Social Graphs and Society

Kellin Pelrine

An interconnected world

Social graphs encode the interactions and relationships between people

Facebook network in 2010
https://paulbutler.org/2010/visualizing-facebook-friends/

Increasingly interconnected

- Where once people were separated by 6 degrees, now around 4

Social Media Network	Degrees of Separation
Twitter Global	4.17
Twitter Brazil	3.78
Twitter Japan	3.89
Twitter U.S.	4.37
Facebook	4.74
MSN Messenger	6.60

S. Aral, The Hype Machine (2020)

Increasingly technological

- Social media an example, naturally

- But all kinds of other interactions are now facilitated by social technologies

And even more in the future...

Top stories

CNN

Facebook is planning to change its name, report says

5 hours ago

Vox

Facebook's name change plan is a reflection of priority on the metaverse

1 hour ago

THE VERGE

Facebook plans to change company name to focus on the metaverse

19 hours ago

So...

- In short, technology is changing who we interact with and the tools to interact with them...
- But why do we care as a society?

- It’s also changing how we interact, and those interactions shape us and our world

“Breaking: Two Explosions in the White House and Barack Obama is injured”

- Associated Press (hacked)

- Over $100 billion gone in seconds from a fake tweet

Data compiled by Bloomberg shows the Dow Jones lost 120 points within minutes of an erroneous report about a bombing at the White House. (Pete Evans/CBC)

Integrating Social Media for Pandemic Response: A “Double-Edged” Path

- Banerjee and Meena (2021)

“Lies spread faster than the truth”

- Science, referencing Vosoughi et al. (2018)

Managing the COVID-19 infodemic: Promoting healthy behaviours and mitigating the harm from misinformation and disinformation

Joint statement by WHO, UN, UNICEF, UNDP, UNESCO, UNAIDS, ITU, UN Global Pulse, and IFRC

WHO, 2020

Political Polarization

- Complex interactions between social media and polarization.

- Simplistic “echo chamber” alone perhaps overstated in popular media… Full effects are probably more complicated.
- E.g. Likes and shares make people more vitriolic (Brady et al. 2020)

Political Polarization

Regardless, understanding and mitigating harms of extreme polarization is crucial:

- Partisan polarization drives political misinformation sharing (Osmundsen et al. 2021)
- Can lead to events like January 6th US capitol riot, politicization of COVID-19...

General Motivation

Huge amount of discussions, interactions, *data* out there.

How can we harness it?

- To understand human behavior, and how it’s changing over time with tech
- To mitigate harmful things, like misinformation or increasing polarization
- To promote beneficial things, like science-based COVID-19 actions

About Me

PhD student, CS/ML

- Complex Data Lab, supervised by Prof. Rabbany

Interests:

- Learning from “broad data:” combining heterogeneous types, sources, tasks, etc.
- Polarization: measuring, understanding, mitigating
- Misinformation: likewise ^
- Empirical evaluation: good performance on paper -> good performance in real world?

kellin.pelrine@mila.quebec

PoliSci-CompSci Collaboration

- Polarization: over time, in response to events, geographically, etc.
  - 2020 US election
  - COVID-19
  - Future: Canada, France...

- PIs (alphabetical): André Blais, Jean-François Godbout, Reihaneh Rabbany
- Other students: Aarash Feizi, Anne Imouza, Gabrielle Desrosiers-Brisebois, Jiewen Liu, Sacha Lévy, Zachary Yang

- Website: https://politicalpolarization.github.io/

The Surprising Performance of Simple Baselines for Misinformation Detection

- Kellin Pelrine, Jacob Danovitch (equal contribution), Reihaneh Rabbany

- Standard text-only models can be competitive with state-of-the-art models.
- Implications for evaluation. Also some flaws in construction of datasets.

- Paper

Overview

- Data
  - Getting/setting up social network data
  - Connecting other types of data
- Modelling
  - Don’t skip the simple approaches
  - Examples
- Understanding
  - Synthetic experiments
  - Annotating
  - Aggregate measures, and going beyond them too
- Conclusion

Getting data

3 main strategies:

- Existing dataset
- API
- Archive

Not the only options…

- Inside connections
- Direct scraping

But most projects at least one of first 3 sufficient, or at least to get started

Twitter Data

- Over 500 million tweets per day[1]

- How can we get data from Twitter to our tools?
- How can we get the data we're interested in?

[1] https://blog.twitter.com/official/en_us/a/2014/the-2014-yearontwitter.html

API prerequisite: developer/academic approval

- https://developer.twitter.com/en/apply-for-access

- https://developer.twitter.com/en/products/twitter-api/academic-research

Then, not too hard to get started

```
app_key = "..."
app_secret = "..."
oauth_token = "..."
oauth_token_secret = "..."

from twython import Twython

api = Twython(app_key=app_key, app_secret=app_secret, oauth_token=oauth_token,
                             oauth_token_secret=oauth_token_secret)
```

```
tweet_id = 1405966038982930437
tweet = api.lookup_status(id=tweet_id, include_entities=True, map=True, tweet_mode='extended')
```

How to filter it?

