Fake Cures: User-centric Modeling of Health Misinformation in Social Media

22 Oct 2018

The 21st ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)

November 3rd-7th, 2018, New York City

Amira Ghenai (Waterloo University), Yelena Mejova (ISI Foundation - Turin, Italy)

UNIVERSITY OF WATERLOO

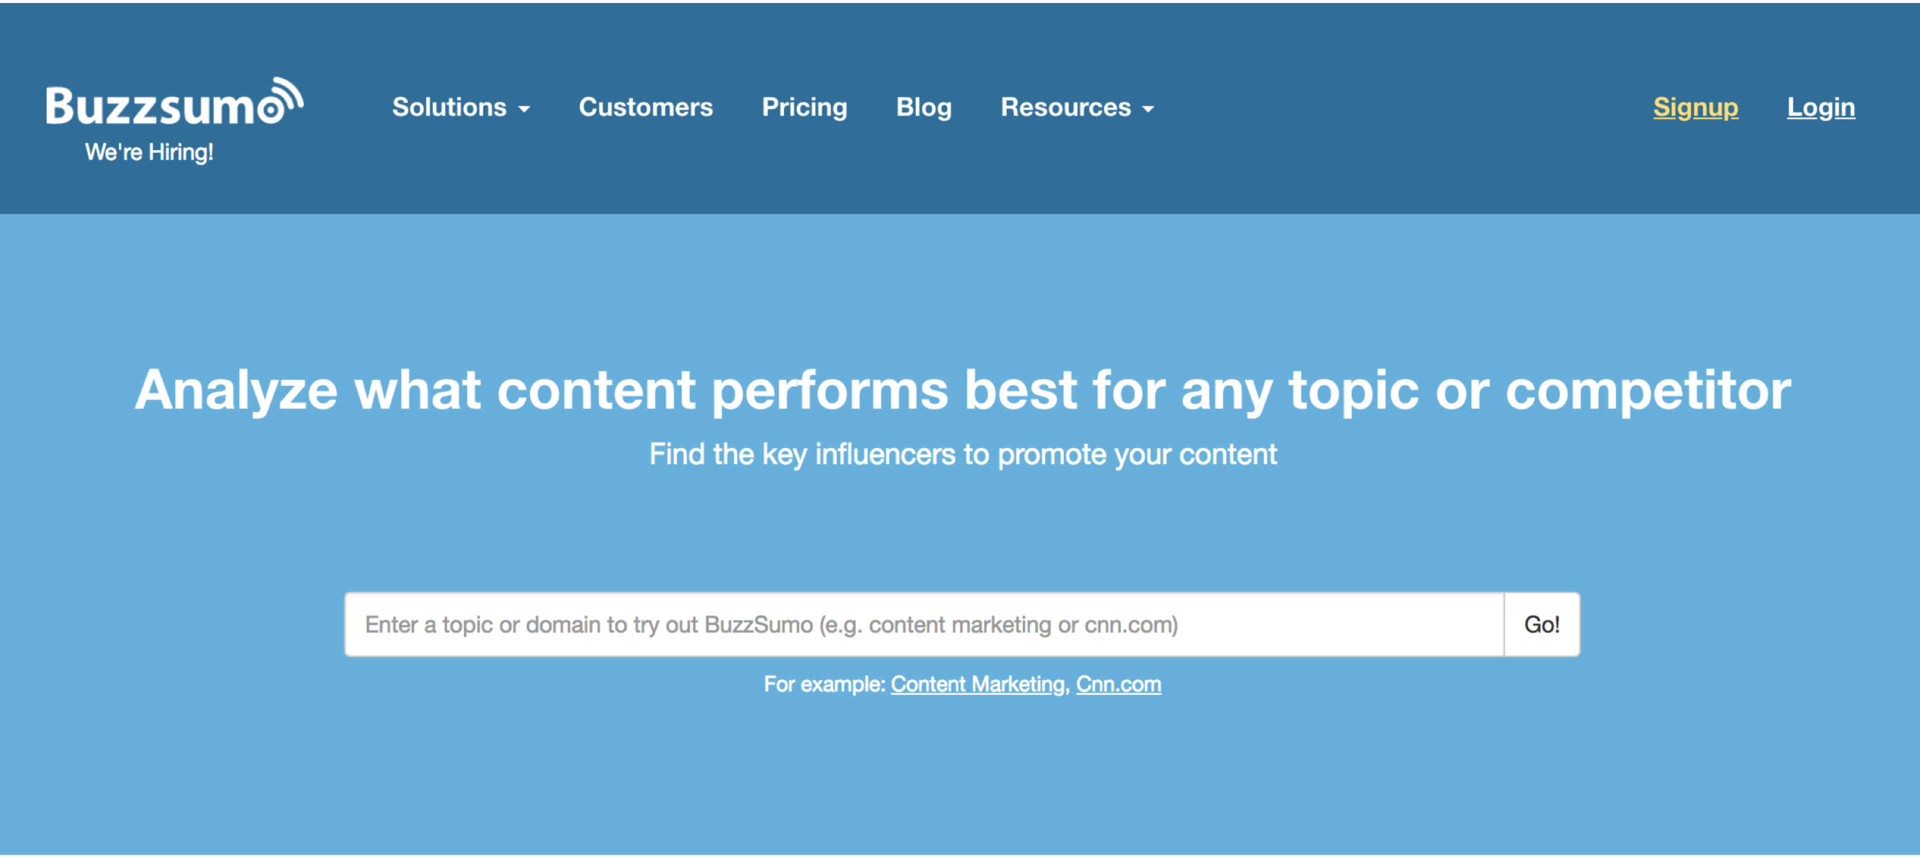
Buzzsumo
We're Hiring!
Solutions
Customers
Pricing
Blog
Resources
Signup
Login
Analyze what content performs best for any topic or competitor
Find the key influencers to promote your content
Enter a topic or domain to try out BuzzSumo (e.g. content marketing or cnn.com)
Go!
For example: Content Marketing, Cnn.com

Topic: “cancer cure”

cancer cure | Search | Save Search | Export | Page 1 of 449

How to run an Advanced Search

Sort by: Total Engagements

Article	Facebook Engagements	Twitter Shares	Pinterest Shares	Reddit Engagements	Number of Links	Evergreen Score	Total Engagements
Hemp Could Free Us From Oil, Prevent Deforestation, **Cure Cancer**, and It's Environmentally Friendly By Truthcommand – Mar 31, 2018 truthcommand.com	294.1K	783	28	1	5	13	294.9K
A YouTuber who claimed being vegan **cured** her **cancer** has died from **cancer** By Harry Shukman – Feb 14, 2018 babe.net	94.6K	3K	2	1.7K	49	7	99.2K
Cancer vaccine that **cured** 97% of tumours in mice will be tested on humans By Alexandra Thompson Health... – Mar 28, 2018 dailymail.co.uk	82.8K	481	0	5.5K	22	12	88.8K
A New Israeli **Cancer** Vaccine May **Cure** 90% of All **Cancer** Types in Just One Shot By Jews News – Nov 25, 2017 jewsnews.co.il	88.1K	24	0	0	1	10	88.2K
Woman '**cured**' her terminal **cancer** with cannabis oil By Alexandra Thompson Health... – Mar 12, 2018 dailymail.co.uk	87.6K	232	0	73	22	15	88K
Ex-Clinton staffer: If Obama **cured cancer**, Trump would try to bring it back By Avery Anapol – Dec 5, 2017 thehill.com	24.2K	7.7K	2	43K	6	6	74.9K
Vegan YouTuber Who Said Veganism **Cures Cancer** Has Died Of **Cancer** By Neelam Tailor – Feb 20, 2018 unilad.co.uk	59.3K	43	0	0	8	4	59.3K
DANDELION **CURES CANCER**, HEPATITIS, LIVER, KIDNEYS, STOMACH ... HERE'S HOW TO USE IT! By Admin – Mar 9, 2018 dailynativenews.site How Article	54.1K	104	143	0	0	4	54.3K

Topic: “cancer cure”

They are all unproven treatments

cancer cure

How to run an Advanced Search

Sort by Total Engagements

Export

Page 1 of 449

Article		Facebook Engagements	Twitter Shares	Pinterest Shares	Reddit Engagements	Number of Links	Evergreen Score	Total Engagements
Hemp Could Free Us From Oil, Prevent Deforestation, **Cure Cancer**, and It's Environmentally Friendly By Truthcommand — Mar 31, 2018 truthcommand.com	Save View Backlinks View Sharers Share	294.1K	783	28	1	5	13	294.9K
A YouTuber who claimed being vegan **cured** her **cancer** has died from **cancer** By Harry Shukman — Feb 14, 2018 babe.net	Save View Backlinks View Sharers Share	94.6K	3K	2	1.7K	49	7	99.2K
Cancer vaccine that **cured** 97% of tumours in mice will be tested on humans By Alexandra Thompson Health... — Mar 28, 2018 dailymail.co.uk	Save View Backlinks View Sharers Share	82.8K	481	0	5.5K	22	12	88.8K
A New Israeli **Cancer** Vaccine May **Cure** 90% of All **Cancer** Types in Just One Shot By Jews News — Nov 25, 2017 jewsnews.co.il	Save View Backlinks View Sharers Share	88.1K	24	0	0	1	10	88.2K
Woman '**cured**' her terminal **cancer** with cannabis oil By Alexandra Thompson Health... — Mar 12, 2018 dailymail.co.uk	Save View Backlinks View Sharers Share	87.6K	232	0	73	22	15	88K
Ex-Clinton staffer: If Obama **cured cancer**, Trump would try to bring it back By Avery Anapol — Dec 5, 2017 thehill.com	Save View Backlinks View Sharers Share	24.2K	7.7K	2	43K	6	6	74.9K
Vegan YouTuber Who Said Veganism **Cures Cancer** Has Died Of **Cancer** By Neelam Tailor — Feb 20, 2018 unilad.co.uk	Save View Backlinks View Sharers Share	59.3K	43	0	0	8	4	59.3K
DANDELION **CURES CANCER**, HEPATITIS, LIVER, KIDNEYS, STOMACH ... HERE'S HOW TO USE IT! By Admin — Mar 9, 2018 dailynativenews.site How Article	Save View Backlinks View Sharers Share	54.1K	104	143	0	0	4	54.3K

