

# The infamous #Pizzagate conspiracy theory: Insight from a TwitterTrails investigation

Panagiotis Metaxas\*

Samantha Finn†

## 1 Introduction

Social media have become part of modern news reporting, used by journalists to both inform and find sources, and by individuals as a news source. The quest for prominence and recognition on social media sites, like Twitter and Facebook, can sometimes eclipse accuracy and lead to the spread of false information. In the last US Congressional elections we saw the rise of the so-called “fake news” phenomenon, a euphemism for “lies in the shape of news articles”, aiming to confuse and anger voters. And this is not just happening in the US: during the recent elections in several European countries including Germany and France, we saw a similar pattern.

As a way of studying and reacting to this trend, we are using TwitterTrails, an interactive, web-based investigative tool ([twittertrails.com](http://twittertrails.com)) that allows users to investigate the origin and propagation characteristics of a rumor and its refutation, if any, on Twitter [1]. Visualizations of burst activity, spreading timeline, and co-retweeted networks help its users investigate the spread of a story. Within minutes TwitterTrails will collect relevant tweets and automatically answer several important questions regarding a rumor: its originator, burst characteristics, propagators and main actors according to the audience. In addition, it will compute and report the rumor’s level of visibility and, as an example of the power of crowdsourcing, the audience’s skepticism towards it, which correlates with the rumor’s credibility.

We envision TwitterTrails as valuable tool for journalists (amateur and professional) investigating recent and breaking stories. Further, its expanding collection of investigated rumors can be used to answer questions regarding the amount and success of misinformation on Twitter [2]. This paper explains how this can be accomplished using as an example a well-known story

**On “Fake News”.** The term “fake news” refers to lies presented as news, that is, falsehoods online formatted and circulated in such a way that a reader might mistake them for legitimate news articles. However, there is some confusion around the term as it has been used more loosely to mean different things to different people. Some use it to denote opinions with which they disagree, reporting errors, or incorrect predictions, though these are not technically fake news.

“Fake news” has been around since ancient times, but technology has made it possible to produce and consume it today on a massive scale. Such articles appear on a variety of little

---

\*Computer Science, Wellesley College; Center for Research on Computation and Society, Harvard University

†Library and Technology Services, Wellesley College

known websites, then turn a profit by competing for clicks as advertisements on social media sites. In order to be successful in attracting user attention, they present a fake story of political nature, religious nature, or anything with strong emotional appeal. Typically, fake news stories are planted on social media sites using provocative titles and images. “Clickbait” attracts the attention of unsuspecting social media users who click on links to these stories thinking they are visiting a legitimate news site. These engaged users are drawn in by the emotional appeal, while fake news providers get a share of advertising money from each click.

These made-up stories, which may or may not have some remote connection to reality, are a form of propaganda, aiming to trick readers into behaving in ways beneficial to the fake news provider, who is essentially a propagandist. The benefit may be political (e.g., persuading readers to vote as the propagandist wants), financial (e.g., persuading readers to click on advertisements and bring money to the propagandist), religious (e.g., persuading readers that a particular religion is good or bad), entertaining (e.g., persuading readers to spread a joke and show how gullible people are), etc. While thousands of these stories exist, the vast majority do not succeed in getting widespread attention. Those that are successful, however, spread for a variety of reasons. One of the main reasons is that most readers are familiar with a historically valid model of trusting news, which is typically edited and printed by some authoritative source to the Web. And these days, search engines and social media algorithms show results that can make anyone’s opinion look equally authoritative.

While we can often identify and avoid being tricked by fake news, the unfortunate fact is that any of us could fall for one of these lies, especially if it is presented in a way that matches our biases and prior beliefs. In order to recognize fake news, diversity is key. It can be easier to recognize fake news if a group of diverse people, with a broad variety of individual biases (and therefore tendencies to believe or be skeptical about different stories), engages with the information together. In contrast, members of a homogeneous group (an “echo chamber”) are easily fooled when presented with lies that conform to their common biases.

Unfortunately, people tend to form echo chambers in social media and in society. We find comfort and safety with others who are similar to ourselves. And when we are presented with evidence that our beliefs are incorrect, we try hard to avoid challenging our belief system. This is when we are most susceptible to lies; when they are presented in a way that confirms our prior beliefs.