- Retrieve past tweets by user
  - Ex: all tweets from US representatives, senators, and presidential candidates
  - -> 150k tweets

- Retrieve real-time tweets by keyword
  - Ex: 1% of all tweets with keywords [JoeBiden, DonaldTrump, Biden, Trump, vote, election, 2020Elections, Elections2020, PresidentElectJoe, MAGA, BidenHarris2020, Election2020]
  - -> 350mil tweets

- Other: go back in time by keyword… followers and friends… retweets… etc.

Parler

- Archives of posts:
  - https://ddosecrets.com/wiki/Parler
  - https://web.archive.org/web/*/https://parler.com
  - -> 6.5mil posts

- A bit spottier than Twitter API, but enough to get some insights

Reddit

- Archive
  - https://files.pushshift.io/reddit/comments/
  - Huge, near comprehensive?
- API

Getting data: Summary

- Lots of data out there

- Not too hard to get started

Choice of graphs

- Even though social interactions seem naturally graph-like, the raw data doesn’t actually come as a graph

- Have to choose:
  - Unipartite vs. bipartite vs. bipartite projection
  - Node attributes (if any)
  - Undirected vs. directed edges (not to mention, weights? More complex features?)
  - Which relations(s) to even use

Unipartite vs. Bipartite

- User-entity: bipartite
  - Ex: user A says “#something”
- User-user: unipartite?
  - Ex: user A mentions user B
- But could also separate mentioning users from mentioned ones
  - Especially if there’s some group structure, e.g. mentions from random users to politicians
- Tip: don’t be like me and forget users aren’t the same as hashtags

Projecting

- To complicate further:
    - Direct connections: edge between A and B if A mentions B
    - Indirect connections: edge between A and B if both mention C

- But it’s worth considering
    - Makes all graphs unipartite again (no more pesky node types to mix up)
    - Can test for benefits empirically
    - Downside: can make graph much much larger or more dense

- Computation: A^TA (note: networkx does the projection super slow)

That aside, which relation to use?

- Big differences in behavior

- Ex: how much conservative and liberal users connect within and between their groups (projected version, average).

Relation	# Cons. Users	# Lib. Users	Cons. Intra-Degree	Cons. Inter-Degree	Lib. Intra-Degree	Lib. Inter-Degree
Retweet	2335	1653	582	29 (4.7%)	294	41 (12.2%)
Mention	3108	2307	1350	760 (36.0%)	863	1024 (54.3%)
Hashtag	1309	998	309	149 (32.5%)	227	195 (46.2%)
Quote	1788	1203	284	81 (22.2%)	164	121 (42.5%)

- Far less inter-group retweet connections than other relations

Connecting to other data

- Geolocation: either straight from platform/API, or extracted from profile
  - E.g. https://developers.arcgis.com/python/ or https://www.openstreetmap.org/

- Not trivial because e.g. multiple places with same name
- But viable nonetheless

	Total	Correlation
US users with state	757,601	0.9748
Users affiliated with a party	161,719	0.9726

Matching users to real-world people

- The (relatively) easy ones: politicians, celebrities, etc.
  - Often have well-known accounts
  - May even find convenient lists online
- The really hard ones: normal users
  - Full name + location (county level) can give a reasonable match
  - Ex: geolocation from before, name from profile, matching to voter registration info
  - Surveys can ask for a twitter ID. Data from existing surveys is typically private, but in some cases may be able to run one's own.

Other social interaction data

- Of course, there's tons of data on interactions besides social media
- Varying difficulty of obtaining it

- More exotic example for inspiration: sociometric badge

Credit: MIT Media Lab

Sociometric badges

- Collects measurements on real-world social interactions

- Much harder than an API call, but sensors like this can collect otherwise inaccessible data
- And give surprising insights
  - Ex: reorganize coffee breaks so people can interact more -> $15mil/year productivity increase

- Pentland, Social Physics, 2014

Data: Summary

- It’s worth putting careful thought and analysis into data collection/setup choices
- As we saw, there’s big variety in the data available and what it reflects
- Creativity also can be as worthwhile here, not just in the modelling stage

Modelling

- So, we have data, what can we do with it?

