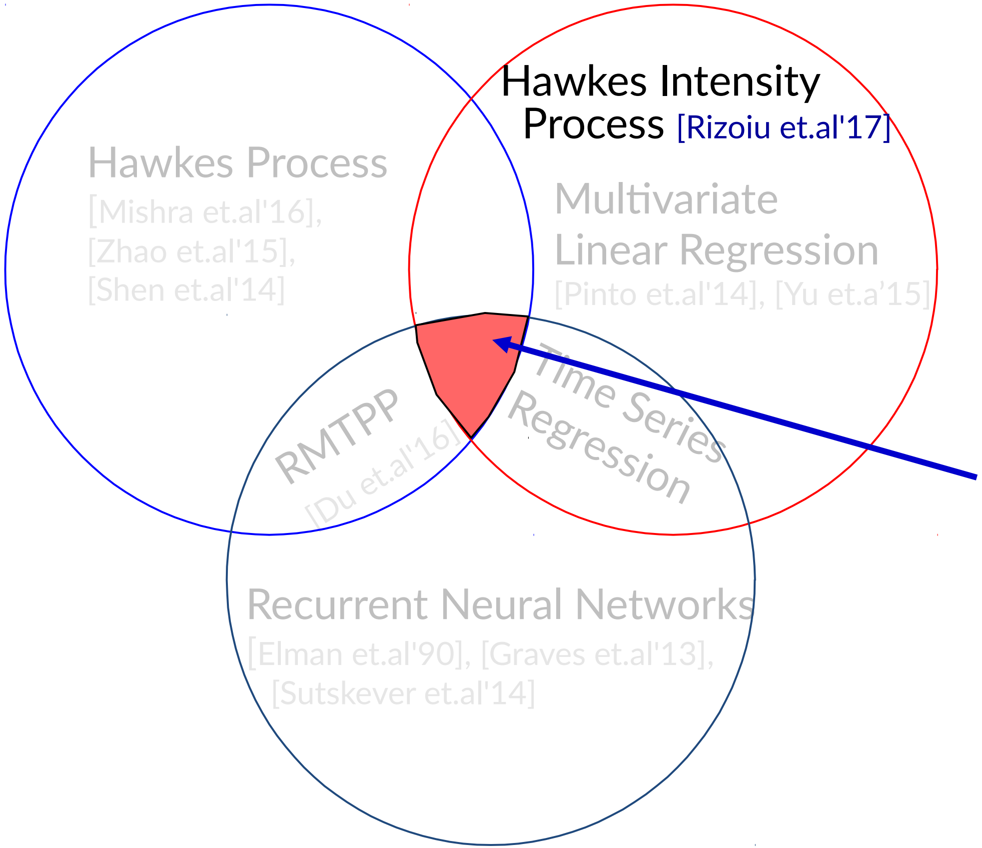

RNN-MAS: popularity modeling using multiple asynchronous streams [Mishra et.al ICWSM'18]

Tweeted Videos dataset

Tweeted Videos: YouTube videos published and tweeted June 2014 until today (5M tweets/day)

Video duration: 4M16S
Visual definition: HD or SD

Video Title:
Shawn Mendes - Treat You Better
Channel Id: UC4-TgOSMJHn-LtY4zCzbQhw
Channel Title:
ShawnMendesVEVO

Freebase topics:
Shawn Mendes; Music; Music video; Pop music

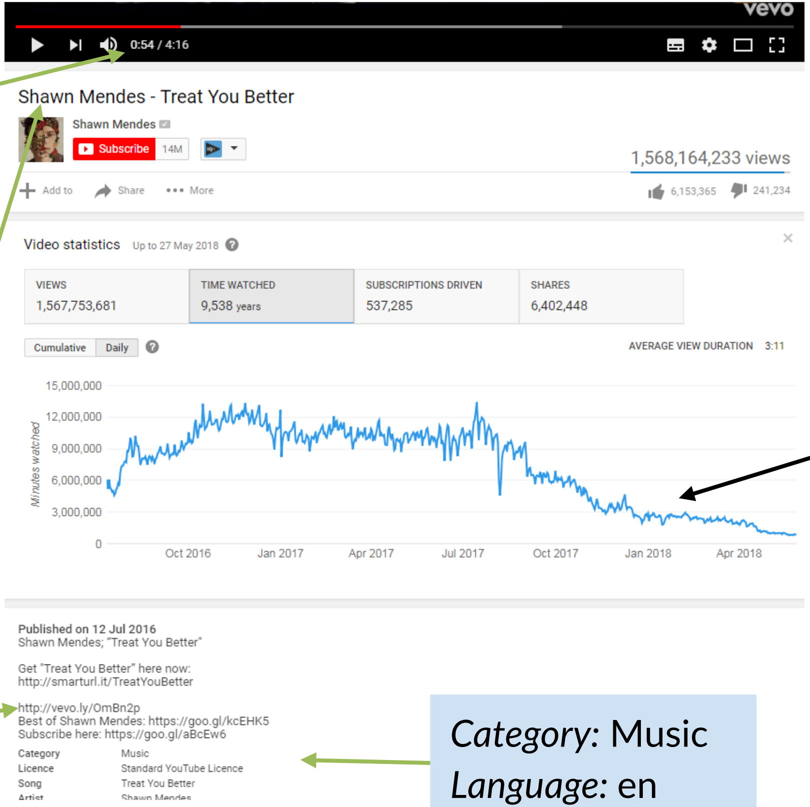

Insight time series:
(a) Daily watch time
(b) Daily view count
(c) Daily share count
(d) Avg watch time

Category: Music
Language: en

Presentation outline

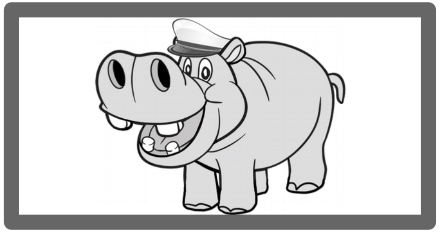

Modeling and predicting popularity using HIP

[Rizoiu et.al WWW'17]

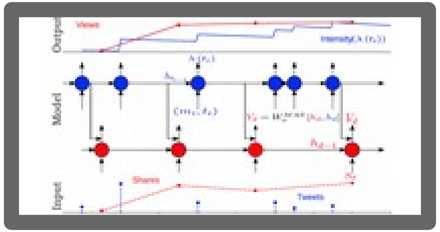

Popularity in Asynchronous Social Media Streams with RNN

[Mishra et.al ICWSM'18]

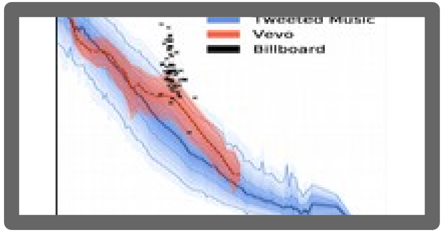

Measuring and Predicting Engagement in Online Videos

[Wu et.al ICWSM'18]

Linking exo-endo popularity

[Rizoiu et.al WWW'17]

exogenous stimuli

endogenous response

observed popularity

Linking exo-endo popularity

[Rizoiu et.al WWW'17]

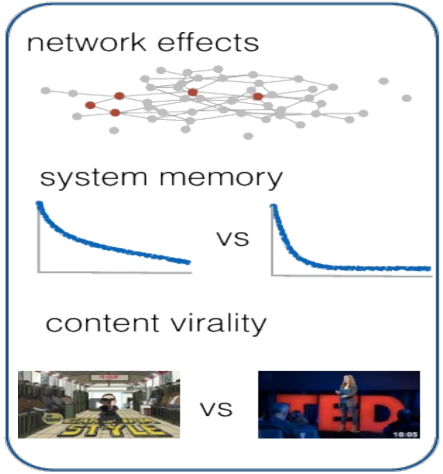

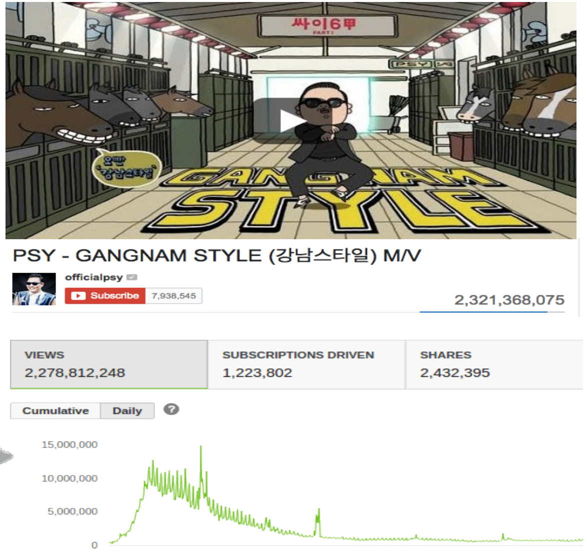

exogenous stimuli

endogenous response

observed popularity

Hawkes Process [Hawkes '71]

$$\lambda(t) = \mu(t) + \sum_{t_i < t} \phi_{m_i}(t - t_i)$$

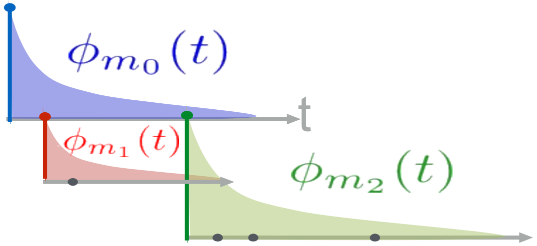

Most state-of-the-art popularity prediction systems require observing individual events.

[Zhao et al KDD'15][Shen et al AAAI'14]
[Farajtabar et al NIPS'15][Mishra et al CIKM'16]

Hawkes Process [Hawkes '71]

$$\lambda(t) = \mu(t) + \sum_{t_i < t} \phi_{m_i}(t - t_i)$$

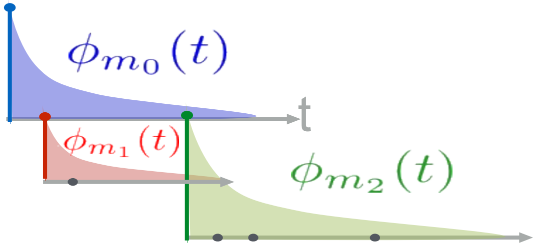

the rate of 'daughter' events — content virality — user influence — memory

$$\phi_m(\tau) = \kappa \, m^{\beta} \, \hat{\tau}^{-(1+\theta)}$$

Most state-of-the-art popularity prediction systems require observing individual events.

[Zhao et al KDD'15][Shen et al AAAI'14]
[Farajtabar et al NIPS'15][Mishra et al CIKM'16]

Hawkes Intensity Process (HIP)

[Rizoiu et.al WWW'17]

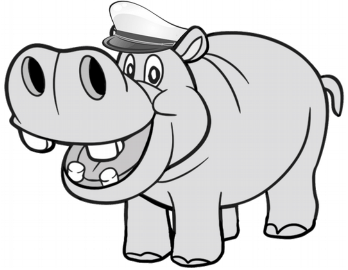

$$\lambda(t) = \mu(t) + \sum_{t_i<t} \phi_{m_i}(t - t_i)$$

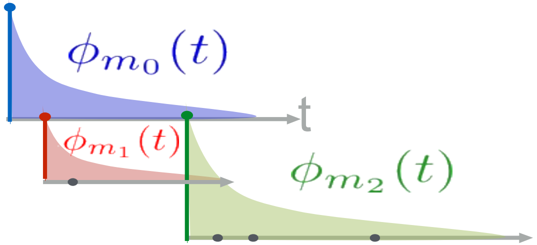

the rate of 'daughter' events — content virality — user influence — memory

$$\phi_m(\tau) = \kappa\, m^\beta \hat{\tau}^{-(1+\theta)}$$

expected number of events

$$\xi(t) = \mu s(t) + C \int_0^t \xi(t-\tau)\hat{\tau}^{-(1+\theta)} d\tau$$

popularity

exogenous stimuli

Hawkes Intensity Process (HIP)

[Rizoiu et.al WWW'17]

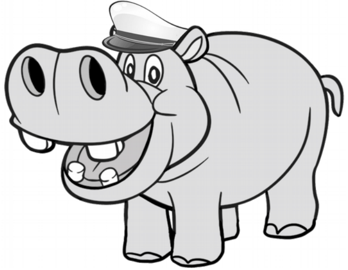

$$\lambda(t) = \mu(t) + \sum_{t_i < t} \phi_{m_i}(t - t_i)$$

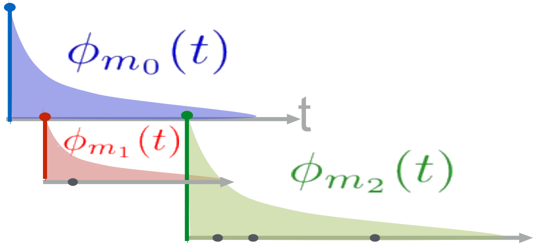

the rate of 'daughter' events | content virality | user influence | memory

$$\phi_m(\tau) = \kappa \, m^{\beta} \, \hat{\tau}^{-(1+\theta)}$$

expected number of events

$$\xi(t) = \mu s(t) + C \underbrace{\int_0^t \xi(t - \tau)\hat{\tau}^{-(1+\theta)} d\tau}_{\text{endogenous reaction}}$$

popularity

exogenous sensitivity | exogenous stimuli | endogenous reaction

Estimating the HIP model

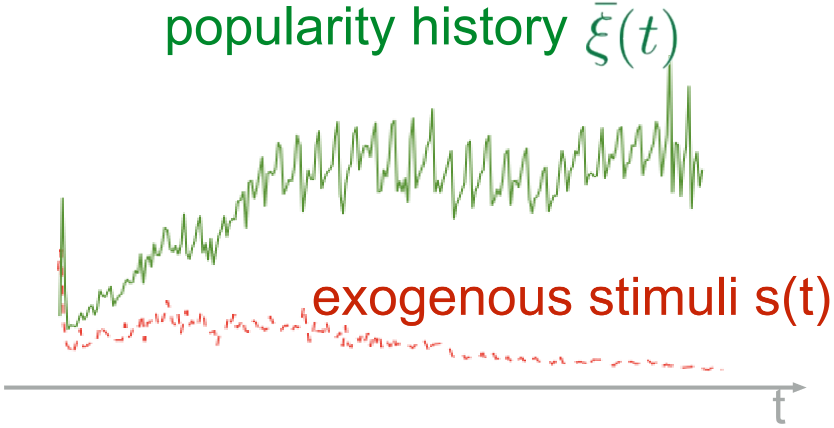

$$\text{find } \{\mu, C, \theta, \ldots\}$$

$$\text{s.t. } \min \sum_t l(\xi(t) - \bar{\xi}(t))$$

$$\xi(t) = \mu s(t) + C \underbrace{\int_0^t \xi(t-\tau)\hat{\tau}^{-(1+\theta)} d\tau}_{\text{endogenous reaction}}$$

popularity

exogenous sensitivity

exogenous stimuli

endogenous reaction

Estimating the HIP model

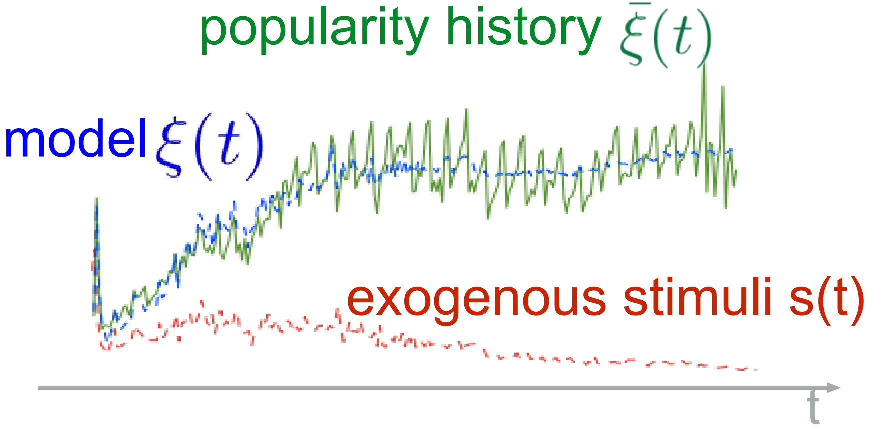

$$\text{find } \{\mu, C, \theta, \ldots\}$$

$$\text{s.t. } \min \sum_t l(\xi(t) - \bar{\xi}(t))$$

$$\underbrace{\xi(t)}_{\text{popularity}} = \underbrace{\mu}_{\text{exogenous sensitivity}} \underbrace{s(t)}_{\text{exogenous stimuli}} + \underbrace{C \int_0^t \xi(t-\tau)\hat{\tau}^{-(1+\theta)} d\tau}_{\text{endogenous reaction}}$$

