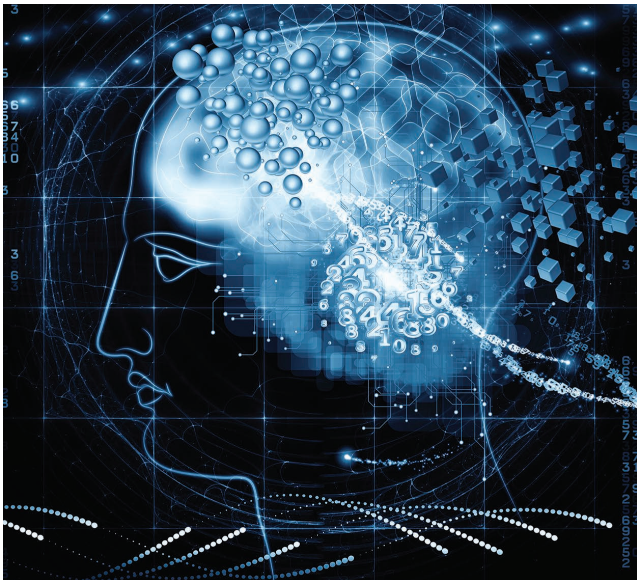

Misinformation Propagation in the Age of Twitter

Fang Jin, Wei Wang, Liang Zhao, Edward Dougherty, Yang Cao, Chang-Tien Lu, and Naren Ramakrishnan, *Virginia Tech*

A quantitative analysis of tweets during the Ebola crisis reveals that lies, half-truths, and rumors can spread just like true news.

Although Ebola isn't a new disease, the current outbreak in West Africa is believed to be more than three times worse than all previous ones in history combined. In addition, public health experts fear massive underreporting in the countries with the most widespread transmission—Guinea, Liberia, and Sierra Leone—due to various social considerations. Even syndromic surveillance strategies, such as social media mining and participatory surveillance, aren't effective due to poor Internet penetration and the lack of roads and communication infrastructure where Ebola is most prevalent.

Mark Twain is credited with the aphorism that a lie can travel halfway around the world while the truth is putting on its shoes. As Ebola rages on, another epidemic being talked about is the rapid spread of misinformation on social media about the deadly virus, its origin and impact, and response strategies. Since social media has become one of the primary means by which people learn about worldwide developments, we sought to characterize the dissemination of both news and rumors on Twitter about Ebola with a view to understanding the prevalence of misinformation.

EBOLA IN THE NEWS

Tweets about Ebola peaked in late September through mid-October 2014, when there was extensive reporting on the disease in the US and Europe.

On 30 September 2014, the Centers for Disease Control and Prevention (CDC) confirmed the first importation of Ebola into the US when Thomas Eric Duncan, a Liberian exposed to the virus in Monrovia, traveled to Dallas, Texas, to visit family. On 6 October in Madrid, Spain, Teresa Romero, a nurse's aide who had cared for a missionary infected with Ebola while treating patients in Sierra Leone, was reported to be the first person to have contracted the disease outside of West Africa.

On 8 October, Duncan succumbed to Ebola at Texas Health Presbyterian Hospital in Dallas. A few days later, a healthcare worker at the hospital who had been attending Duncan tested positive for the disease. On 14 October, a second healthcare worker at the hospital reported a low-grade fever and was isolated; she subsequently also tested positive for Ebola.

Many states and cities in the US began making contingency plans and issuing travel advisories and guidelines. Some lawmakers also called for screening passengers and proposed travel bans for Ebola-stricken countries. On 23 October, Craig Spencer, a doctor returning from volunteer work in Guinea, was rushed to Bellevue Hospital Center in New York City with a 100.3°F fever. The following day, when blood tests

Published by the IEEE Computer Society

Table 1. Top 10 Ebola-related rumors by Tweet volume from 28 September to 18 October 2014.

