Aalto University
School of Science

Diversity maximization in social networks

Aristides Gionis

Department of Computer Science
Aalto University

International Workshop on Decision Making and Recommender Systems
Bolzano, November 29, 2018

social media

- people use social media to
  - share information, express opinion, comment, interact, discuss, get personalized news feed

- majority of EU citizens get their news from social media

PEW RESEARCH CENTER 2018

social media : good and bad sides

advantages

- no information barriers
- citizen journalism
- social connectivity
- democratization
- …

social media : good and bad sides

advantages

- no information barriers
- citizen journalism
- social connectivity
- democratization
- …

disadvantages

- harassment
- fake news
- echo chambers
- polarization
- …

polarization in US politics

1994

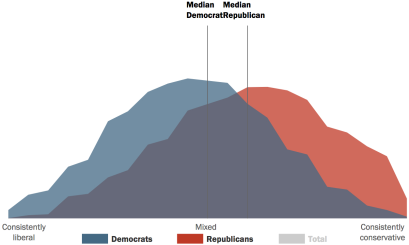

2014

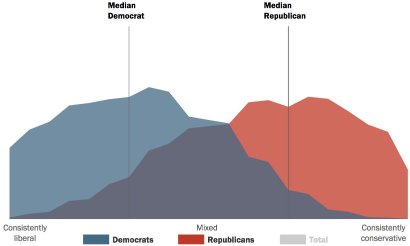

PEW RESEARCH CENTER

The biggest threat to democracy? Your social media feed

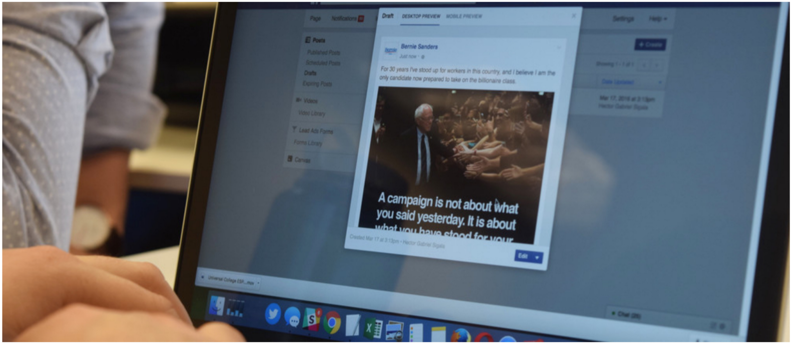

The internet was meant to spread democracy. Could it be having the opposite effect?

Image: REUTERS/Melissa Fares

echo chambers

- a situation in which information, ideas, or beliefs are amplified or reinforced by communication and repetition inside a defined system

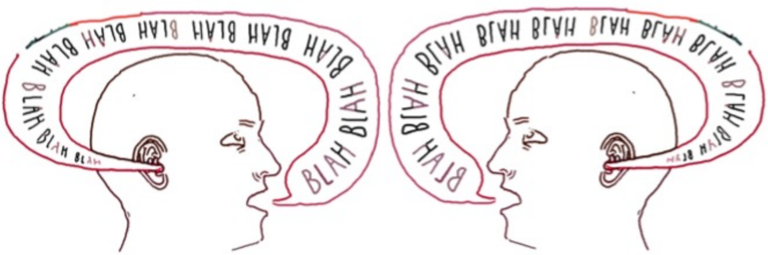

what may cause echo chambers?

- individual biases
  - homophily, confirmation bias,
    cognitive dissonance, selective exposure
- group biases
  - social identity, group polarization,
    in-group favoritism
- system biases
  - algorithmic filtering, algorithmic personalization,
    media bias

the polarization cycle

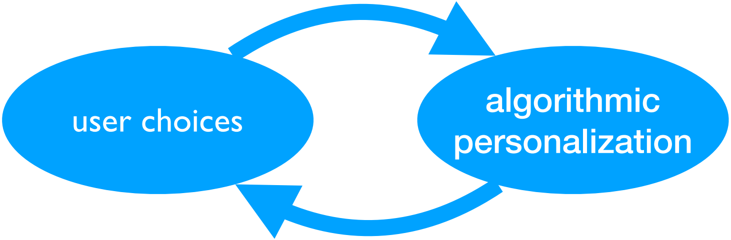

research questions

- do echo chambers exist?
- can we identify polarized discussions in social media?
- can we design algorithms to help reduce polarization?
- can we design algorithms to maximize diversity?

research questions

do echo chambers exist?

what is the interplay between content and network?

who are the key players?