HIP as a Linear Time-Invariant system

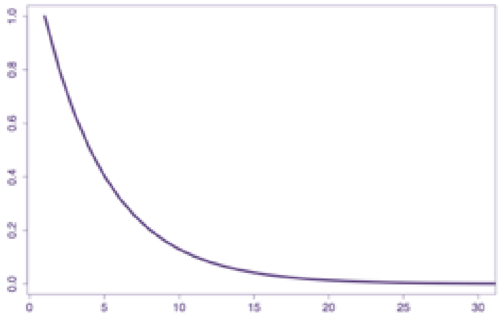

Impulse
response

$$\xi(t) = \mu s(t) + C \underbrace{\int_0^t \xi(t-\tau)\hat{\tau}^{-(1+\theta)} d\tau}_{}$$

popularity

exogenous sensitivity

exogenous stimuli

endogenous reaction

HIP as a Linear Time-Invariant system

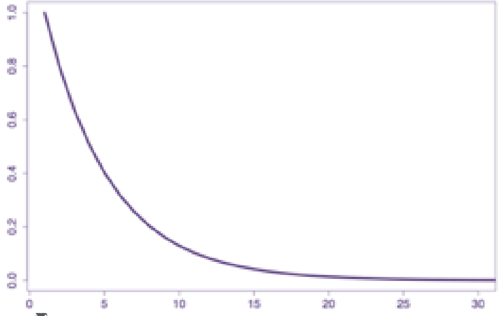

scale, shift,
add

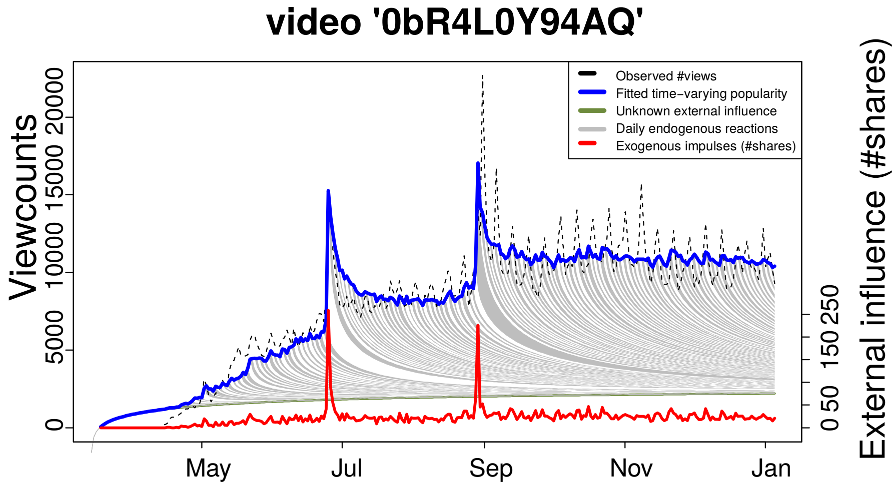

$$\xi(t) = \mu s(t) + C \underbrace{\int_0^t \xi(t-\tau)\hat{\tau}^{-(1+\theta)} d\tau}$$

popularity

exogenous sensitivity

exogenous stimuli

endogenous reaction

HIP as a Linear Time-Invariant system

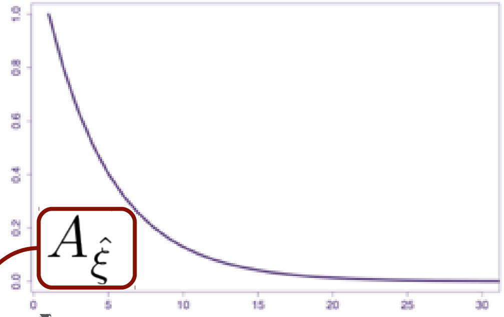

endogenous response

scale, shift, add

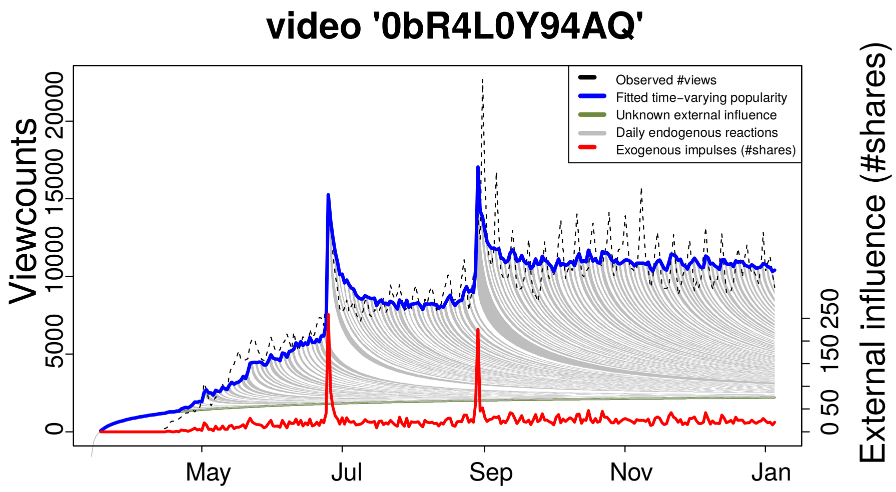

$$\xi(t) = \boxed{\mu} s(t) + C \underline{\int_0^t \xi(t-\tau)\hat{\tau}^{-(1+\theta)} d\tau}$$

popularity

exogenous sensitivity

exogenous stimuli

endogenous reaction

The “endo-exo” map

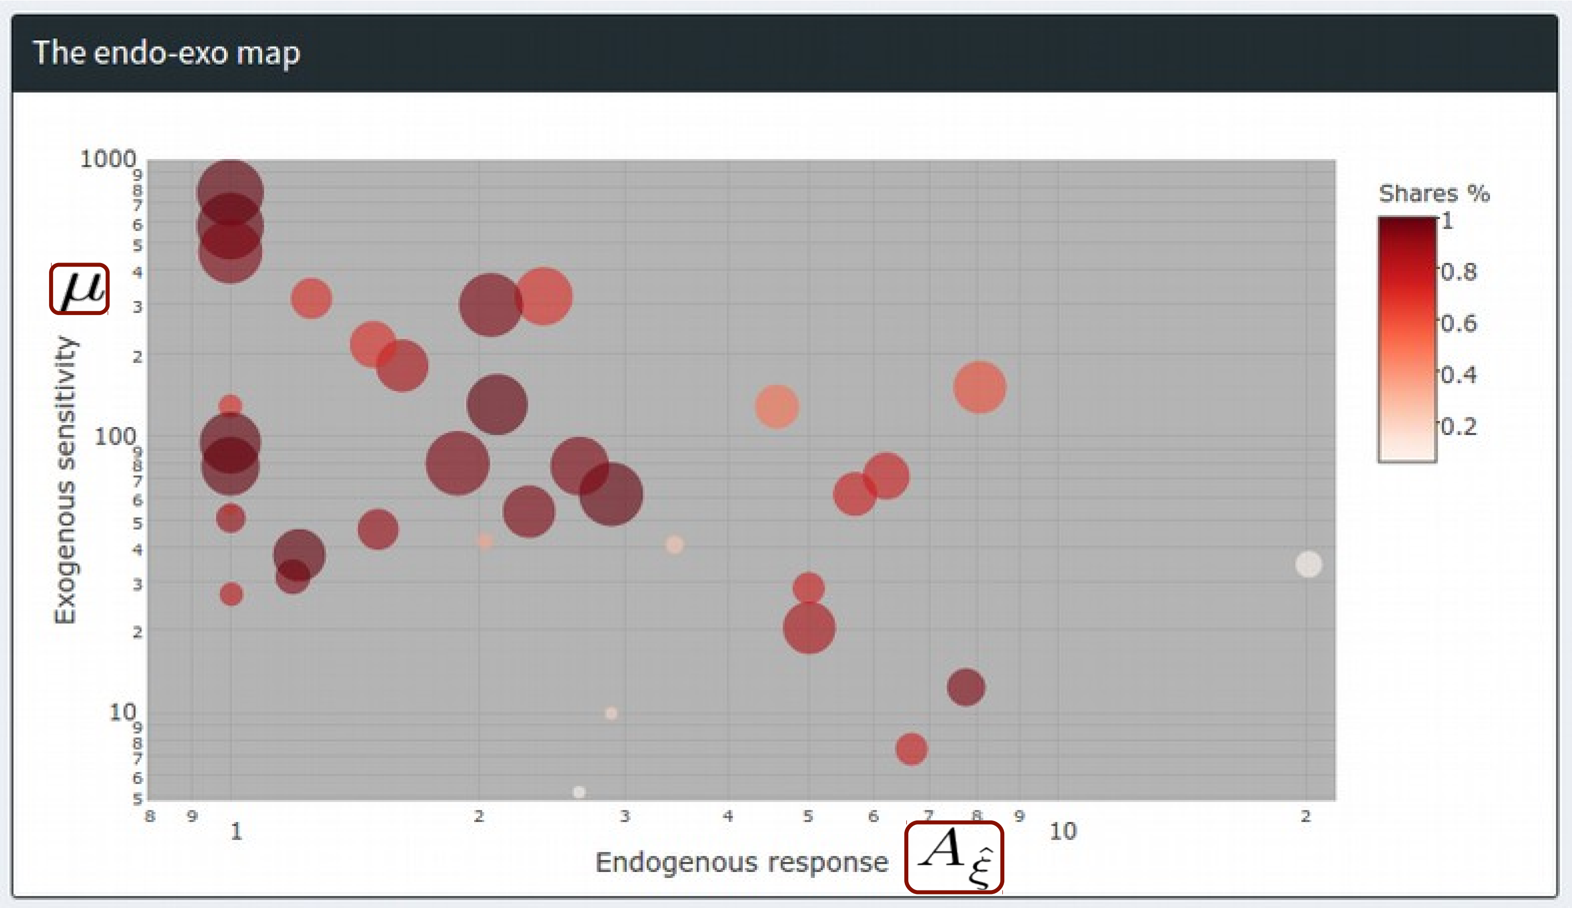

Explain popularity dynamics [Kong et.al WWW'18]

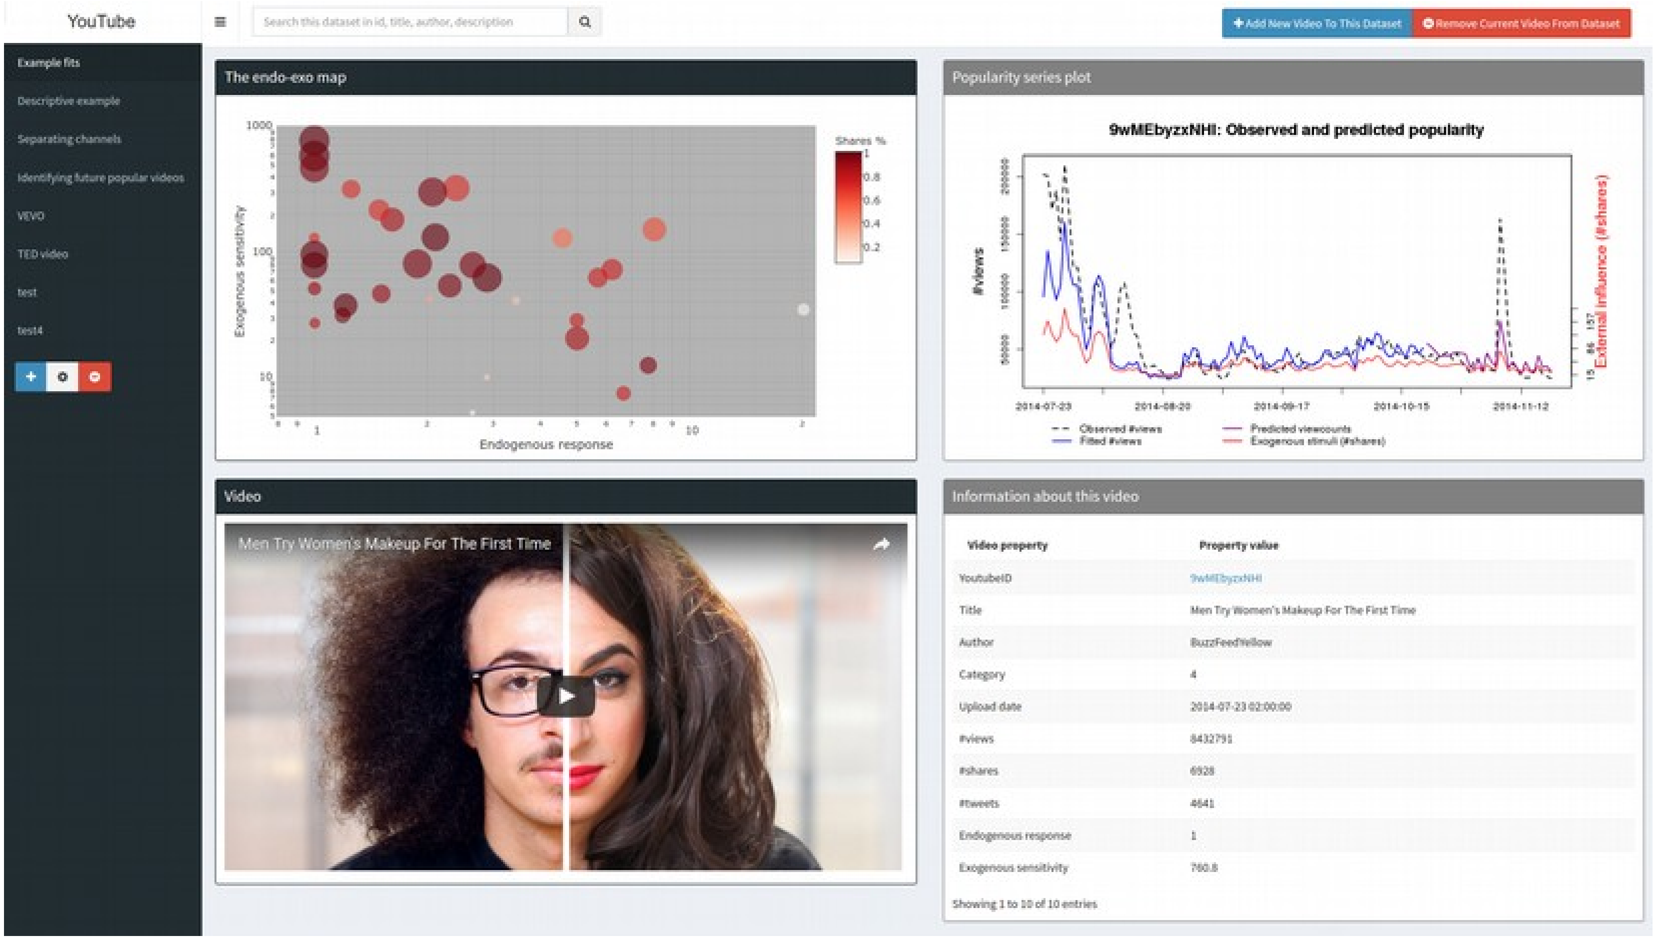

Presentation outline

Modeling and predicting popularity using HIP

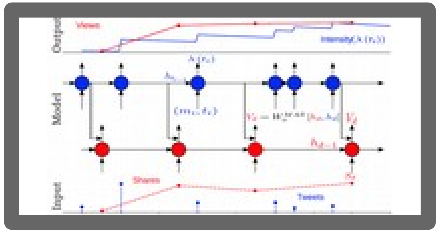

Popularity in Asynchronous Social Media Streams with RNN

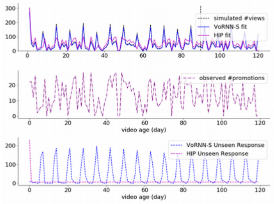

Response to unseen influence

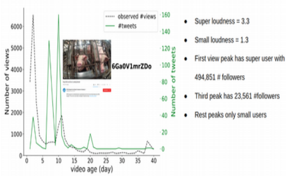

Loudness of User(s)