Topic: "cancer cure"

They are all unproven treatments

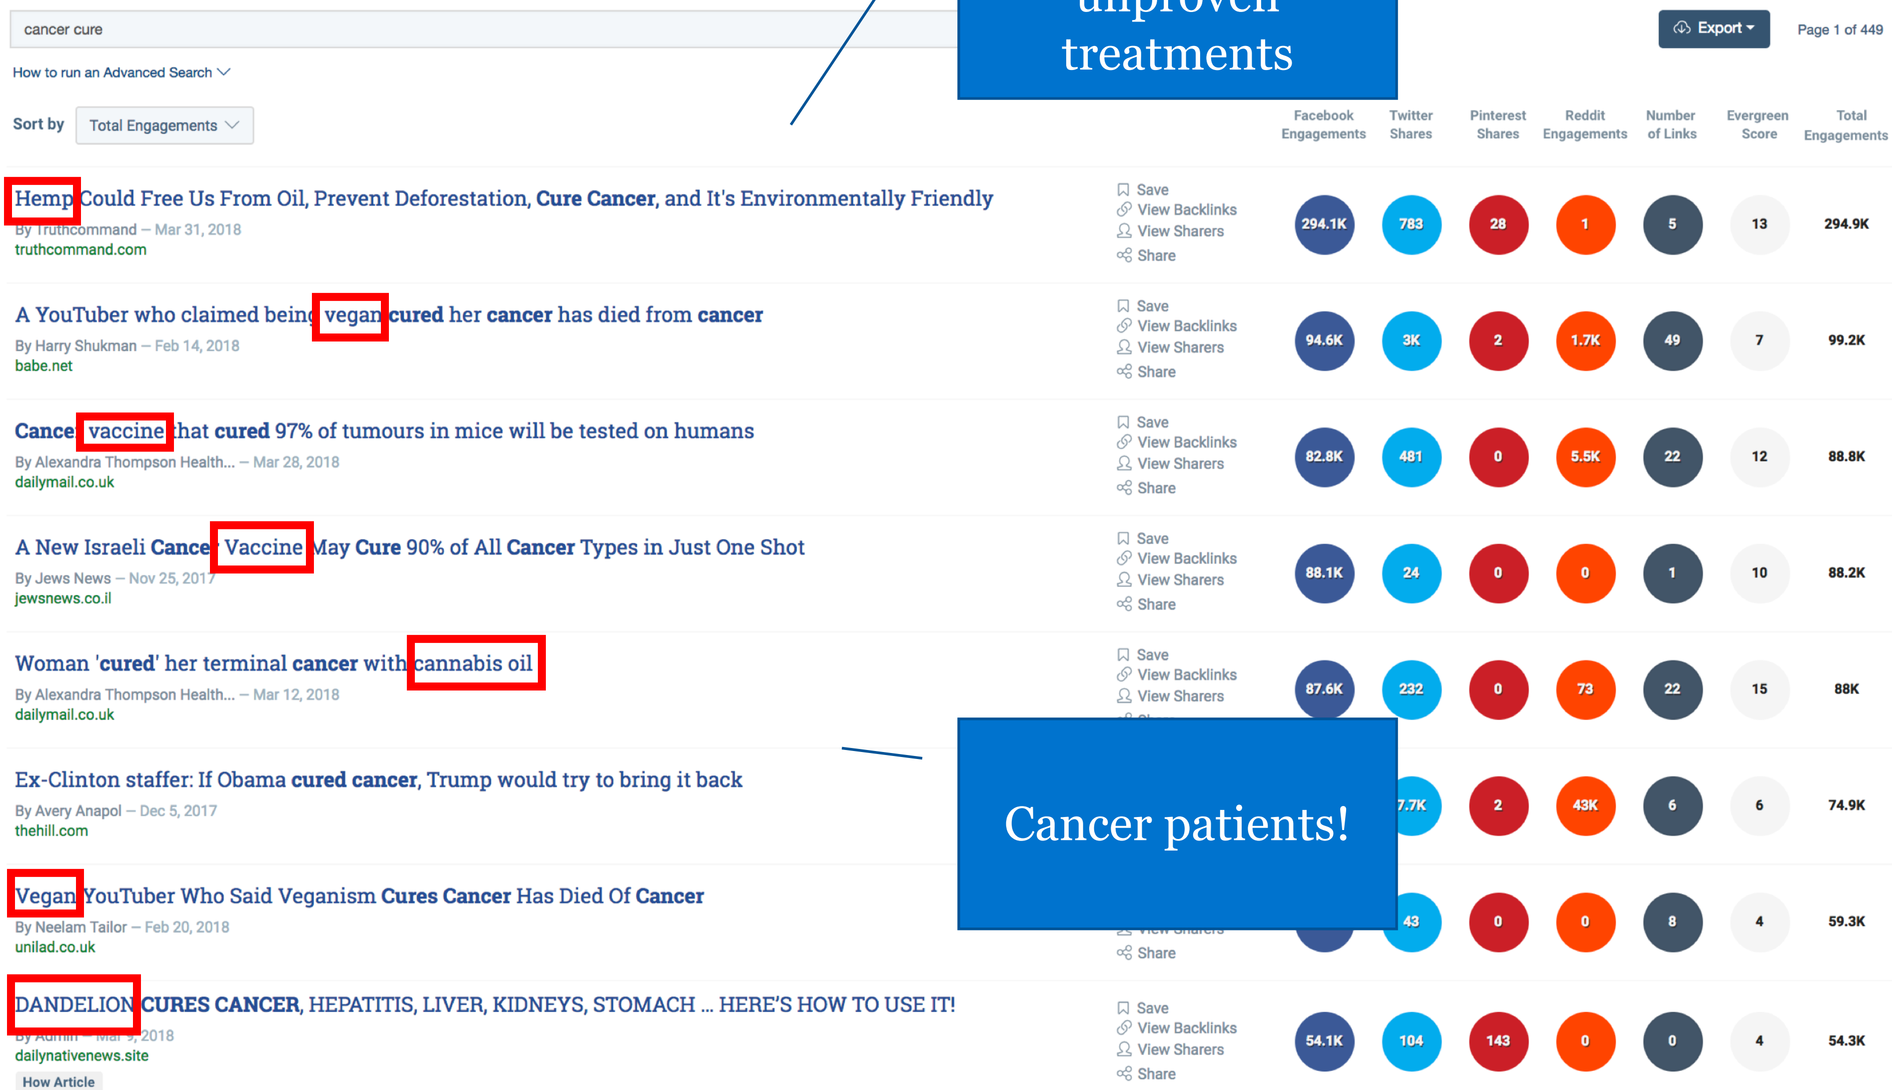

Cancer patients!

Problem Statement

- Social media use for health management is growing
  - 62% of internet users in U.S. use social networking sites for health related topics
- Accountability, quality and confidentiality issues
- Perfect medium for propagating possible medical misinformation
  - Serious threat to public health

Proposed Solution

"Fake cancer treatments" topic

- Method: user modeling
- Aim: determine characteristics of users propagating unverified "cures" of cancer on Twitter
- Benefits: allow public health officials to
  - Detect potential sources of misinformation
  - Monitor social media communications
  - Identify current limitations and improve them
  - Detect new misinformation before it causes harm

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

1. Control Group
   - General cancer topics

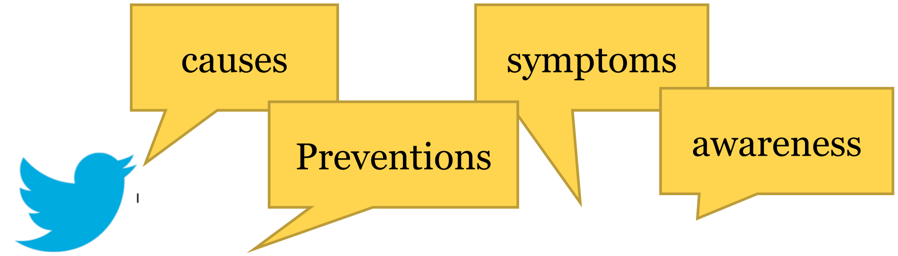

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

1. Control Group

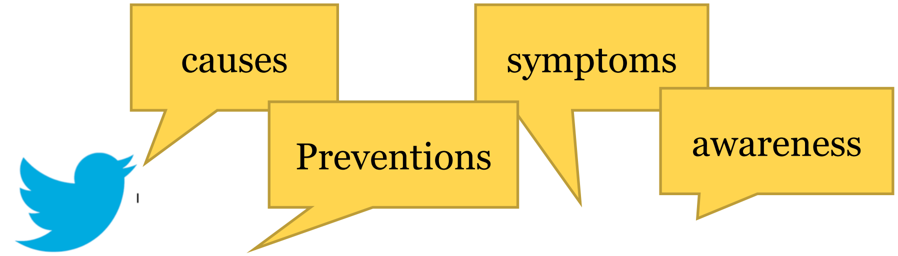

- General cancer topics
- We use Paul and Dredze [1] dataset
  - 144 million tweets related to health topics
  - Dataset time period between 01 August 2011 - 28 February 2013
  - Cancer topic has 676,236 users who posted 969,259 tweets

[1] Michael J Paul and Mark Dredze. 2014. Discovering health topics in social media using topic models. *PloS one* 9,8 (2014), e103408.

