

Meme Extraction and Tracing in Crisis Events

Saike He¹, Xiaolong Zheng¹, Jiaoqiao Wang¹, Zhijun Chang¹, Yin Luo², Daniel Zeng^{1,3}

¹The State Key Laboratory of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China

²School of Management and Economics, Beijing Institute of Technology, Beijing 100081, China

³University of Chinese Academy of Sciences, Beijing, China

{saike.he, xiaolong.zheng, jiaoqiao.wang, zhijun.chang, dajun.zeng}@ia.ac.cn, vflyvv@bit.edu.cn

Abstract—The proliferation of social media has increased the competition among different memes, which can be free texts, trending catchphrases, or micro media. As human attention is limited, these memes compete with each other, and go in and out of popularity at a rapid pace, sometimes even faster than we can recognize. Popular memes often shape the mindsets of online communities, and also shed light on their future tendencies. Considering the huge volume of memes generated and their continuous mutations, extracting and tracing online memes automatically is rather challenging. In this paper, we propose an automatic meme extraction algorithm. The proposed algorithm extracts massive memes based on phrases independency, and clusters phrase variants of a single meme efficiently. Evaluation on measles outbreak in the USA in 2015 indicates that the proposed algorithm could extract typical memes reflecting the fierce campaign between the pro-vaccination community and the anti-vaccination community. In both communities, memes are power-law distributed, and popular ones have many variants that appear more frequently. By tracing the evolution of online memes, we uncover that popular memes converge and generate peaks at times. Though the pro-vaccination community and the anti-vaccination community may focus on similar memes, they comprehend memes from totally different perspectives and deliver opposing opinions of measles vaccination.

Keywords—meme extraction; meme tracing; accessor variety

I. INTRODUCTION

The advent of social media has lowered the cost of information production and diffusion, boosting the potential reach of each meme to online users. Meme is a concept first coined by Dawkins [1] that refers to a cultural gene passed among person. In social media, it can be free text, trending catchphrase, or micro media. Nowadays, the abundance of online memes is exceeding human capability to consume it. Thus, memes compete with each other, and they go in and out of popularity at a rapid pace, sometimes even faster than we can recognize [2, 3]. During this process, some memes prevail and shapes the mindsets in online communities. These popular memes also shed light on the future tendencies of public sentiments, and inform potential strategies for emergency management in crisis events [4].

Considering the large volume of memes and their continuous mutations [5], extracting and tracing online memes proves to be rather challenging. Existing studies on meme

diffusion can be roughly categorized into two strands. One strand of studies try to investigate how a set of predefined memes evolve and mutate in online communities [6]. These work can help to track the diffusion process of co-evolving memes and measure their effects on online users, yet they are incapable to examine unprecedented or newly emerging memes, thus limiting their scale in real world applications. Another strand of studies try to detect memes automatically. The complexity of these approaches is often extremely high, and the flow chart of phrases mutation is not readily understandable sometimes [5]. Thus, there demands an approach to extract massive memes automatically and efficiently.

In this paper, we propose an efficient algorithm to automatically extract massive memes in real world dataset. This algorithm primarily identifies memes based on a dependency criterion, and then clusters variants to a single meme. Evaluation on measles outbreak in the USA in 2015 indicates that the proposed algorithm could extract typical memes reflecting the fierce campaign between the pro-vaccination community and the anti-vaccination community. In both communities, memes are power-law distributed, and popular ones have many variants that appear more frequently. By tracing the evolution of online memes, we uncover that popular memes converge and generate peaks at times. Though the pro-vaccination community and the anti-vaccination community may focus on similar memes, they comprehend memes from totally different perspectives and deliver opposing opinions of measles vaccination. Findings in this paper have practical significance for those who design and maintain online communities, especially in crisis events.

II. MEME EXTRACTION

In this section, we present the algorithm for extracting memes. From an algorithmic perspective, textual variants of distinctive phrases act as “genetic signatures” for different memes. In analogy to genetic signatures, these phrases undergo continuously mutation while keep majorly constant. Thus, meme extraction ends up to two subtasks: (a) independent phrase identification, and (b) clustering of mutational variants of the phrase (Fig. 1). These two subtasks will be described in turn in the following subsections.

