Discovering Health Beliefs in Twitter

Sanmitra Bhattacharya, Hung Tran, and Padmini Srinivasan
Department of Computer Science, The University of Iowa, Iowa City, IA 52246
{sanmitra-bhattacharya, hung-trv, padmini-srinivasan}@uiowa.edu

Abstract

Social networking websites such as Twitter have invigorated a wide range of studies in recent years ranging from consumer opinions on products to tracking the spread of diseases. While sentiment analysis and opinion mining from tweets have been studied extensively, surveillance of beliefs, especially those related to public health, have received considerably less attention. In our previous work, we proposed a model for surveillance of health beliefs on Twitter relying on the use of hand-picked probe statements expressing various health-related propositions. In this work we extend our model to *automatically* discover various probes related to public health beliefs. We present a data driven approach based on two distinct datasets and study the prevalence of public belief, disbelief or doubt for newly discovered probe statements.

Introduction

Social networking websites and social media are an integral part of our daily life now-a-days. Our views and opinions on a specific topic or the world in general are largely molded by not only traditional information sources (e.g. news, literature, etc.) but also by social media (e.g. Twitter, Facebook, blogs, etc.). Recent survey shows that almost 13% of online adults use Twitter[1], which generates over 1 billion tweets[2] per week from over 500 million users around the globe[3]. Online presence of individuals may be active or passive, where one can contribute to and/or seek information from various web sources. In the United States, 74% of adults use the internet with 61% of them looking online for health information and 6% of users posting health-related information on the internet[4]. Hence the use of social media for tracking and using health information is as important as traditional approaches for tapping into various biomedical issues.

There is a long-standing recognition that social and biobehavioral scientists and policy makers need accurate and up-to-date information about the broad spectrum of beliefs and opinions voiced in the population (Cummings et al. 2004). As an example, having an accurate estimate of the frequency of people who believe the HPV increases the risk of cervical cancer (Mosavel and El-Shaarawi 2007) or that using deodorant increases the risk of cancer (Gansler et al. 2005) allows public health scientists to decide whether there is a need to mount a special health campaign to correct those beliefs. Large-scale survey approaches using mail, telephone, and special websites can provide useful data, but such approaches, by definition, having already formulated the content of their questions, do not tap the naturally occurring opinions or beliefs expressed by people. Moreover, typically there are always time-delays between preparation of the survey questions and administration. The development of a methodology that assesses the prevalence of the naturally occurring expression of beliefs and opinions would be immensely useful. Motivated by this, we recently proposed, in a research note, the novel function of belief surveillance and demonstrated how this could be done using Twitter (Bhattacharya et al. 2012).

The surveillance methods we proposed involve specific propositions that we call probes. A probe is a statement presenting a directed, binary relationship between two key concepts. An example is *smoking causes cancer*. In our prior work we studied belief surveillance for 32 probes and showed, for example, that although factual probes (e.g. smoking causes cancer) generally garner high degree of belief, there is still considerable doubt regarding some false probes (e.g. honey treats allergies). Quite alarmingly, we find several debatable (e.g. Actos causes bladder cancer) and false statements also generate high level of belief among Twitter users.

Our prior work was mostly limited to the belief analysis of manually selected probe statements. We did not fully explore an automatic approach for identifying health beliefs in Twitter. In this paper, we extend our prior work with analysis of new beliefs for probes mined automatically from Twitter using two data-driven approaches. Automation is necessary to be able to scale our methodology to handle surveillance of beliefs as they arise. In summary, we ask the following new questions in this paper.

- What kinds of health beliefs are revealed by the naturally occurring discussions on Twitter? In particular we mine beliefs related to a set of health-hashtags and also a set

[1] http://bit.ly/mwmzOp (links to PewInternet.org)
[2] http://blog.twitter.com/2011/03/numbers.html
[3] http://twopcharts.com/twitter500million.php
[4] http://bit.ly/3b8Np4 (links to PewInternet.org)

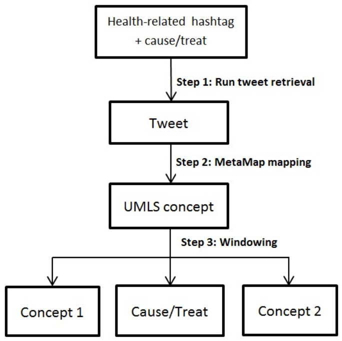

Figure 1: Flowchart of Health-related Tweet Retrieval and Concept Extraction (e.g. examples of Concepts 1 and 2 may be 'smoking' and 'cancer')

30 diseases and 20 drugs. Thus we are able to ask: What is the public perceptions on a health belief X? What are the public perceptions of known effects (or side-effects) of drugs? What beliefs are observed regarding cures of diseases using prescription and OTC drugs? For this, we extend our earlier methods for mining new beliefs from Twitter.