- Not going to go too in depth here
  - Depends a lot on the application
  - You’re probably aware of many tools already, and many more out there

- So just one general suggestion, and a few examples

Don’t visit the model store when you’re starving

- It’s like going to the grocery store - if I’m starving and everything looks good, who knows how much I’ll spend on unnecessary stuff

- Try some simple baselines first
  - Helps understand the data and get stuff running
  - Useful for evaluating and reporting results later anyways

- If you prepare them well, might find they’re pretty good

Misinfo baselines

- Tried to make complex multimodal model for misinfo detection

- Struggled to beat our baseline… but turns out other models did too

	PHEME9 T/F	PHEME5 R/NR	PHEME5 3-way	PHEME9 4-way	PHEME5 Lc	Average Rank
SOTA	82.5 [96]	87.6 [16, 92]	66.7 [16]	75.3 [97]	**51.3** [16]	5.6
CT-BERT	92.0 ± 0.9	89.0 ± 0.8	84.6 ± 1.5	79.0 ± 2.6	27.9	3.4
Funnel	86.7 ± 3.2	87.3 ± 0.6	79.4 ± 3.7	71.4 ± 3.3	28.7	6.4
RoBERTa	**93.2 ± 0.9**	**89.4 ± 0.3**	**87.7 ± 1.9**	**82.5 ± 3.3**	29.0	**2.0**
BERT	89.9 ± 1.1	87.2 ± 0.4	81.2 ± 1.4	76.8 ± 2.7	24.2	5.8
BERTweet	89.8 ± 0.6	87.3 ± 0.6	81.8 ± 0.9	76.6 ± 4.1	29.0	5.0
DeCLUTR	90.2 ± 0.8	88.3 ± 0.4	83.7 ± 2.1	77.8 ± 3.5	30.2	3.2
ELMo	81.7 ± 2.4	84.2 ± 0.8	65.8 ± 1.8	64.3 ± 4.0	30.3	9.4
ALBERT	85.3 ± 2.9	84.2 ± 2.7	71.1 ± 2.2	65.7 ± 3.1	29.4	7.2
BERT-Tiny	81.6 ± 2.0	84.7 ± 0.8	67.3 ± 2.0	61.0 ± 2.5	36.5	7.6

More results

- When test data content is sufficiently similar to train, transformer-based language models compete with or even beat state of the art models.

	PolitiFact	GossipCop	Twitter15	Twitter16	Twitter15 T/F	Twitter16 T/F	WNUT-2020	Average Rank
SOTA	**92.8** [77]	85.0 [28]	**91.0** [30]	**92.4** [30]	82.5 [50]	75.9 [50]	**91.0** [43, 58]	4.4
CT-BERT	86.0 ± 3.2	90.6 ± 0.2	83.5 ± 2.8	83.9 ± 0.9	93.8 ± 1.6	94.0 ± 3.5	90.6	2.8
Funnel	86.4 ± 3.2	– [a]	66.9 ± 3.0	69.6 ± 2.9	83.2 ± 3.8	90.8 ± 2.2	88.5	–
RoBERTa	86.7 ± 1.2	**92.8 ± 0.5**	81.8 ± 1.5	84.8 ± 1.9	**94.4 ± 0.8**	**95.7 ± 2.8**	90.5	**2.3**
BERT	81.8 ± 3.0	89.8 ± 0.4	77.5 ± 3.3	78.2 ± 4.1	89.7 ± 1.6	91.6 ± 4.5	88.5	5.3
BERTweet	88.5 ± 1.2	92.6 ± 0.6	76.7 ± 2.9	77.7 ± 2.7	86.7 ± 1.8	92.0 ± 3.7	88.8	4.4
DeCLUTR	36.6 ± 1.4 [b]	43.3 ± 0.4 [b]	80.4 ± 2.6	80.5 ± 1.7	91.7 ± 1.5	94.5 ± 2.5	89.1	5.1
ELMo	83.1 ± 1.6	92.0 ± 0.5	53.7 ± 2.7	55.5 ± 4.9	74.4 ± 3.5	83.3 ± 5.0	82.4	8.0
ALBERT	80.1 ± 2.9	88.2 ± 0.9	63.4 ± 4.0	68.0 ± 3.5	83.3 ± 1.8	88.9 ± 4.3	86.8	7.7
BERT-tiny	85.3 ± 2.8	86.5 ± 0.6	54.6 ± 3.4	48.8 ± 3.9	77.8 ± 4.8	77.8 ± 4.9	79.9	8.9

Simple models can work well

- Those models had text-only baselines… but not more recent, strongest ones
- Another example: MLP + Label Propagation > GNN (Huang et al. 2021)
- Optimizing something relatively simple might work pretty well, even compared to more complicated approaches
    - Especially if your goal/contribution is more solving an application than a new model itself
    - The ideal model may be super complicated… but harder to find it