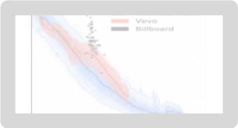

Measuring and Predicting Engagement in Online Videos

Modelling Popularity in Asynchronous Social Media Streams with RNNs

Swapnil Mishra, Marian-Andrei Rizoiu, Lexing Xie

ComputationalMedia @ANU: http://cm.cecs.anu.edu.au

ICWSM '18, Stanford, CA, USA

ICWSM talk Tuesday, 1:15PM

RNN-MAS: Accounting for tweets

Output

Model

Input

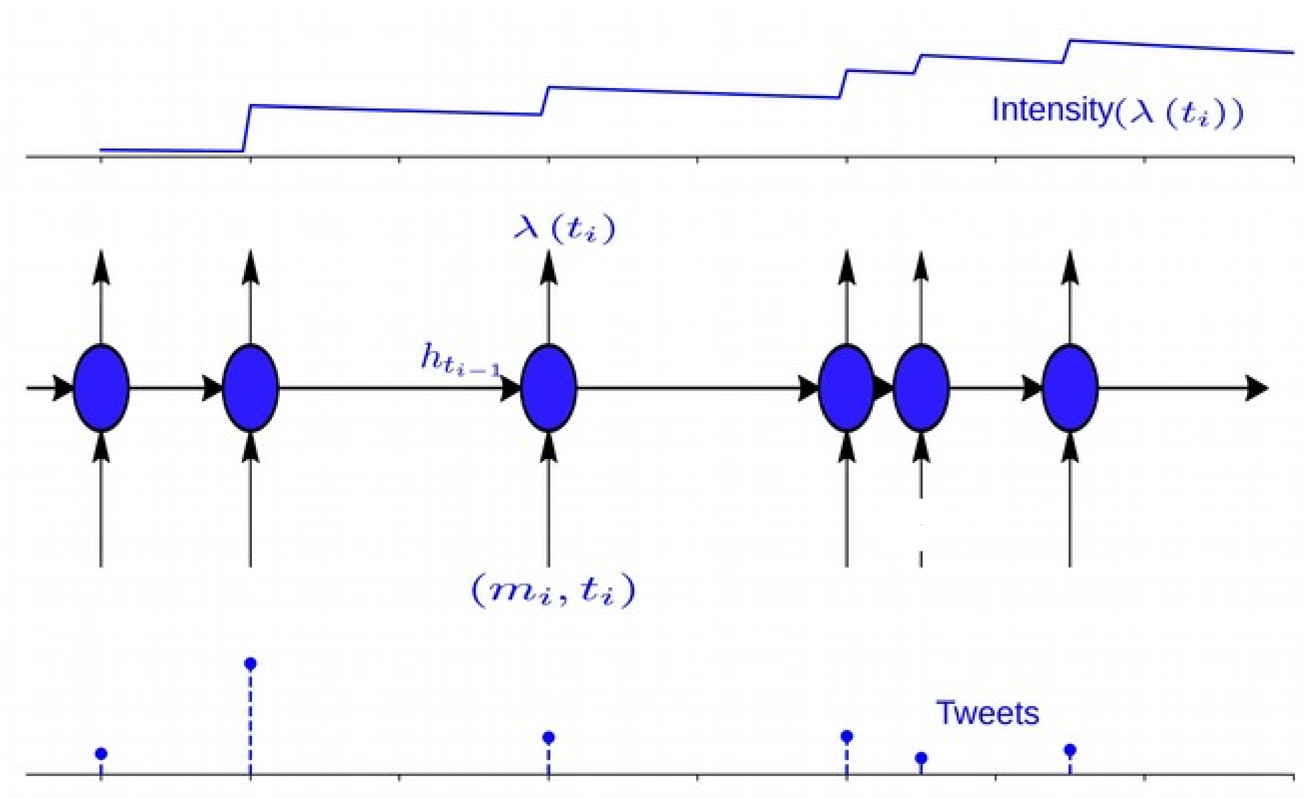

RNN-MAS: Daily Aggregated Shares

Output

Model

Input

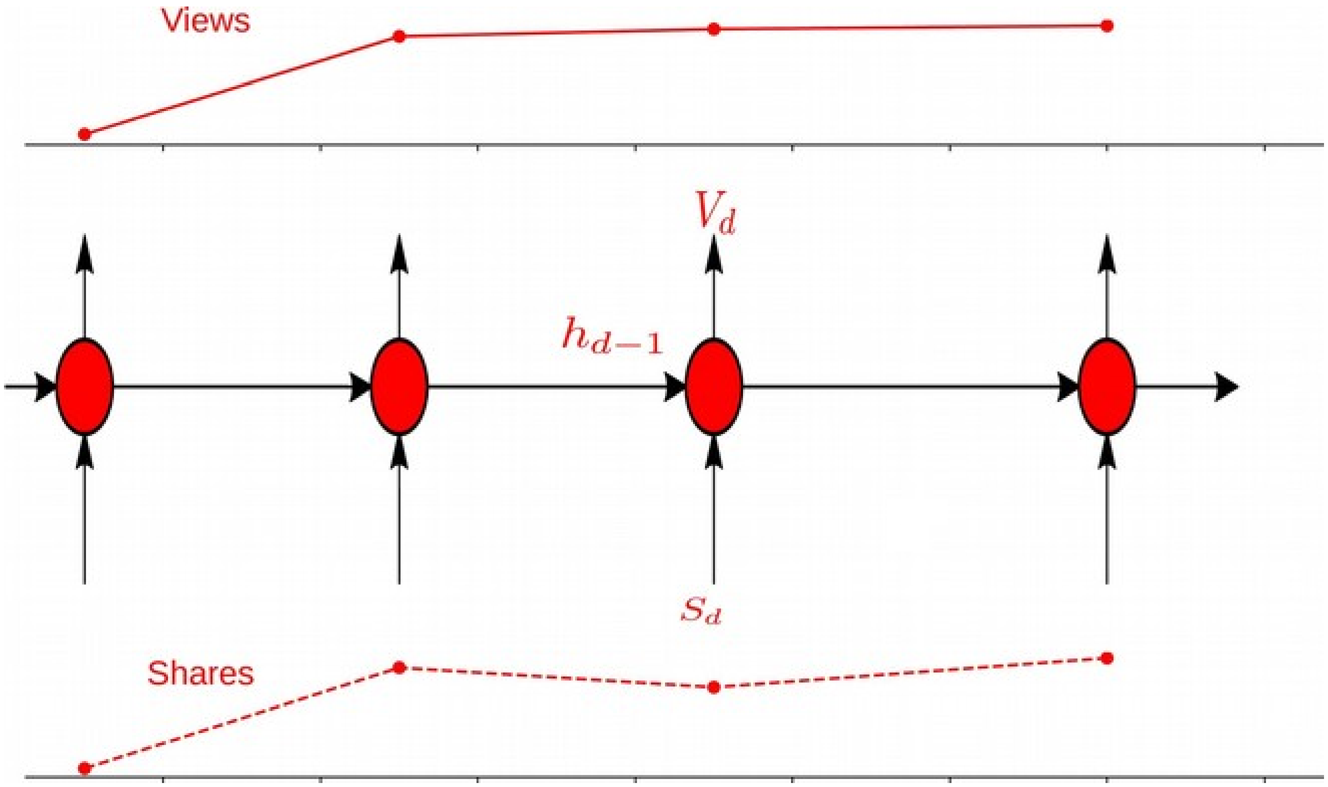

RNN-MAS: Multiple asynchronous streams

[Mishra et.al ICWSM'18]

Output

Model

Input

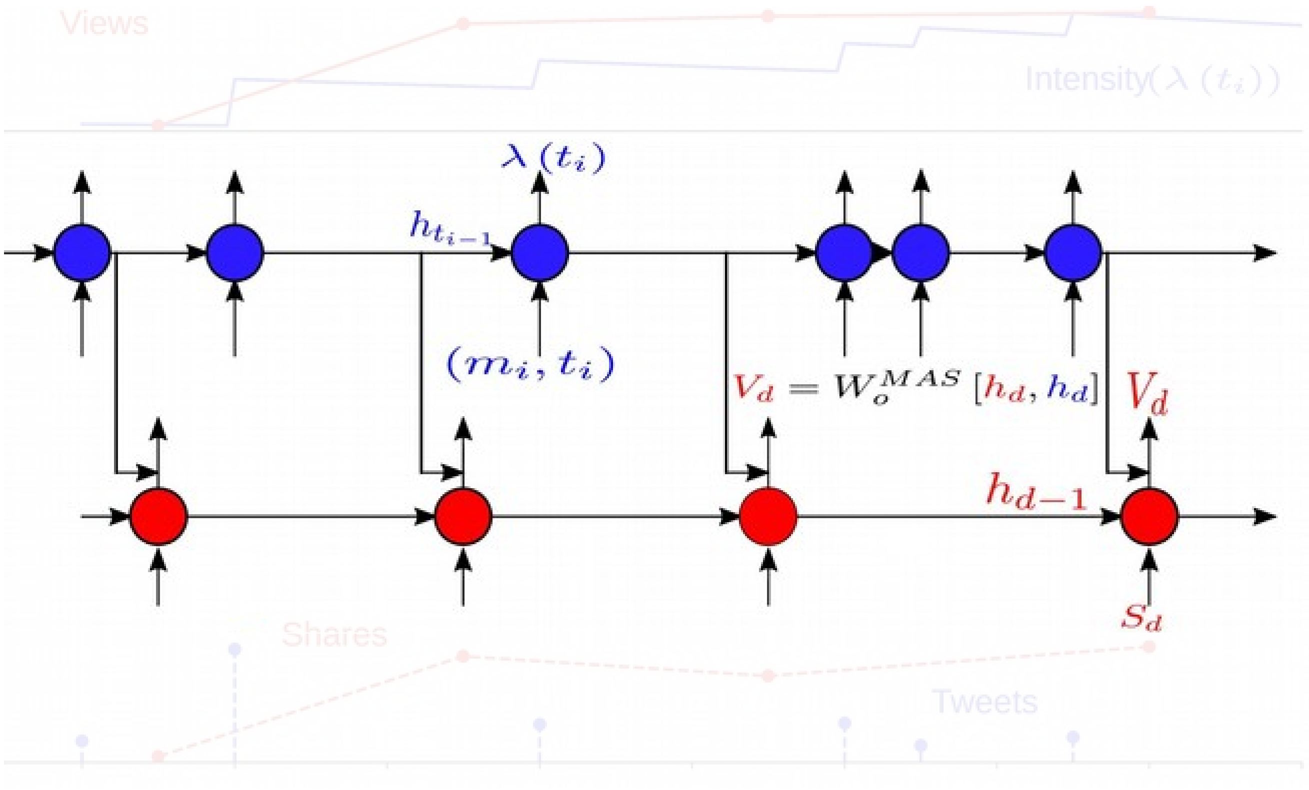

RNN-MAS: Multiple asynchronous streams

[Mishra et.al ICWSM'18]