- Which health beliefs are most prevalent in Twitter conversations? Thus we will be able to determine if the discovered beliefs are more or less common in this population.

Related Research

Recent studies in biomedicine and healthcare informatics have seen an increasing use of Twitter. The scope of such studies can be either broad or narrow. Using Twitter for disease surveillance (Signorini, Segre, and Polgreen 2011) or analyzing topics relevant to public health (Paul and Dredze 2011) takes a much broader scope compared to disease specific studies. Studies to monitor health information dissemination for specific health-related issues like dental pain (Heaivilin et al. 2011) or concussion (Sullivan et al. 2011) show the spectrum of applicability of Twitter to get insight into the specific health problems.

Our goal of belief surveillance is inspired by work on disease surveillance. There has been an increasing use of social media in disease surveillance. Popular social networking websites have been shown to be important sources for monitoring real-time data for disease outbreaks. Twitter-based influenza epidemics detection (Aramaki et al. 2011) shows the importance of mining social media for early stage detection of a disease outbreak. A similar study also reinforces the timeliness of outbreak detection using tweets and found high correlation between patterns gathered from Twitter messages and CDC statistics (Culotta 2010). The use of twitter traffic and tweets has also been shown for not only gauging public interest about a particular disease (H1N1), but also for tracking disease outbreak in real-time (Signorini, Segre, and Polgreen 2011).

In comparison to such disease specific studies, few studies have taken a much broader scope. For example, Paul and Dredze (2011) propose a topic model based approach to explore Twitter for public health research. In this study they identify various public health topics that can be studied using Twitter such as, syndromic surveillance, behavioral risk factors, and symptoms and medications taking into account the geographic distribution for each such topic. Similarly, Prier et al. (2011) study methods for mining general health topics from Twitter. Using an LDA topic model they identify health-related issues that garner a lot of attention on Twitter, such as, weight loss programs, Obama's health reform policy, marijuana uses, etc.

In contrast to these studies, we use Twitter to identify and estimate public health beliefs. Towards this end, we use a novel surveillance framework and show its applicability in mining and guaging naturally occurring health beliefs in the population.

Belief Surveillance

In our previous work (Bhattacharya et al. 2012), we proposed a novel framework for *belief surveillance*. We handpicked a set of 32 probe statements[5] from various websites (CDC, FDA, etc.), news sources (Google News) and physicians. Each probe statement was categorized as either *true*, *false* or *debatable* by a physician. We then used the Twitter Search API[6] for collecting tweets for each probe statement. A subset of these tweets were manually annotated using the crowdsourcing platform oDesk[7]. For each tweet, annotators had to judge the *relevance* as well as the *position* (i.e. supporting, opposing or questioning) of the tweet vis-à-vis a probe statement. For each probe we then calculated the degrees of belief, disbelief and doubt[8] using equation (1) and its analogues. These estimates were then aggregated as averages to find the extents of belief, disbelief and doubt (equation (2) and analogues) for the true, false and debatable probe categories.

$$Degree\_Belief(S_i) = \frac{\#\ relevant\ tweets\ supporting\ S_i}{\#\ tweets\ relevant\ to\ S_i} \quad (1)$$

$$Belief(S_1, S_2, ..., S_N) = \frac{1}{N}\sum_{i=1}^{N} Degree\_Belief(S_i) \quad (2)$$

In summary, we presented our belief estimates for the 32 probes some of which represented factual information, others represented false information and the remaining represented debatable propositions. We found high belief for true as well as false statements and low disbelief for false and debatable statements. Finally, we showed preliminary evidence of feasibility for automating this estimation process by using off-the-shelf classifiers, namely, SVM with PolyKernel and Bagging using unigram features.