This work was supported in part by the following grants: the National Natural Science Foundation of China under Grant Nos. 71402177, 71602184, 71603253, 71472175, 61172106, 91546112, 91224008, 61671450, 71621002, and 71103180, the National Major Research Program of China under Grant No. 2016YFC1200702, the Early Career Development Award of SKLMCCS under Grant Nos. Y6S9011F4F and Y3S9021F37.

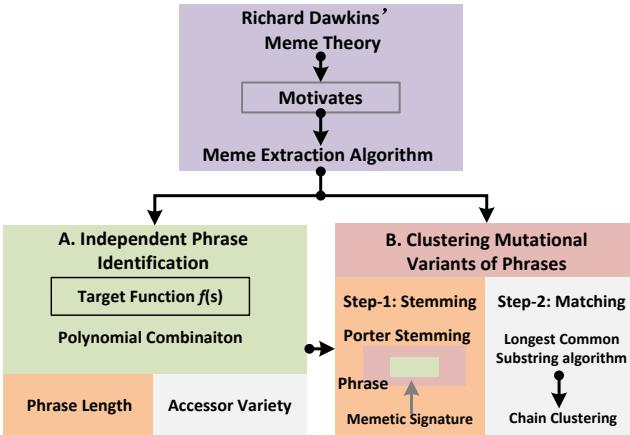


Fig. 1. Illustration of meme extraction algorithm.

A. Independent Phrase Identification

Identifying independent phrases in massive dataset proves to be a challenging task since new phrases (e.g. slangs and neologisms [7]) continuously emerge and die out of popularity beyond human recognition. To tackle with this issue, we propose a scalable algorithm to extract independent phrases automatically. Specifically, we try to extract phrases based on its independency of the context. To quantify the independency degree, we calculate the *accessor variety* [8] value of a string. This criterion is a straightforward formulation of Harris's motivation in determining a morph in terms of the number of contexts in which it can appear: the larger the number is, the more likely it is a morph. For a succinct illustration, we will use an example to describe the calculation of accessor variety. Suppose we have the four free texts:

- **Text A:** Nearly all states allow *religious exemptions* for vaccinations.
- **Text B:** California would remove personal *religious exemptions* for vaccination.
- **Text C:** Hundreds of schools have high rates of *religious exemptions* from vaccinations.
- **Text D:** *Religious exemptions* to vaccination endanger us all.

In the above four texts, the string “*religious exemptions*” has four distinct prefixes, i.e., “allow”, “personal”, “of”, “S” (“S” denotes the start of a text), and three distinct suffixes, i.e., “for”, “from”, “to”. This means the string “*religious exemptions*” can be used in at least three different contexts, and thus might carry some semantic meanings independent from the rest of the words in the four texts. In this case, $\text{three} = \min\{\text{three, four}\}$ is the accessor variety value of the string “*religious exemptions*”.

In phrase extraction, we use the criterion accessor variety (AV for short) to evaluate how independent a string is in a corpus. The formal definition of accessor variety is given by:

$$AV(s) = \min\{L_{av}(s), R_{av}(s)\} \quad (1)$$

where, $L_{av}(s)$ is the *left accessor variety*, and is defined as the number of distinct words (predecessors) excepts “S” that precede s plus the number of distinct texts in which s appears at the beginning; while $R_{av}(s)$ is the *right accessor variety*, and is defined as the number of distinct words (successors) except “E” that succeed s plus the number of distinct texts in which s appears at the end.

Under this criterion, the words “S” and “E” will be counted repeatedly to account for the phrases that usually appear at the beginning or the end of a text, yet seldomly in the middle of a text. Such phrases include “what’s worse”, “In addition”, and more. Though there might be very few words preceding or succeeding them, these phrases can be readily identified as independent meaningful phrases.