## 2 TwitterTrails investigation of #Pizzagate

TwitterTrails is a system designed to help us investigate the spread and skepticism of particular rumors or memes on Twitter. It is initiated by a keyword search and produces automatically, within minutes, a report that answers several key questions regarding the spreading of the rumor, such as: who started the rumor on Twitter, who made it widely known, who are the main groups participating, who are the main actors in its spreading, and whether there is any skepticism in the spreading or does the rumor spread in an echo chamber.

We describe here a particular investigation of a rumor that has gotten great attention in the media, the spreading of the #Pizzagate rumor. Searching the Web one finds hundreds of thousands of news articles and analyses related to this discredited conspiracy theory<sup>1</sup>, and it is a reasonable question to ask whether TwitterTrails can add anything to what we already know. In fact, we demonstrate that TwitterTrails can offer facts and insights that were not previously known.

---

<sup>1</sup>See, e.g., related articles on the New York Times or snopes

TwitterTrails’ investigation of #PizzaGate can be found in: <http://bit.ly/TTpizzagate> an online interactive page. Interested readers are encouraged to visit the page and interact with it on their own. Beware that the data set is large and it may take a couple minutes to load on a computer. Examining the page reveals several points worth noting. Here are a few of them:

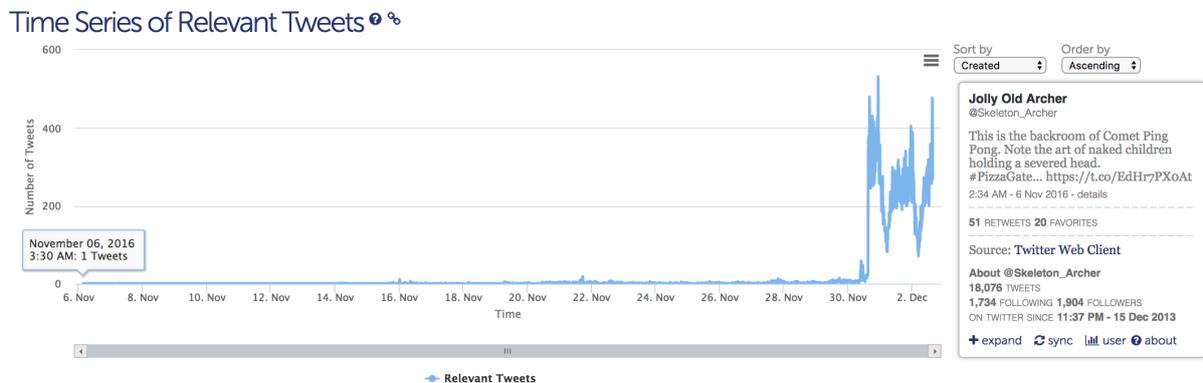


Figure 1: The timeline reveals that #Pizzagate was a rumor that had not gotten attention on Twitter for much of November, 2016, until groups formed an echo chamber to discuss it.

1. It has been widely reported in the news (see, e.g., “Pizzagate: From rumor, to hashtag, to gunfire in D.C.”, Dec. 6, 2016, The Washington Post, <http://wapo.st/2h9EXgi>) that the hashtag #PizzaGate appeared on Twitter on Nov. 7, 2016, a day before the US Presidential elections. There is no mention of who first used the hashtag. TwitterTrails data reveals that the hashtag appeared earlier than that (Nov. 6, 2016 at 3:30 AM ET), and was created by a trolling account that is still present on Twitter but has not tweeted since Feb. 13, 2017. See Figure 1.