[5]A probe statement is a sentence representing a specific hypothesis or idea. e.g. 'vaccine causes autism', 'lemon treats cancer'. The rational for selection of probe statements is presented in (Bhattacharya et al. 2012).

[6]https://dev.twitter.com/docs/api/1/get/search

[7]https://www.odesk.com/

[8]'belief', 'disbelief' and 'doubt' reflect the levels of public support, opposition and questioning, respectively, for a probe state-

Table 1: Health Related Hashtags

Hashtags			
#disease	#medicine	#doctor	#patient
#doctors	#patients	#health	#pharma
#healthcare	#pharmacy	#hospital	#physician
#hospitals	#physicians	#medical	#therapy

Data Driven Probe Statements

One of the limitations of our previous research is that our work has revolved around 32 hand-picked probe statements. Though important, this set ignores other spontaneous and naturally-occurring discussions on health topics in Twitter. We now propose two approaches to identify discussions related to causes and treatments of illnesses in Twitter. In the first approach we use a set of health-related hashtags as hooks to fetch probes from Twitter. As a second approach we propose the use of a variety of drug and disease names to identify more specific probe statements for health belief surveillance.

Hashtag-based approach

In this approach, we first build a dataset (*HashtagDataset*) using a strategy independent of the probe statements. We identify a set of 16 general health-related hashtags (Table 1) from Fox's ePractice Healthcare Hashtag Project[9]. We do not select country (e.g. #cdnhealth), organization (e.g. #FDA) or technology-specific (e.g. #HealthIT) hashtags. We also choose not to explore disease-specific (e.g. #hepatitis) hashtags; though relevant we are leaving these for future research. Each hashtag is then coupled with causes and treats verbs as search terms. These pairs are searched using the Twitter Search API (as in previous work). Using this API users can get at most 1,500 tweets per query within a time frame of the past 7 days. The *HashtagDataset* was built on October 13, 2011 and contains 1,313 non-unique tweets. After removal of user mentions (@ prefixed), retweet mentions (@RT) and URLs we get 613 unique tweets.

Our goal here is to identify naturally occurring discussions on causes and treatments of illnesses. With this in mind we process these 613 tweets using the procedure outlined in Figure 1. We process each tweet with National Library of Medicine's MetaMap (Aronson 2001) program. MetaMap can identify coherent words or phrases from a particular sentence (tweet in our case) and map them to Unified Medical Language System[10] (UMLS) metathesaurus concepts. UMLS provides a standard set of health and biomedical vocabularies. After the mapping step, each tweet is annotated with several concepts and their semantic types (categories). We then extract pairs of concepts belonging to key semantic types ('Disease or Syndrome', 'Finding', 'Pharmacologic Substance', 'Manufactured Object', etc.) appearing within a specific window size of 4 words[11].

Table 2: *HashtagDataset* mined probes (T: True; F: False; D: Debatable)

Mined Probe Statements	#Tweets	
smoking causes death	1181	T
skin product causes aging	673	F
chemotherapy treats breast cancer	392	T
oral sex causes throat cancer	272	T
marijuana treats PTSD	235	D
smokeless tobacco causes cancer	150	T
antidepressant causes depression	149	T
stress causes sickness	129	T
medication causes hair loss	84	T
milk causes acne	52	T
milk causes osteoporosis	29	F
magic mushroom causes personality change	23	D
nasal polyp causes nasal block	17	T
tea tree oil treats infection	17	D
cialis treats enlarged prostrate	16	T
diet causes bad breath	16	T
listeria causes miscarriage	14	T

This procedure results in 49 new concept pairs linked by 'causes' or 'treats'. Now with these pairs as probes we use the Twitter Search API for retrieving tweets for each new probe. This was done on November 1, 2011; retrieved tweets dated back to the previous 7 days as per the API. Table 2 lists the 17 pairs that retrieved at least 10 tweets and that did not appear in our initial probe statement set. So these are the probes we have discovered in our data. They represent the conversations (of at least 10 tweets) around causes and treatments that were occurring in Twitter. Here too each probe statement was judged as either true, false or debatable by a physician.