K. Garimella, G. De Francisci Morales, A. Gionis, M. Mathioudakis, *“Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship”*, The Web Conference (WWW) 2018

studying echo chambers

- working definition

  the political leaning of the content that users receive from the network agrees with that of the content they share

- consider the two components of the phenomenon
  - echo : the opinion shared (content)
  - chamber : the place it is shared (network)

methodology

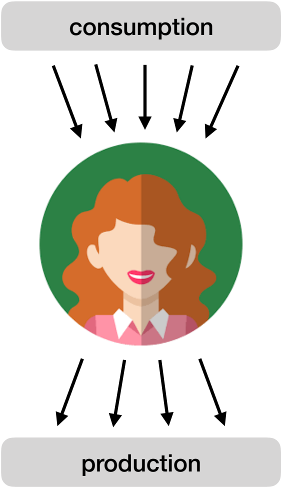

methodology

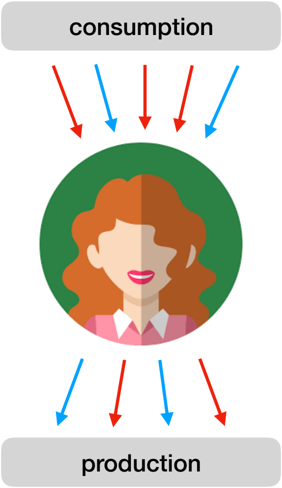

datasets

Topic	#Tweets	#Users	Event
`guncontrol`	19M	7506	Democrat filibuster for gun-control reforms (June 12–18, 2016)[6]
`obamacare`	39M	8773	Obamacare subsidies preserved in us supreme court ruling (June 22–29, 2015)[7]
`abortion`	34M	3995	Supreme court strikes down Texas abortion restrictions (June 27–July 3, 2016)[8]
`combined`	19M	6391	2016 US election result night (Nov 6–12, 2016)
`large`	2.6B	676 996	Tweets from users retweeting a U.S. presidential/vice presidential candidate (from [4], 2009–2016)
#ff	4M	3204	
#gameofthrones	5M	2159	
#love	3M	2940	filtering for these hashtags
#tbt	28M	12 778	
#foodporn	8M	3904	

content

- focus on news outlets e.g., NYT, BBC, CNN, etc.
- assign content polarity score at each outlet

  0 : liberal — 1 : conservative
- obtain ground-truth scores for top-500 outlet

[Bakshy et al., Science, 2015]

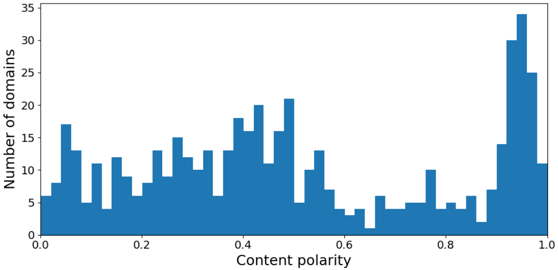

characterize users based on

- production polarity : avg polarity of shared content
- consumption polarity : avg polarity of followees' content