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

2. Rumor Group

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

2. Rumor Group

- 139 total unproven cancer treatments from 3 different sources
- Validated by trained oncologist

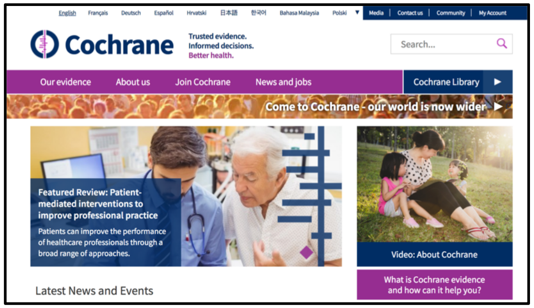

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

2. Rumor Group

- 139 total unproven cancer treatments from 3 different sources
- Validated by trained oncologist
- Collect tweets about treatments:
  - Same time period as control group
  - Hand craft query & query expansion
  - 39,675 users with 215,109 tweets

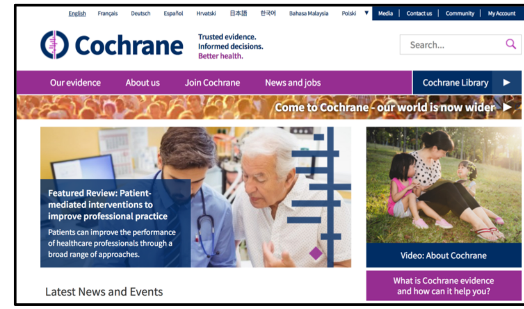

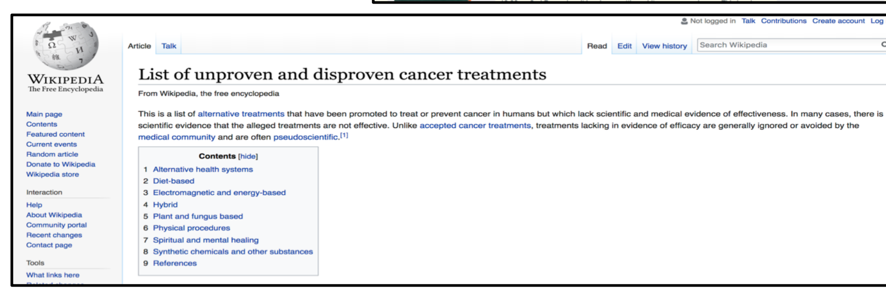

Data Collection

Rumor/Control data collection → User Selection → Relevance Refinement

Topic*	**Expanded Query**	**Example Tweet**
Soursop	(Soursop:OR:Graviola:OR:guyabano:OR:guanabana:OR:"Annona:muricata":OR:"Annona:crassiflora":OR:"Guanabanus:muricatus":OR:"Annona:bonplandiana":OR:"Annona:cearensis":OR:"Annona:muricata"):AND:cancer	"[...] *University show that the* ***soursop*** *fruit kills* ***cancer*** *cells effectively, particularly prostate cancer cells, pancreas and lung.*"
Ginger	ginger:AND:cancer	"*Can* ***ginger*** *help cure ovarian* ***cancer****? Since 2007, the University of [...] has been studying GINGER...* <url>"
Antineoplaston	(antineoplaston:OR:burzynski):AND:cancer	"*RT Dr.* ***Burzynski*** *He has the cure for* ***cancer****, the FDA want to shut him down <url>*"

* The topics (along with the keyword queries) are available at https://tinyurl.com/y78mkg6s

Data Collection

Rumor/Control data collection → **User Selection** → Relevance Refinement

Data Collection

Rumor/Control data collection → **User Selection** → Relevance Refinement

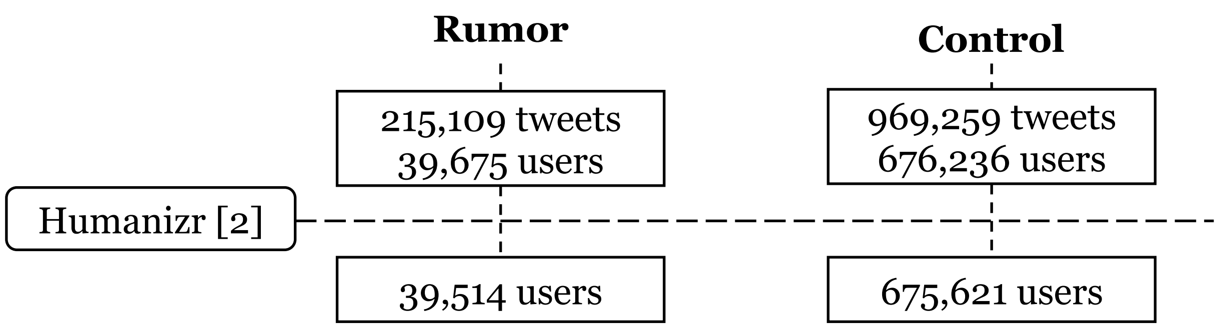

[2] James McCorriston, David Jurgens, and Derek Ruths. 2015. Organizations Are Users Too: Characterizing and Detecting the Presence of Organizations on Twitter. In ICWSM. 650–653.

Data Collection

Rumor/Control data collection → **User Selection** → Relevance Refinement

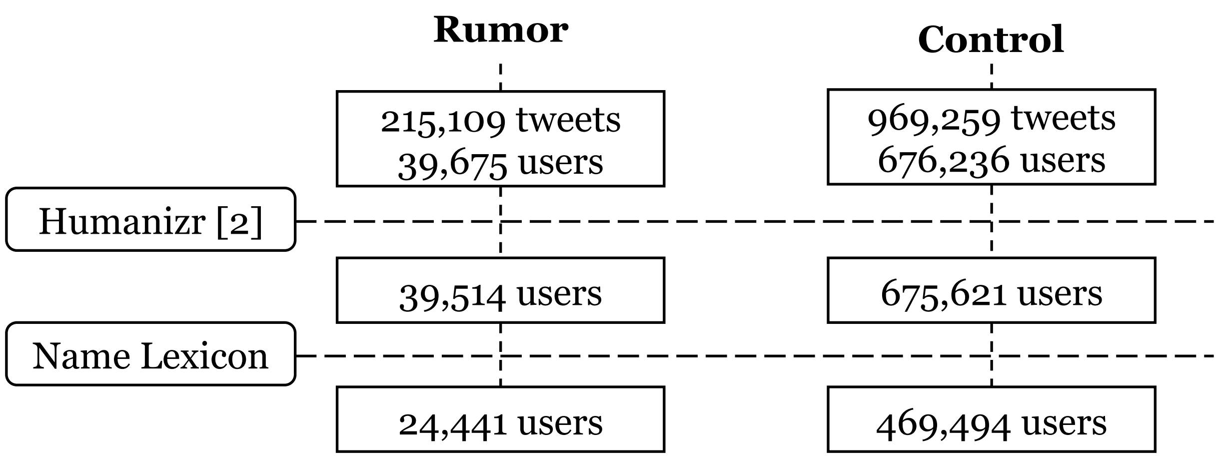

[2] James McCorriston, David Jurgens, and Derek Ruths. 2015. Organizations Are Users Too: Characterizing and Detecting the Presence of Organizations on Twitter. In ICWSM. 650–653.

Data Collection

Rumor/Control data collection → **User Selection** → Relevance Refinement

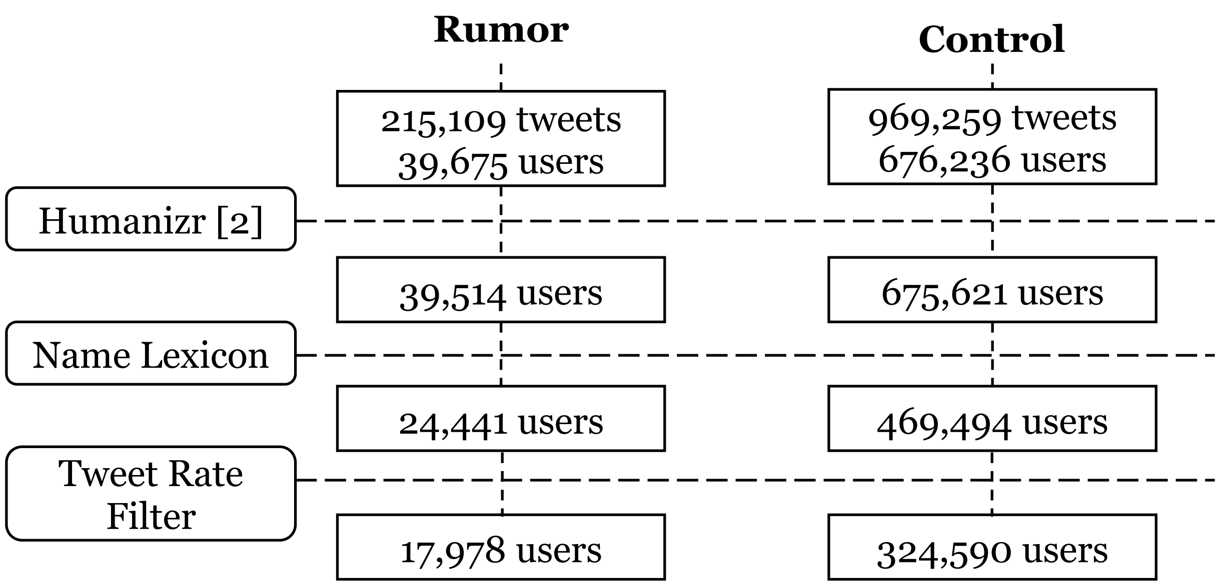

[2] James McCorriston, David Jurgens, and Derek Ruths. 2015. Organizations Are Users Too: Characterizing and Detecting the Presence of Organizations on Twitter. In ICWSM. 650–653.