Output

Model

Input

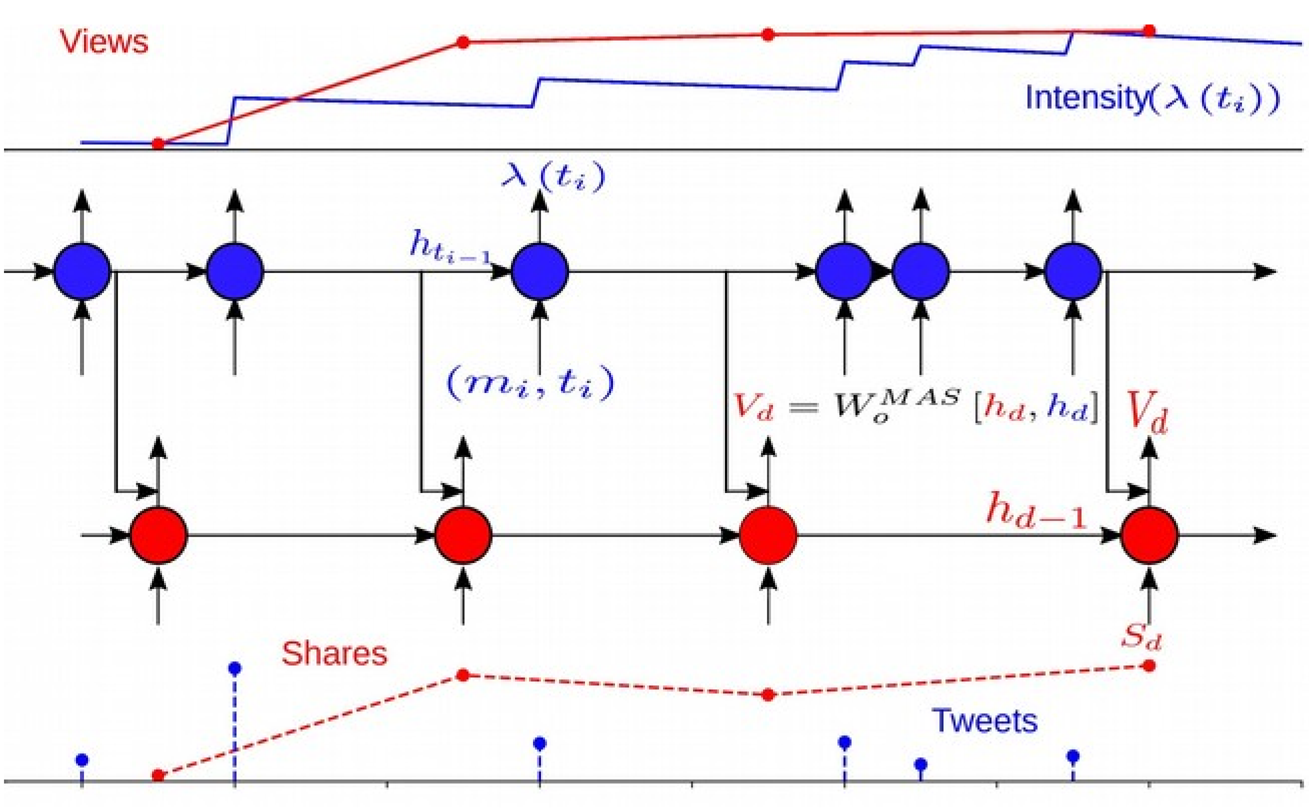

RNN-MAS example fittings

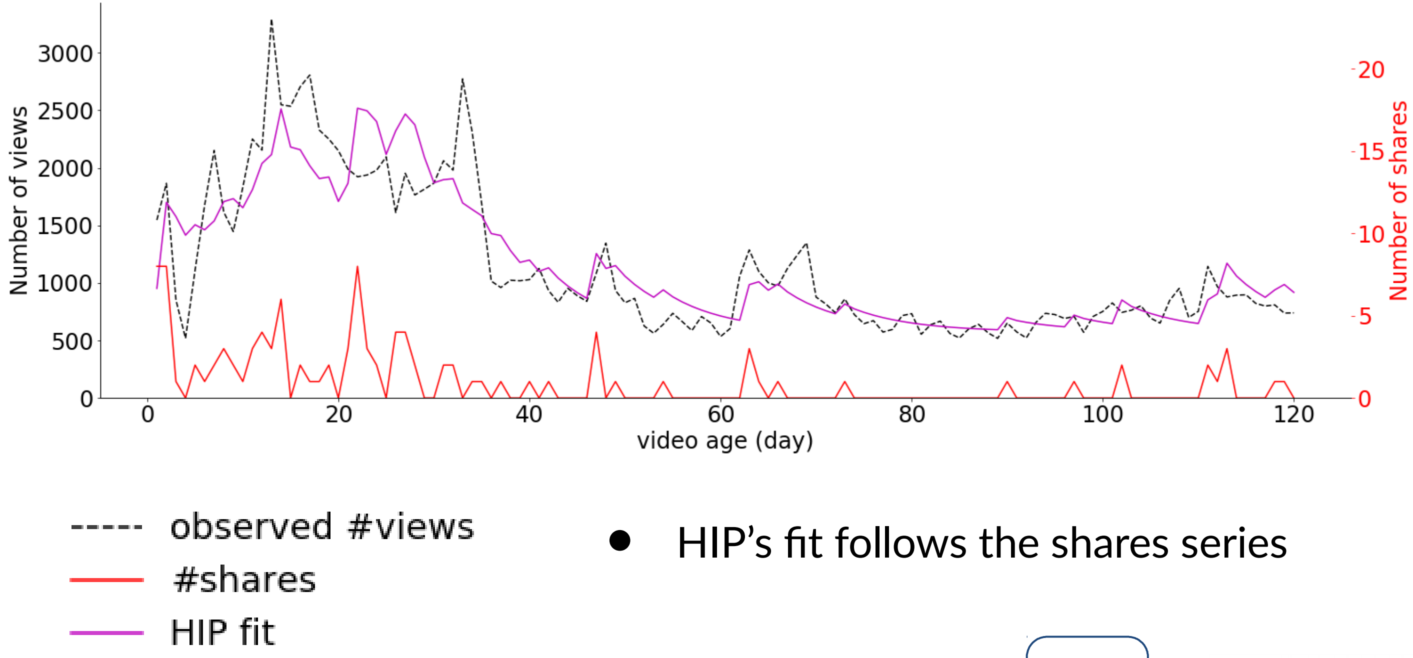

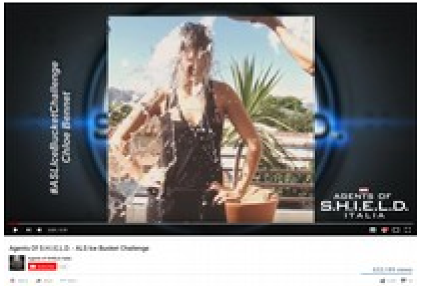

- HIP's fit follows the shares series

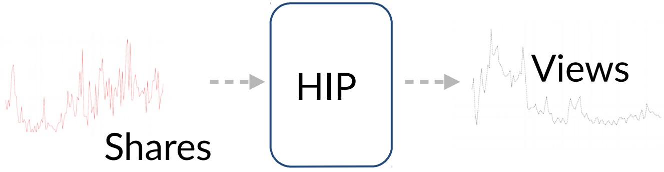

RNN-MAS example fittings

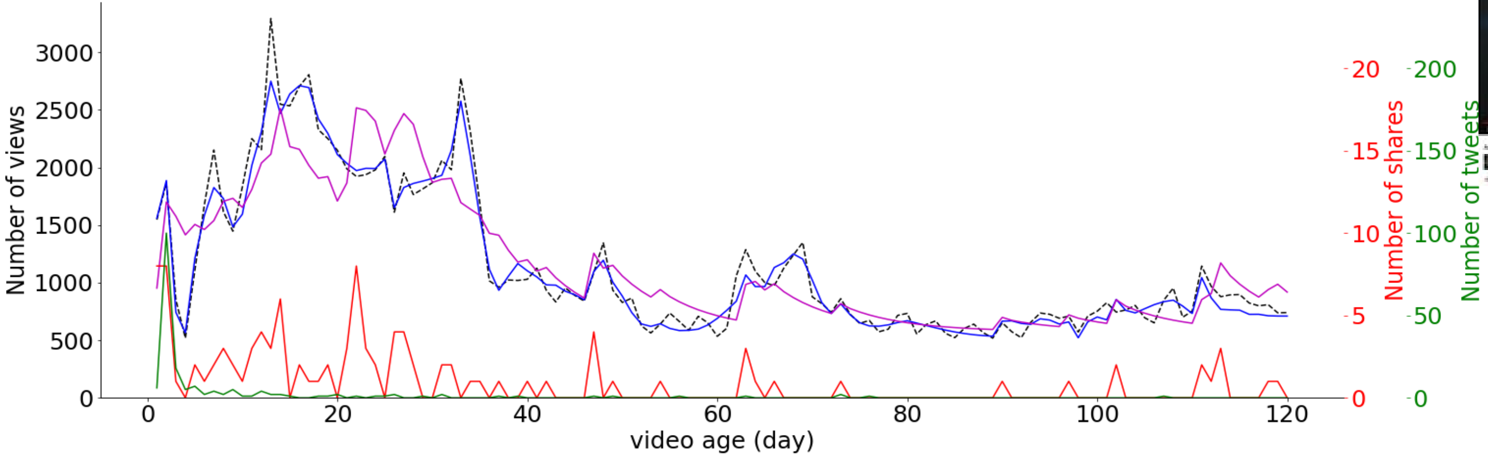

- HIP’s fit follows the shares series
- RNN-MAS handles multiple series with different granularities
- RNN-MAS follows view series closely
- RNN-MAS outperforms HIP by 17% on HIP’s dataset

Presentation outline

Modeling and predicting popularity using HIP

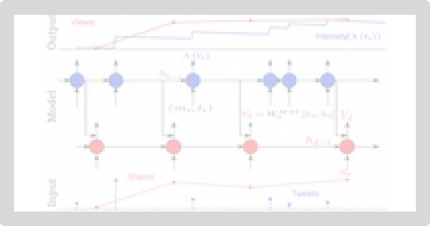

Popularity in Asynchronous Social Media Streams with RNN

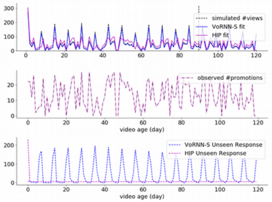

Response to unseen influence

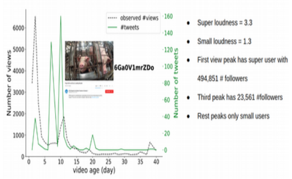

Loudness of User(s)