- Combining 5 modalities to measure polarization

- Core: builds on basic GCN, basic LM

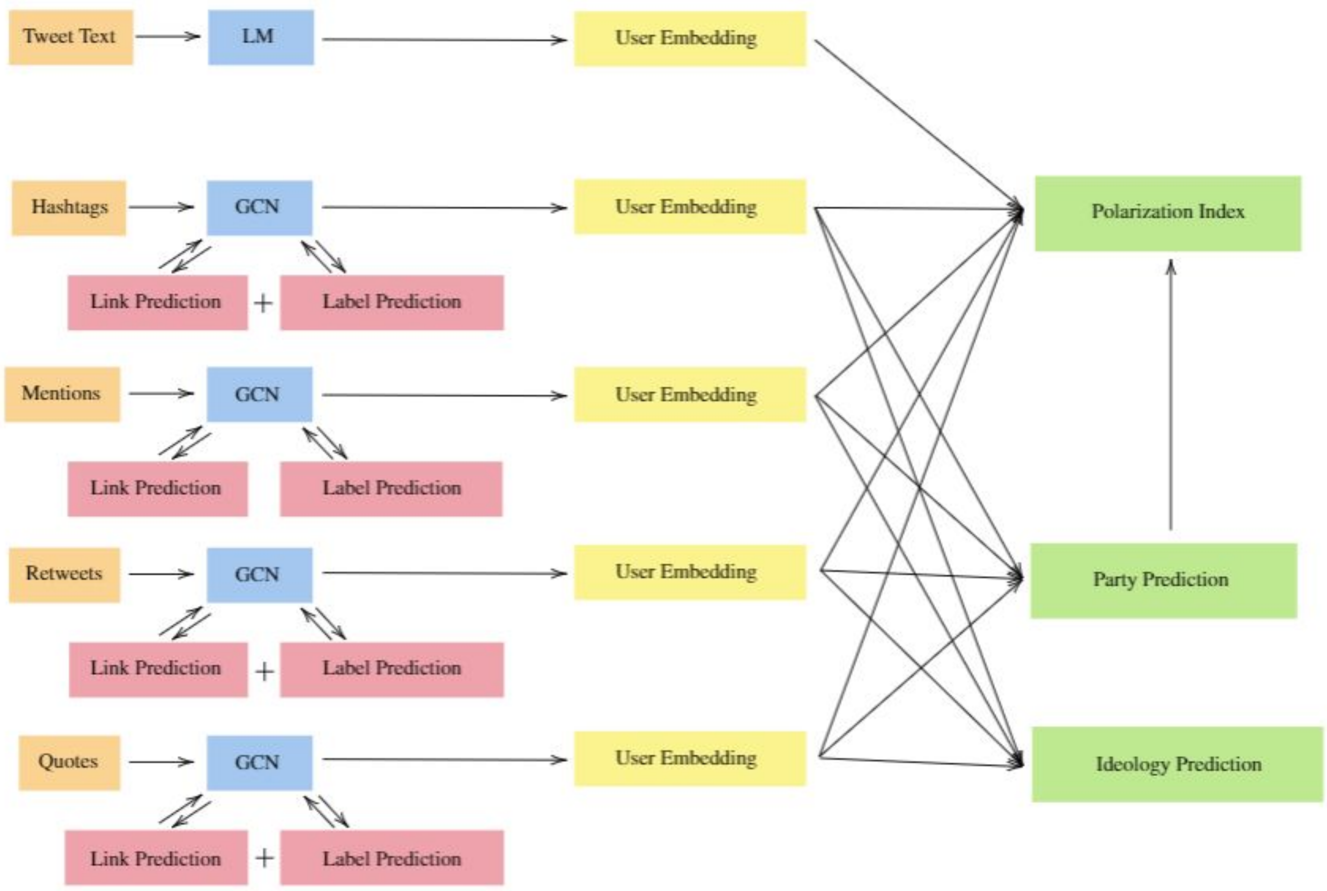

Modeling summary

- Not a blanket statement…
    - Sometimes a really complicated model can be necessary or perform much better
    - Or your inspiration might be to try some modeling idea itself, e.g. GAN invention
- But often a reasonable starting point
    - Especially with social graph data, where there's often tons of options and it's hard to know right away which approach will be best
- And sometimes may find surprises

Understanding

- We’ve got data, we’ve got at least some model… How do we know if it’s working correctly? And if it is, interpret what it’s saying?

- I’ll discuss 3 subtopics here (briefly):
    - Synthetic experiments
    - Data annotation/labeling
    - Going beyond simple aggregate metrics

Synthetic Experiments

- Real data is complicated and it's hard to isolate different factors

- So, can use much simpler synthetic data to better understand how a model/pipeline will behave in various situations

- Construction depends on application and goal
    - Would like some similarity to real world data, without complicating the synthetic model so much it becomes similarly uninterpretable
    - Social graphs often have a community structure, so SBM may be a reasonable starting point

Example

- We found our polarization measure varied drastically depending on community size ratio (left). Fixed with upsampling to balance sizes (right).