Rumor no.	Content	Label
1	Ebola vaccine only works on white people	White
2	Ebola patients have risen from the dead	Zombie
3	Ebola could be airborne in some cases	Airborne
4	Health officials might inject Ebola patients with lethal substances	Inject
5	There will be no 2016 election and complete anarchy	Vote
6	The US government owns a patent on the Ebola virus	Patent
7	Terrorists will purposely contract Ebola and spread it around	Terrorist
8	The new iPhone 6 is infecting people with Ebola	iPhone
9	There is a suspected Ebola case in Kansas City	Kansas
10	Ebola has been detected in hair extensions	Hair

confirmed he had Ebola, the governors of New York and New Jersey jointly announced automatic quarantines for medical workers returning from Ebola-stricken countries.

RUMORS ON TWITTER

In conjunction with news reports about Ebola, conspiracy theories, innuendo, and rumors about the disease began to propagate wildly on Twitter. We gathered tweets from late September through late October 2014 and filtered them by mention of the keyword "Ebola" or relevant hashtags such as #ebola, #EbolaVirus, #EbolaOutbreak, #EbolaWatch, #EbolaEthics, #EbolaChat, #nursesfightebola, #ebolafacts, #StopEbola, #FightingEbola, and #UHCRevolution.

From the gathered tweets, we removed stopwords for further processing and constructed word clouds for specific days. On 29 September, the day prior to CDC's confirmation that Duncan tested positive for Ebola, the most common words included "Liberia," "Africa," "virus," "outbreak," and "exposed." On 30 September, words specifically related to news about Duncan became prominent: "case," "diagnosed," "patient," "first," "Dallas," "CDC," and so on.

A simple frequency plot of specific keywords in Ebola-related tweets during this period likewise highlighted significant upticks in words such as "Dallas," "Texas," "CDC," "Africa," and "enfermera" (Spanish for "nurse," referring to Romero and possibly other health professionals) after 29 September. By mid-October, mentions of President Obama also noticeably increased, likely due to proposed mitigation and response strategies by the US government.

Next, we studied information cascades in our tweet collection with a view toward identifying misinformation and the spread of falsehoods. We identified several widespread rumors circulating on Twitter, the top 10 of which are shown in Table 1. (Our study focused on rumors in English only.) Most are self-explanatory as to their intent and interpretation.

Two Ebola-related rumors not listed in the table are noteworthy. According to the Snake rumor, which originated at least as early as late summer 2014, the disease crossed the border from Guinea to Sierra Leone via a snake in a bag. The Maldives rumor pertains to an uncorroborated report that Ebola patients had been quarantined in that island nation.

We geocoded tweets spreading these rumors to better understand their geographical scope. The rumor distribution changed over time; for example, in the US, on 1 October the Zombie and Airborne rumors were most prevalent but by 8 October the White and Inject rumors were more common. In African countries like Ghana, Nigeria, and Kenya, the most dominant rumor was the Inject rumor. Other rumors were scattered across other parts of the world.

We next employed a dynamic query expansion model[1] to study the rumors in greater detail. DQE begins with a seed set of keywords (for example, "Ebola," "rumor"), identifies tweets that mention these keywords, and iteratively expands them to identify a larger set of keywords. By conducting a modularity-based optimization over the underlying network of expanded tweets connected by shared keywords, it can identify specific localized instantiations of rumors. As Figure 1 shows, on 29 September 2014—when there was no reported incidence of Ebola in the US—the Zombie rumor was dominant. By 6 October, rumors that Ebola can be airborne and is a potential terrorist weapon gained hold.

Although Figure 1 might suggest that Ebola-related rumors were quite rampant, it's important to keep in mind that they were a small fraction of information propagated on Twitter. We compared the time-indexed