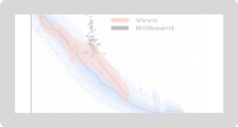

Measuring and Predicting Engagement in Online Videos

Response to unseen influence

Shares and tweets are two of the factors influencing popularity

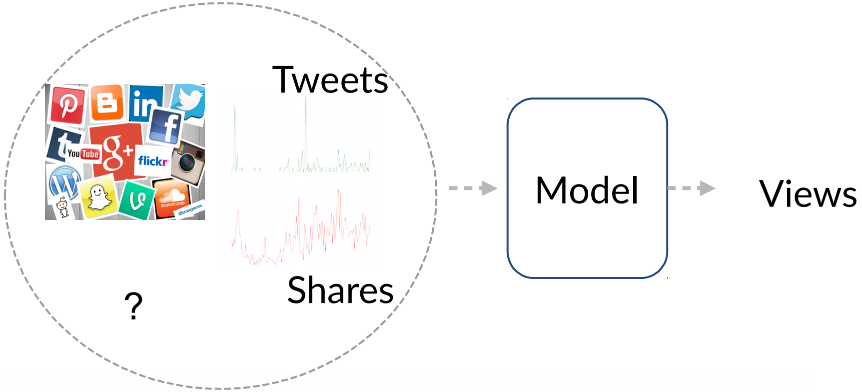

Seasonality is important

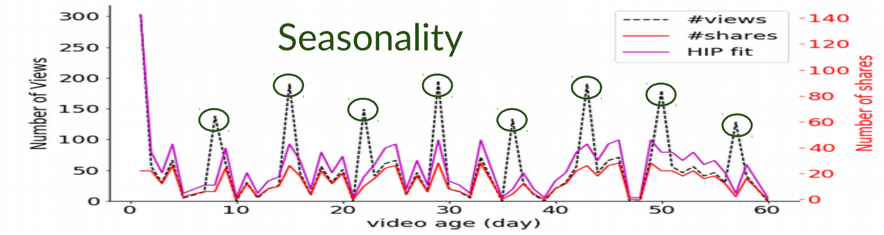

Response to unseen influence

Shares and tweets are two of the factors influencing popularity

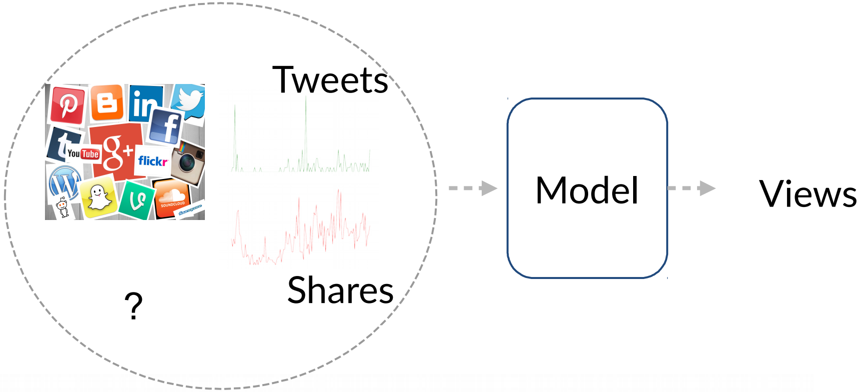

Seasonality is important

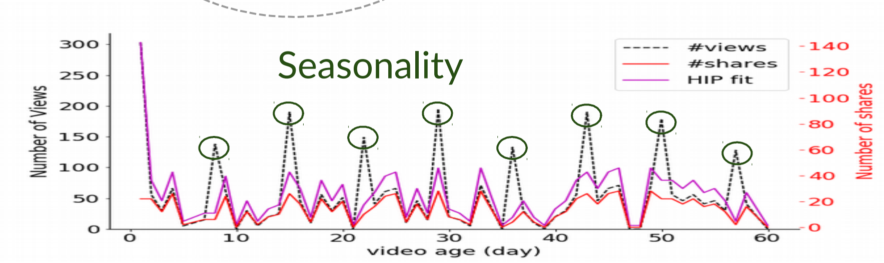

New metric:
Total response of RNN-MAS with zero promotion

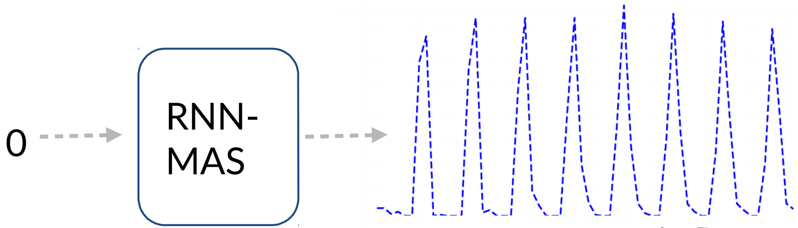

Response to unseen influence

Response to unseen influence: Results

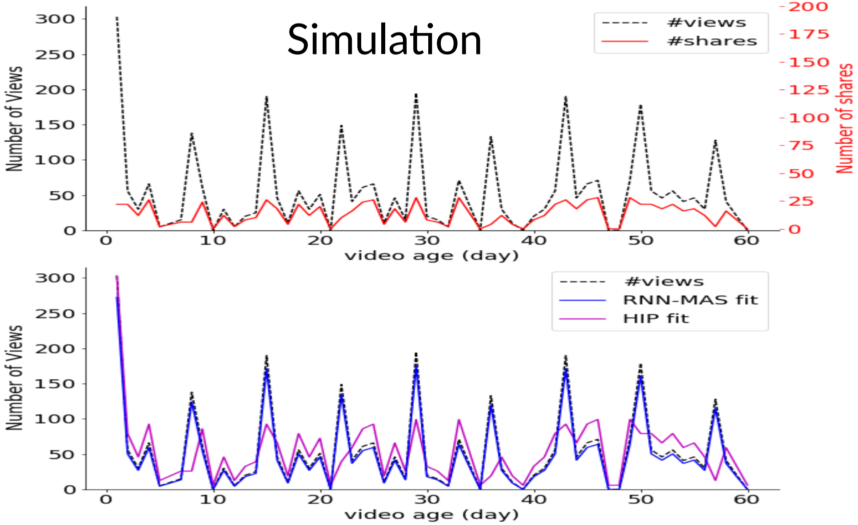

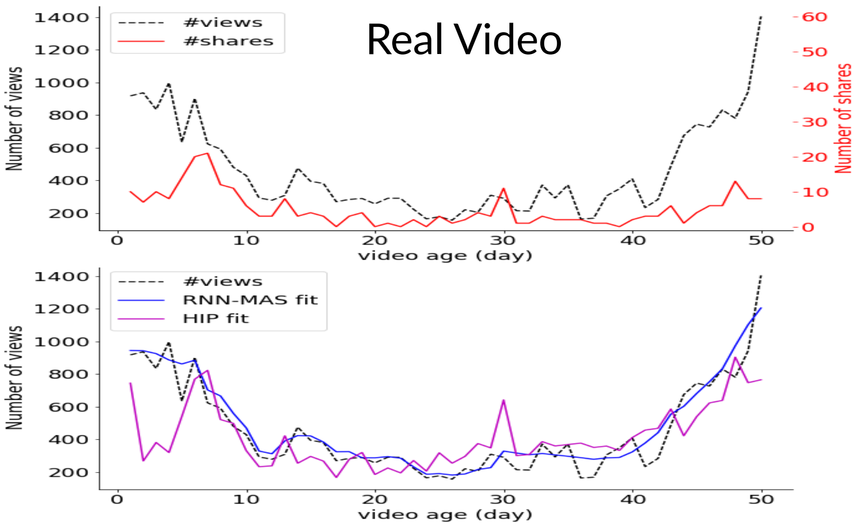

Response to unseen influence: Results

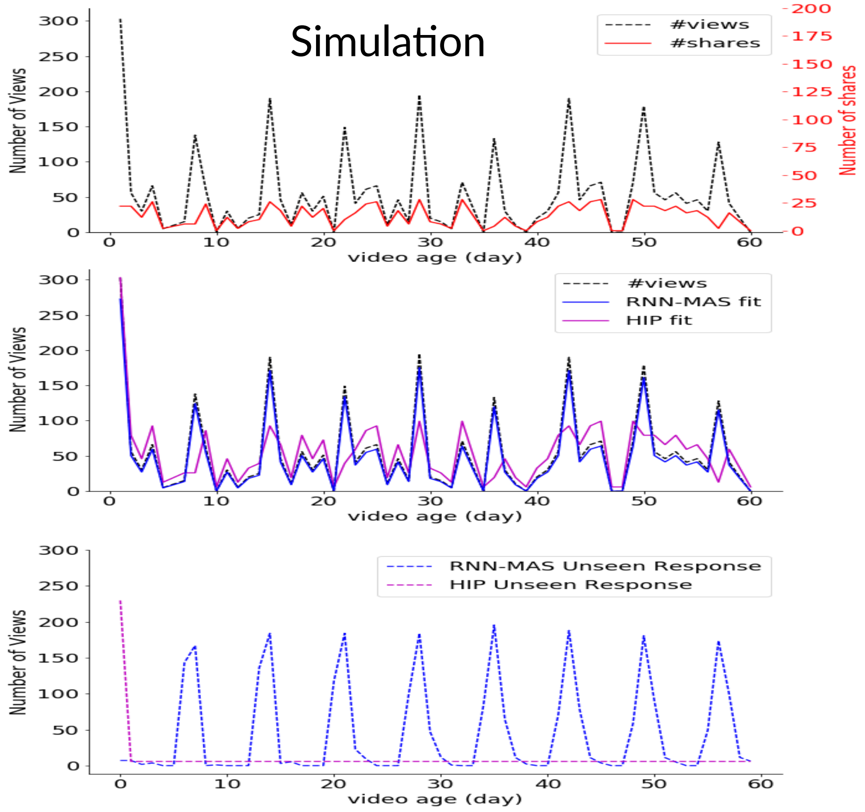

Latent response has a seasonal behavior

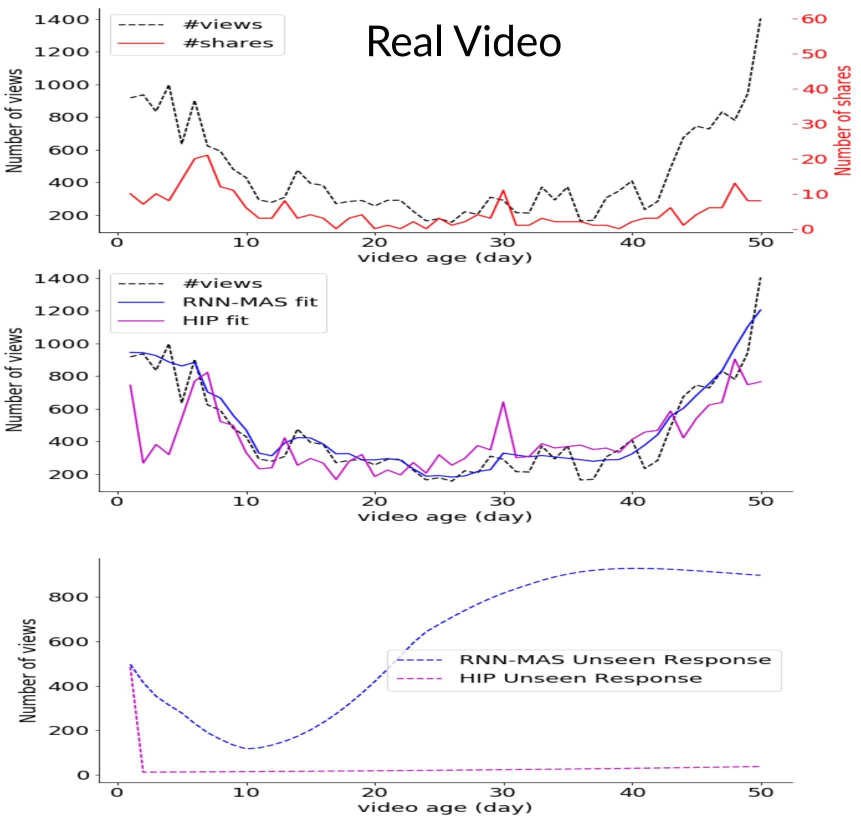

Latent response starts after a delay

Loudness of Users

Which promotion will gather more views for the video?

- [Bakshay et.al'11]
- [Budak et.al'12]

Top 1% most followed

Super user

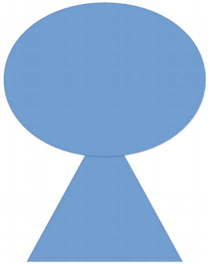

Cohort of median users

Small user

Loudness of Users

Which promotion will gather more views for the video?

- [Bakshay et.al'11]
- [Budak et.al'12]

~38K Followers

Super user loudness = log(sum(#views (

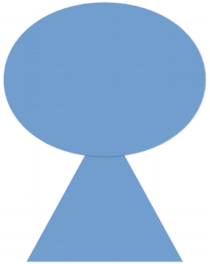

)))

RNN-MAS

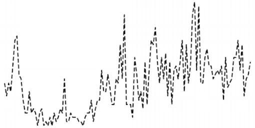

Small user

Cohort of median users

Loudness of Users

Which promotion will gather more views for the video?

- [Bakshay et.al'11]
- [Budak et.al'12]

~38K Followers

Super user loudness = log(sum(#views (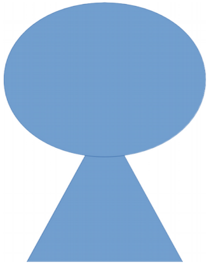)))

RNN-MAS

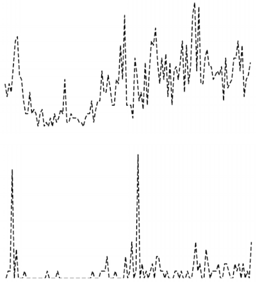

Small user loudness = log(sum(#views ()))

156 Followers

Disproportionate Influence

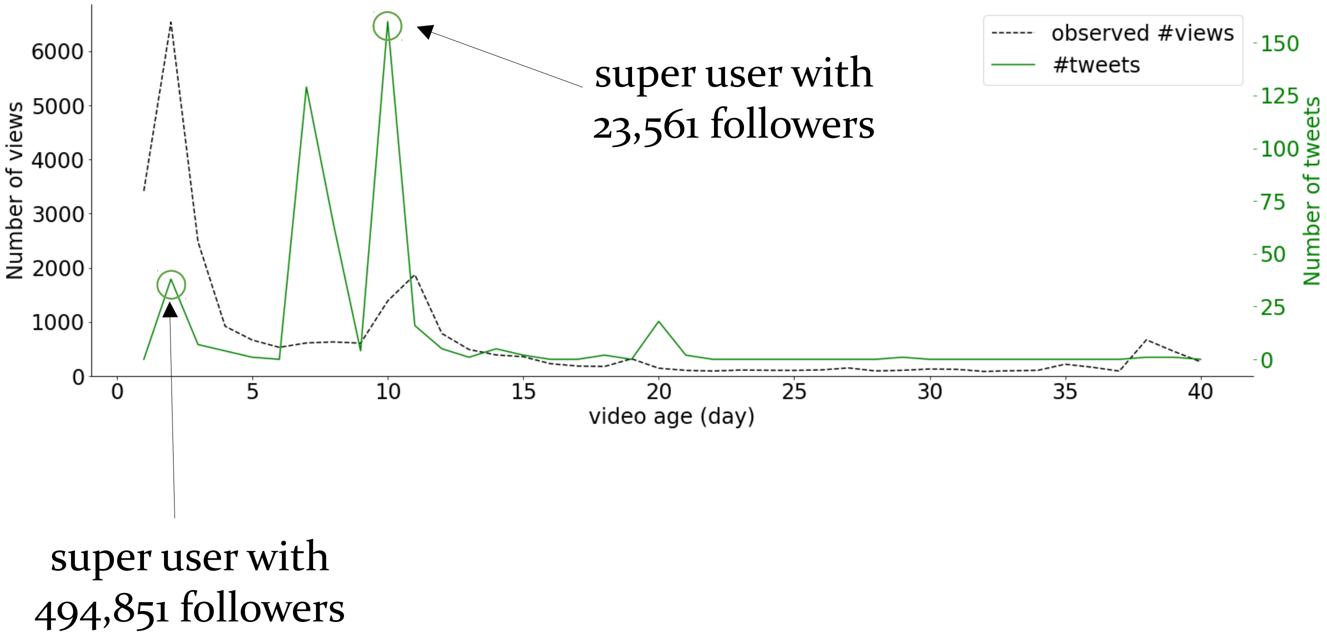

Super user loudness = 3.3 > Small user loudness = 1.3

Disproportionate Influence

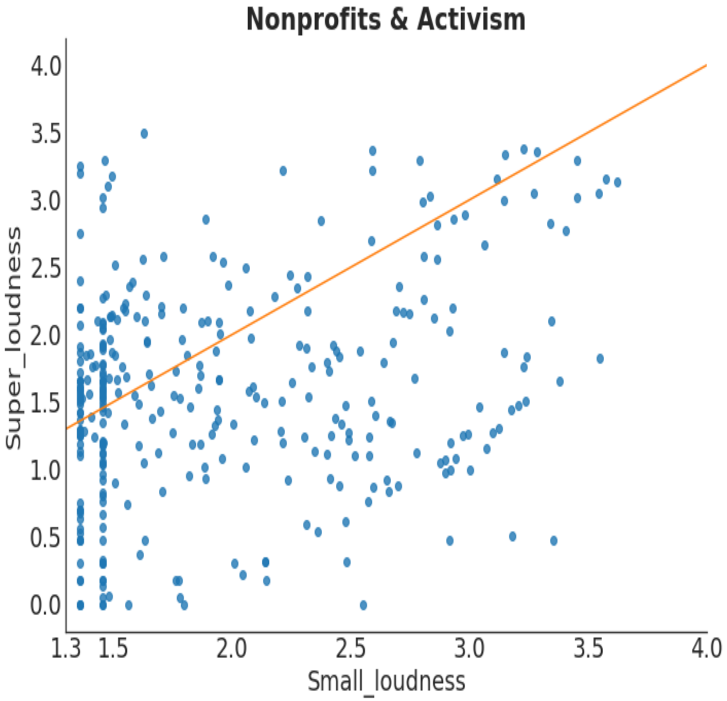

Best promotion:

small: 63%, Super: 37%

Disproportionate Influence

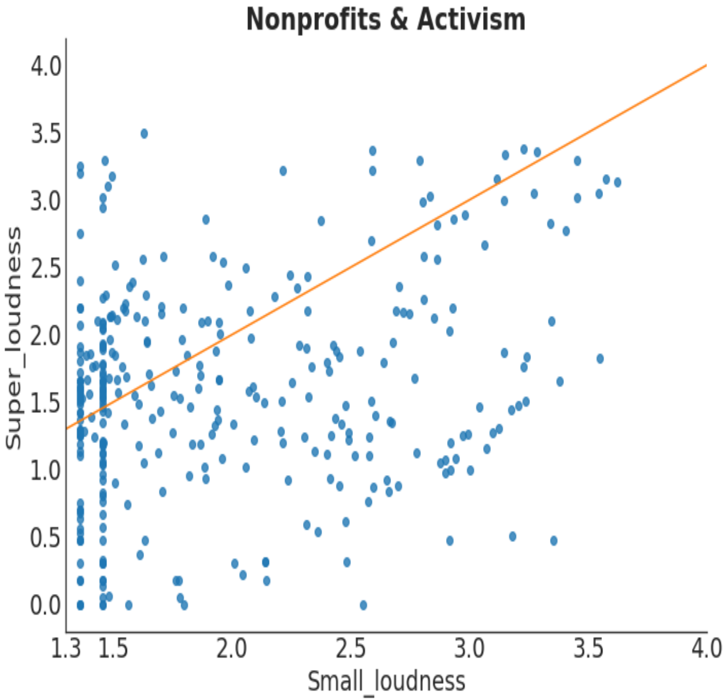

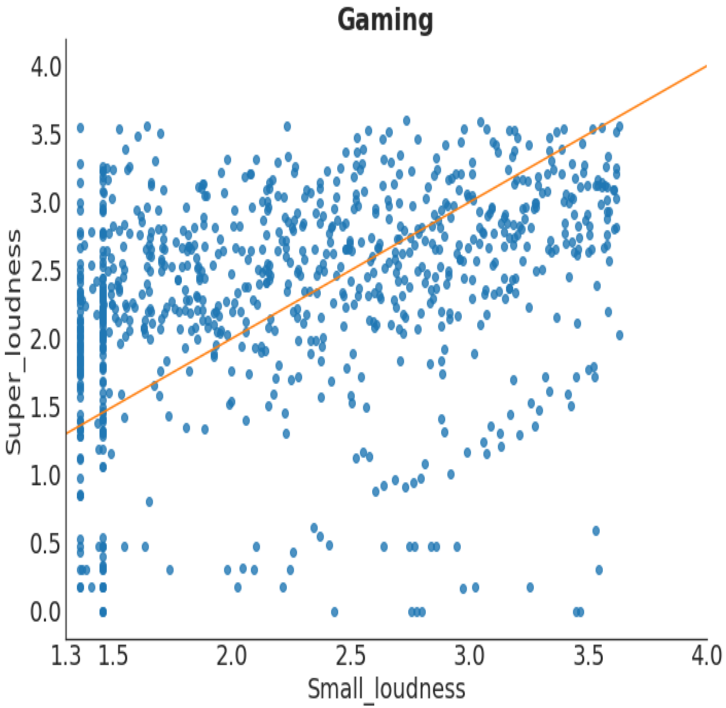

Best promotion:

small: 63%, Super: 37%

small: 42%, Super: 58%

RNN-MAS

Get code and data from
https://github.com/computationalmedia/rnn-mas

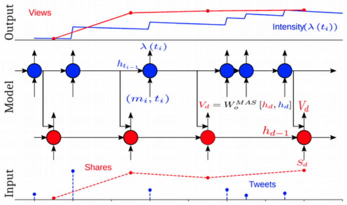

1. RNN-MAS: Joint Model for Asynchronous heterogeneous Stream
Models multiple asynchronous streams of different time granularity
Outperforms state of the art by 17%.

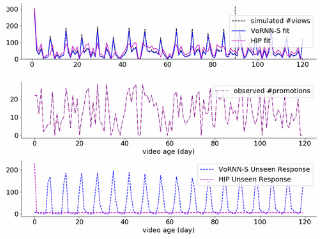

2. New Metric: Response to unseen influence
Explains model behaviour including seasonality, uncovers latent influences

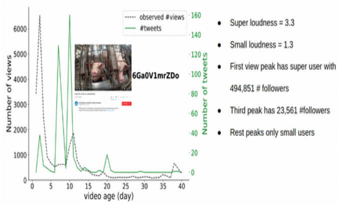

3. New Metric: Loudness of User(s)
Quantifies user influence across network boundaries. Compares effects of celebrity versus grass-root users.

Presentation outline

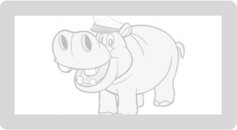

Modeling and predicting popularity using HIP

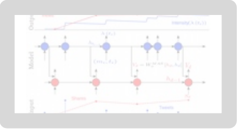

Popularity in Asynchronous Social Media Streams

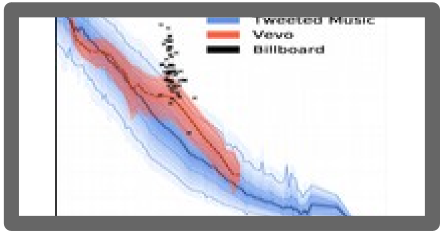

Measuring and Predicting Engagement in Online Videos

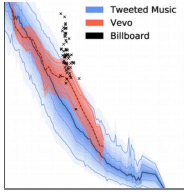

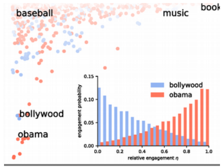

Does engagement relate to content quality?

Can aggregate engagement be predicted?

Beyond Views: Measuring and Predicting Engagement in Online Videos

Siqi Wu, Marian-Andrei Rizoiu, Lexing Xie

ComputationalMedia @ANU: http://cm.cecs.anu.edu.au
ICWSM '18, Stanford, CA, USA

ICWSM talk Thursday, 4:15PM

View count does NOT translate to watch time

View count: 3,917,179 < View count: 7,833,595

Watch time: 62 years > Watch time: 32 years

Research questions on online video engagement

User-specific engagement: the key for video recommendation [Covington et al. *RecSys* '16][Park et al. *ICWSM* '16]

Aggregate engagement: open data available to researchers

Applications: better recommender systems, mitigate information overload, etc.

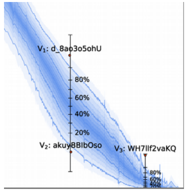

1. How to measure aggregate engagement?

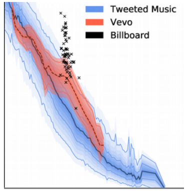

2. Characteristics of aggregate engagement

(a) Does engagement relate to content quality?

(b) How does engagement evolve over time?