user roles : partisan

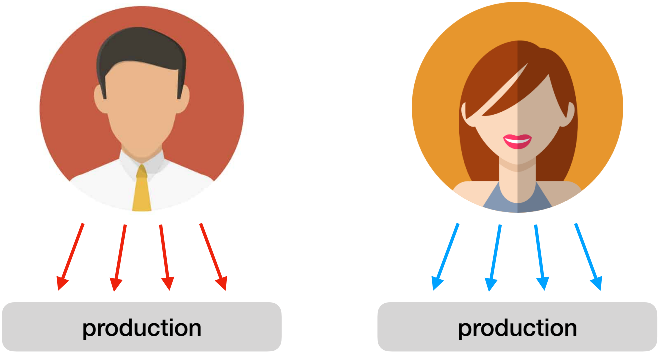

user roles : bi-partisan

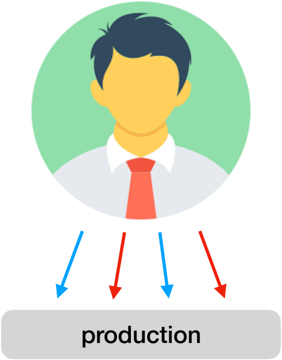

user roles : consumer

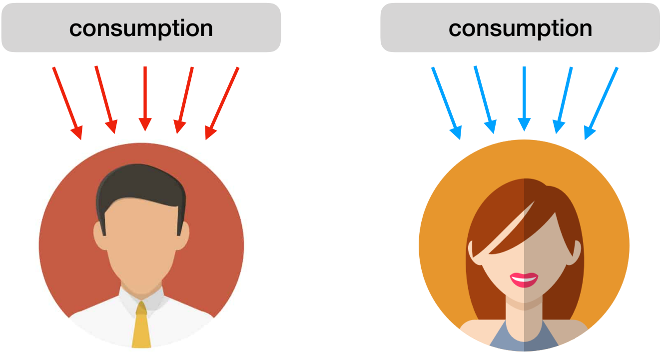

user roles : gatekeeper

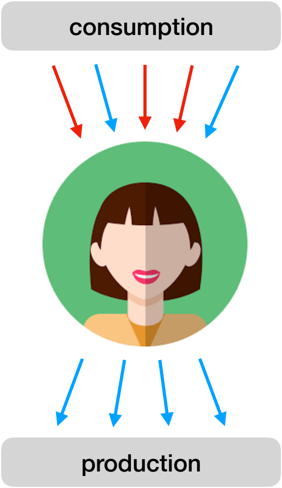

users — production-polarity distribution

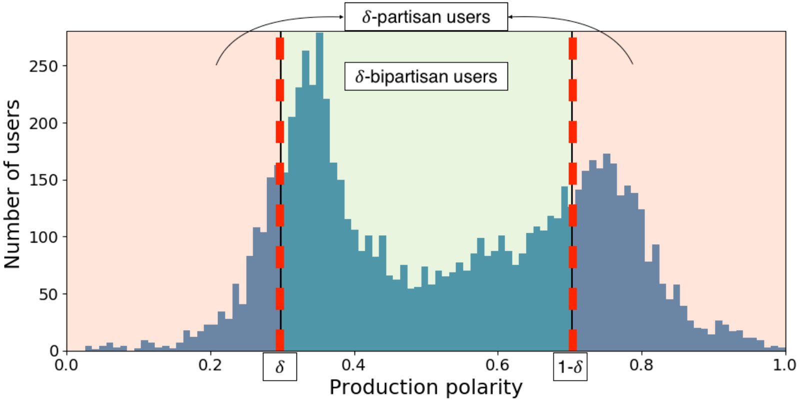

network features

- user polarity (democrat vs. republican)

  [Barberá et al., Psychological Science, 2015]
- network centrality : PageRank, in-degree
- clustering coefficient
- retweet ratio
- retweet volume

echo chambers

content production and consumption

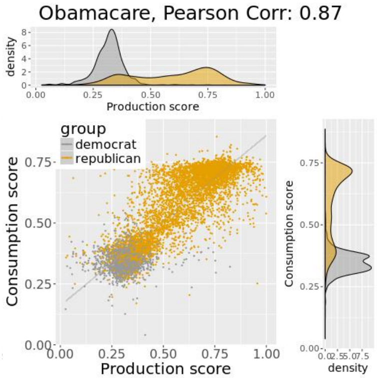

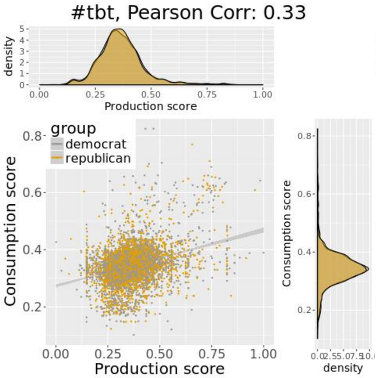

partisans vs. bi-partisans

gatekeepers vs. non gatekeepers

Features	Partisans	Gatekeepers
PageRank	✓	✓
clustering coefficient	✓	✓ (-)
user polarity	✓	✓ (-)
degree	✓	✓
retweet rate	✓	✗
retweet volume	✓	✗
favorite rate	✓	✗
favorite volume	✓	✗
# followers	✗	✗
# friends	✗	✗
# tweets	✗	✗
age on Twitter	✗	✗

partisans vs. bi-partisans

gatekeepers vs. non gatekeepers

Features	Partisans	Gatekeepers
PageRank	✓	✓
clustering coefficient	✓	✓ (-)
user polarity	✓	✓ (-)
degree	✓	✓
retweet rate	✓	✗
retweet volume	✓	✗
favorite rate	✓	✗
favorite volume	✓	✗
# followers	✗	✗
# friends	✗	✗
# tweets	✗	✗
age on Twitter	✗	✗

there is a price to be bi-partisan

prediction

- tweet features
  - n-grams with tf $\cdot$ idf weights
- profile features
  - number of tweets / followers / friends, age on twitter
- network features
  - PageRank, degree, clustering coefficient

predicting partisans (accurasy $\approx$ 0.81)

is easier than

predicting gatekeepers (accurasy $\approx$ 0.68)

summary of findings

- echo chambers observed in politically contentious topics
- echo chambers not observed in non-contentious topics
- bi-partisan users pay a price in terms of network centrality, community connection, and endorsements
- gatekeepers : who are they and what is their role?

  e.g., open-minded citizens or “soldiers” of one side?

research question

can we identify and quantify polarization ?