Data Collection

Rumor/Control data collection → User Selection → **Relevance Refinement**

- We check whether every tweet is relevant to the topic of interest, we define users as follows:
  - **Rumor group** - users who claim a cure is <u>helpful</u> for treating cancer and **not** users who talk about other topics such as <u>prevention</u> or <u>debunking</u>
  - **Control group** - users who post at <u>least once about cancer</u> symptoms, awareness, prevention, cause or personal experience etc. but **not** about a <u>cancer cure</u>
  - To make our users follow these definitions, we use:
    - Crowdsourcing & Classification – machine learning

Data Collection

Rumor/Control data collection → User Selection → **Relevance Refinement**

1. Crowdsourcing

   a) Sample the tweets (4,000 tweets from rumor and control groups)

   b) Label the sampled tweets:

Rumor group - whether the tweet is about:	Control group - whether the tweet is about:
i. cancer treatment **helps** with treating cancer ii. cancer treatment **does not help** with treating cancer (debunks the claim) iii. cancer treatment **prevents** cancer iv. No potential cancer remedy	**i. cancer**, and has **personal** (or friend/family) experience ii. about cancer **treatment** **iii. other** cancer-related information (symptoms, awareness, prevention, causes, etc.) iv. No information about cancer

(Note: participants did not access the veracity of the tweets!)

   c) 184 CrowdFlower annotators contributed to the task

   d) A minimum of three labels collected per tweet

Data Collection

Rumor/Control data collection → User Selection → **Relevance Refinement**

2. Classification

- We train several classifiers on the labeled tweets using 1,2,3-grams as features
- We train the classifiers on the labeled tweets, which we then apply to the rest to characterize each user's behavior
- For every label in every group, we build a binary logistic regression classifier
  - *Example: from the crowdsourcing task of rumor group: 2,564 were cancer cure tweets and 1,587 were not. We build the classifier and apply it to the rest of (non-labeled) rumor tweets which results in 12,685 tweets about cancer cure and 7,872 not*
- 7,221 rumor user and 433,883 control users

Modeling Rumormongering

- We observe the behavioral statistics of three different users:
  - Rumor group users
  - Control group personal experience users
  - Control group non-personal experience users
- The different groups are compared using:
  - Mann-Whitney U test (a non-parametric test that is more appropriate for highly skewed data for which normality cannot be assumed)
  - p-value level

Modeling Rumormongering

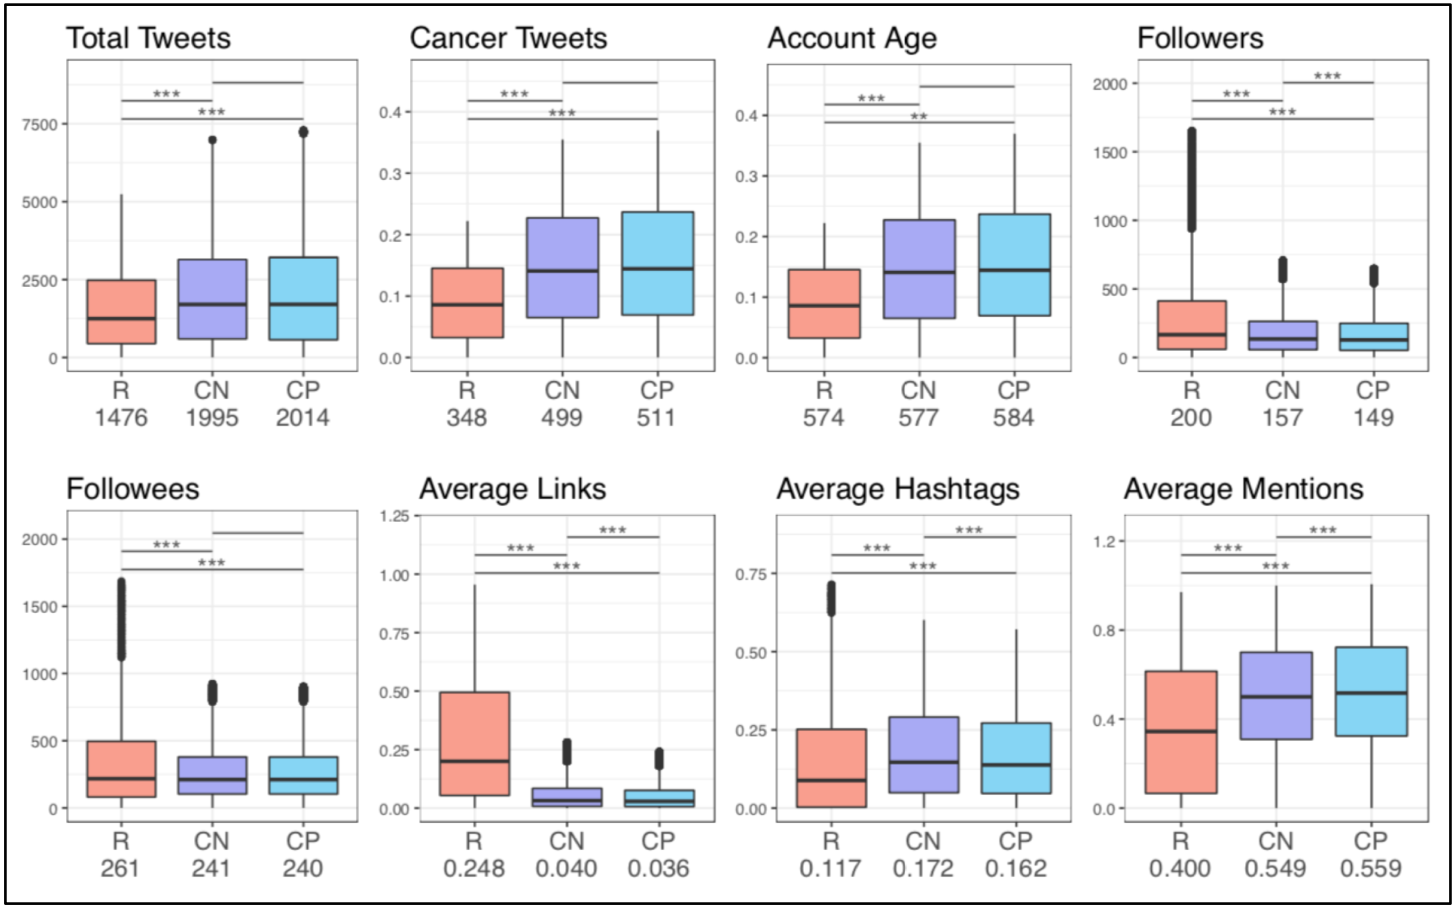

Figure 1: For each characteristic a box plot (excluding outliers outside 90th percentile) is shown with median values under the title. Differences in medians are tested using Mann-Whitney U test, for which p-values: p < 0.0001 ***, p < 0.001 **, p < 0.01 *.

Modeling Rumormongering

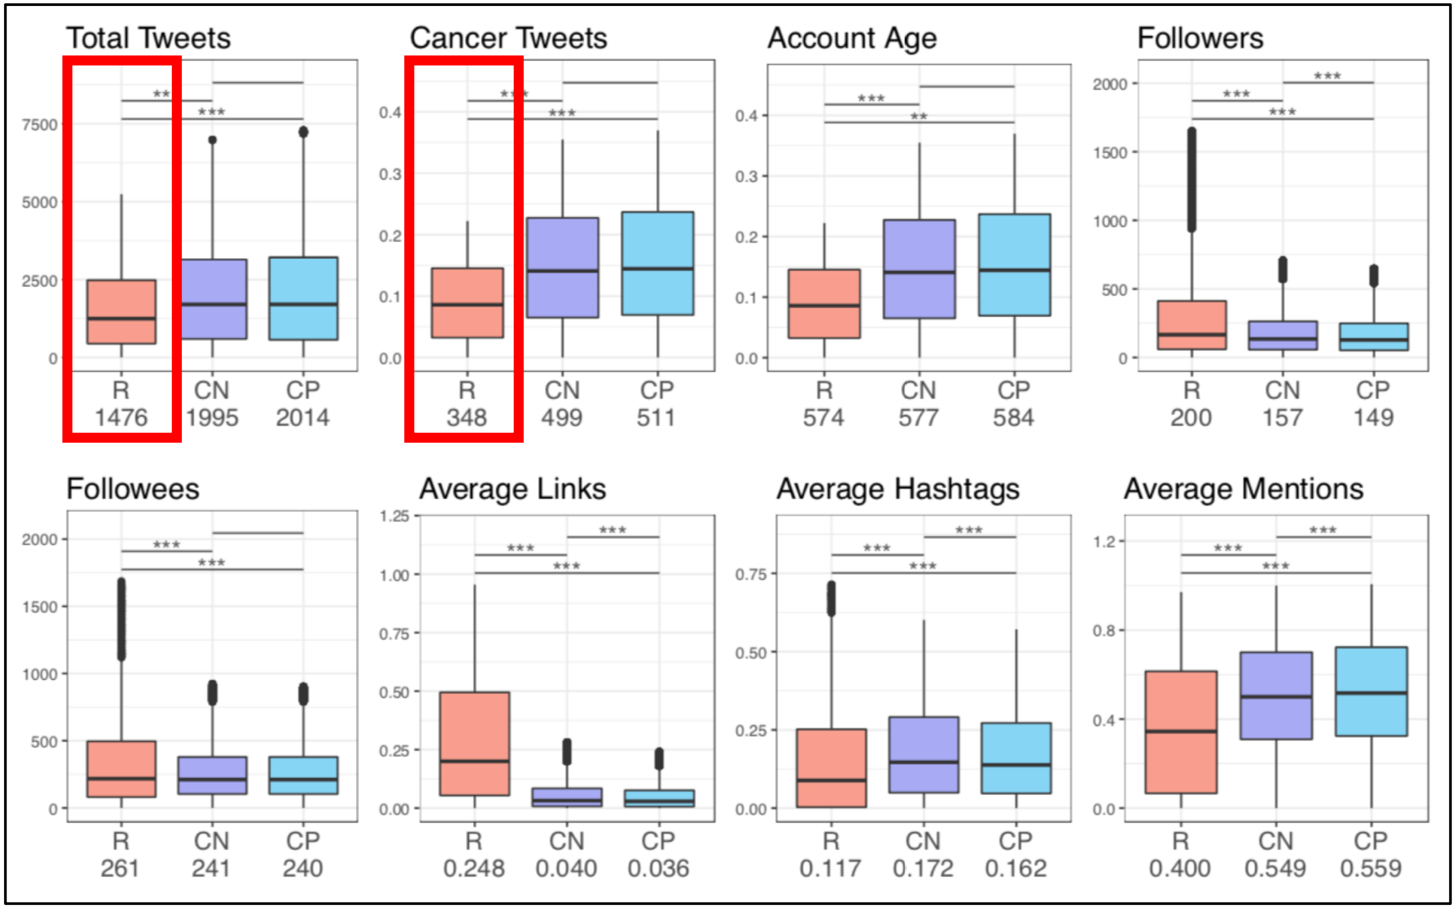

Figure 1: For each characteristic a box plot (excluding outliers outside 90th percentile) is shown with median values under the title. Differences in medians are tested using Mann-Whitney U test, for which p-values: $p < 0.0001$ ***, $p < 0.001$ **, $p < 0.01$ *.