Before using accessor variety directly to extract independent phrases, we should consider another important factor, i.e., the length of a phrase. Without this factor, AV is prone to extract shorter phrases since the AV value of a shorter phrase is generally higher than that of a longer one that contains or partly contains it (e.g., “religious exemptions” versus “religious exemptions for vaccinations”). To guarantee that longer yet meaningful phrases are extracted, we have designed a target function $f(s)$ to balance the influence of phrase length and its AV value.

Essentially, $f(s)$ can be designed in many types, representative functions include: polynomial, exponential, logarithm, and a combination of them. From the perspective of computation efficiency and effectiveness, we design $f(s)$ with polynomial functions:

$$f(s) = |s|^a \times AV(s)^b \quad (2)$$

where, $|s|$ and $AV(s)$ are respectively the length (in terms of word number) and the AV value of phrase s ; a and b are integer parameters that define the target function. Theoretically, a and b can be assigned with any possible numbers. In the experiment section, we set them by fine-tune the extraction results.

B. Clustering Mutational Variants of Phrases

To make our extraction algorithm scale up to massive texts, instead of employing traditional top-down or bottom-up hierarchical approaches [9], we design the phrase clustering procedure in a two-step manner, i.e., stemming and matching. In the first step (stemming), phrases variants are stemmed to retain their essential meaning. In the second step (matching), stemmed phrases are linked together according to their identical meaning. These two steps are detailed below.

Step-1: Stemming. The stemming approach is designed based on the Porter stemming algorithm [10]. Primarily, we divide each phrase into a word list. Then for each word, we remove its commoner morphological and inflexional endings. After this process, each phrase is truncated to its memetic signature.

According to the porter stemmer, we use a variable m to measure the length of a word or part of a word. Specifically, we use character C to denote a sequence of one or more consonants, and use character V to denote a sequence of one or more vowels. Then, each word can be represented as:

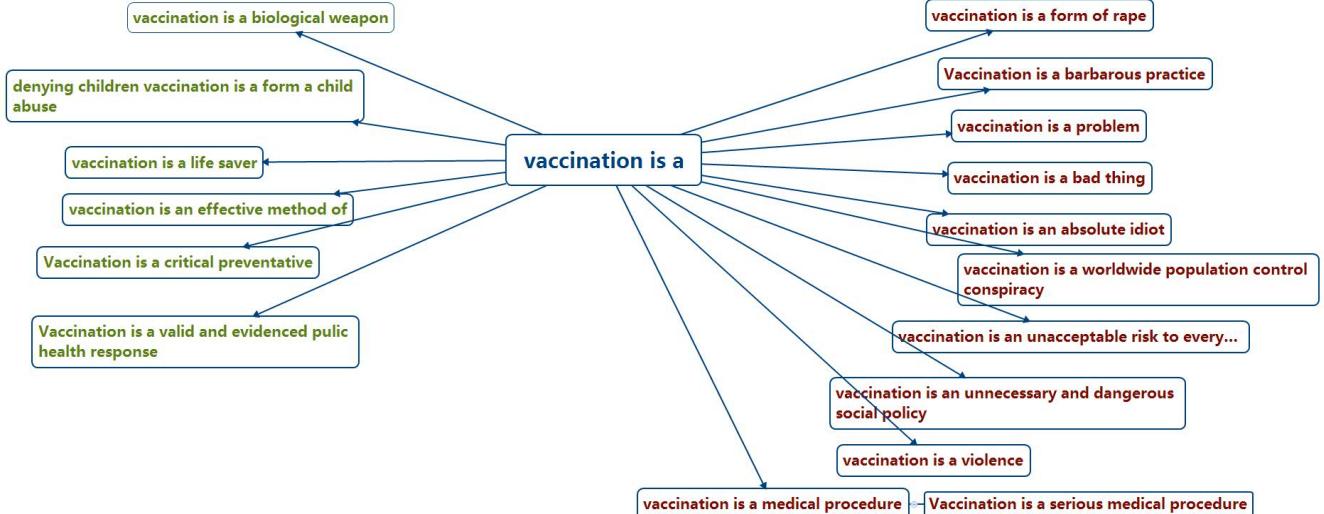


Fig. 2. A small portion of the clustering of variants of phrase “vaccination is a”. Positive phrases are colored in green, while negative phrases are colored in red.