K. Garimella, G. De Francisci Morales, A. Gionis, M. Mathioudakis, *"Quantifying controversy in social media"*, ACM WSDM 2016

how can we identify polarization ?

ideas

- content
  - do opposing sides say different things ?
- sentiment
  - do polarized topics exhibit wider range of emotions ?
- interactions
  - do people interact more with their own side ?

high-level approach

- build an interaction graph
- is the interaction graph polarized?
- output polarization score

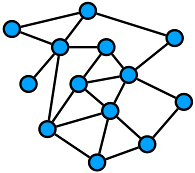

non polarized

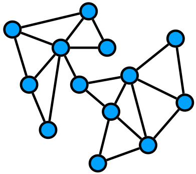

polarized
two sides well separated

many different options

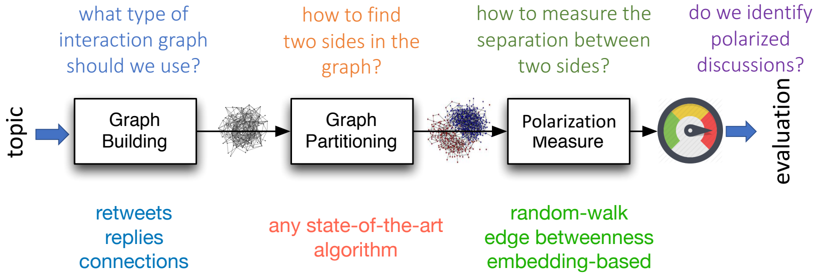

random-walk controversy score (RWC)

- how likely a random user on either side to be exposed to authoritative content from the opposing side
- assume graph is partitioned in two sides, A and B
- consider a random walk that started at a random node and finished in a hub in $Y \in \{A, B\}$
- probability that random walk started in $X \in \{A, B\}$

$$P_{XY} = \Pr(\text{r.w. started in } X \mid \text{r.w. finished in } Y)$$

- random-walk controversy score (RWC)

$$\textsc{rwc} = P_{AA}P_{BB} - P_{AB}P_{BA}$$

does not depend on cluster sizes and relative in-degrees

evaluation

- annotate polarized and non-polarized topics
- polarized
  - indian beefban, nemtsov protests, netanyahu US congress speech, baltimore riots, ukraine
- non-polarized
  - germanwings plane crash, sxsw, mother's day, jurassic world movie, national kissing day
- evaluate different settings on ground truth