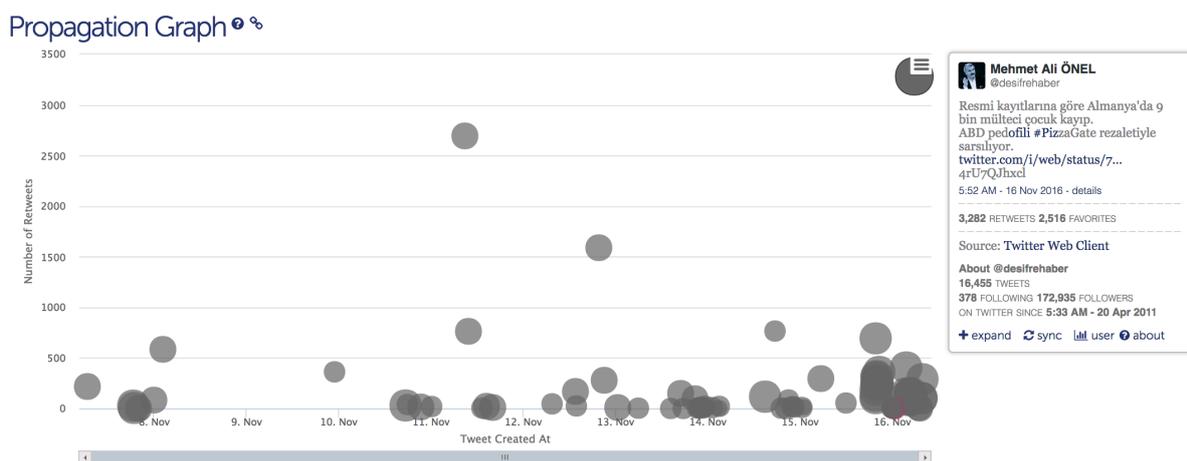


Figure 2: The propagation graph shows who and when the story “broke” on Twitter. In the upper right corner the tweet that received the most retweets early on appears as a (partially covered) gray circle. Moving the cursor over the circle shows the actual tweet on the right. Under this tweet is a barrage of tweets sent a few hours earlier by a troll, “informing” and provoking the Turkish journalist.

2. It has been reported that Turkish journalists promoted #PizzaGate (see, e.g., “The saga of ‘Pizzagate’: The fake story that shows how conspiracy theories spread”, BBC Trending, Dec. 2, 2016: <http://bbc.in/2v5wLmX>) but no direct evidence has been given on how large was that role. Our data identify a Turkish journalist that played a major role in the internationalization of the rumor earlier than the news articles reported. Importantly, TwitterTrails point out who prompted him about the rumor. He was prompted by a bilingual troll that bombarded with at least 118 tweets the Turkish Twittersphere with conspiracy theories a few hours before the journalist picked it up. See This trolling account was created just minutes before tweeting for the first time and within a two week period acquired about 14 thousand followers. Figure 2.
3. One might wonder whether there was any skepticism during the spreading of the rumor. The answer is no, because the rumor spread in a dense echo chamber, creating a perfect environment for growing the conspiracy theory. See Figure 3. TwitterTrails can easily show what the echo chamber looked like, what the main keywords in the profiles of the group that populated the echo chamber were, and who the main actors in its spreading were. It turns out that a pair of identical twins aspiring writers, played a major role in spreading the rumor in the US. See Figure 4.