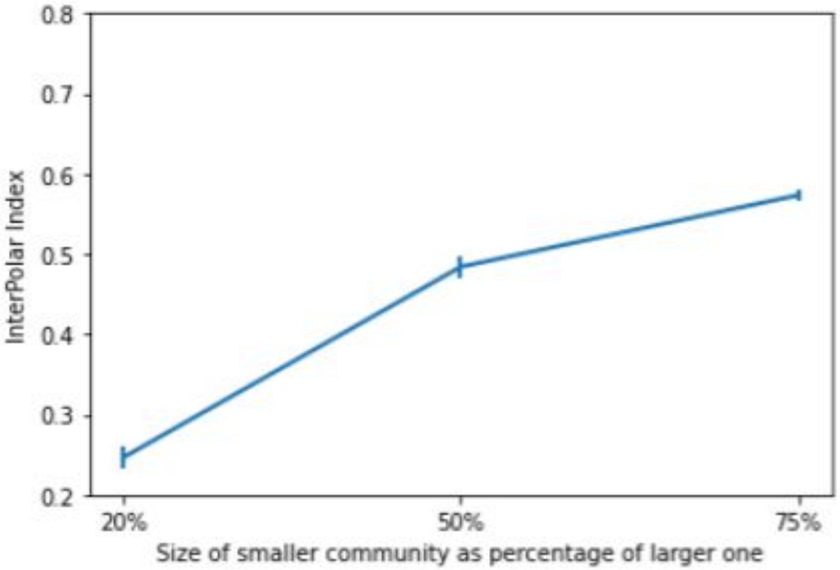

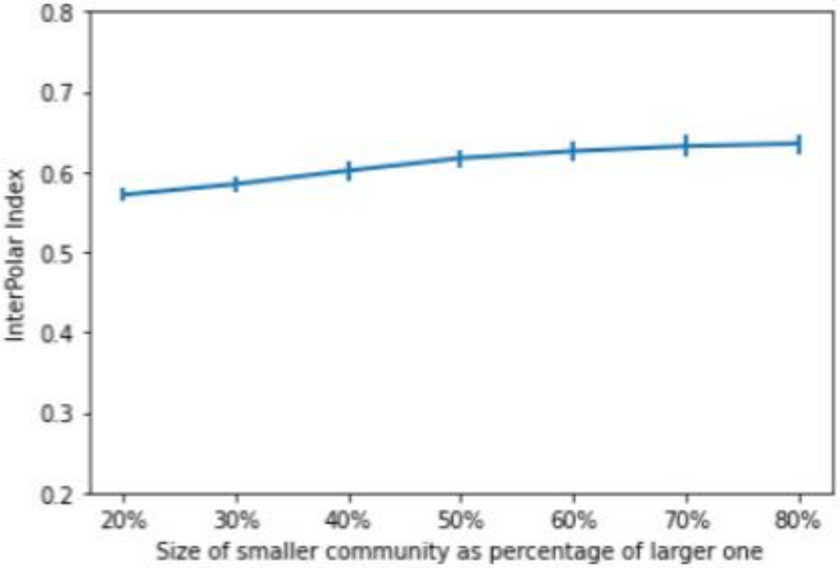

Annotation

- For evaluation, as well as training, often want labeled data
- Existing data may come with nice labels, but for new data, what to do?
- Tradeoff between quality and quantity

Annotation

- Expert labeling
  - Most accurate, most challenging to get at scale

- Crowd labeling
  - Human, but (typically) not expert
  - Usually not free, but much less expensive and much more scalable than expert labels
  - Quality may vary; accuracy measurements advised

Annotation

- Keywords
  - Can use profiles or post content
  - Very scalable
  - Fits naturally with social media data: can be used in collection pipeline, correspondence with hashtags, etc.
  - Accuracy needs verification