Now applying our two general classifiers (SVM with PolyKernel and Bagging for relevance and for position, developed in our previous work) we classified the tweets retrieved for the 17 probes to produce the belief chart shown in Figure 2. The chart is clustered into 3 groups; the leftmost represents factual statements (T), the rightmost represents fictional statements (F) and the middle statements are debatable (D). The y-axis conveys the degree of belief, disbelief, doubt and other (estimated using Equation (1) and its analogues). Within each group the statements are ordered by the degree of belief. First we observe that the ideal situation would be if the tweets convey 100% belief in factual statements and 100% disbelief in false statements. For the debatable probe statements, we can expect a mixture of belief, disbelief and some degree of doubt via questioning.

Instead we see high levels of belief in false and debatable probe statements. For example there is almost no disbelief in the false notion that *skin products cause aging* and little disbelief in the debatable notion that *tea tree oil treats infection*. Aggregating belief, disbelief and doubt using equations (2) and its analogues gives us some very alarming results (Table 3). Not surprising, there is high belief (0.80) in true statements. However, there is almost equal belief in both false and debatable statements. Disbelief is surprisingly low in false (0.05) statements where the deviation from ideal is 0.84. These results are even more striking compared to the findings from our previous research. There is considerably

ment (S_i)

[9]http://www.foxepractice.com/healthcare-hashtags/

[10]https://uts.nlm.nih.gov/

[11]This parameter was tuned on a training set of 30 MetaMapped tweets

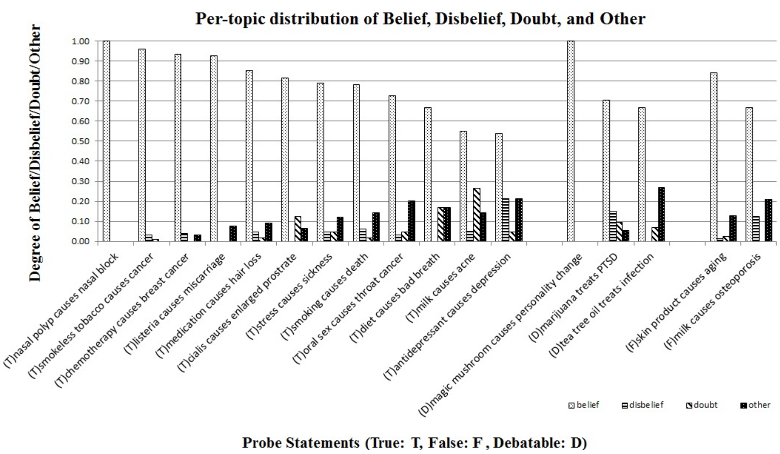

Figure 2: Belief Plot for *HashtagDataset*

Table 3: Belief, Disbelief and Doubt Measures for *Hashtag-Dataset*

	True	False	Debatable
Belief	**0.80**	0.79	0.75
Disbelief	0.04	0.05	**0.07**
Doubt	**0.06**	0.05	0.01
Other	0.10	0.11	0.17

Table 4: Drugs and Diseases Explored

Top 10 DTC Drugs	Lipitor, Cialis, Advair, Abilify, Cymbalta Symbicort, Pristiq, Plavix, Chantix, Lyrica
OTC Drugs	Aspirin, Advil, Prilosec, Centrum, Robitussin Tylenol, Nyquil, Dramamine, Zantac, Benadryl
Chronic Diseases	Diabetes, Asthma, Arthritis, Schizophrenia, Cardiac failure Glaucoma, Haemophilia, Hypertension, Multiple sclerosis, Parkinson's disease, Osteoporosis, Psoriasis, Obesity, Epilepsy
Infectious Diseases	HIV/AIDS, Dengue, Malaria, Anthrax, Cholera Bubonic plague, Influenza, Typhoid, Smallpox, Pneumonia Tuberculosis, Yellow fever, Bird flu, Ebola, Leprosy, Hepatitis

less disbelief and doubt for both false and debatable statements. As discussed before, these results may be used by health educators to make health campaign decisions or even to identify angles to probe further in the general population through more formal surveys.

Drug and disease name-based approach

In contrary to the more general approach of mining illness related probe statements, in our second approach we try to gather the public perceptions of known effects (or side-effects) of drugs, and beliefs observed regarding cures of diseases using prescription and OTC drugs. Here we select a list of drugs and diseases (*DrugDiseaseDataset*) as hooks to automatically fetch naturally-occurring probe statements in Twitter (Table 4). Though our primary interests are in beliefs related to specific diseases and drugs, we are also interested in entities that co-occur with drugs and diseases. These might be food, drinks, recreational drugs, animals, etc. For drugs, we selected the top 10 most Direct-to-Consumer (DTC) advertised drugs[12] for heart diseases, neuropathic pain, etc. We also selected 10 of over-the-counter (OTC) or generic drugs[13] that are less advertised but frequently used for common problems like fever, pain, heartburn,etc. For diseases we selected chronic[14] and infectious diseases[15] from the World health Organization's (WHO) fact sheets for such diseases (Table 4).