best performing setting

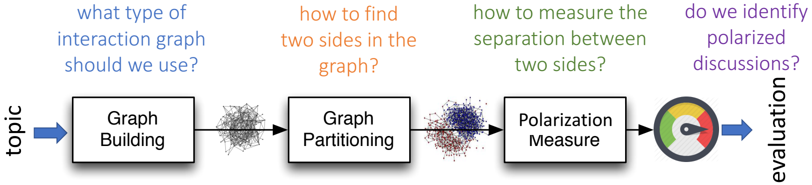

- retweet graph
- RWC

other good settings: edge betweenness score

sentiment variance

example of results

high RWC

low RWC

polarized topics

nemtsov protests

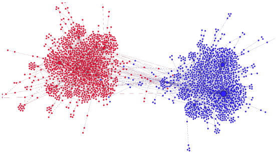

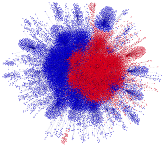

germanwings plane crash

indian beef ban

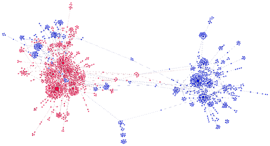

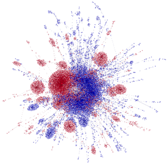

sxsw conference

non-polarized topics

using retweet graph

example of results

nemtsov protests

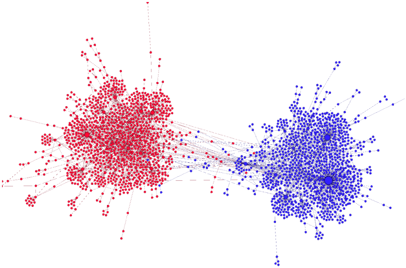

retweets

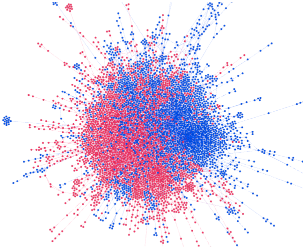

replies

research questions

design algorithms to help reducing polarization

design algorithms to maximize diversity

mitigation action I

improve awareness

P. Lahoti, K. Garimella, A. Gionis, *"Joint non-negative matrix factorization for learning ideological leaning on twitter"*, ACM WSDM 2018

improve awareness

- develop tools for users to perceive their "news diet"
- visualize/navigate in the underlying ideology space, their position, the accounts they follow, the news they read
- offer functionalities such as

  "find a high-quality article on the same topic from the opposing viewpoint"

learning of ideological leanings

- infer ideological stances of users and content
  e.g., liberal–conservative space
- common latent space for users and content
- e.g., substitute ground-truth polarities in previous study with learned polarities
- joint non-negative matrix-factorization task

intuition

- map users and content in a joint latent ideology space

such that

- similar users are more likely to follow each other
- similar users are more likely to share similar content
- similar content is more likely to be shared by similar users

*similar means close in the latent ideology space

the problem setting

- social network $G = (V, E)$
  - adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$
- user–content matrix $\mathbf{C} \in \mathbb{R}^{m \times n}$
- latent matrix representing user ideology $\mathbf{U} \in \mathbb{R}^{n \times k}$
- latent matrix representing content ideology $\mathbf{V} \in \mathbb{R}^{m \times k}$
- decompose

$$\mathbf{A} \approx \mathbf{U}\mathbf{H}_u\mathbf{U}^T \quad \text{and} \quad \mathbf{C} \approx \mathbf{U}\mathbf{H}_v\mathbf{V}^T$$

subject to orthonormal $\mathbf{U}$ and $\mathbf{V}$ and graph-regularization

evaluation

- twitter data from 2011 to 2016, focusing on controversial topics (gun control, abortion, obamacare)
- 6 391 users and 19 million tweets
- gather ground-truth polarity scores
  - content polarity [Bakshy et al., 2015]
  - user polarity [Barberá et al., 2015]