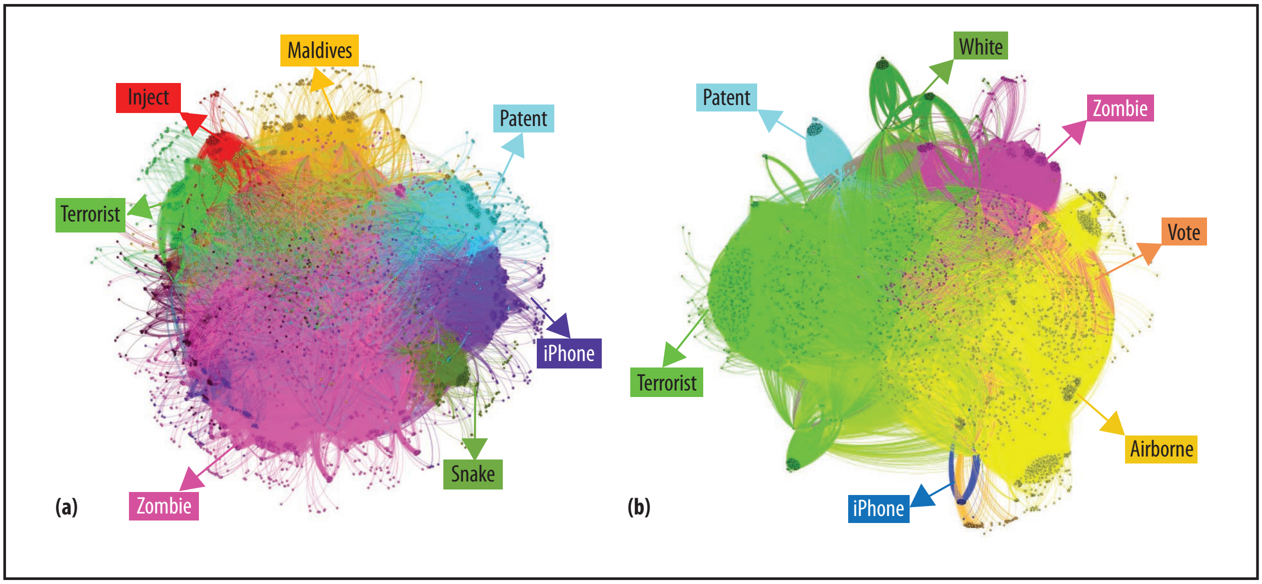

Figure 1. **Clustering of Ebola-related rumors on (a) 29 September 2014 and (b) 6 October 2014. Rumors are color-coded consistently across the two projections.**

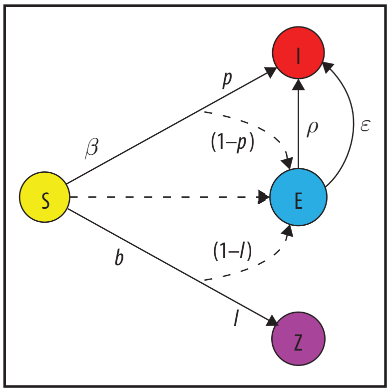

Figure 2. **SEIZ compartmental model applied to information propagation on Twitter: susceptible (S), exposed (E), infected (I), and skeptical (Z) users.**

spread of rumors versus true news stories and found that rumors are more localized, distributed, and comparatively smaller in permeation than news stories.

EPIDEMIOLOGICAL MODELING OF RUMORS

Another way to study the spread of rumors versus news is through epidemiological modeling. Such a model helps capture the likelihood of an individual getting infected with a virus or, in this case, of adopting an idea he or she has been exposed to.

In earlier work,[2] we demonstrated how we can accomplish this objective using the SEIZ compartmental model, which was originally proposed to study the adoption of ideas.[3] This model is particularly suited to studying rumor propagation, as it captures distinctions in how people respond to ideas: whether they adopt it readily or are initially skeptical.

As applied to information propagation on Twitter, the SEIZ model compartmentalizes users into four categories: *susceptible* (S) users haven't received the information; *exposed* (E) users have received the information via Twitter but have delayed tweeting about it; *infected* (I) users have received and tweeted the information; and *skeptical* (Z) users have received the information but chosen not to tweet it. Figure 2 shows the transitions between these states. Note that the model doesn't capture the underlying truth or falsehood of the information but simply whether users readily accept or adopt it.

We fitted SEIZ to the different rumors and plotted the time-course data for each state variable. As Figure 3a shows, the model can capture quasilinear (Patent), sigmoidal (White), and other nonlinear information-spread patterns (Zombies and Airborne). Time-course results, shown in Figure 3b, depict broadly similar patterns: S initially has a value close to the total Twitter population size (N) and then rapidly decreases with a comparable rapid increase in Z, and I gradually increases as E gradually decreases.