Analyzing COVID-19 Discussions

Topic	Sentiment	Keywords
Lockdown	Neutral	quarantine, secondlockdown, lockdownDC, californialockdown, 2ndLockdown, Lockdown3, lockdowns, coronavirusshutdown, covidlockdown, covidshutdown, shutdowns, coronaviruslockdown, lock down, Wuhanlockdown, lockdownextension, home-quarantine, lockdownTrump
	Positive	Stayhome, StayHomeStaySafe, lockdownlife,StayHomeSaveLives,nationallockdown, TogetherAtHome, Prolockdown, Proshut-down, LockdownWorks, AvoidGatherings, stay home challenge, safe at home, stay at home, stay home, sheltering in place, quarantine life, 14DayQuarantine, inmyquarantinesurvivalkit, quarantine shelter, shelteringinplace, stay home challenge, stayathome, stay_home_safe, stayhometosavelives, workfromhome, stayhome!, saferathome, safe at home
	Negative	endthelockdown, endlockdowns, NoShutdown, NoMoreLockdown, NoMoreShutdown, ReopenAmerica, OpenAmericaNow , Antilockdown, LockdownsKill, Breakthelockdown, LockdownsAreNotACure, nolockdown2, BreakTheLockdowns, Lockdown-Chaos, LockdownsDontWork, StopThelockdowns, LockdownFraud, Bidenlockdown, NoMoreLockdowns, CancelTheLock-down, freetheUSA2020, NoLockdowns, NoLockdown, Antishutdown, Anti shutdown, endtheshutdown
Mask	Neutral	mask, masks, facemask, facemasks, ppe, n95, kn95, CoronavirusMask, surgicalmasks, clothmasks, n95facemask, kn95facemask, facecover, cloth mask, ffp2mask, ffp3mask, ffp3, ffp1, ffp2, kn95 mask, n95 mask, surgical mask, faceshield
	Positive	masksSaveLives, wearamask, maskup, wearadamnmask, MaskYourKids, MaskMandates, WearMaskProtectLife, Wear a mask protect a life, WearAMaskSaveALife, maskon, Doublemasking, Doublemask, MaskOnAmerica, MaskSelfie, MasksWork, MaskWorks, mandatorymask, GetMePPE, masks4all, wear face mask, CoronavirusCoverup
	Negative	maskdontWork, Nomasks, Nomask, MasksOff, MaskOff, antimasker, antimaskers, NoMaskMandate, Nomoremasks, Un-MaskAmerica, Maskless, IWillNotWearAMask, SheepNoMore, unmask, MaskOffAmerica, Talesoftheunmaskedpatriot, take-offthemask, maskburning, Burnyourmask, Burnyourmaskchallenge, nomaskonme, nomaskselfie, maskhoax, nomaskEVER, facefreedom, masksmakemesweaty, MasksAreDangerous, TakeMaskOff, Stopforcingmaskonme, takeoffyourmask, refusemask, NeverMasker, StopWearingMask, StopWearingtheDamnMasks, MasksdontMatter, stopmasking, stopthestupidmask, mask-ingchildrenischildabuse, MomsAgainstMasks, MasksRUnhealthy, SheepWearMasks, MasksAreForSheep, RefuseToWearMasks, MasksAreMurderingMe, maskshoax, TakeOffTheMask

Analyzing COVID-19 Discussions

- Found 88%+ accuracy on topic

- Wildly varying accuracy on stance (as low as 6%!)

Annotation

- Simple classifier can pick up differences between e.g. “I love x” and “I hate x”
  - Can improve performance a lot compared to keywords, without sacrificing scalability
  - Usually needs train (and test) data labeled some other way

- Example: improved on profile keyword label accuracy ~70% -> ~90% using 2000 hand-labeled examples (left).

Dataset	Counts		Accuracy	
	Cons.	Lib.	Cons.	Lib.
Politicians	1,174	1,068	97.7%	96.8%
Election	183,207	176,271	87.0%	90.5%
Parler	31,966	808	93.1%	82.9%

Aggregate measures, and beyond

- Once you have some labels (or even without them, in unsupervised setting), can compute standard metrics
- Amour et al. (2020), among others, highlights need to evaluate models beyond simple top-level metrics
    - Interesting paper, 40 authors

- Particularly important with this data where labeling is often challenging and human behavior is varied and evolving
- At the same time, extensive data means lots of opportunities for better understanding