content ideology scores

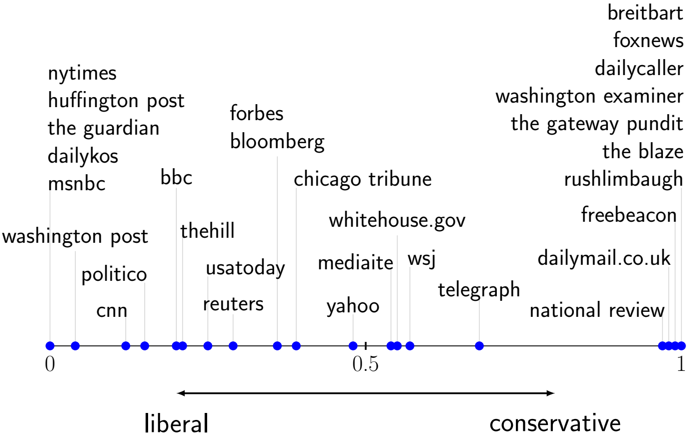

correlation with ground-truth scores 0.82

audience ideology scores

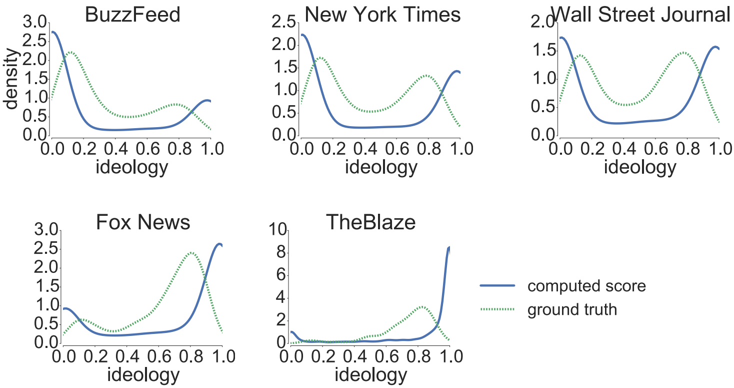

correlation of user ideology scores with ground-truth 0.90

visualizing the information bubble

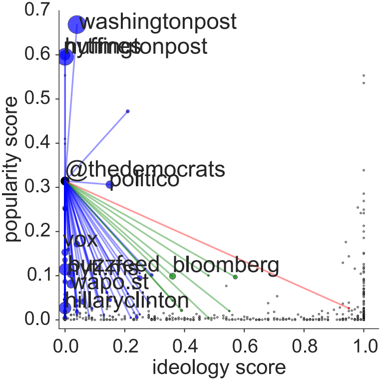

@thedemocrats

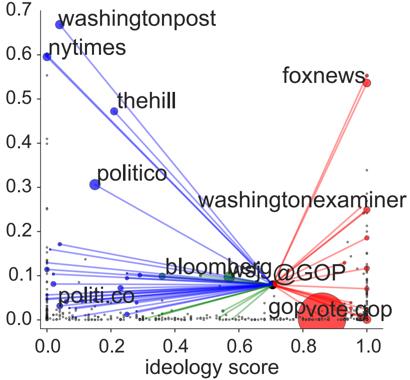

@gop

mitigation action II

user-to-user recommendation

K. Garimella, G. De Francisci Morales, A. Gionis, M. Mathioudakis, *"Reducing controversy by connecting opposing views"*, ACM WSDM 2017

user-to-user recommendation

- social network has clustered structure

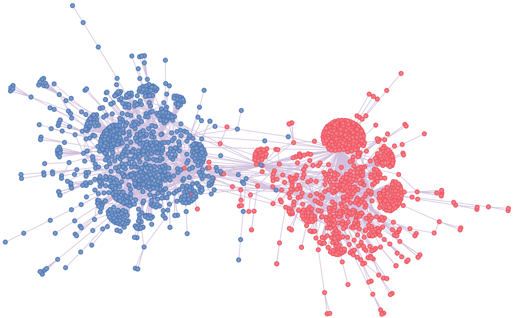

- user-to-user recommendation to reduce clustered structure
- e.g., minimize average shortest path length, maximize conductance, etc.

reducing polarization

how can we bridge the divide?

- assuming
  - polarization score measured by RWC

$\Rightarrow$ we want to reduce RWC

- problem
  - add k edges that maximally reduce RWC

reducing polarization

- greedy algorithm
  - find the single best edge to reduce RWC
  - repeat k times
- inefficient
  - computing RWC requires $\mathcal{O}(\text{MMULT}(n))$
    faster in practice with iterative computation
  - still, greedy requires $\mathcal{O}(n^2 \cdot k \cdot \text{MMULT}(n))$
- improvements
  - consider adding edges only between hubs
  - incremental RWC computation using
    Sherman-Morrison formula

reducing polarization

- what does it mean "add k edges"?
- answer: recommendations
- but many recommendations are unlikely to be materialized

  no point recommending D. Trump to retweet H. Clinton

- incorporate probability of accepting a recommendation
  - compute user polarity, and
  - acceptance probability as a function of user polarity

reducing polarization : real example

polarity=-.99

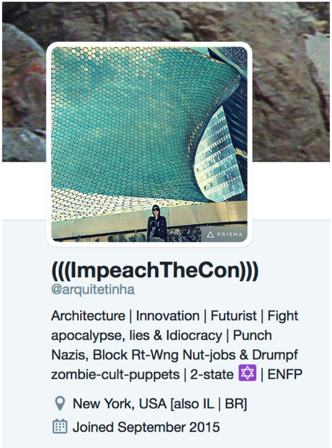

polarity=.95

reducing polarization : real example

polarity=-.99

polarity=.15

reducing polarization : results

	obamacare		guncontrol	
	node1	node2	node1	node2
ROV	mittromney	barackobama	ghostpanther	barackobama
	realdonaldtrump	truthteam2012	mmflint	robdelaney
	barackobama	drudge_report	miafarrow	chuckwoolery
	barackobama	paulryanvp	realalexjones	barackobama
	michelebachmann	barackobama	goldiehawn	jedediahbila
ROV-AP	kksheld	ezraklein	chuckwoolery	csgv
	lolgop	romneyresponse	liamkfisher	miafarrow
	irritatedwoman	motherjones	csgv	dloesch
	hcan	romneyresponse	jonlovett	spreadbutter
	klsouth	dennisdmz	drmartyfox	huffpostpol

mitigation action III

maximize diversity

A. Matakos, A. Gionis, *"Tell me something my friends do not know: Diversity maximization in social networks"*, ICDM 2018