where $[]$ means the object enclosed is optional. Take “*religious*” for example, its m value is 2, as calculated below:

$$\begin{array}{c}
 \text{r} \quad \text{e} \quad \text{l} \quad \text{i} \quad \text{g} \quad \text{i} \quad \text{o} \quad \text{u} \quad \text{s} \\
 | \quad | \quad | \quad | \quad | \quad | \\
 [\text{C}] \text{ V } \text{ C } \text{ V } \text{ C } [\text{V}] \\
 1 \quad 2
 \end{array} \tag{4}$$

The suffix “ous” will be removed according to rules based on the value of m , i.e., removing “ous” if $m > 0$. After the stemming process, “*religious*” changes to “*religi*”. These stems, derived from the procedure are not necessarily valid words, but rather common sub-strings. We consider phrases combined with such stems already manifest enough meanings, and can be used to trace meme diffusion, as will be shown later in the *Results* section.

Step-2: Matching. In order to cluster mutational variants of phrases together, we need to match phrases with similar stems together. The matching strategy is based on the Longest Common Substring (LCS) algorithm [11]. To guarantee high recall rate of all phrase variants while avoid missing matching, we cluster phrases based on the relative length LDS, i.e., phrases with the length of LDS proceeds a predefined threshold will be clustered. To accelerate clustering procedure among huge amount of candidate phrases, we utilize chain clustering procedure [12]. Finally, for each clustering, we choose the phrase with the highest AV value to represent the cluster. Fig. 2 present an example of clustering mutational variants of phrase “vaccination is a”.

III. EXPERIMENTAL SETUP

A. Data Description

In 2015, the United States experienced a large, multi-state measles outbreak. Totally, 189 people from 24 states and the District of Columbia were reported to have measles. This outbreak has incurred fierce discussion on Twitter. To study

how memes manifested different emotions spread on Twitter, we have collected all tweets that contain any one of the keywords listed in Table I¹. The relevance of the collected data was checked manually. The time span of the collected data is from Jan. 1st 2015 to Aug. 1st 2015. Totally, we have collected 206,688 tweets written by 100,987 Twitter users. After data collection, SentiStrength [13] is employed to label tweets as positive, negative, or neutral with respect to the intent of getting vaccinated against measles. Statistics of the dataset is shown in Table II.

TABLE I. KEYWORDS USED FOR DATA COLLECTION.

Category	Keywords
Measles	measle, morbilli, rubeola, rubeolla
vaccination	vacinate, vaccinate, vaccination

TABLE II. DATA STATISTICS.

Sentiments	Tweet Nr.	Retweet Nr.	Reply Nr.	Total
POS	24481	4312	2260	31053
NEU	88933	16868	6311	112112
NEG	51949	7367	4206	63522
Total	165363	28547	12777	206687

B. Settings

In the following experiments, we set $a=1$, $b=2$ in the target function defined in (2). The threshold for clustering is set as 0.7, which means phrase stems having similarity higher than 0.7 is clustered together. In the following experiments, we mainly focus on the data samples classified as positive or negative, and trace how memes belong to these two extremes evolve and compete with each other.

¹ Keywords are selected according to CHV Wiki:
<http://consumerhealthvocab.chpc.utah.edu/CHVwiki/>

IV. RESULTS

A. Meme Distribution

Fig. 3 shows the complementary cumulative distribution function (CCDF) of meme volume. For each volume x , we plot the number of memes with volume $\geq x$. If the meme volume is power-law distributed with exponent α , $p(x) \sim x^{-\alpha}$, then when plotted on log-log axes the CCDF will be a straight line with slope $-\alpha$. In Fig. 3, we superimpose two kinds of memes, i.e., positive memes (blue circles) and negative memes (green circles). We find that the two kinds of memes are all power-law distributed with some cut-offs. In addition, the volume of positive memes and negative memes respectively decay with exponent $\alpha_{POS} = 1.187$ and $\alpha_{NEG} = 1.018$. This fitting is based on the maximum likelihood approach proposed by Clauset et al. [14]. We find that the tails are heavy as the power-law distributions start to have finite variances only when $\alpha > 3$ [5]. In fact, for $\alpha < 2$ powerlaws have infinite expectations. This means that popular memes have many variants and each of them appears more frequently than an “average” phrase.