Modeling Rumormongering

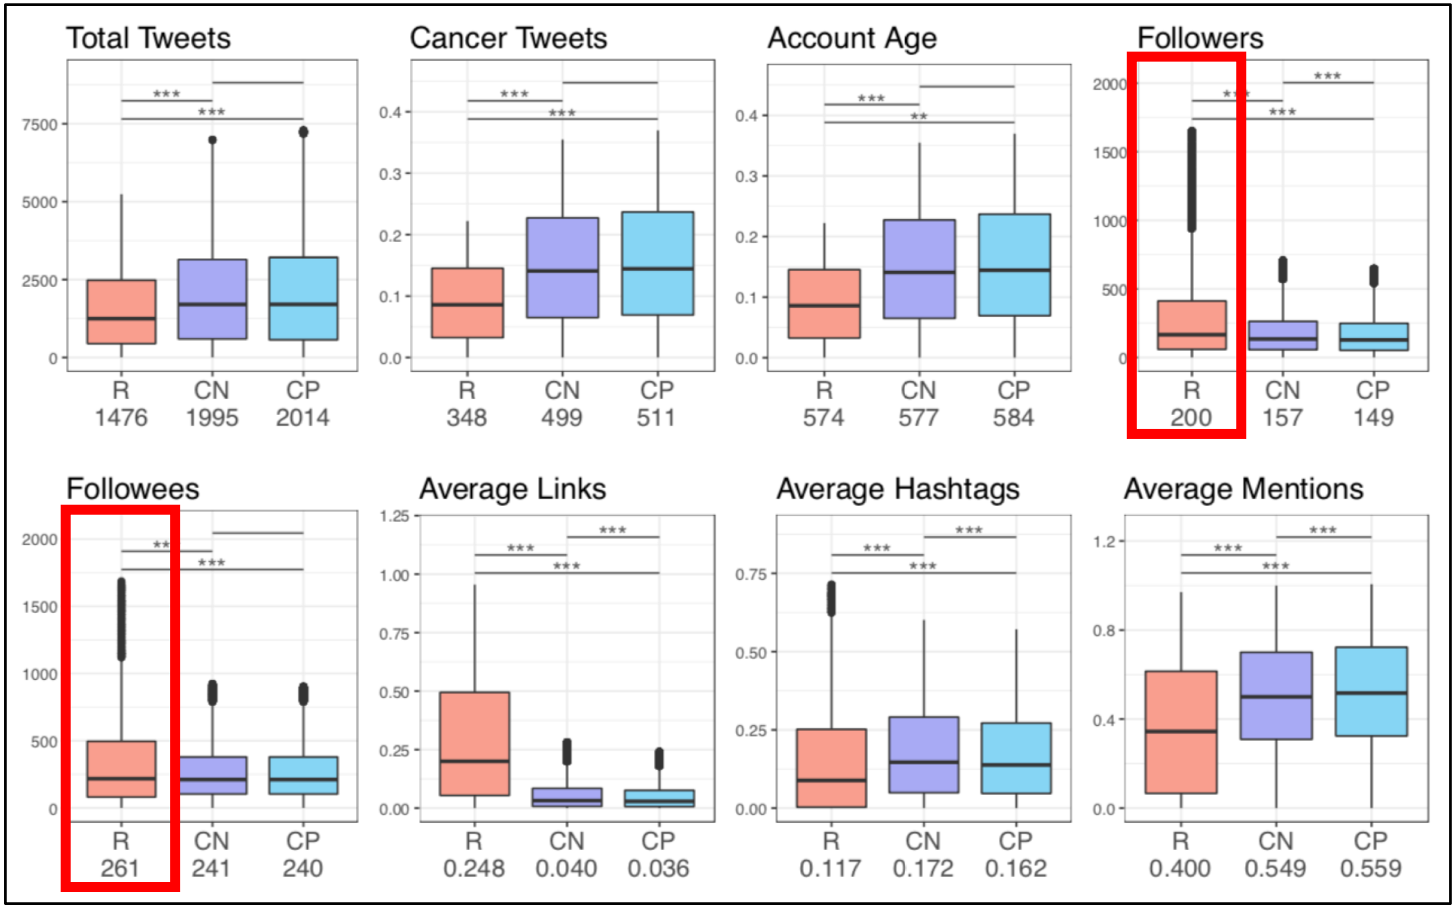

Figure 1: For each characteristic a box plot (excluding outliers outside 90th percentile) is shown with median values under the title. Differences in medians are tested using Mann-Whitney U test, for which p-values: $p < 0.0001$ ***, $p < 0.001$ **, $p < 0.01$ *.

Modeling Rumormongering

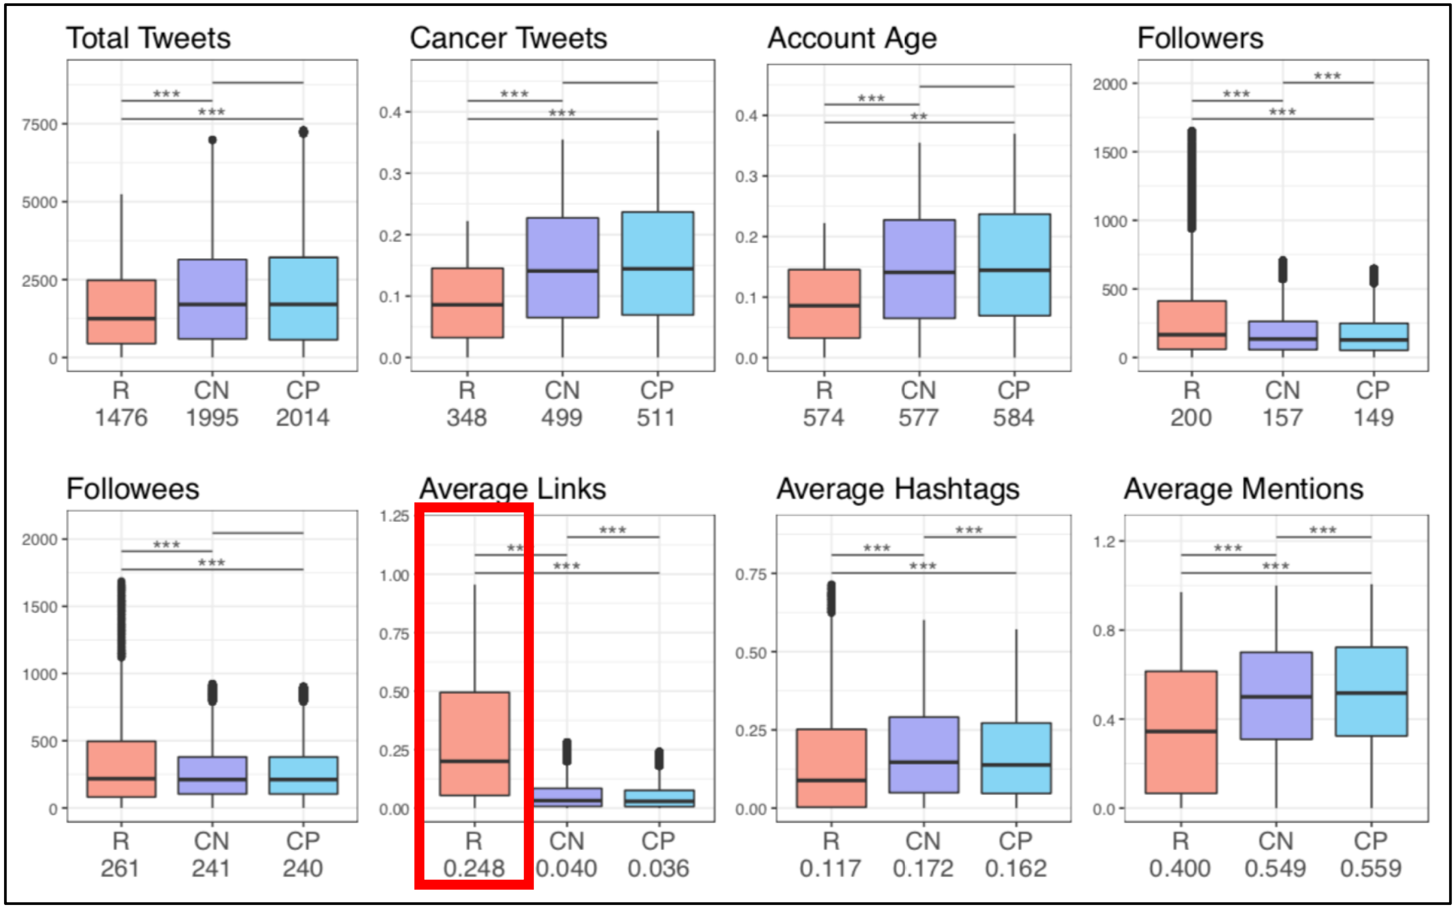

Figure 1: For each characteristic a box plot (excluding outliers outside 90th percentile) is shown with median values under the title. Differences in medians are tested using Mann-Whitney U test, for which p-values: $p < 0.0001$ ***, $p < 0.001$ **, $p < 0.01$ *.