We combined each drug and disease term from Table 4 with relationship terms causes and treats for search on Twitter. This was done on February 24, 2012 using the Twitter Search API. Similar to the previous approach, here also we removed URLs, user mentions and re-tweet mentions and took the unique instances of the remaining tweets. For the drugs and diseases this resulted in 942 and 3722 unique tweets respectively.

These tweets were then processed using an approach similar to the one shown in Figure 1. However given the considerably large number of tweets retrieved using this approach and the potential variability in tweet expressions in this set of tweets, we considered a wider array of semantic type to fetch the concepts for probe statements. For reasons stated above, each tweet was examined for the presence of drug and disease specific semantic types appearing together in the same tweet. Such tweets are expected to portray drug-disease relationships being discussed by Twitter users. For drug related semantic types, we considered [Organic Chemical, Pharmacologic Substance], [Amino Acid, Peptide, or Protein,Pharmacologic Substance], etc. We also consider certain drug-unrelated semantic types which may have important association with diseases (for example, [Food], [Mammal], etc.). For diseases we considered semantic types such as [Disease or Syndrome], [Mental or Behavioral Dysfunction], etc. Using this strategy we identified 361 and 978 term pairs from the drug and disease sets respectively. Manual inspection of these pairs revealed certain anoma-

[12] http://gaia.adage.com/images/bin/pdf/WPpharmmarketing_revise.pdf
[13] http://www.uihealthcare.com/pharmacy/OTCmedications.html
[14] http://www.who.int/topics/chronic_diseases/factsheets/en/
[15] http://www.who.int/topics/infectious_diseases/factsheets/en

Table 5: Drugs and Diseases Datasets

Dataset: ***Drugs***	
Number of tweets retrieved for Disease set	1226
Number of unique tweets	942
Number of pairs (before filtering)	361
Number of pairs (after filtering)	209
Dataset: *Diseases*	
Number of tweets retrieved for Drugs set	8679
Number of unique tweets	3722
Number of pairs (before filtering)	978
Number of pairs (after filtering)	556

Table 6: Select probes from *DrugsDiseaseDataset* (probe status: True: T, False: F, Debatable: D)

Mined Probe Statements	
Advil treats hangover (T)	Ginkgo treats diabetes (F)
Viagra treats anxiety (F)	Viagra treats hypertension (D)
Advil causes stomach bleeding (T)	Marijuana causes schizophrenia (T)
Armadillos causes leprosy (T)	Methotrexate treats cancer (T)
Video causes seizure (T)	Water treats hangover (T)
Benadryl causes itching (F)	Nigella sativa treats diabetes (D)
Bilberry treats diabetes (D)	Nyquil causes coma (F)
Weed treats AIDS (D)	Weed treats asthma (F)
BPA causes obesity (D)	Overeating causes memory loss (D)
Cannabis treats bronchitis (F)	Seroquel causes diabetes (D)
Weed treats cancer (D)	Weed treats depression (D)
Cialis treats impotency (T)	Stress causes schizophrenia (T)
Coffee causes diabetes (F)	Viagra causes hearing loss (D)
Weed treats glaucoma (D)	

lies in the identified term pairs. For example, for the tweet "The girl in my class is giving a speech and said weed causes schizophrenia", MetaMap identifies the verb 'said' as 'said (Simian Acquired Immunodeficiency Syndrome) [Disease or Syndrome]'. Other examples of frequently appearing but incorrectly mapped terms are 'I' (identified as 'I NOS (Blood group antibody I)'), dnt (abbreviated don't) (identified as 'DNT (Dysembryoplastic neuroepithelial tumor)'), etc. On the other hand, we also identified some very interesting probes such as 'armadillos[16] cause leprosy' because of the use of wider range of semantic types. After manually filtering out pairs containing incorrectly mapped terms we got 209 and 556 pairs from the drugs and diseases datasets respectively. Table 5 summarizes these dataset characteristics. We note that disease-related search terms retrieve more tweets than drug-related terms. Consequently, the number of probes mined from disease-related tweets is also greater than that of drug-related tweets dataset.