The exception to this pattern is the Patent time-course profile. In this case, S begins with less than half of the total population size and is only slightly higher than the initial Z and E values. In addition, Z is essentially constant, meaning that the number of skeptics doesn't change throughout propagation. Finally, the decrease in S doesn't correspond to a change in Z as in the other rumor profiles; rather, the drop in S is accompanied by a nearly identical increase in E.

In our earlier work on characterizing rumors,[2] we defined the notion of a *response ratio* that provides a relative measure of the population influx into E versus the efflux from

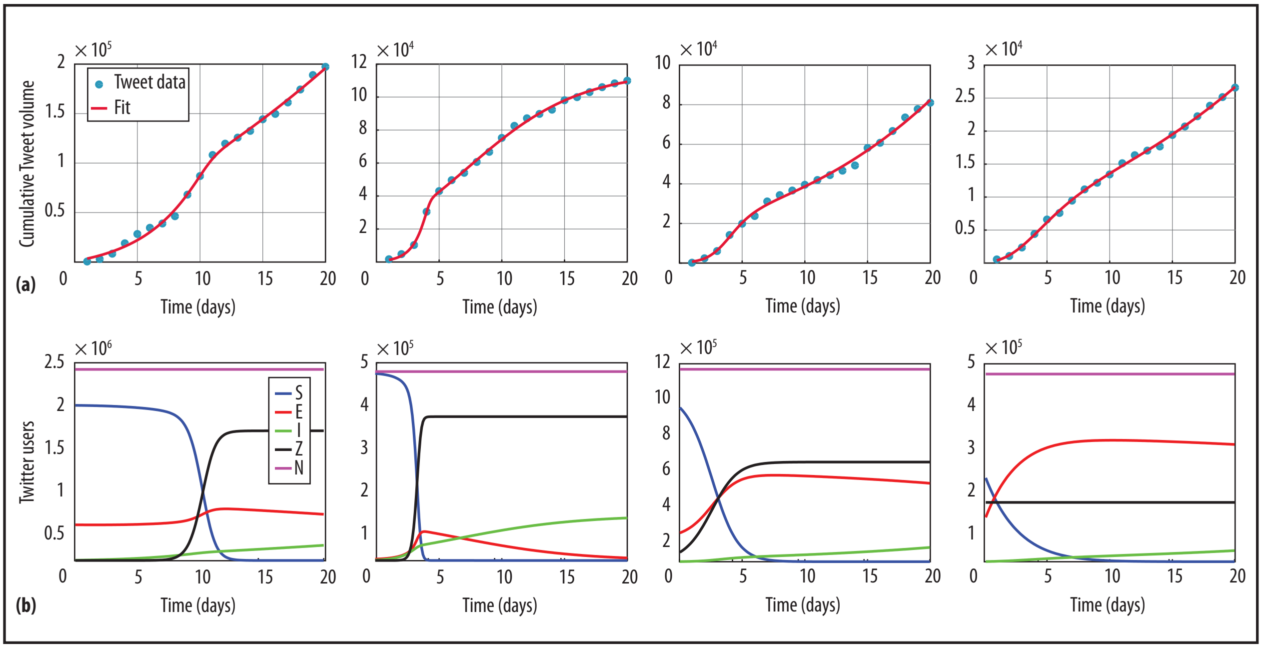

Figure 3. Model fits of SEIZ to (a) different rumors and (b) time-course data for each state variable (N denotes the total number of Twitter users). From left to right: White, Zombie, Airborne, and Patent rumors.

this compartment. The Patent time-course findings suggest that a large influx of users into E without a corresponding efflux to I, combined with a stagnant Z group, will produce an elevated response ratio—in other words, a large exposure to the rumor without significant change in skepticism.