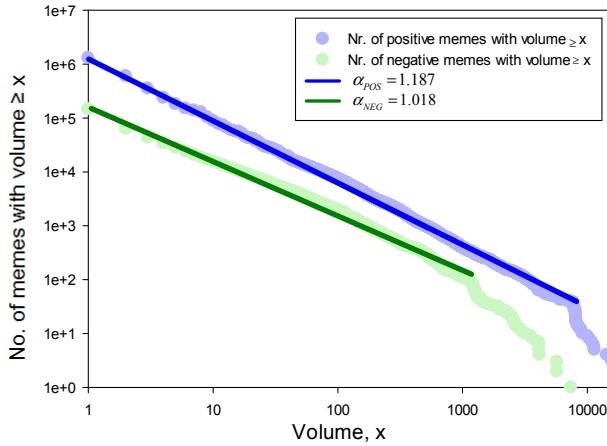


Fig. 3. Phrase volume distribution.

B. Meme Clouds

Fig. 4 shows the clouds of memes manifesting positive sentiments and negative sentiments respectively. The clouds indicate that distinctive sentiments drive different meme diffusions in online communities. For the positive memes (Fig. 4-a), the most widely spread memes are those related to children vaccination against measles, including “childhood vaccination”, “vaccinate their children”, “please vaccinate your”, et al. Positive memes also suggest that the benefits of vaccination is discussed and propagandized in online communities, related memes include “protect your”, “MMR vaccine”, and “vaccines save”. As a result, some online users concern the necessity of mandatory vaccination, as reflected by the memes of “mandatory vaccination” and “vaccination

obligatoire”. In addition, online users who support vaccination also mention Jimmy Kimmel (a doctor promoting vaccination) to deliver vaccination mandatory.

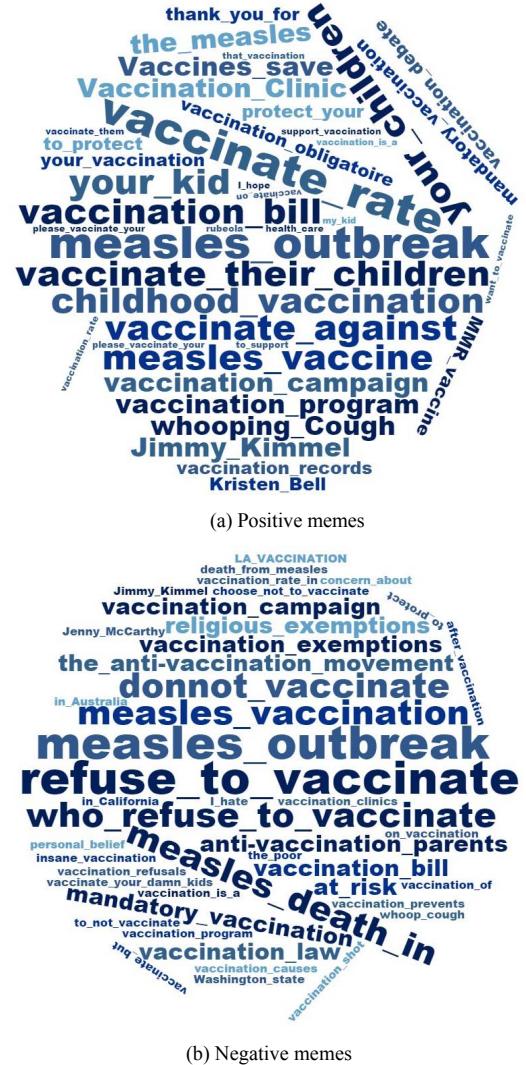


Fig. 4. Clouds of (a) positive memes and (b) negative memes. The size of each meme corresponds to its number of variant phrases.