Measuring polarization

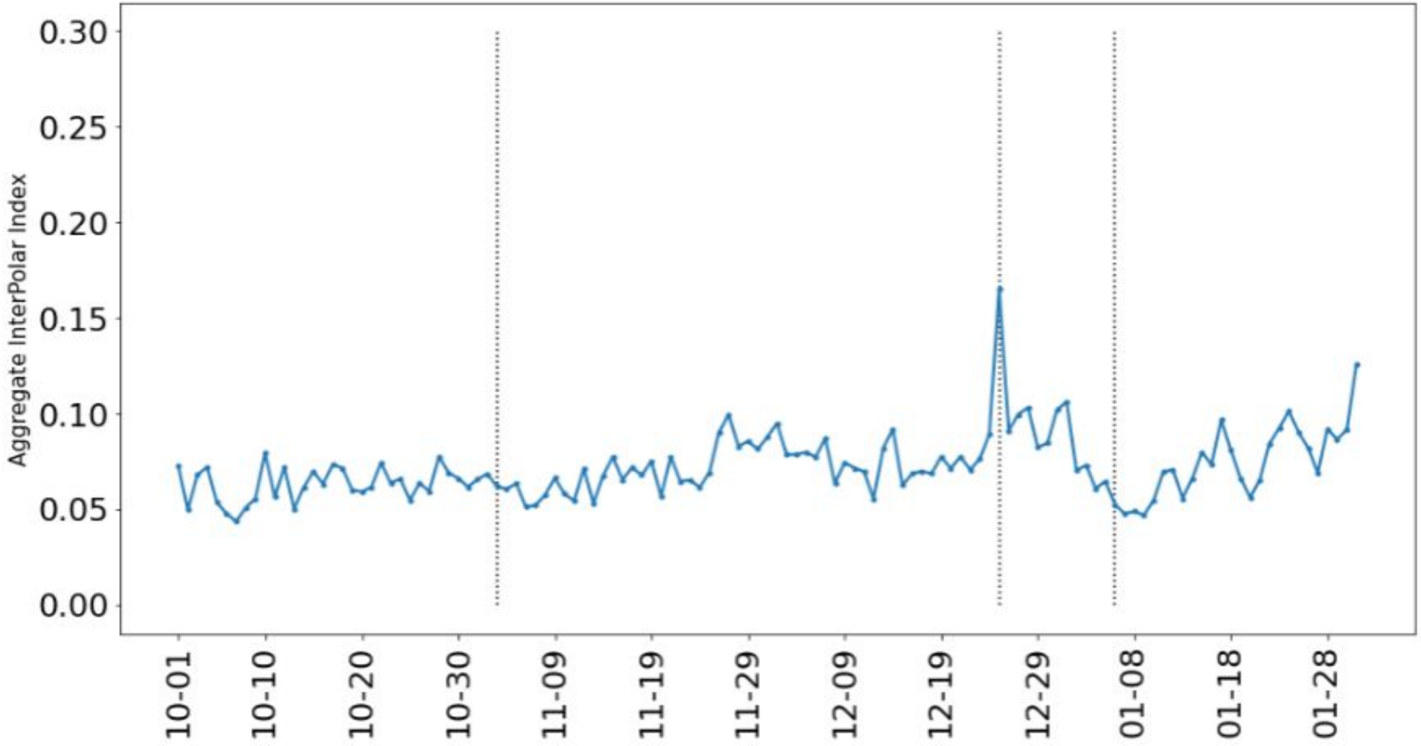

Measuring polarization

- Why is polarization high on December 25th?

- Mostly driven by quote interactions

Dec. 25				Jan. 6			
Liberal		Conservative		Liberal		Conservative	
realDonaldTrump	102 (6.4%)	realDonaldTrump	406 (13.8%)	LindseyGrahamSC	68 (2.7%)	ElijahSchaffer	22 (2.0%)
JoeBiden	38 (2.4%)	TimRunsHisMouth	171 (5.8%)	ElijahSchaffer	38 (1.3%)	TheLeoTerrell	21 (1.9%)
donwinslow	33 (2.1%)	marklevinshow	85 (2.9%)	atrupar	35 (1.2%)	JoeBiden	16 (1.4%)
michaelluo	29 (1.8%)	NewDayForNJ	79 (2.7%)	Phil_Lewis_	32 (1.1%)	JackPosobiec	12 (1.1%)
kylegriffin1	25 (1.6%)	JoeBiden	55 (1.9%)	igorbobic	32 (1.1%)	TomiLahren	11 (0.9%)

Understanding misinfo data

- Showed results on Twitter15/16 in previous slide
- But what are we really classifying?