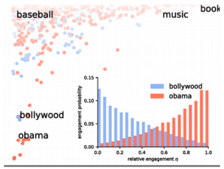

3. Can aggregate engagement be predicted?

Popularity and engagement for web content

Domain	Popularity metrics	Engagement metrics
Webpages	Visit number [Li and Moore, *JMLR* '08]	Click-through-rate [Richardson et al. *WWW* '07]
Search ads	Display number [He et al. *ADKDD* '14]	Conversion rate [Barbieri et al. *WWW* '14]
Songs	Listening count [Bellogin et al. *ICWSM* '13]	Download number [Salganik et al. *Science* '06] [Krumme et al. *PloS* '12]
Videos	View count [Pinto et al. *WSDM* '13] [Szabo and Huberman *Com.ACM* '10] [Rizoiu et al. *WWW* '17]	Watch time [Guo et al. *L@S* '14] [Park et al. *ICWSM* '16]

Popularity and engagement for web content

Domain	Popularity metrics	Engagement metrics
Webpages	Visit number [Li and Moore, *JMLR* '08]	Click-through-rate [Richardson et al. *WWW* '07]
Search ads	Display number [He et al. *ADKDD* '14]	Conversion rate [Barbieri et al. *WWW* '14]
Songs	Listening count [Bellogin et al. *ICWSM* '13]	Download number [Salganik et al. *Science* '06] [Krumme et al. *PloS* '12]
Videos	View count [Pinto et al. *WSDM* '13] [Szabo and Huberman *Com.ACM* '10] [Rizoiu et al. *WWW* '17]	Watch time [Guo et al. *L@S* '14] [Park et al. *ICWSM* '16]

★ **No browser extension**
★ **New metric**
★ **Cold-start prediction**

The engagement maps

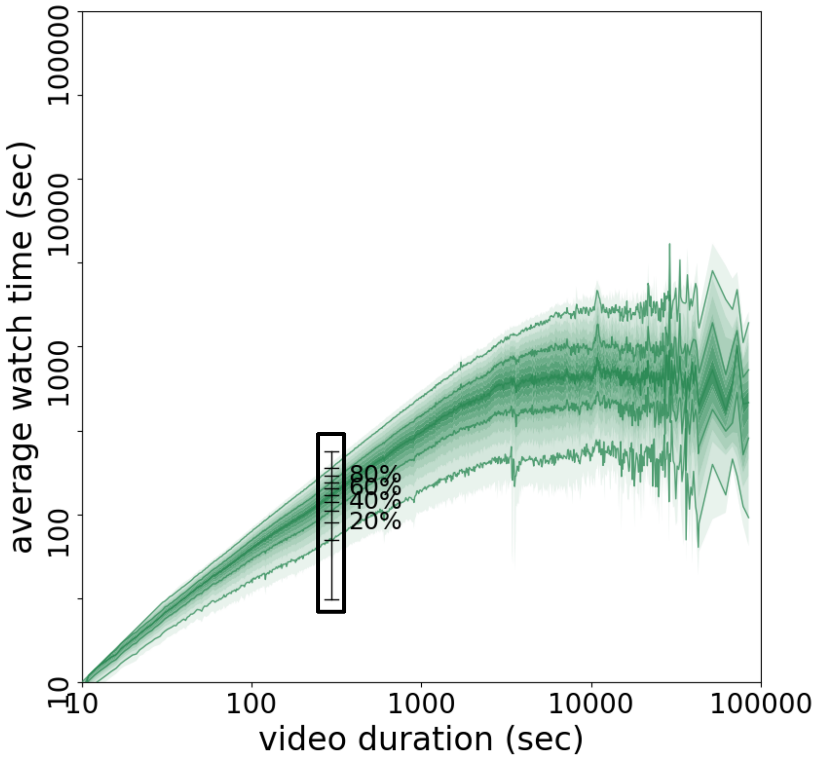

The engagement maps

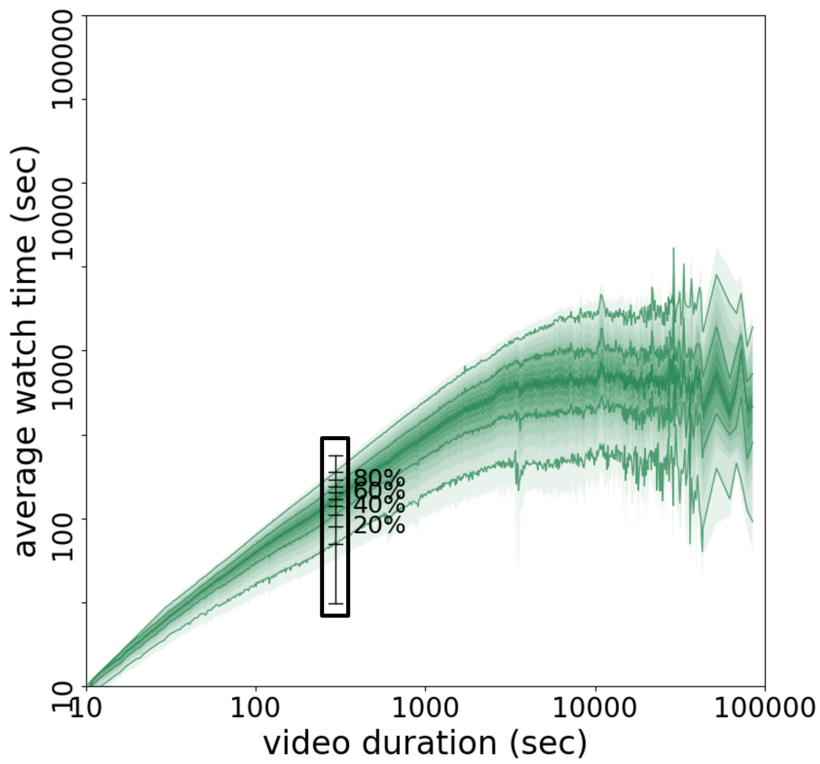

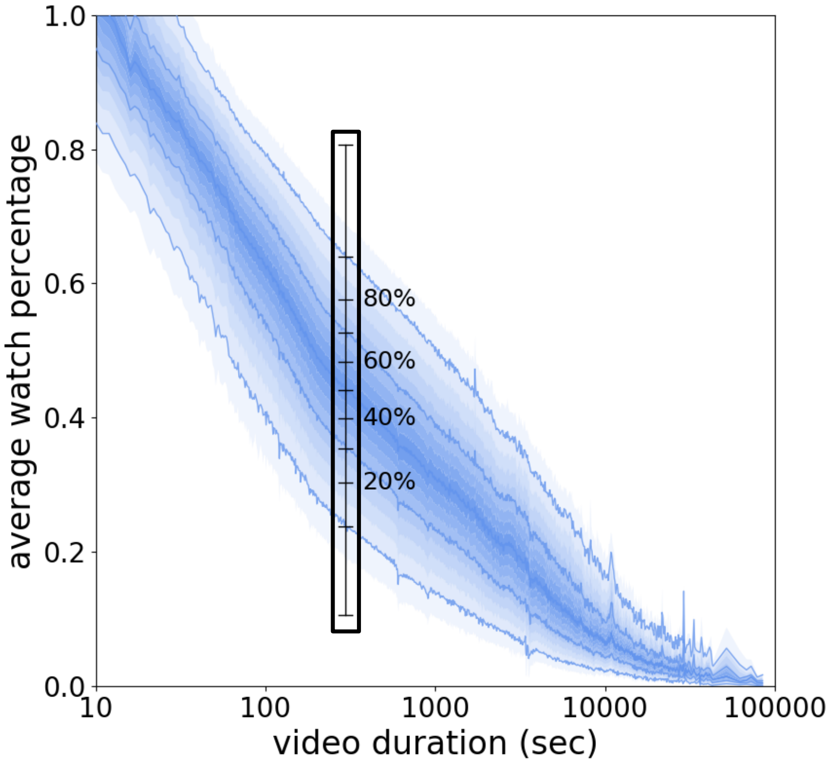

The engagement maps

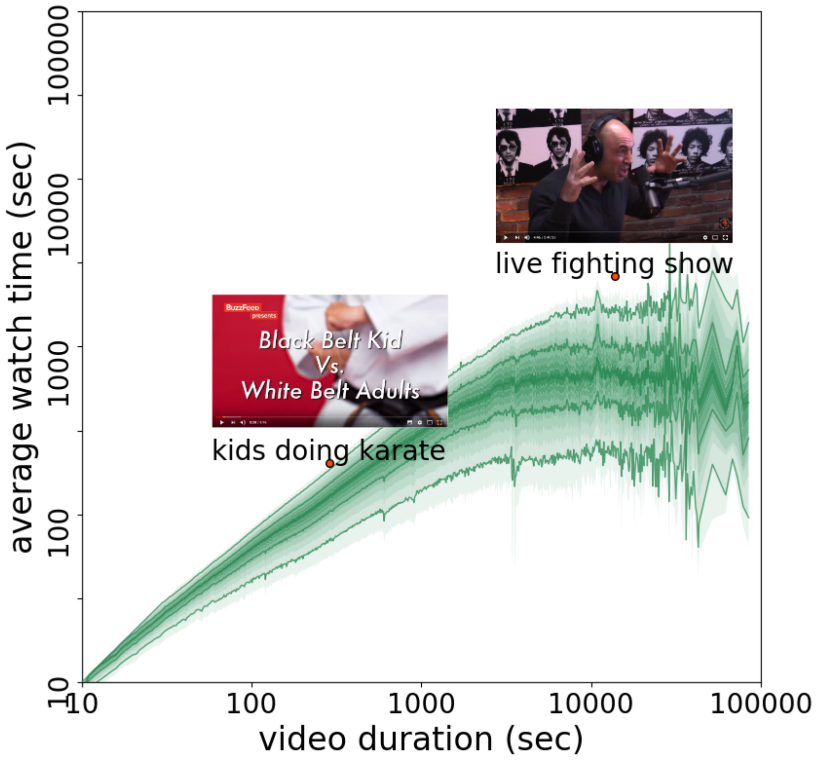

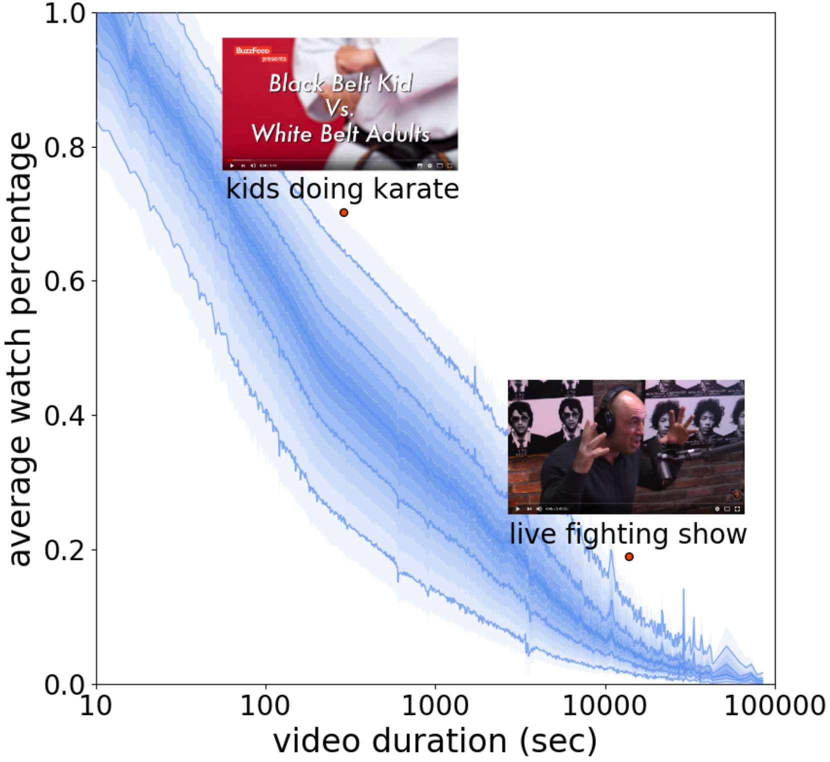

New metric: ***relative engagement*** [Wu et.al ICWSM'18]

Relative engagement

Rank percentile of average watch percentage among videos with similar lengths

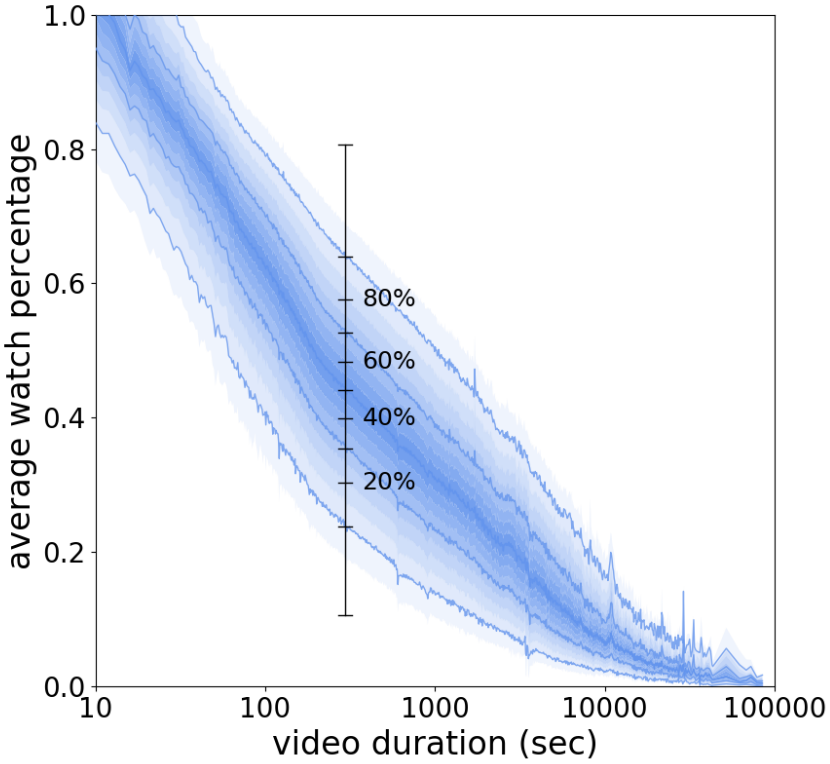

New metric: ***relative engagement*** [Wu et.al ICWSM'18]

Relative engagement

Rank percentile of average watch percentage among videos with similar lengths

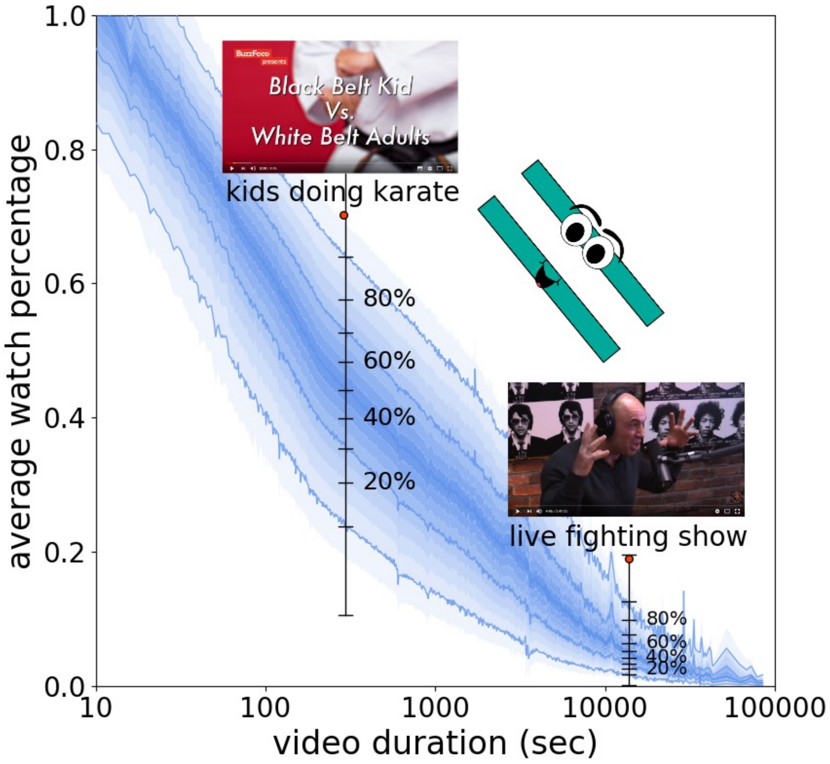

Online video engagement

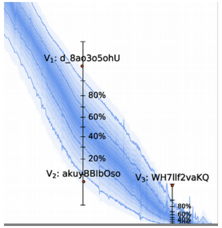

1. How to measure aggregate engagement?

Relative engagement - a new metric invariant *wrt* video duration

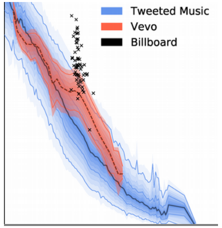

2. Characteristics of aggregate engagement

(a) Does engagement relate to content quality?