While for the negative memes (Fig. 4-b), the frequently mentioned argument is anti-vaccination, such as “refuse to vaccinate”, “donnot vaccinate”, “vaccination refusals” and “the anti-vaccination movement”. Online anti-vaccine activist also express extreme sentiments, such as “i hate” and “vaccinate your damn kid”. Some also make inappropriate analogies, such as “vaccination is a form of rape”, “vaccination is a barbarous practice”, “vaccination is an absolute idiot” and “a vaccination is a violence”. In order to exempt from vaccination, online users cite Jenny McCarthy as an example, who is an American model, television host, and most importantly, anti-vaccine activist. Besides, they also use “personal beliefs” or “religious exemptions” in excuse of

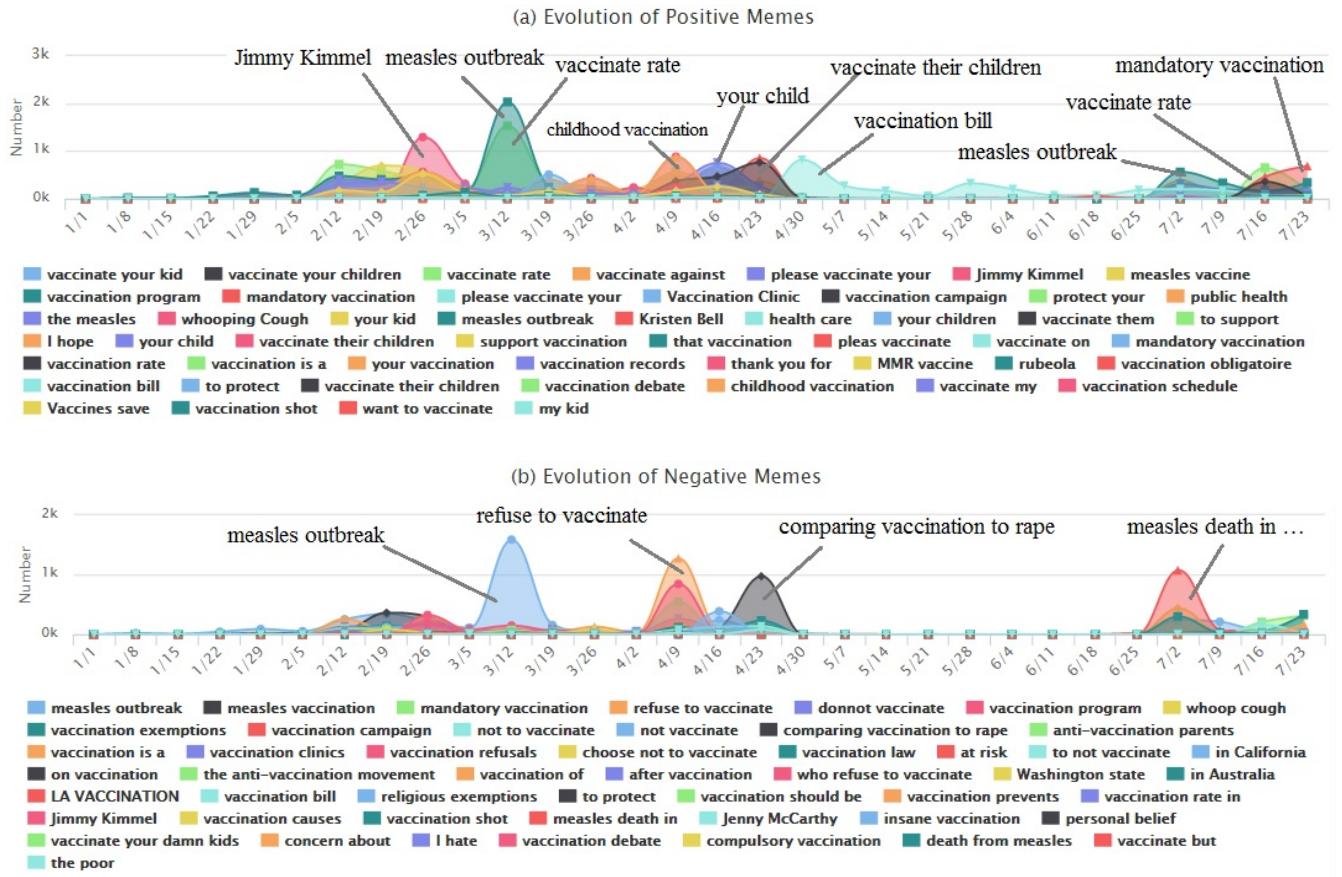


Fig. 5. Evolution of top 50 memes. (a). Evolution of positive memes; (b). Evolution of negative memes.

vaccination exemptions. These typical mindsets are accurately captured by our meme extraction algorithm.

C. Meme Tracing

Having extracted memetic phrases, we now track how different memes evolve over time, considering both their individual temporal dynamics as well as their interactions with one another. Using our clustering approach, we automatically created and labeled the plot in Fig. 5, which presents the 50 largest memes for the period of Jan. 1st 2015 – Jul. 23rd 2015. It is drawn as a stacked plot, in which the thickness of each strand corresponds to the volume of the memetic phrases over time. The rising and falling patterns could tell the tendency by which online users successively focus on and defocus on specific memes. The periods when the upper envelope of the curve are high correspond to times when there is a convergence on popular memes, while the low periods mean that meme popularity is low among online users.

During the evolution of positive memes, there are several peaks (Fig. 5-a). The first peak occurs after Feb. 5th, 2015, when an American doctor, Jimmy Kimmel, began to promote vaccination against measles, as indicated by the meme “Jimmy Kimmel”. The second peak comes with two crests, one for “measles outbreak”, and another for “vaccinate rate”. The prevalence of these two memes indicates that online users supporting vaccination attributed this measles outbreak to the