Modeling Rumormongering

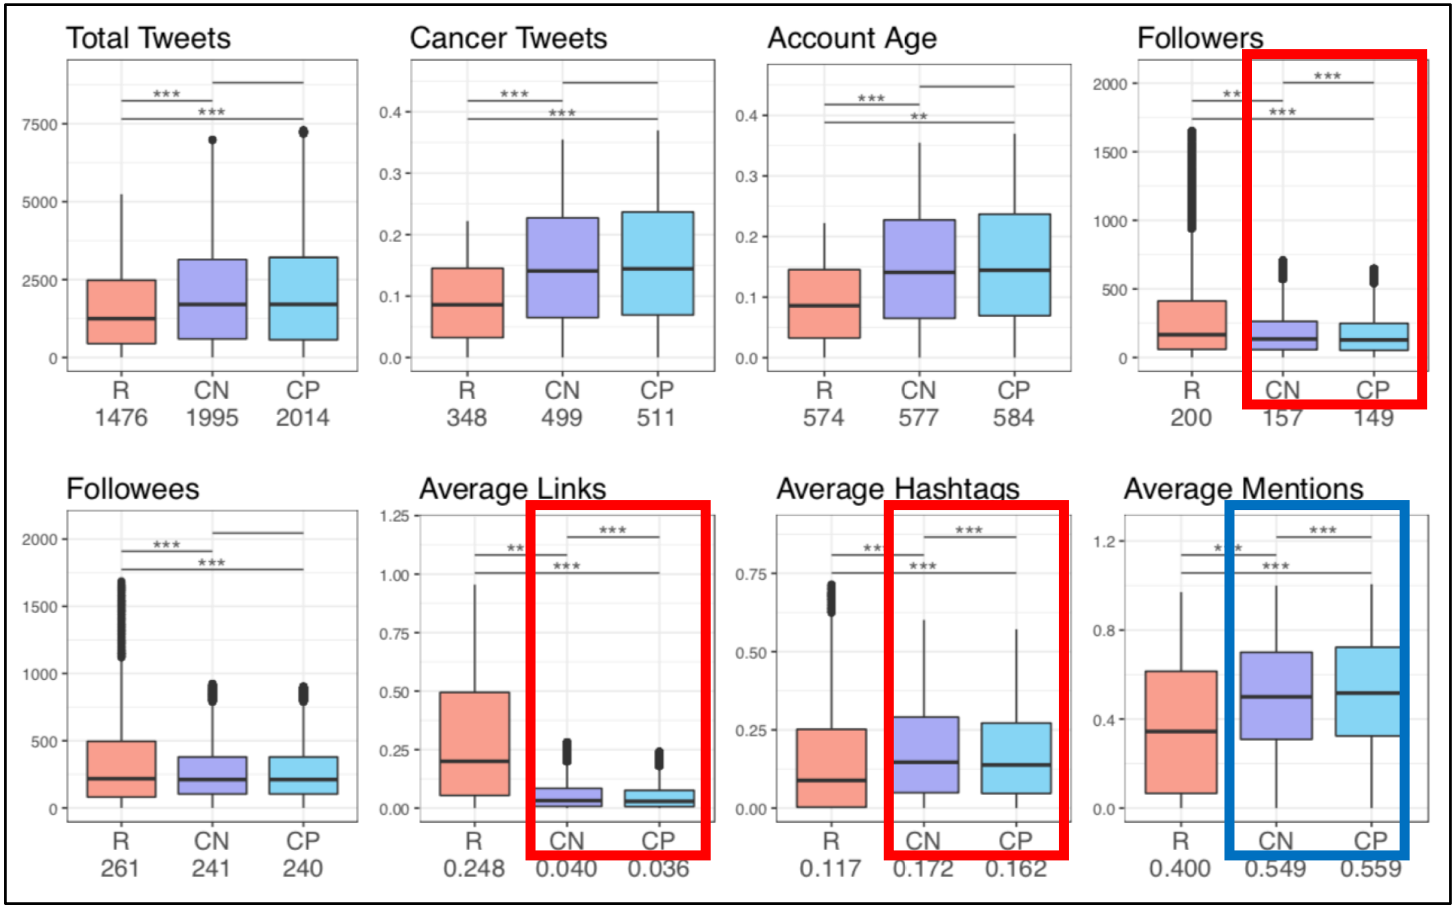

Figure 1: For each characteristic a box plot (excluding outliers outside 90th percentile) is shown with median values under the title. Differences in medians are tested using Mann-Whitney U test, for which p-values: $p < 0.0001$ ***, $p < 0.001$ **, $p < 0.01$ *

Modeling Rumormongering

- We are interested in examining whether we can **predict** the "rumor spreading" behavior before users spread the rumors
- We look at all the tweets before a user posts a tweet about the rumor (not necessarily claims the rumor)
  1. We collect all tweets timeline of every user –>get more information about users online behavior/content
  2. We keep only tweets **before** the rumor tweet -> only behavior before posting a rumor tweet