We hypothesized that the response ratio could be a useful factor in discriminating rumors from true news, with higher response ratios associated with people's greater belief in the underlying information. To compare response ratios across rumors and news, we selected three breaking stories pertaining to Ebola: the story of the first Ebola patient identified in the US (Dallas), the first confirmation of an Ebola patient in New York City (NYC), and the symptoms and travel activities of Doctor Spencer in the days before he was diagnosed. As Figure 4 shows, all three news stories have response ratios higher than 25, with a mean of approximately 38, while 8 of the 10 rumors have a response ratio less than or equal to 6.4, with a mean of only 3.33. Patent and Airborne have elevated response values, suggesting greater belief in these rumors than in the other eight.

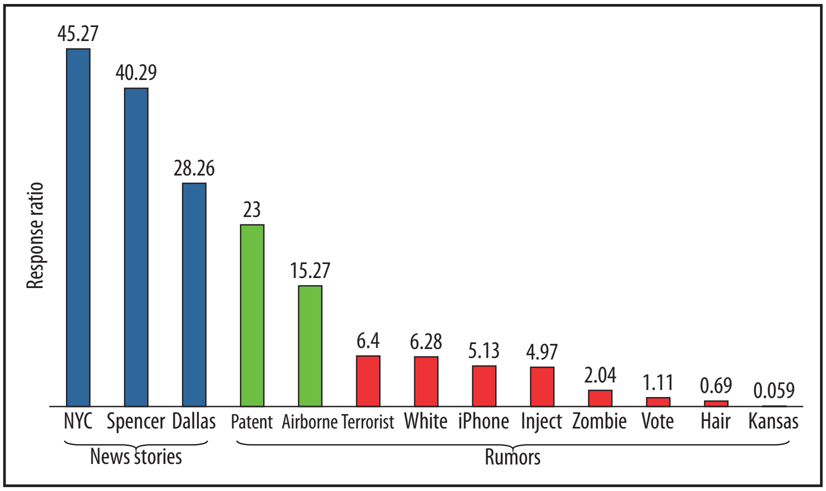

Figure 4. Response ratios for 3 news stories and 10 rumors related to Ebola.

The propagation of misinformation on Twitter can sometimes resemble that of genuine newsworthy events. Given that many consumers now receive news from real-time social media platforms, it's important to have quantitative methods to distinguish news from rumors.

References

1. L. Zhao et al., "Unsupervised Spatial Event Detection in Targeted Domains with Applications to Civil Unrest Modeling," *PLOS ONE*, vol. 9, no. 10, 2014; doi:10.1371/journal.pone.0110206.

2. F. Jin et al., "Epidemiological Modeling of News and Rumors on Twitter," *Proc. 7th ACM SIGKDD Workshop Social Network Mining and Analysis* (SNAKDD 13), 2013; doi:10.1145/2501025.2501027
3. L.M.A. Bettencourt et al., "The Power of a Good Idea: Quantitative Modeling of the Spread of Ideas from Epidemiological Models," *Physica A: Statistical Mechanics and its Applications*, vol. 364, 2006, pp. 513–536.

Fang Jin *is a PhD student in the Department of Computer Science at Virginia Tech. Contact her at jfang8@vt.edu.*

Wei Wang *is a PhD student in the Department of Computer Science at Virginia Tech. Contact him at tskatom@vt.edu.*

Liang Zhao *is a PhD student in the Department of Computer Science at Virginia Tech. Contact him at liangz8@vt.edu.*

Edward Dougherty *is a PhD student in the Genetics, Bioinformatics, and Computational Biology (GBCB) program at Virginia Tech. Contact him at edougherty@vt.edu.*

Yang Cao *is an associate professor in the Department of Computer Science at Virginia Tech. Contact him at ycao@cs.vt.edu.*

Chang-Tien Lu *is an associate professor in the Department of Computer Science at Virginia Tech. Contact him at ctlu@vt.edu.*

Naren Ramakrishnan, *Discovery Analytics column editor, is the Thomas L. Phillips Professor of Engineering at Virginia Tech and director of the university's Discovery Analytics Center. Contact him at naren @cs.vt.edu.*

DISCLAIMER

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via DoI/NBC contract number D12PC000337, the US Government is authorized to reproduce and distribute reprints of this work for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the US Government.