(b) How does engagement evolve over time?

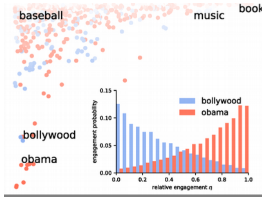

3. Can aggregate engagement be predicted?

Quality Videos datasets: Music and News

Music

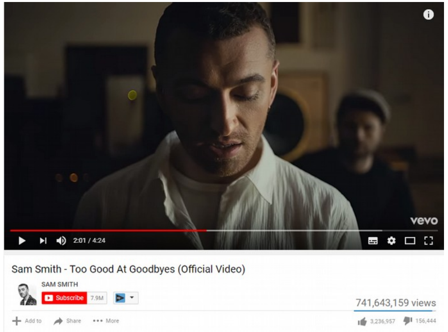

Random music clip
449,314 videos

Professional Vevo video
67,649 videos

Billboard top hit
63 videos

News

Random news clip
459,728 videos

Top News video
28,685 videos

Relative engagement is correlated with video quality

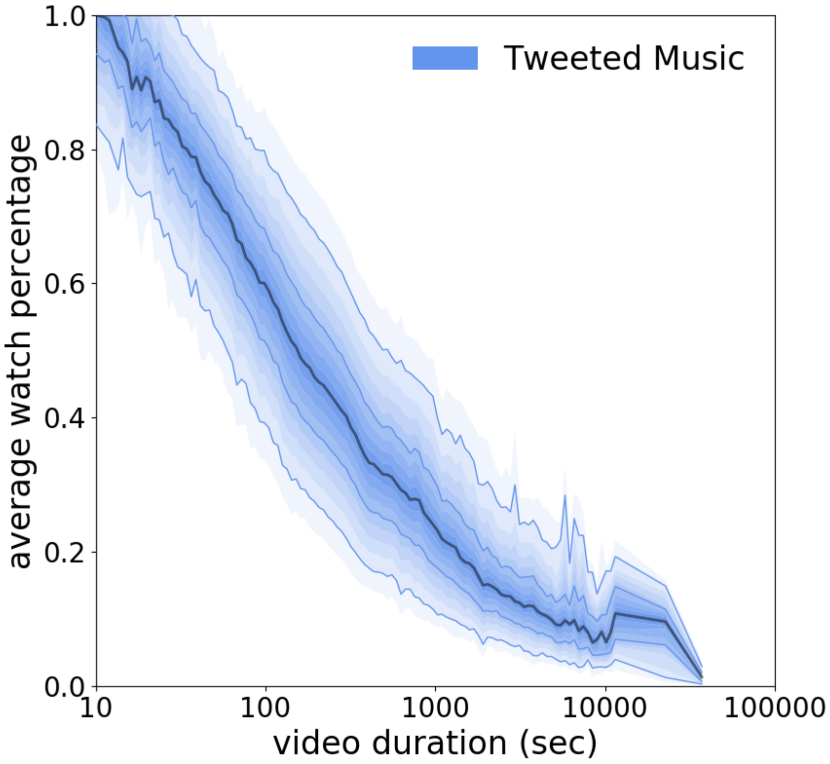

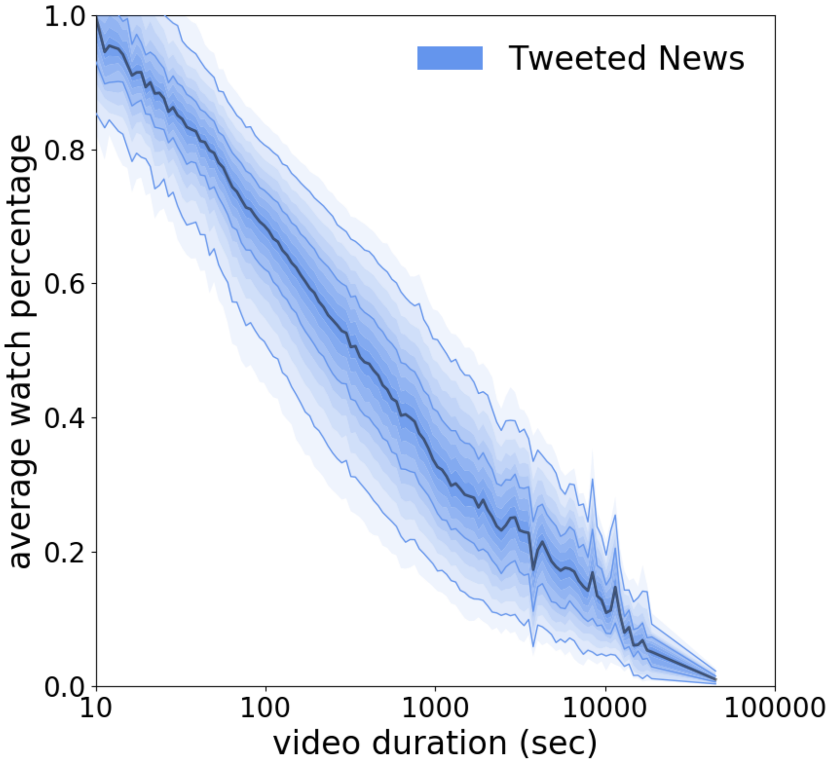

Relative engagement is correlated with video quality

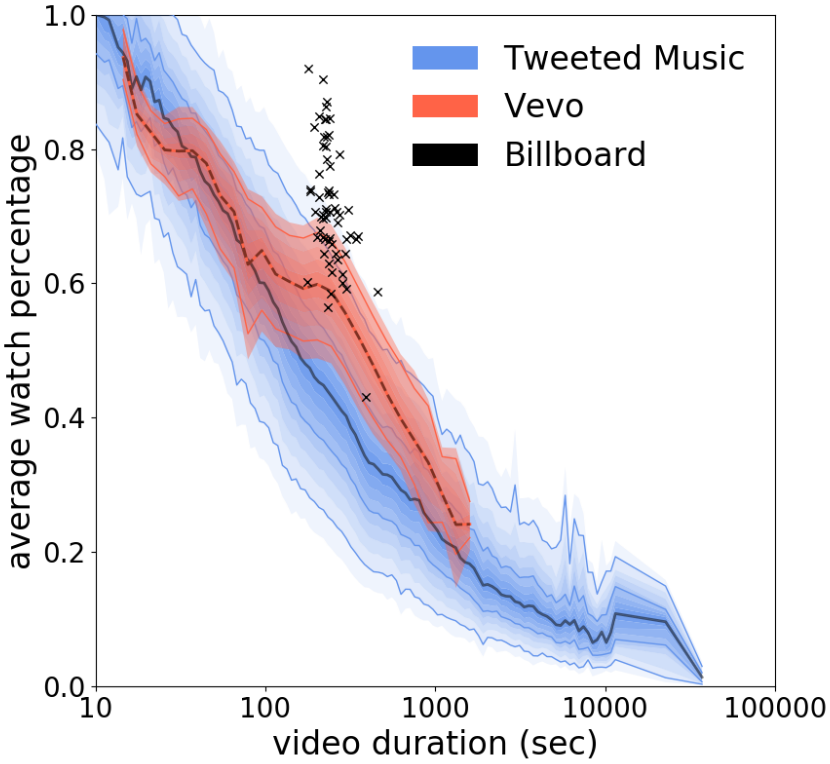

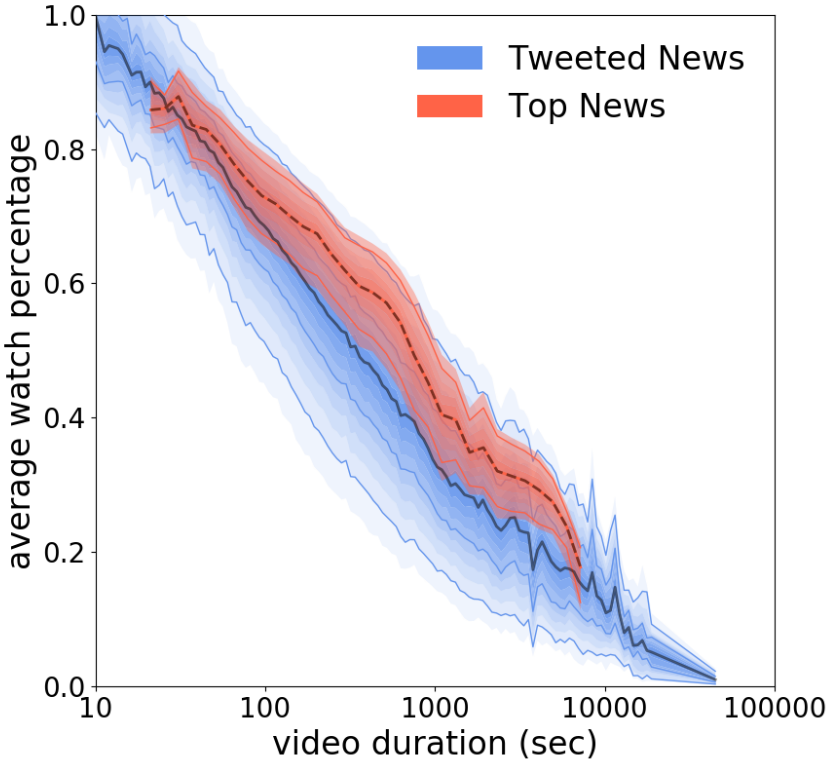

Relative engagement is stable over time

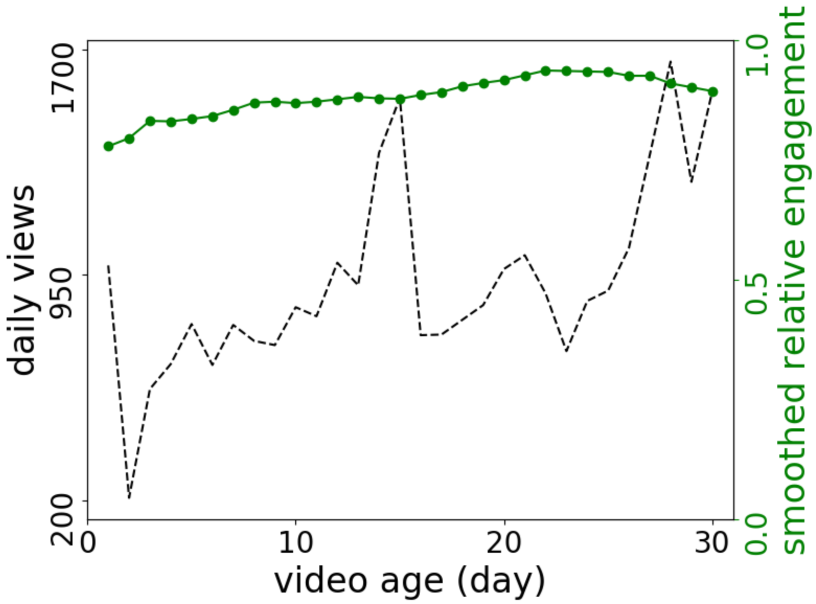

Video Id: XIB8Z_hASOs
Video Title: DC Super Hero Girls S02E10

Relative engagement is stable over time

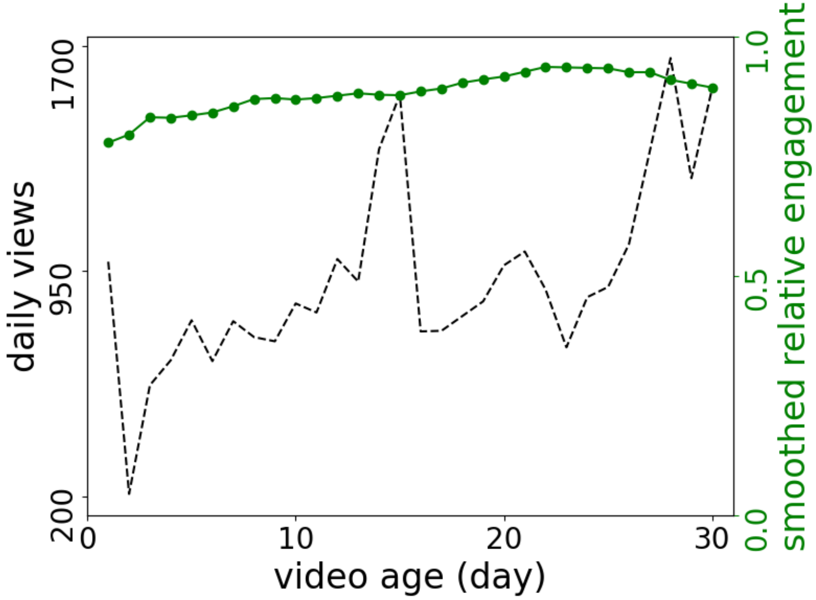

Video Id: XIB8Z_hASOs
Video Title: DC Super Hero Girls S02E10

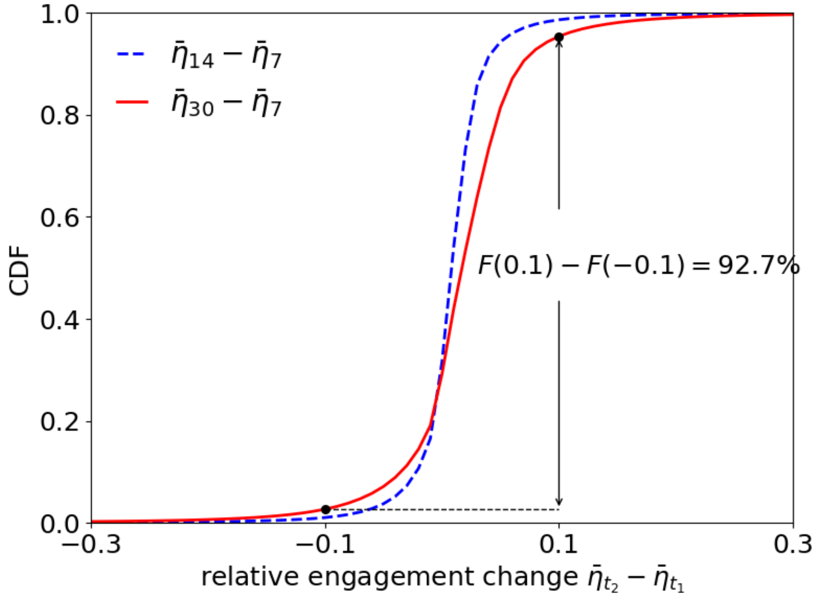

93% of videos stay within 0.1 in relative engagememt

Online video engagement

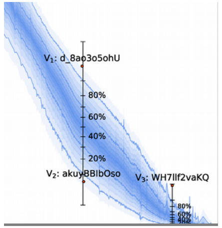

1. How to measure aggregate engagement?

Relative engagement - a new metric invariant *wrt* video duration

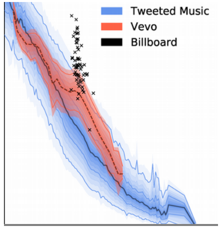

2. Characteristics of aggregate engagement

(a) Relative engagement is correlated with content quality

(b) Relative engagement is stable over time

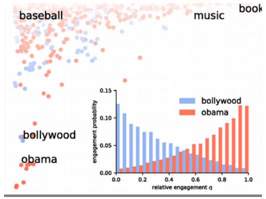

3. Can aggregate engagement be predicted?

Prediction task setup

Video duration: 4M16S

Channel activity level:
Daily upload number
Channel past engagement:
Summary of past performance