For the mined probes from *DrugDiseaseDataset* we follow the same retrieval strategy shown before. The Twitter search was performed on February 28, 2011; retrieved tweets dated back to the previous 7 days as per the API. We collected 88048 tweets for over 700 filtered probes (Table 5). 117 probes of the drugs dataset and 324 probes of the diseases dataset retrieved more than 10 tweets. We selected a subset of 27 from these probes to analyze further. Table 6 shows the selected probes that retrieved at least 10 relevant tweets. As mentioned earlier, note that the mined probes include some entities that are not about drugs or diseases (e.g. water, coffee, weed, armadillos, etc.) but co-occur with a drug or disease terms. For each probe in our subset a physician provided decisions on whether the belief is true, false or debatable. While our hand picked set might not be representative of the wide range of probes collected from the *Drugs-DiseaseDataset* (Table 5), it gives us a sample which is feasible to study within the scope of this paper. 26 of the 66 of selected probes pertain to effects/side-effects of therapeutic drugs. Quite a few refer to recreational drugs. There were a number of probes in alternative medicine (herbal therapy, homeopathy, etc.) and dietary substances with causal or curative relationships with diseases. Not surprisingly, we find a large number of probes relate to recent studies with animal models that might have generated a buzz. Naturally probes mined depend upon current events and developments. This is because social media often correlates to current events in news media or even pop-culture[17]. Overall, our probe mining outcome supports our intuition that health beliefs are discussed in twitter and may be identified.

[16]Armadillos (*Armadillo officinalis*) are mammals primarily found in Central and South America.

Table 7: Belief, Disbelief and Doubt Measures for *DrugDiseaseDataset*

	True	False	Debatable
Belief	**0.82**	0.80	0.76
Disbelief	0.06	**0.13**	0.03
Doubt	0.02	**0.04**	0.02
Other	0.10	0.03	0.19

Mined Probes

Figure 3 shows the plots of belief, disbelief, doubt, and other for the factual, fictional and debatable probes (Table 6) having at least 10 relevant tweets. Here we notice that there is surprisingly low belief in some true probes such as *Cialis treats impotency* which is a known prescribed medication for impotency. Additionally a high level of belief in debatable probes, especially those pertaining to alternative medicines and recreational drugs (e.g. weed). We also notice that a significant proportion of tweets retrieved for probes like *weed treats glaucoma* contain large number of spam tweets (advertisements) which are classified in the 'other' category. These were apparently quite challenging for our relevance classifier. We see high belief in false probes related to alternative medicine, recreational drugs and even therapeutic drugs. For example, tweets like "*smokin weed helps people wit asthma #fact*" emphasize such beliefs in false probes.

The aggregate calculations of belief, disbelief, doubt, and other is shown in the third set of columns in Table 7. Similar to the belief plot, we see high belief (0.82) in true statements. Moreover, there is almost equally high belief (0.80) in false probes and debatable probes (0.76). Low disbelief (only 0.13) in false statements is alarming and emphasizes the need to embark on public health campaigns to correct such misinformed beliefs.

The findings for the automatically mined probes mirror our earlier findings using the 32 probes in several aspects. With both the *HashtagDataset* and *DrugDiseaseDataset* we have shown that belief in false and debatable probes is quite high and sometimes even comparable to belief in true

[17]Several tweets related to the probe "video causes seizure" refer to a popular music video that might cause epileptic seizure and contains a related disclaimer