Modeling Rumormongering

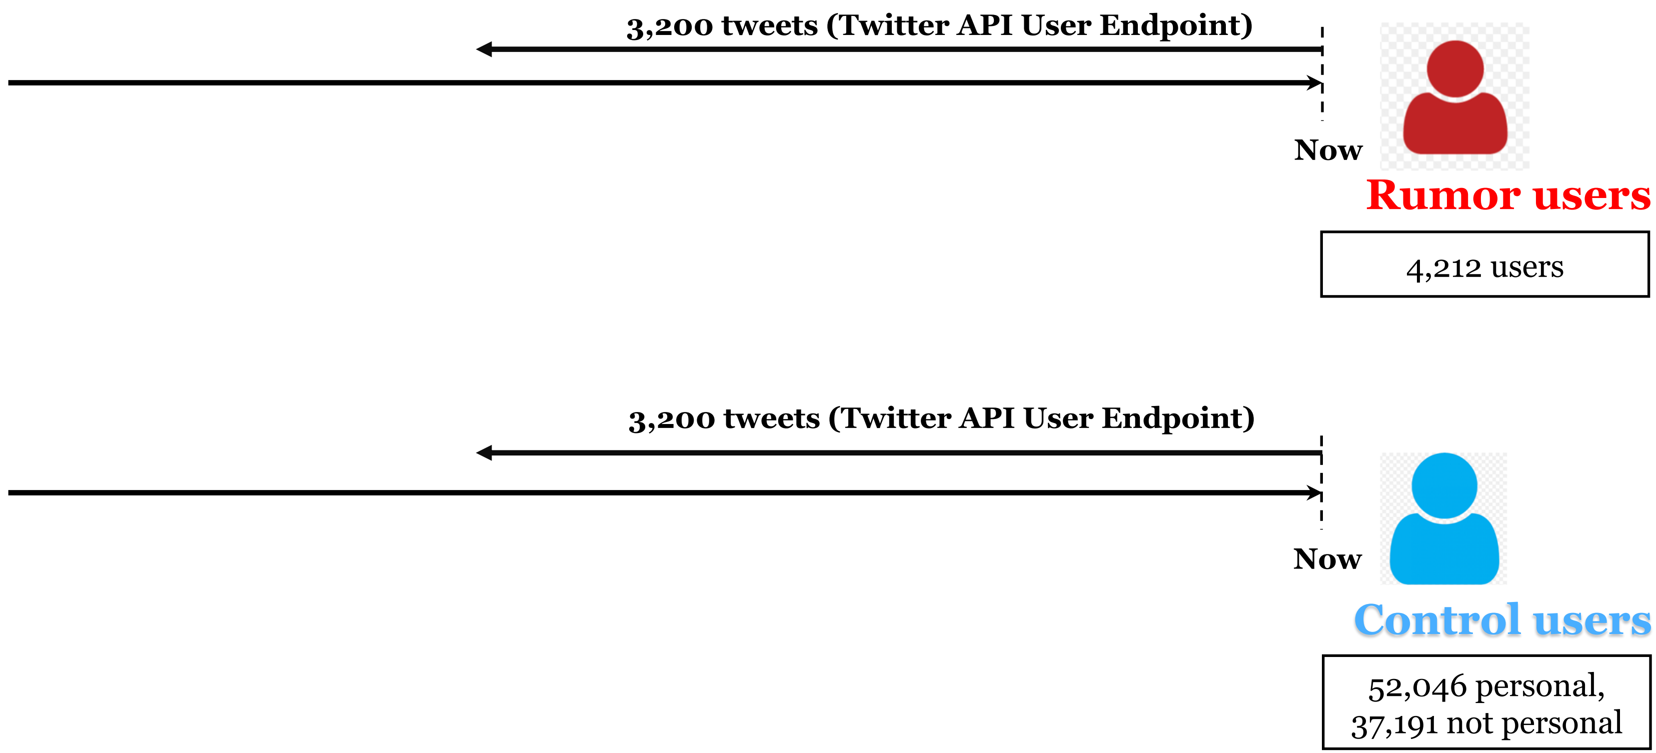

Modeling Rumormongering

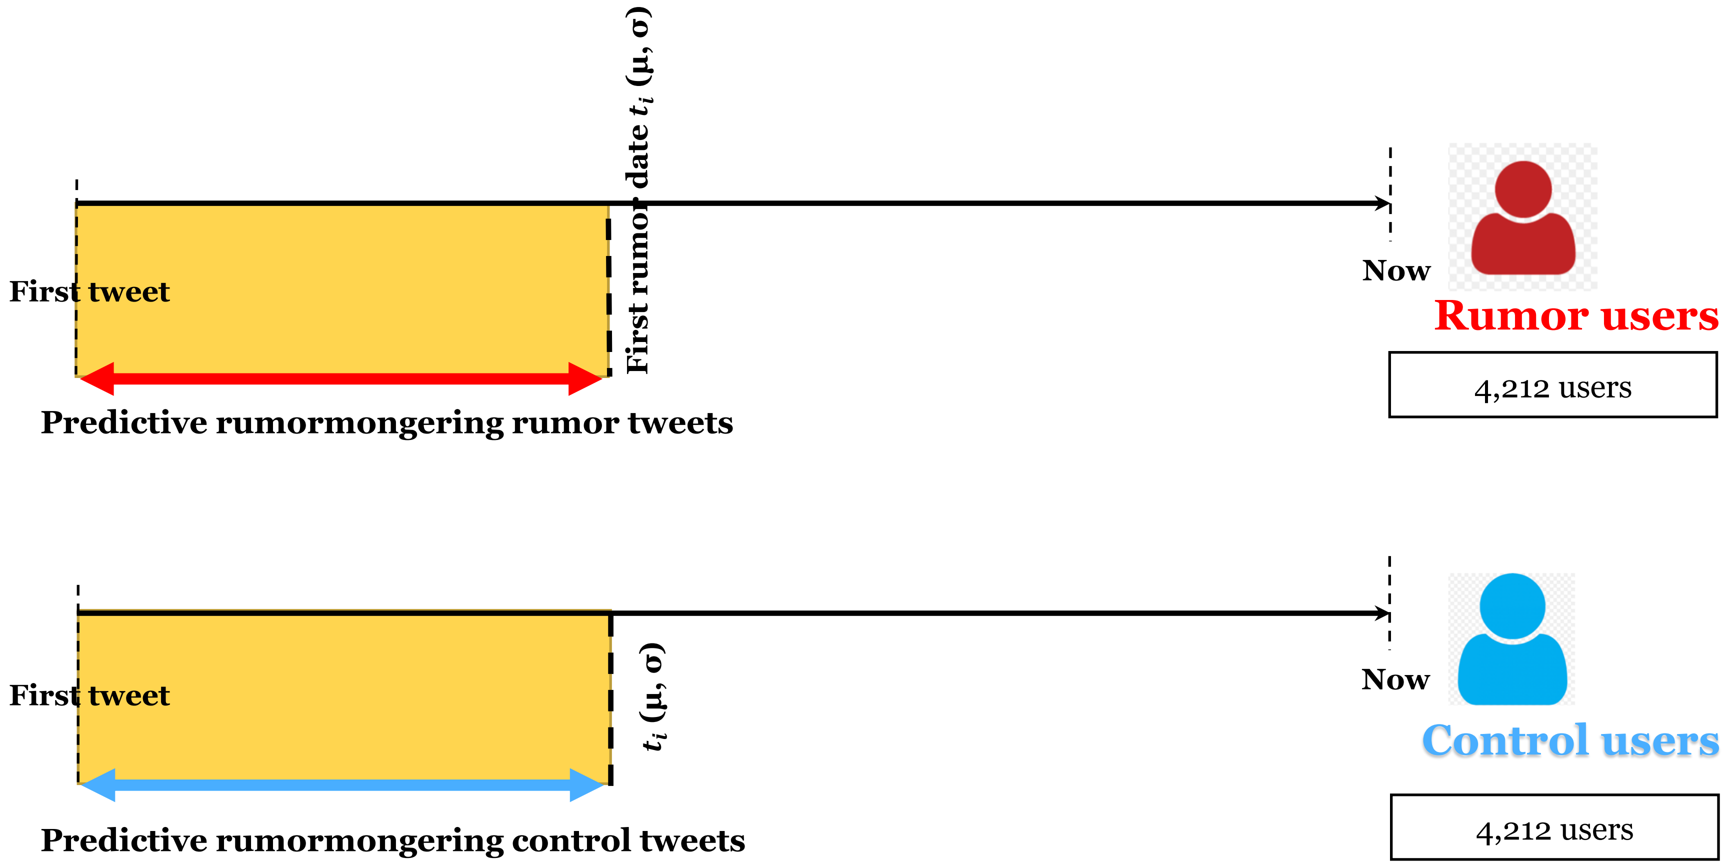

Modeling Rumormongering

- Based on our previous work, we use behavior and content features to access the credibility content in Twitter
  - User features[3]: encompass proxies of popularity (#followers, #followees), as well as productivity (# tweets up to date).
  - Tweet features[3]: linguistic and semantical forms of the tweet averaged for every user (sentiment, characters, domains etc…)
  - Entropy: the intervals between posts to measure the predictability of retweeting patterns
  - LIWC (Linguistic Inquiry and Word Count): psycholinguistic measures shown to express user mindset

[3] Amira Ghenai, Yelena Mejova, 2017, January. Catching Zika Fever: Application of Crowdsourcing and Machine Learning for Tracking Health Misinformation on Twitter. The Fifth IEEE International Conference on Healthcare Informatics (ICHI 2017), Park City, Utah.

Results – Modeling Rumormongering

- We apply logistic regression with LASSO regularization

Results – Modeling Rumormongering

- We apply logistic regression with LASSO regularization
- We use Akaine Information Criterion (AIC) for feature selection

Results – Modeling Rumormongering

- We apply logistic regression with LASSO regularization
- We use Akaine Information Criterion (AIC) for feature selection
- Results of the model shows McFadden R^2 is 0.925

Results – Modeling Rumormongering

- We apply logistic regression with LASSO regularization
- We use Akaine Information Criterion (AIC) for feature selection
- Results of the model shows McFadden R^2 is 0.925
- Instead of randomly sampling the control, we apply a matched experiment:
  - For each rumor user, select the control user with closest match of number of followers

Results – Modeling Rumormongering

- We apply logistic regression with LASSO regularization
- We use Akaine Information Criterion (AIC) for feature selection
- Results of the model shows McFadden R^2 is 0.925
- Instead of randomly sampling the control, we apply a matched experiment:
  - For each rumor user, select the control user with closest match of number of followers
- Results of the regression model with new matched samples shows McFadden R^2 is 0.906

Results – Modeling Rumormongering

Figure 2: Logistic regression with LASSO regularization model, predicting whether a user posts about a rumor, with forward feature selection. For each feature, coefficient (unstandardized), standard error, and accompanying p-value are shown. Significance levels: p < 0.0001 ***, p < 0.001 **, p < 0.01 *, p < 0.05 .

variable	coefficient	std. error	p-value
(Intercept)	-6.160	1.405	***
Avg syllables per word	17.120	0.660	***
Is verified	-40.310	42310	
Percentage uppercase / lowercase	-0.201	0.018	***
Word count	1.491	0.131	***
SMOG readability score	-0.753	0.123	***
Percentage uppercase	0.191	0.019	***
Character count	-0.163	0.024	***
Number of cancer tweets	0.001	1.9E-04	***
LIWC48: ingest	1.839	0.722	*
Negative word count	-1.460	0.262	***
URL count	3.364	0.505	***
Is retweet	4.947	0.790	***
word2vec count	-0.634	0.165	***
LIWC55: focuspast	-1.636	0.567	**
LIWC37: tentat	2.531	0.859	**
Number of sentences	-0.610	0.205	**
LIWC32: male	-1.820	1.000	
Interval entropy	0.508	0.105	***
Account age	-0.001	2.7E-04	***
LIWC23: posemo	-0.490	0.384	
LIWC61: time	-1.431	0.378	***
LIWC13: adverb	1.758	0.536	**
LIWC20: number	2.936	1.317	*
Statuses count	7.1E-05	2.6E-05	**
LIWC42: hear	-4.742	1.799	**
Has 1st person pronoun	-1.504	0.662	*
LIWC62: work	1.591	0.665	*
LIWC40: percept	1.217	0.754	

Results – Modeling Rumormongering

Readability

Figure 2: Logistic regression with LASSO regularization model, predicting whether a user posts about a rumor, with forward feature selection. For each feature, coefficient (unstandardized), standard error, and accompanying p-value are shown. Significance levels: p < 0.0001 ***, p < 0.001 **, p < 0.01 *, p < 0.05 .

variable	coefficient	std. error	p-value
(Intercept)	6.160	1.405	***
Avg syllables per word	17.120	0.660	***
Is verified	-40.310	42310	
Percentage uppercase / lowercase	-0.201	0.018	***
Word count	1.491	0.131	***
SMOG readability score	-0.753	0.123	***
Percentage uppercase	0.191	0.019	***
Character count	-0.163	0.024	***
Number of cancer tweets	0.001	1.9E-04	***
LIWC48: ingest	1.839	0.722	*
Negative word count	-1.460	0.262	***
URL count	3.364	0.505	***
Is retweet	4.947	0.790	***
word2vec count	-0.634	0.165	***
LIWC55: focuspast	-1.636	0.567	**
LIWC37: tentat	2.531	0.859	**
Number of sentences	-0.610	0.205	**
LIWC32: male	-1.820	1.000	
Interval entropy	0.508	0.105	***
Account age	-0.001	2.7E-04	***
LIWC23: posemo	-0.490	0.384	
LIWC61: time	-1.431	0.378	***
LIWC13: adverb	1.758	0.536	**
LIWC20: number	2.936	1.317	*
Statuses count	7.1E-05	2.6E-05	**
LIWC42: hear	-4.742	1.799	**
Has 1st person pronoun	-1.504	0.662	*
LIWC62: work	1.591	0.665	*
LIWC40: percept	1.217	0.754	

Results – Modeling Rumormongering

LIWC

Figure 2: Logistic regression with LASSO regularization model, predicting whether a user posts about a rumor, with forward feature selection. For each feature, coefficient (unstandardized), standard error, and accompanying p-value are shown. Significance levels: p < 0.0001 ***, p < 0.001 **, p < 0.01 *, p < 0.05 .

variable	coefficient	std. error	p-value
(Intercept)	-6.160	1.405	***
Avg syllables per word	17.120	0.660	***
Is verified	-40.310	42310	
Percentage uppercase / lowercase	-0.201	0.018	***
Word count	1.491	0.131	***
SMOG readability score	-0.753	0.123	***
Percentage uppercase	0.191	0.019	***
Character count	-0.163	0.024	***
Number of cancer tweets	0.001	1.9E-04	***
LIWC48: ingest	1.839	0.722	*
Negative word count	-1.460	0.262	***
URL count	3.364	0.505	***
Is retweet	4.947	0.790	***
word2vec count	-0.634	0.165	***
LIWC55: focuspast	-1.636	0.567	**
LIWC37: tentat	2.531	0.859	**
Number of sentences	-0.610	0.205	**
LIWC32: male	-1.820	1.000	
Interval entropy	0.508	0.105	***
Account age	-0.001	2.7E-04	***
LIWC23: posemo	-0.490	0.384	
LIWC61: time	-1.431	0.378	***
LIWC13: adverb	1.758	0.536	**
LIWC20: number	2.936	1.317	*
Statuses count	7.1E-05	2.6E-05	**
LIWC42: hear	-4.742	1.799	**
Has 1st person pronoun	-1.504	0.662	*
LIWC62: work	1.591	0.665	*
LIWC40: percept	1.217	0.754	

Results – Modeling Rumormongering

Cancer topic

Figure 2: Logistic regression with LASSO regularization model, predicting whether a user posts about a rumor, with forward feature selection. For each feature, coefficient (unstandardized), standard error, and accompanying p-value are shown. Significance levels: p < 0.0001 ***, p < 0.001 **, p < 0.01 *, p < 0.05 .

variable	coefficient	std. error	p-value
(Intercept)	-6.160	1.405	***
Avg syllables per word	17.120	0.660	***
Is verified	-40.310	42310	
Percentage uppercase / lowercase	-0.201	0.018	***
Word count	1.491	0.131	***
SMOG readability score	-0.753	0.123	***
Percentage uppercase	0.191	0.019	***
Character count	-0.163	0.024	***
Number of cancer tweets	0.001	1.9E-04	***
LIWC48: ingest	1.839	0.722	
Negative word count	-1.460	0.262	***
URL count	3.364	0.505	***
Is retweet	4.947	0.790	***
word2vec count	-0.634	0.165	***
LIWC55: focuspast	-1.636	0.567	**
LIWC37: tentat	2.531	0.859	**
Number of sentences	-0.610	0.205	**
LIWC32: male	-1.820	1.000	
Interval entropy	0.508	0.105	***
Account age	-0.001	2.7E-04	***
LIWC23: posemo	-0.490	0.384	
LIWC61: time	-1.431	0.378	***
LIWC13: adverb	1.758	0.536	**
LIWC20: number	2.936	1.317	*
Statuses count	7.1E-05	2.6E-05	**
LIWC42: hear	-4.742	1.799	**
Has 1st person pronoun	-1.504	0.662	*
LIWC62: work	1.591	0.665	*
LIWC40: percept	1.217	0.754	