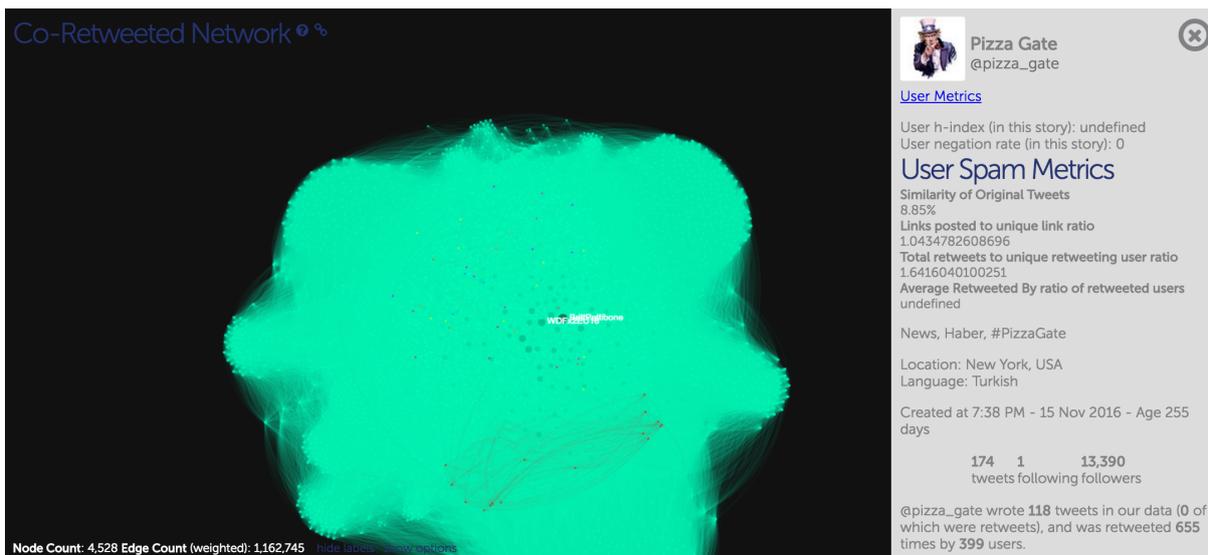


Figure 3: The co-retweeted network of the echo chamber formed by the group discussing #Pizzagate. Typically, co-retweeted networks of stories that are of political nature show at least two polarized groups. For example, in US politics one typically sees two groups one by liberals and the other by conservatives. In this story, however, there is a single network representing an echo chamber wither claims are accepted without strong doubt.

There are many more discoveries that one can make interacting with the TwitterTrails page <http://bit.ly/TTpizzagate> such as the pictures posted by the groups, but given the limited space for this publication we will leave it for the interested reader.

An investigation starts with data collection. Using appropriate keywords that should guarantee as great precision and recall as possible, we collect all relevant tweets. For our case study, we performed a keyword search on Twitter for a single keyword: “#PizzaGate” on Dec. 2, 2016. After that step, TwitterTrails is fully automatic. Throughout the process we also col-

## Co-Retweeted Network Statistics

There are a total of 17 communities of similar users in the Co-Retweeted Network. The largest community has 4474 users in it, and the smallest has 2, with an average of 266 users, and median of 2 users. Nodes are colored based on their community. To view aggregation statistics about any of the communities, you can either click on a node in the graph, or select a community from the panel below.



Figure 4: There was one major group discussing #Pizzagate effectively forming a dense echo chamber. The keywords in the word cloud are those that appear more often in the profiles of the participants in the group, effectively describing the group members. In the large group the keywords are trump, #maga, truth, love and god.

lect relevant data, such as the pictures shared with these tweets and the URLs of any websites mentioned.

## 3 Conclusion

Even though #Pizzagate is a topic that has been extensively examined, TwitterTrails easily points out many facts and insights that were previously missed. This is but one of the over 500 publically visible investigations we have conducted using our system. There are many others that have been conducted by journalists and are visible only to them. We welcome and work with journalists or researchers that want to make use of our system. If interested, please contact us and follow us on Twitter (Trails Research @tweet\_trails)

In addition to the findings we describe above, it is worth knowing that this investigation reveals the technique that propagandists used to create a conspiracy theory on Twitter: Find a community that is emotionally charged on some issue; create fake accounts that become members of the community; then launch the conspiracy theory and sit back and watch while the rest of the community promotes the conspiracy without skepticism [3]

## References

- [1] P. T. Metaxas, S. Finn, and E. Mustafaraj, "Using twittertrails.com to investigate rumor propagation," in Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing, ser. CSCW'15 Companion. New York, NY, USA: ACM, 2015, pp. 69–72. [Online]. Available: <http://bit.ly/2uGeoEc>
- [2] S. Finn, P. T. Metaxas, and E. Mustafaraj, "Spread and skepticism: Metrics of propagation on twitter," in Proceedings of the ACM Web Science Conference, ser. WebSci '15. New York, NY, USA: ACM, 2015, pp. 39:1–39:2. [Online]. Available: <http://bit.ly/2n7o2Ml>
- [3] E. Mustafaraj and P. T. Metaxas, "The fake news spreading plague: Was it preventable?" in Proceedings of the 2017 ACM on Web Science Conference, ser. WebSci '17. New York, NY, USA: ACM, 2017, pp. 235–239. [Online]. Available: <http://bit.ly/2sehUCv>