Figure 1: False vs. Other Classes

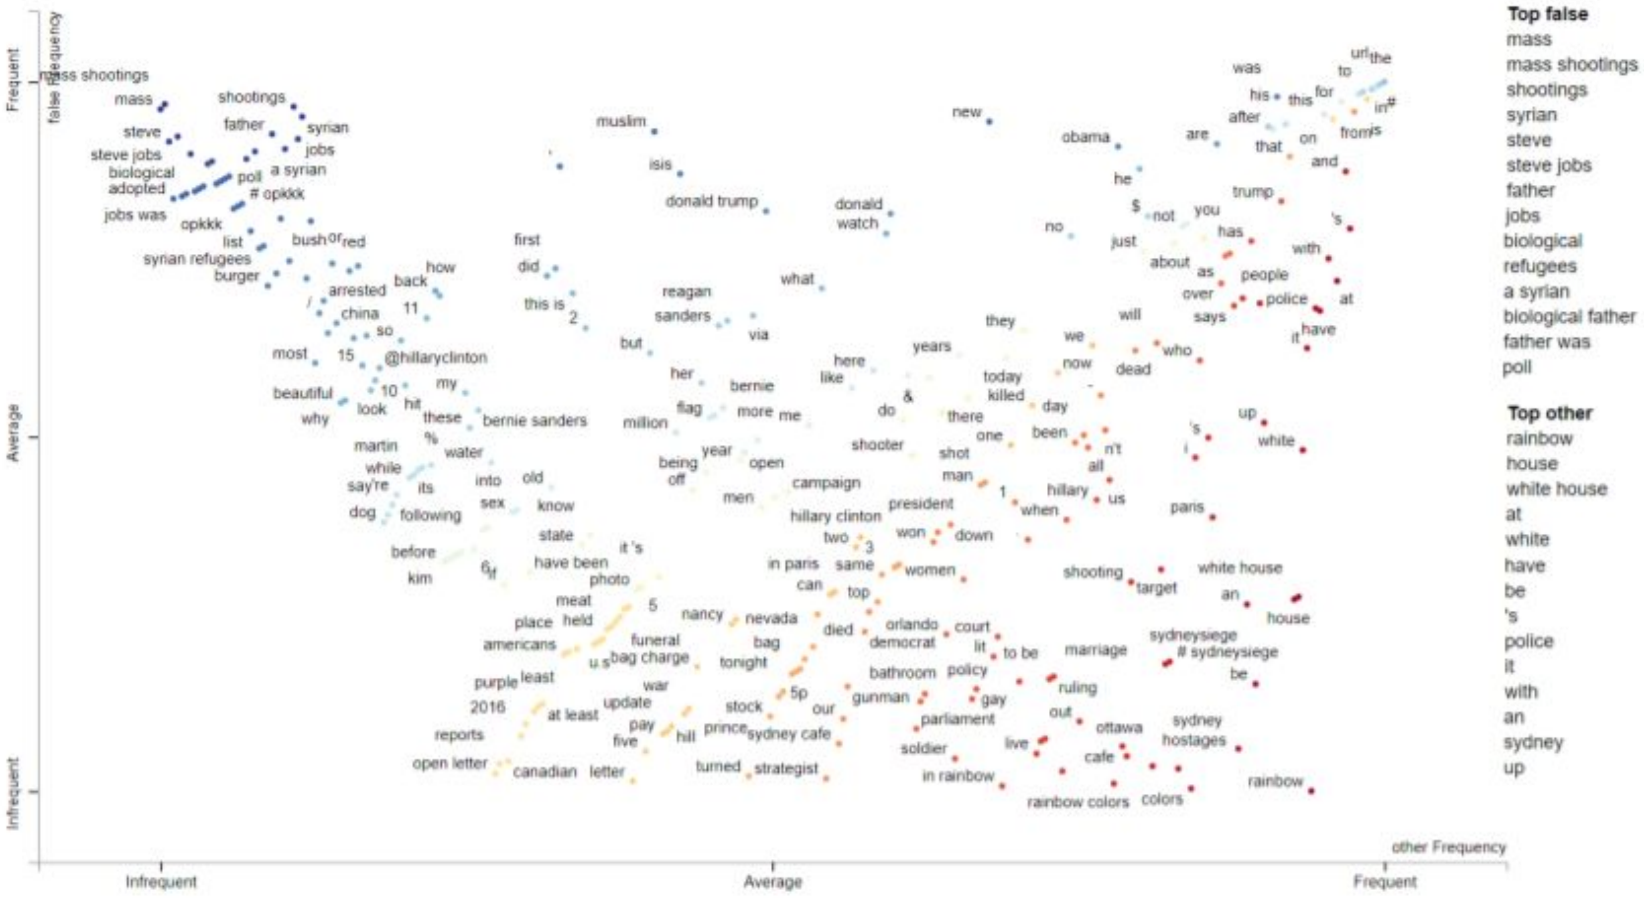

What are we really classifying?

- In this data different classes were collected at different dates
- Leading to good performance if you know something that gives the date (e.g. tweet ID, above) or unrealistically informative keywords (e.g. no true/false tweets from Clinton/Trump, below)

Table 4: Evaluating tweet ID classification (Macro F1 score)

Twitter15	False	True	Unverified	Non-rumor	Macro Avg.
SOTA [30]	92.9	90.5	85.4	95.3	91.0
2-digit RF	62.4	65.6	61.1	99.4	72.1
3-digit RF	73.0	69.5	79.7	98.2	80.1
Twitter16	False	True	Unverified	Non-rumor	Macro Avg.
SOTA [30]	91.3	94.7	89.9	93.5	92.4
2-digit RF	83.5	87.6	82.1	90.7	86.0
3-digit RF	90.7	95.3	84.4	92.9	90.8

Table 5: Label counts of tweets containing "Clinton" and "Trump"

Twitter15	Clinton	Trump	Twitter16	Clinton	Trump
True	0	0	True	0	0
False	0	0	False	17	18
Unverified	22	30	Unverified	17	39
Non-rumor	6	14	Non-rumor	8	6

Summary

- Complicated data -> need careful evaluation

- Synthetic experiments can help understand what model alone is doing
- A combination of aggregate and zoomed in analysis can help understand what the data says

Conclusion

- We discussed some ideas and approaches for three key parts of working with social graphs: data collection/construction, modelling, and evaluation/understanding

- Examples here focused on polarization and misinfo, but there’s tons of ways data like this can be used

- Technology and data already shape our interactions, and will play an even bigger role in the future

Questions?

Thank you!