low vaccination rate. As for the third peak, several memes caught public attentions, with the theme of childhood vaccination. Related memes include “childhood vaccination”, “your child”, “vaccinate their children”, and “vaccination bill”. Among these memes, “vaccination bill” prevailed for a long time, from Apr. 23rd, 2015 to Jun. 25th, 2015. The last peak comes with the reported one death due to measles on Jul. 2nd, 2015. After this death, online users began to concern about the relation between “measles outbreak” and “vaccinate rate”, some of them even promoted “mandatory vaccination”, as shown in Fig. 5-a from Jun. 25th, 2015 to Jul. 23rd, 2015.

As for the evolution of negative memes, there are mainly four peaks (Fig. 5-b). The first peak occurs when online user focused on the meme of “measles outbreak”, which also focused by those supporting vaccination (Fig. 5-a). However, anti-vaccination activists comprehended this outbreak from a totally different perspective, as reflected by the second peak “refuse to vaccinate”. In the nearest peak that follows, some users also compared vaccination to rape (as reflected by the dominated meme “comparing vaccination to rape”). In contrast to this radical comparison, positive memes at the same time period delivered persuasion about “childhood vaccination” (as shown in Fig. 5-a). This direct contrast suggests that online users hold two extreme sentiments toward measles vaccination, supervising and tracing their campaign may benefit public health. The last peak nearly concurred with the rising of

positive memes, i.e., after with the reported one death due to measles on Jul. 2nd, 2015. This shock event did not greatly affect those refusing vaccination, since they just mentioned this sad news by spreading the meme “measles death in ...” without commenting too much.

V. RELATED WORK

The original work regarding meme traces back to a theory proposed by Dawkin [1], who first coined the concept of meme. This concept is utilized to describe the potential process of information diffusion among online users, in analogy with gene in genetics.

Existing studies focuses on meme diffusion can be roughly categorized into two strands. One strand of studies try to investigate how a set of predefined memes evolve and mutate in online communities [6, 15]. These studies usually propose various diffusion models to depict meme dynamics. One major drawback of these models is that they could only characterize predefined memes, yet can do nothing about newly emerging memes. This drawback severely limits the wide scale of the models’ applications.