Visual definition: HD or SD
Category: Music
Language: en

Freebase topics:
Shawn Mendes; Music; Music video; Pop music

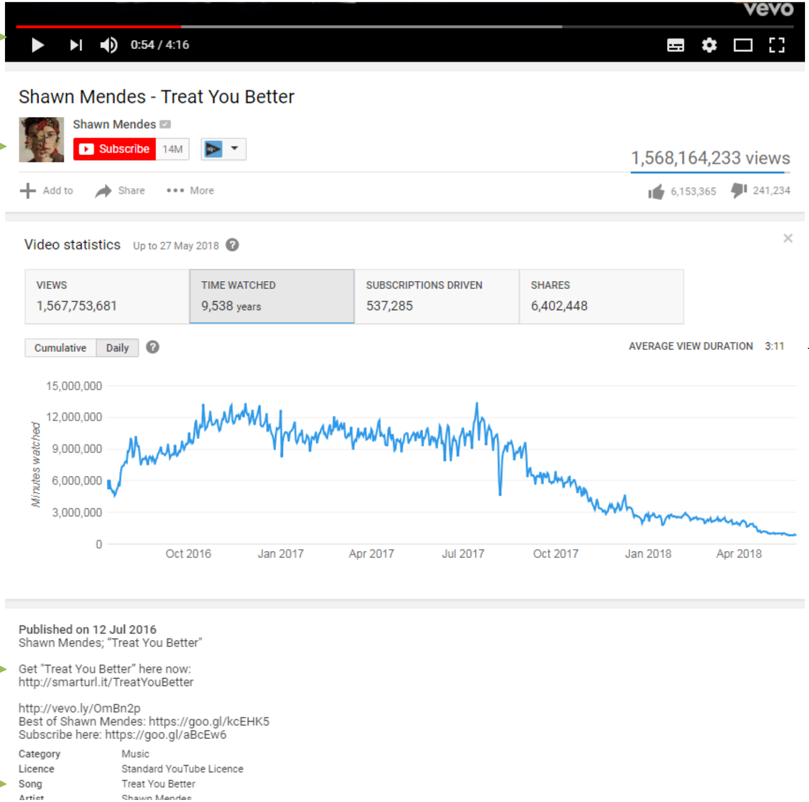

Prediction targets:
(a) Relative engagement
(b) Avg watch percentage

Prediction method:
Ridge regression

Evaluation metric:
R2

Prediction results

Predict relative engagement

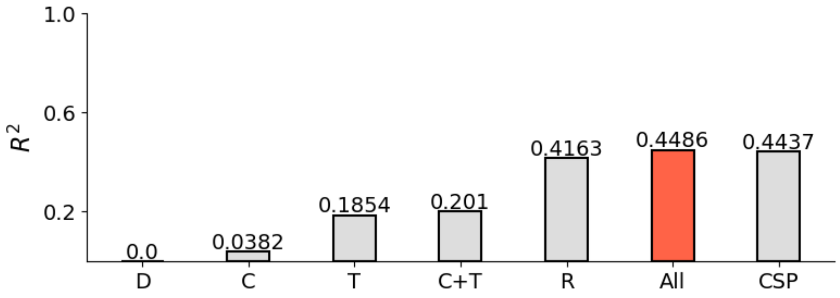

Predict average watch percentage

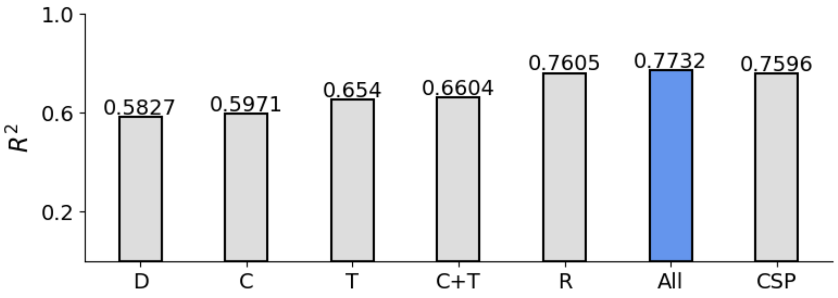

D: duration; **C:** context; **T:** topic; **C+T:** context+topic;
R: channel past reputation; **All:** all features; **CSP:** channel specific predictor

- R2 up to 0.45 for relative engagement and 0.77 for average watch percentage.

Prediction results

Predict relative engagement

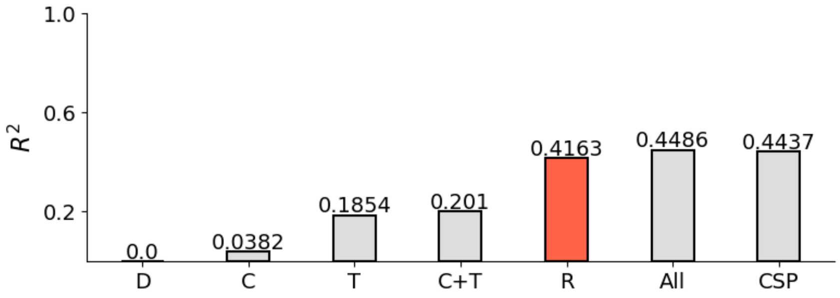

Predict average watch percentage

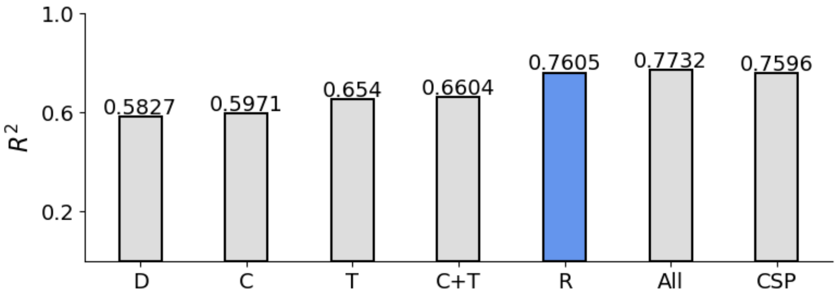

D: duration; **C:** context; **T:** topic; **C+T:** context+topic;
R: channel past reputation; **All:** all features; **CSP:** channel specific predictor

- R2 up to 0.45 for relative engagement and 0.77 for average watch percentage.
- Channel related features are the most predictive, consistent with [Cheng et al. *WWW* ’14]

Prediction results

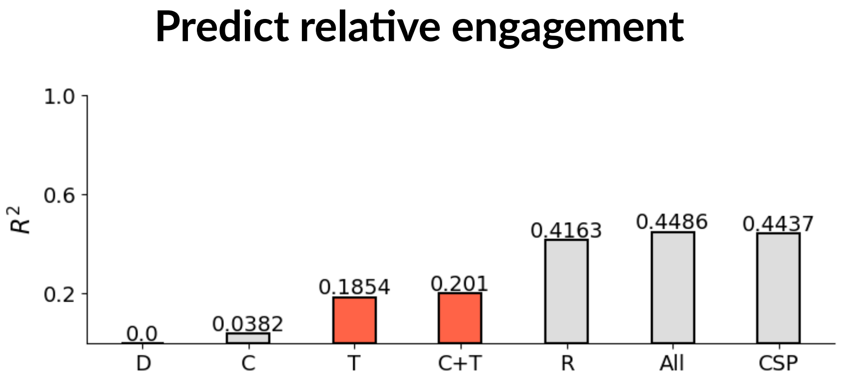

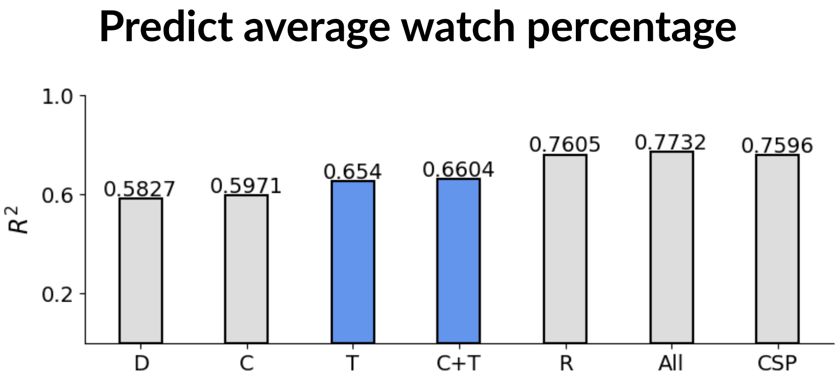

D: duration; **C:** context; **T:** topic; **C+T:** context+topic;
R: channel past reputation; **All:** all features; **CSP:** channel specific predictor

- R2 up to 0.45 for relative engagement and 0.77 for average watch percentage.
- Channel related features are the most predictive, consistent with [Cheng et al. *WWW* '14]
- Topic features are somewhat predictive, contrasting to [Martin et al. *WWW* '16]

What are engaging topics?

Conditional entropy: $H(Y|X_i = 1) = -\sum_{y \in Y} P(y|x_i = 1) \log_2 P(y|x_i = 1)$

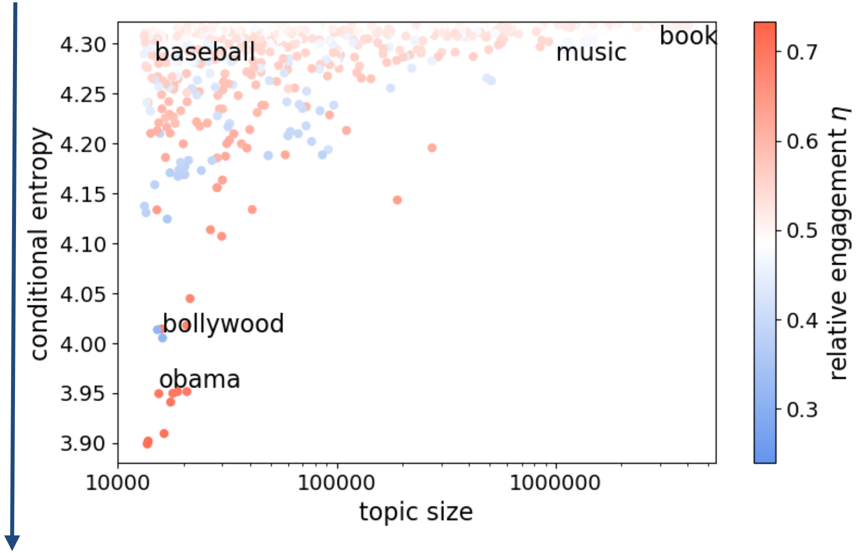

500 most frequent topics

What are engaging topics?

Conditional entropy: $H(Y|X_i=1) = -\sum_{y \in Y} P(y|x_i=1)\log_2 P(y|x_i=1)$

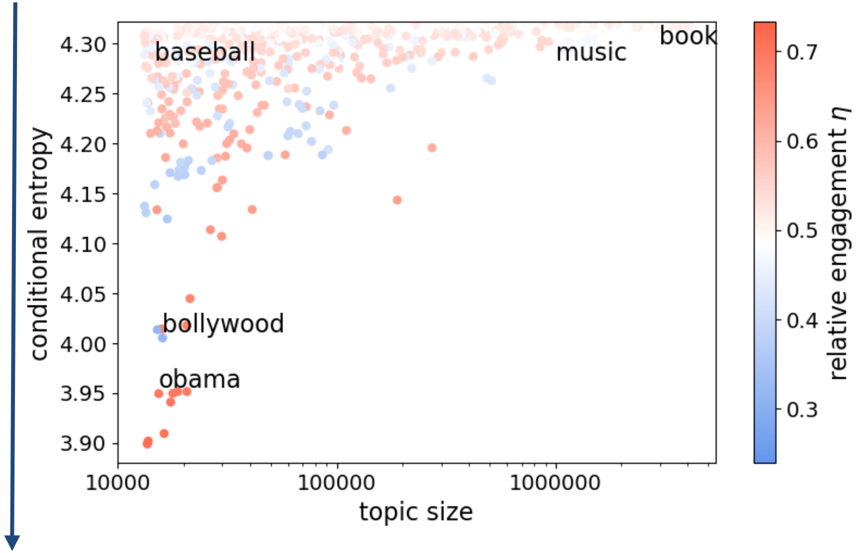

500 most frequent topics

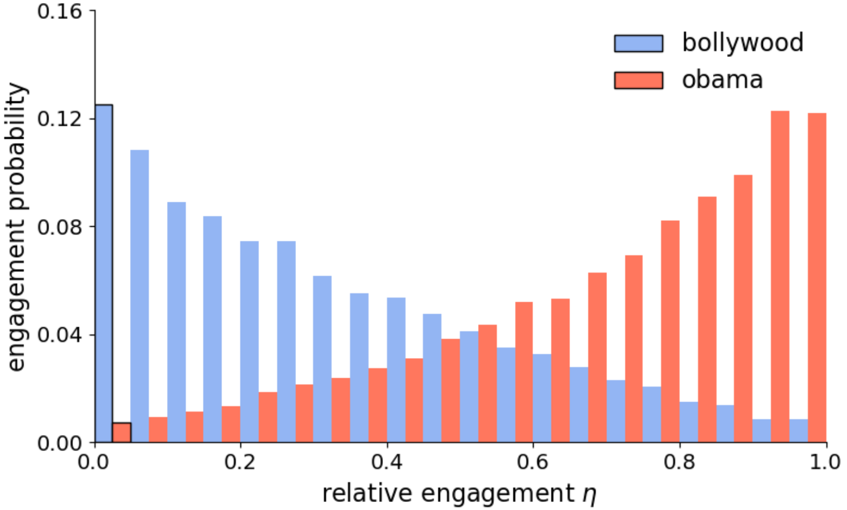

Engagement distribution of "*Obama*" and "*Bollywood*" videos

Online engagement

Get code and data from
https://github.com/avalanchesiqi/youtube-engagement

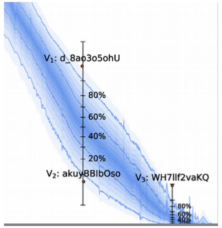

1. How to measure aggregate engagement?

Relative engagement - a new metric invariant *wrt* video duration

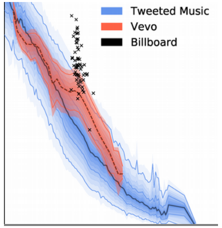

2. Characteristics of aggregate engagement

(a) Relative engagement is correlated with content quality

(b) Relative engagement is stable over time

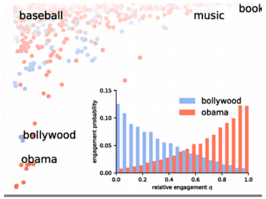

3. Can aggregate engagement be predicted?

Engagement can be predicted before video's upload, R2=0.77

Thank you!

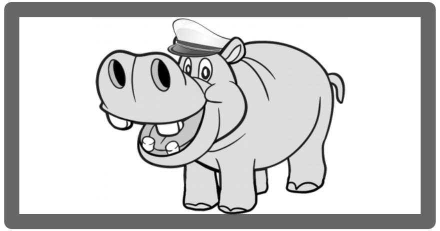

Modeling and predicting popularity using HIP

[Rizoiu et.al WWW'17]

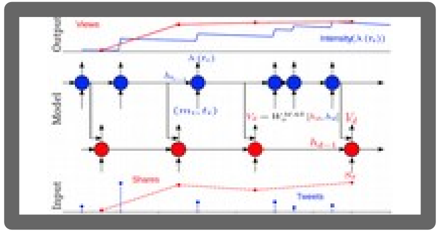

Popularity in Asynchronous Social Media Streams with RNN

[Mishra et.al ICWSM'18]

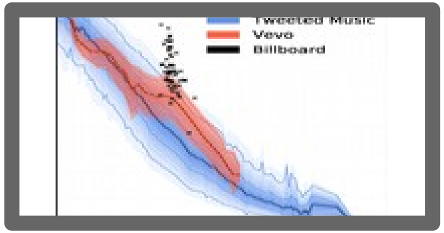

Measuring and Predicting Engagement in Online Videos

[Wu et.al ICWSM'18]