Results – Modeling Rumormongering

Tentative lang

Figure 2: Logistic regression with LASSO regularization model, predicting whether a user posts about a rumor, with forward feature selection. For each feature, coefficient (unstandardized), standard error, and accompanying p-value are shown. Significance levels: p < 0.0001 ***, p < 0.001 **, p < 0.01 *, p < 0.05 .

variable	coefficient	std. error	p-value
(Intercept)	-6.160	1.405	***
Avg syllables per word	17.120	0.660	***
Is verified	-40.310	42310	
Percentage uppercase / lowercase	-0.201	0.018	***
Word count	1.491	0.131	***
SMOG readability score	-0.753	0.123	***
Percentage uppercase	0.191	0.019	***
Character count	-0.163	0.024	***
Number of cancer tweets	0.001	1.9E-04	***
LIWC48: ingest	1.839	0.722	*
Negative word count	-1.460	0.262	***
URL count	3.364	0.505	***
Is retweet	4.947	0.790	***
word2vec count	-0.634	0.165	***
LIWC55: focuspast	-1.636	0.567	**
LIWC37: tentat	2.531	0.859	**
Number of sentences	-0.610	0.205	**
LIWC32: male	-1.820	1.000	
Interval entropy	0.508	0.105	***
Account age	-0.001	2.7E-04	***
LIWC23: posemo	-0.490	0.384	
LIWC61: time	-1.431	0.378	***
LIWC13: adverb	1.758	0.536	**
LIWC20: number	2.936	1.317	*
Statuses count	7.1E-05	2.6E-05	**
LIWC42: hear	-4.742	1.799	**
Has 1st person pronoun	-1.504	0.662	*
LIWC62: work	1.591	0.665	*
LIWC40: percept	1.217	0.754	

Results – Modeling Rumormongering

Tentative lang

Figure 2: Logistic regression with LASSO regularization model, predicting whether a user posts about a rumor, with forward feature selection. For each feature, coefficient (unstandardized), standard error, and accompanying p-value are shown. Significance levels: p < 0.0001 ***, p < 0.001 **, p < 0.01 *, p < 0.05 .

variable	coefficient	std. error	p-value
(Intercept)	-6.160	1.405	***
Avg syllables per word	17.120	0.660	***
Is verified	-40.310	42310	
Percentage uppercase / lowercase	-0.201	0.018	***
Word count	1.491	0.131	***
SMOG readability score	-0.753	0.123	***
Percentage uppercase	0.191	0.019	***
Character count	-0.163	0.024	***
Number of cancer tweets	0.001	1.9E-04	***
LIWC48: ingest	1.839	0.722	*
Negative word count	-1.460	0.262	***
URL count	3.364	0.505	***
Is retweet	4.947	0.790	***
word2vec count	-0.634	0.165	***
LIWC55: focuspast	-1.636	0.567	**
LIWC37: tentat	2.531	0.859	**
Number of sentences	-0.610	0.205	**
LIWC32: male	-1.820	1.000	
Interval entropy	0.508	0.105	***
Account age	-0.001	2.7E-04	***
LIWC23: posemo	-0.490	0.384	
LIWC61: time	-1.431	0.378	***
LIWC13: adverb	1.758	0.536	**
LIWC20: number	2.936	1.317	*
Statuses count	7.1E-05	2.6E-05	**
LIWC42: hear	4.742	1.799	**
Has 1st person pronoun	-1.504	0.662	*
LIWC62: work	1.591	0.665	*
LIWC40: percept	1.217	0.754	

Results – Modeling Rumormongering

Entropy

Figure 2: Logistic regression with LASSO regularization model, predicting whether a user posts about a rumor, with forward feature selection. For each feature, coefficient (unstandardized), standard error, and accompanying p-value are shown. Significance levels: p < 0.0001 ***, p < 0.001 **, p < 0.01 *, p < 0.05 .

variable	coefficient	std. error	p-value
(Intercept)	-6.160	1.405	***
Avg syllables per word	17.120	0.660	***
Is verified	-40.310	42310	
Percentage uppercase / lowercase	-0.201	0.018	***
Word count	1.491	0.131	***
SMOG readability score	-0.753	0.123	***
Percentage uppercase	0.191	0.019	***
Character count	-0.163	0.024	***
Number of cancer tweets	0.001	1.9E-04	***
LIWC48: ingest	1.839	0.722	*
Negative word count	-1.460	0.262	***
URL count	3.364	0.505	***
Is retweet	4.947	0.790	***
word2vec count	-0.634	0.165	***
LIWC55: focuspast	-1.636	0.567	**
LIWC37: tentat	2.531	0.859	**
Number of sentences	-0.610	0.205	**
LIWC32: male	-1.820	1.000	
Interval entropy	0.508	0.105	***
Account age	-0.001	2.7E-04	
LIWC23: posemo	-0.490	0.384	
LIWC61: time	-1.431	0.378	***
LIWC13: adverb	1.758	0.536	**
LIWC20: number	2.936	1.317	*
Statuses count	7.1E-05	2.6E-05	**
LIWC42: hear	-4.742	1.799	**
Has 1st person pronoun	-1.504	0.662	*
LIWC62: work	1.591	0.665	*
LIWC40: percept	1.217	0.754	

Results – Modeling Rumormongering

Control History				Rumor History				Rumor Misinformation			
love	1.95%	night	0.66%	good	1.01%	video	0.54%	cancer	1.43%	cells	0.50%
good	1.55%	life	0.63%	health	1.00%	food	0.54%	juice	0.81%	out	0.48%
day	1.34%	happy	0.60%	day	0.96%	back	0.50%	RT	0.77%	healthy	0.45%
time	1.22%	ill	0.59%	love	0.85%	free	0.46%	breast	0.73%	diabetes	0.44%
people	1.00%	hope	0.58%	time	0.78%	work	0.45%	risk	0.61%	prostate	0.44%
lol	0.99%	feel	0.55%	great	0.73%	diet	0.44%	help	0.58%	antioxidant	0.42%
today	0.96%	haha	0.51%	people	0.71%	healthy	0.40%	health	0.55%	pain	0.40%
back	0.94%	follow	0.51%	today	0.68%	post	0.38%	helps	0.54%	chronic	0.37%
great	0.73%	home	0.49%	news	0.62%	weight	0.38%	cure	0.54%	patients	0.37%
work	0.70%	man	0.47%	life	0.57%	blog	0.36%	treatment	0.53%	study	0.36%

Figure 3: Word frequency tables summarizing the top 20 most popular terms, excluding stopwords, in all historical tweets by control users (left), all historical tweets of rumor users (center), and only rumor tweets (right).

Results – Modeling Rumormongering

Control History				Rumor History				Rumor Misinformation			
love	1.95%	night	0.66%	good	1.01%	video	0.54%	cancer	1.43%	cells	0.50%
good	1.55%	life	0.63%	health	1.00%	food	0.54%	juice	0.81%	out	0.48%
day	1.34%	happy	0.60%	day	0.96%	back	0.50%	RT	0.77%	healthy	0.45%
time	1.22%	ill	0.59%	love	0.85%	free	0.46%	breast	0.73%	diabetes	0.44%
people	1.00%	hope	0.58%	time	0.78%	work	0.45%	risk	0.61%	prostate	0.44%
lol	0.99%	feel	0.55%	great	0.73%	diet	0.44%	help	0.58%	antioxidant	0.42%
today	0.96%	haha	0.51%	people	0.71%	healthy	0.40%	health	0.55%	pain	0.40%
back	0.94%	follow	0.51%	today	0.68%	post	0.38%	helps	0.54%	chronic	0.37%
great	0.73%	home	0.49%	news	0.62%	weight	0.38%	cure	0.54%	patients	0.37%
work	0.70%	man	0.47%	life	0.57%	blog	0.36%	treatment	0.53%	study	0.36%

Figure 3: Word frequency tables summarizing the top 20 most popular terms, excluding stopwords, in all historical tweets by control users (left), all historical tweets of rumor users (center), and only rumor tweets (right).

Results – Modeling Rumormongering

Control History				Rumor History				Rumor Misinformation			
love	1.95%	night	0.66%	good	1.01%	video	0.54%	cancer	1.43%	cells	0.50%
good	1.55%	life	0.63%	health	1.00%	food	0.54%	juice	0.81%	out	0.48%
day	1.34%	happy	0.60%	day	0.96%	back	0.50%	RT	0.77%	healthy	0.45%
time	1.22%	ill	0.59%	love	0.85%	free	0.46%	breast	0.73%	diabetes	0.44%
people	1.00%	hope	0.58%	time	0.78%	work	0.45%	risk	0.61%	prostate	0.44%
lol	0.99%	feel	0.55%	great	0.73%	diet	0.44%	help	0.58%	antioxidant	0.42%
today	0.96%	haha	0.51%	people	0.71%	healthy	0.40%	health	0.55%	pain	0.40%
back	0.94%	follow	0.51%	today	0.68%	post	0.38%	helps	0.54%	chronic	0.37%
great	0.73%	home	0.49%	news	0.62%	weight	0.38%	cure	0.54%	patients	0.37%
work	0.70%	man	0.47%	life	0.57%	blog	0.36%	treatment	0.53%	study	0.36%

Figure 3: Word frequency tables summarizing the top 20 most popular terms, excluding stopwords, in all historical tweets by control users (left), all historical tweets of rumor users (center), and only rumor tweets (right).

Discussion

- The model exemplifies a tool to monitor misinformation on large scale
  - Automatically detect users more likely to post questionable facts
  - Use *persuasive technologies* to change users' attitudes
  - Timely identification of new potential rumor topics
- Useful dataset to explore other research topics
  - Understand the emotional and mental state of susceptible users