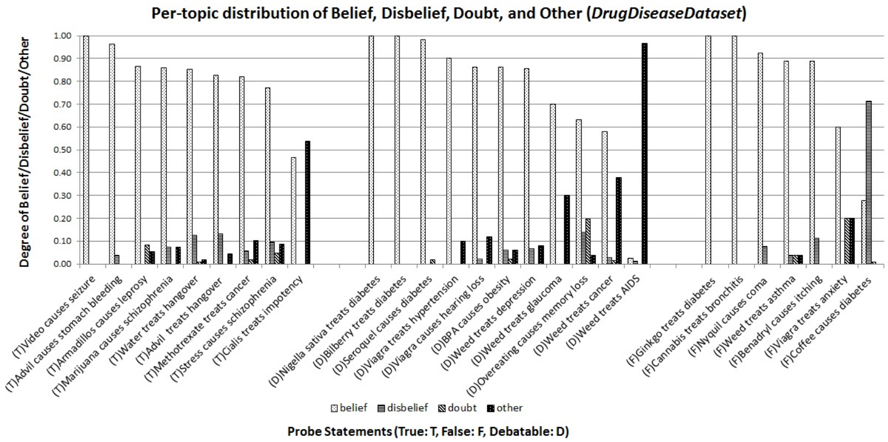

Figure 3: Belief Plot for *DrugDiseaseDataset*

probes. However, the levels of disbelief in false and debatable probes is much lower for the mined tweets compared to the 32 probes from our previous study. The low levels of disbelief in certain probes like *vaccine causes autism* (Bhattacharya et al. 2012) may be deemed more worrisome compared to *weed treats asthma* based on the reach and impact of such beliefs. Similarly, probes that generate a lot of doubt in the form of questioning such as *chocolate causes acne* or *Viagra treats anxiety* may be identified as equally important for disseminating public health knowledge in the population.

Conclusions

In our previous research we developed methods for gauging public belief, i.e., for belief surveillance, with Twitter as an exemplar social medium. We proposed a novel approach built around *using* statements representing specific propositions as probes and measuring belief, disbelief and doubt for such propositions. In this paper we propose methods for *mining* probe statements automatically from twitter and thereby discovering naturally occurring health beliefs on Twitter. This is an exciting aspect as it points to being able to proactively conduct belief surveillance with Tweet data. This also shows that our method is robust in the sense that it does not rely solely on the predefined probes to measure a population's beliefs, but it can automatically discover and monitor new beliefs from social media. From the *HashtagDataset* we show that there is a high level of belief in false and debatable probes. From the *DrugDiseaseDataset* we find an large number of probes related to not just therapeutic drugs in our study, but also recreational drugs which demonstrate the prevalence of discussions on such topics in Twittersphere. Other than the known effects or side-effects of drugs, were able to uncover several probes containing novel information (e.g. *Cialis causes anxiety*). In future research we plan to test the validity of such information including potential adverse drug effect reports. Most interesting, especially from the perspective of public health education, is that we show that some fictional statements also garner a high degree of belief. These results potentially offer an informed basis for targeting educational strategies. Overall we show that Twitter offers valuable signals that can be used for belief surveillance.

References

Aramaki et al., E. 2011. Twitter catches the flu: Detecting influenza epidemics using twitter. In *EMNLP*, 1568–1576.

Aronson, A. R. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. *Proc AMIA Symp* 17–21.

Bhattacharya, S.; Tran, H.; Srinivasan, P.; and Suls, J. 2012. Belief surveillance with twitter. In *Proceedings of the Fourth ACM Web Science Conference (WebSci2012)*.

Culotta, A. 2010. Towards detecting influenza epidemics by analyzing twitter messages. In *Proceedings of SOMA '10*.

Cummings et al., K. M. 2004. Are smokers adequately informed about the health risks of smoking and medicinal nicotine? *Nicotine Tob. Res.* 6 Suppl 3:S333–340.

Gansler et al., T. 2005. Sociodemographic determinants of cancer treatment health literacy. *Cancer* 104(3):653–660.

Heaivilin et al., N. 2011. Public health surveillance of dental pain via twitter. *J Dent Res*.

Mosavel, M., and El-Shaarawi, N. 2007. ''I have never heard that one'': young girls' knowledge and perception of cervical cancer. *J Health Commun* 12(8):707–719.

Paul, M. J., and Dredze, M. 2011. You are what you tweet: Analyzing twitter for public health. In *ICWSM*. The AAAI Press.

Prier et al., K. W. 2011. Identifying health-related topics on twitter: an exploration of tobacco-related tweets as a test topic. SBP'11, 18–25.

Signorini, A.; Segre, A. M.; and Polgreen, P. M. 2011. The use of Twitter to track levels of disease activity and public concern in the U.S. during the influenza A H1N1 pandemic. *PLoS ONE* 6(5):e19467.

Sullivan et al., S. J. 2011. 'what's happening?' a content analysis of concussion-related traffic on twitter. *Br J Sports Med*.