K. Garimella, A. Gionis, N. Parotsidis, N. Tatti, *"Balancing information exposure in social networks"*, NIPS 2017

C. Aslay, E. Galbrun, A. Matakos, A. Gionis, *"Maximizing the diversity of exposure in a social network"*, ICDM 2018

maximizing diversity

- goal : make recommendations to maximize diversity
- what is diversity and how to measure it?
- user level : recommend diverse content
- network level : make recommendations so that friends see different content
  - motivation : friends can discuss / debate
- combinations
- another consideration : propagation effects, or not

maximizing diversity

"*tell me something my friends do not know*"

- goal : recommendations to maximize network diversity
- make a small number of recommendations (k)
  - why ? intervene as little as possible
- a simple formulation that captures the essence of setting
  - graph $G = (V, E)$ where nodes have values $+1$ or -1
  - corresponds to what kind of content they see
  - select k nodes to swap their values so as to maximize the number of edges having different values at their endpoints, i.e., edges with values (+1,-1)

Matakos et al., ICDM 2018

maximizing diversity

“*tell me something my friends do not know*”

toy example in “karate club”

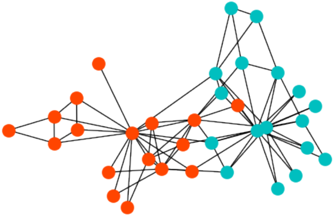

(a) Echo-chamber graph

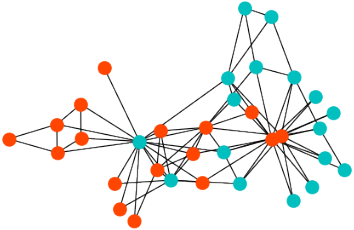

(b) Graph with diversified exposure

optimal solution for $k = 4$

maximizing diversity

"*tell me something my friends do not know*"

- problem **NP**-hard (generalization of MAX-CUT)
  - also **NP**-hard to approximate
- problem formulation non convex 0-1 quadratic problem
  - an instance of quadratic knapsack QK
- proposed solutions
  - SDP relaxation + rounding, inspired by QK solutions
  - Glover's linearization, solve LP, round
  - greedy (extremely scalable and high-quality)
  - exact solution (not scalable)

maximizing diversity
"*tell me something my friends do not know*"

- many future directions
- study more realistic problem formulations
  - continuous user leaning scores
  - continuous content leaning scores
  - probability of accepting a recommendation

balancing information exposure

- setting inspired by viral marketing
  - a social network and two campaigns
  - seed nodes I_1 and I_2 for the two campaigns
  - a model of information propagation
- the problem of balancing information exposure
  - find additional seeds S_1 and S_2, with $|S_1| + |S_2| \leq k$
  - s.t. minimize # of users who see only one campaign
    or maximize # of users who see both or none

Garimella et al., NIPS 2017

illustration

social discussion on fracking

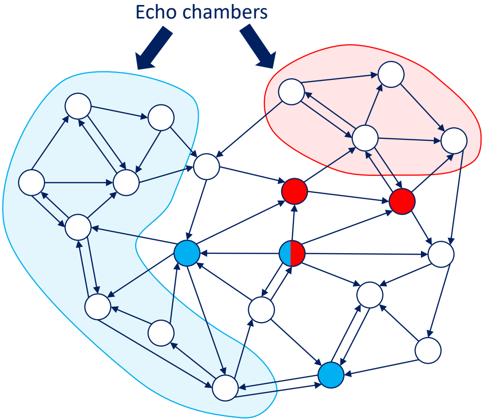

balancing information exposure — results

- optimization problem is **NP**-hard
- minimization problem is **NP**-hard to approximate
- maximization problem: objective function non monotone and non submodular
- different models of how the two campaigns propagate
- approximation guarantee $\frac{1}{2}(1 - \frac{1}{e})$

balancing information exposure — example

Side 1
Pro-Choice

Side 2
Pro-Life

Hedge

Pro-Remain

Pro-Leave

maximizing diversity of exposure

- goal : recommendations to maximize user diversity
- consideration : recommended content may be shared among users, creating possible cascades
- make a small number of recommendations
  - why ? intervene as little as possible
  - make at most k recommendations in total
  - make at most k_i recommendations to user i

Aslay et al., ICDM 2018

maximizing diversity of exposure

- problem formulation inspired by influence maximization
- item propagation modeled by independent cascade
  - influence prob. depend on user and item leanings
- we want to recommend k items to k users,
  so as to maximize the diversity score

$$\sum_{v \in V} \left(\max_{i \in E(v)} \ell(i) - \min_{i \in E(v)} \ell(i) \right),$$