Another strand of studies focuses on developing algorithms that detect memes automatically. These studies are mainly based on graph model [5] or network analysis [16]. The computational complexity of these approaches are often comparatively high, thus could not be directly applied to analyze real world datasets.

Unlike these previous studies, the algorithm proposed in this paper could extract massive memes automatically and efficiently.

VI. CONCLUSIONS AND FUTURE WORK

We have proposed an efficient and automatic algorithm for meme extraction. Evaluation on measles outbreak in the USA in 2015 indicates that the proposed algorithm could extract typical memes reflecting the fierce campaign between the pro-vaccination community and the anti-vaccination community. By tracing the evolution of online memes, we uncover that popular memes converge and generate peaks at times.

In our future work, we plan to study how meme competition affects the popularity of different memes and the diversity of information we are exposed to.

REFERENCES

- [1] R. Dawkins, *The selfish gene*: Oxford university press, 2006.
- [2] L. Weng, A. Flammini, A. Vespignani, and F. Menczer, “Competition among memes in a world with limited attention,” *Scientific reports*, vol. 2, 2012.
- [3] S. He, X. Zheng, and D. Zeng, “A model-free scheme for meme ranking in social media,” *Decision Support Systems*, vol. 81, pp. 1-11, 2016.
- [4] S. He, X. Zheng, X. Bao, H. Ma, D. Zeng, B. Xu, G. Tian, and H. Hao, “Ranking Online Memes in Emergency Events Based on Transfer Entropy,” in Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint, 2014, pp. 236-239.
- [5] J. Leskovec, L. Backstrom, and J. Kleinberg, “Meme-tracking and the dynamics of the news cycle,” in Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, 2009, pp. 497-506.
- [6] X. Wei, N. C. Valler, B. A. Prakash, I. Neamtiu, M. Faloutsos, and C. Faloutsos, “Competing memes propagation on networks: A network science perspective,” *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 6, pp. 1049-1060, 2013.
- [7] X. Lou, Y. Chai, H. Zan, R. Xu, Y. Han, and K. Zhang, "Research on micro-blog sentiment analysis," *Chinese Lexical Semantics*, pp. 466-479: Springer, 2013.
- [8] Z. S. Harris, *Papers in structural and transformational linguistics*: Springer, 2013.
- [9] J. Cong, and M. L. Smith, “A parallel bottom-up clustering algorithm with applications to circuit partitioning in VLSI design,” in Proceedings of the 30th international Design Automation Conference, 1993, pp. 755-760.
- [10] M. F. Porter, “An algorithm for suffix stripping,” *Program*, vol. 14, no. 3, pp. 130-137, 1980.
- [11] L. Bergroth, H. Hakonen, and T. Raita, “A survey of longest common subsequence algorithms,” in String Processing and Information Retrieval, 2000. SPIRE 2000. Proceedings. Seventh International Symposium on, 2000, pp. 39-48.
- [12] S. He, X. Zheng, D. Zeng, C. Luo, and Z. Zhang, “Exploring Entrainment Patterns of Human Emotion in Social Media,” *PloS one*, vol. 11, no. 3, pp. e0150630, 2016.
- [13] M. Thelwall, K. Buckley, G. Paltoglou, D. Cai, and A. Kappas, “Sentiment strength detection in short informal text,” *Journal of the American Society for Information Science and Technology*, vol. 61, no. 12, pp. 2544-2558, 2010.
- [14] A. Clauset, C. R. Shalizi, and M. E. Newman, “Power-law distributions in empirical data,” *SIAM review*, vol. 51, no. 4, pp. 661-703, 2009.
- [15] J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, S. Patil, A. Flammini, and F. Menczer, “Truthy: mapping the spread of astroturf in microblog streams,” in Proceedings of the 20th international conference companion on World wide web, 2011, pp. 249-252.
- [16] J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, A. Flammini, and F. Menczer, “Detecting and Tracking Political Abuse in Social Media,” in ICWSM, 2011.