$E(v)$: items that v is exposed, considering also cascades

$\ell(i)$: leaning score of item i

maximizing diversity of exposure — results

- diversity function is submodular
- greedy algorithm provides $\frac{1}{2}$ approximation
  - maximizing a submodular function under partition matroid constraints
- but computation prohibitively expensive
  - Monte-Carlo simulations
- adapt recent techniques to obtain highly scalable algorithm
  - generalize the idea of reverse-reachable sets
  - sample-size estimation using martingales

mitigation action IV

clustering with non-polarized representatives

B. Ordozgoiti, A. Gionis, *“Reconciliation k-median”*, work in progress

clustering with non-polarized representatives

- clustering formulations are used to select representatives
- *k*-median: select set of representatives R to

$$\text{minimize} \quad \text{cost}(S) = \sum_{x \in X} \min_{r \in R} d(x, r)$$

clustering with non-polarized representatives

- clustering formulations are used to select representatives
- *k*-median: select set of representatives R to

$$\text{minimize} \quad \text{cost}(S) = \sum_{x \in X} \min_{r \in R} d(x,r) \; + \; \underbrace{\lambda \sum_{r \in R} \sum_{s \in R} d(r,s)}_{\text{consensus } k\text{-median}}$$

clustering with non-polarized representatives

- clustering formulations are used to select representatives
- k-median: select set of representatives R to

$$\text{minimize} \quad \text{cost}(S) = \sum_{x \in X} \min_{r \in R} d(x,r) \; + \; \underbrace{\lambda \sum_{r \in R} \sum_{s \in R} d(r,s)}_{\text{consensus } k\text{-median}}$$

- application: select a set of articles to summarize an event
- local-search algorithm yields
  - $\mathcal{O}(k)$ approximation
  - $\mathcal{O}(1)$, if clusters are *large enough*

clustering with non-polarized representatives

consensus k-median on twitter US presidential election dataset

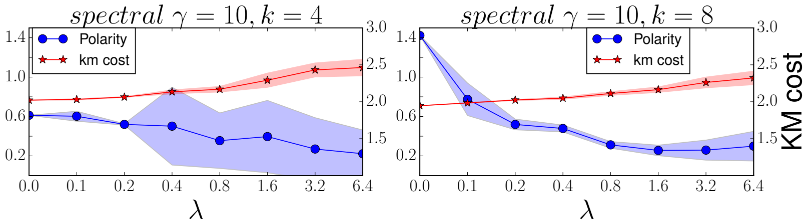

summary

- evidence of echo chambers in social networks
  - price of bi-partizanship
- quantifying polarization in social media
  - random-walk controversy score
- actions to mitigate echo chambers
  - improve awareness
  - connecting opposing sides
  - maximize diversity via recommendations
  - clustering with non-polarized representatives

discussion, limitations, future work

- models use mostly network structure
  - language-independent, but
  - incorporating language can help
- simple models
  - two-sided controversies
  - external influence is ignored
  - “follow” does not imply content consumption
  - simple propagation models
- evaluation is challenging, done on few topics
- analysis limited to twitter

reflections

- attempts to reduce polarization may backfire
  - user interface is very important
- perhaps people do not want to break their filter bubbles
  - still, it should be possible to increase their awareness and transparency of the content they receive

thank you

Q & A

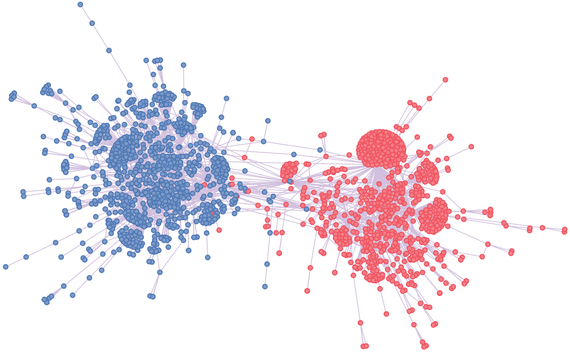

credits

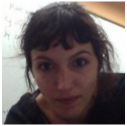

Cigdem Aslay
Aalto Univ

Gianmarco De Francisci Morales
ISI Foundation

Esther Galbrun
Aalto Univ

Kiran Garimella
EPFL

Preethi Lahoti
MPI

Antonis Matakos
Aalto Univ

Michael Mathioudakis
Univ of Helsinki

Bruno Ordozgoiti
UP Madrid

Nikos Parotsidis
Univ of Rome

Nikolaj Tatti
F-Secure