

Disinformation on Facebook – A case study from Finland*

Vili Ketonen
Aalto University
vili.ketonen@aalto.fi

Kiran Garimella
MIT
garimell@mit.edu

Aristides Gionis
Aalto University
aristides.gionis@aalto.fi

Abstract

The spread of disinformation in online media and its effects on politics has been a major issue of concern recently. In this paper, we look at the specific case of disinformation on Facebook, in Finland. We find that unlike the US, disinformation is not widespread in the Finnish media landscape. We analyze the differences in the reach and the social feedback for the disinformation content and the users who spread this content. We show that the top-100 biased articles are shared up to 2 times more than credible articles and have up to 1.7 times more reactions. Given that most large-scale analysis on the disinformation landscape to date has been US-centric and mostly on twitter, our findings provide a look into a less frequently studied landscape. Our results also highlight the importance of cultural differences while characterising the effects of disinformation. We release our code and dataset to promote further investigation into this area.

Introduction

Social media has become a major source for information consumption by users. Social media is used not only for sharing content with friends but also for consuming news and learning what is happening in the world. A recent Pew research study showed that over 60% of users receive their news from social media (Gottfried and Shearer, 2016). On the other hand, the ease with which news can be curated and shared has caused an expanding spread of disinformation, which is developing into a major problem. For example, during the latest US presidential election, it was reported that the most viral fake-news stories were shared more than the top stories published by the mainstream news media (Silverman, 2016). Disinformation has been used as a weapon to manufacture polarizing stories and create propaganda (Graham-Harrison and Cadwalladr, 2018).

Meanwhile, scandinavian countries, including Finland, have remained fairly resistant to this disinformation problem, due to the historic trust in traditional media sources. Nevertheless, the discussion around fake news has become more common in Finland in the past two years, following the public debate on the disinformation propaganda that took

place during the 2016 US election.¹ Typical topics of fake news in Finland have mostly been fake stories about anti-immigration and far-right hate speech, along with content that is suspected to have ties to disinformation propaganda sponsored by the Russian government (Koponen, 2018).

In this paper, we collect and analyze a large dataset from facebook of news articles from several Finnish biased news media and websites known for spreading unconfirmed rumors, disinformation, and fake news. In addition, we collect a similar dataset of credible articles by reputable Finnish news media in order to compare the differences between the datasets. Using these datasets, we study how disinformation spreads on facebook in Finland in contrast with the information spread of reputable news media.

While many previous studies have explored this topic, they have either been conducted using twitter data or are focused on US politics. Those aspects bring a number of limitations: (i) facebook has almost an order of magnitude higher audience than twitter;² (ii) facebook represents better the mass population of users while twitter is somewhat of an ‘elitist’ social network; (iii) with different levels of trust in media across different countries, we hypothesize that the effects of disinformation vary among countries;³ and (iv) facebook has not been extensively studied in the context of disinformation, particularly in Europe (Fletcher et al., 2018).

Our results provide interesting insights, which differ from the existing literature on analyzing disinformation on twitter, mostly in the US. In particular, our main findings are the following:

- The prevalence and spread of disinformation in Finland is limited. There is a small number of popular biased media, which reach a large audience, but in general, credible news sources still have the dominance.
- Biased news get more extreme reactions (“angry” and “haha”) than credible news.

¹In this paper we study *disinformation*, which we define as deliberate propagation of false information. At places we use the popular term *fake news* to refer to disinformation.

²For instance, a mere 5% of Finns use twitter compared to 82% who use facebook: <https://www.adweek.com/digital/finland-social-networks/>

³<http://www.digitalnewsreport.org/interactive/>

*Not peer reviewed. Please do not cite.
Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

- The top 100 biased articles are shared 2 times more than credible articles and have 1.7 times more reactions.
- The users who share biased articles are also ‘rewarded’ socially, by receiving almost 4 times more reactions compared to users who share credible articles.
- In terms of network properties, users who share biased articles have high in-degree but low out-degree.

To facilitate further research in this area, we release our dataset and our data-collection code.⁴

Related work

The overloaded term ‘fake news’ has been used quite often in different contexts (Morgan (2018)): (i) disinformation — deliberately propagated false information with the intention to malign a person, social group, or organization; (ii) misinformation — unintentionally propagated false information, (iii) propaganda — factually correct information, but packaged in a way so as to malign a person, group, or organization, (iv) other — conspiracy, rumors, etc. In this paper, we deal with disinformation and propaganda, i.e., deliberate attempts to mislead users with false or biased information.

Understanding the effects of disinformation. With the unexpected results from recent high profile events like the US elections and Brexit, interest in studying the phenomenon of fake news has peaked. Most studies conflate the different types of fake news mentioned above into a single term. Silverman (2016) analyze engagement data obtained by fake and real news during the 2016 US elections and conclude that fake news reached a larger audience than real news. Guess, Nyhan, and Reifler (2018) consider the actual consumption of fake news using web-traffic data from a representative sample of US users. They estimate that approximately 1 in 4 users in US visited a fake-news website from October 7 to November 14, 2016. They also find that facebook was a key source for exposure to fake news. Similar analysis is also conducted for the 2017 French elections by Ferrara (2017), who conclude that disinformation by twitter bots did not play a big role in the elections since it mostly reached a non-French audience. Fletcher et al. (2018) measure the reach of disinformation by news sources in two different European countries, France and Italy. They find that the average reach of fake-news websites was lower than the reach of the most popular credible news websites.

Understanding properties of fake news. Research has also identified that fake-news stories have different properties compared to credible news stories. These differences have been studied along various dimensions: (i) *Textual content*: Fake-news stories usually have certain unique textual characteristics, e.g., they contain clickbait headlines with false information to catch reader’s attention, they put emphasis on extreme opinions and emotional topics (Gu, Kropotov, and Yarochkin, 2017; Janze and Risius, 2017), etc. (ii) *Visual content*: Fake-news stories often comprise of shocking or even staged images to further trigger the reader’s emotional response (Shao et al., 2017). (iii) *Network*: Looking

at who spreads fake news, Zhao et al. (2018) found that, unlike for real news stories, the original creator of the story is not usually the one who spreads the story, since these fake-news stories often go viral. In a study related to our paper, Vosoughi, Roy, and Aral (2018) analyze the spread of rumors in twitter and find that news containing rumors spreads faster, and farther in the social network. They also find that false-news stories activate emotions, such as fear, disgust, and surprise to the users, whereas real stories bring anticipation, sadness, joy, and trust.

Based on the papers discussed above, we make the following observations:

- (i) In this paper we look at disinformation on facebook at scale, a domain in which there is not much prior work. Most existing work relies on small-scale surveys. For our study we develop an automated data-collection strategy to track facebook social data for tens of thousands of news articles.
- (ii) We aim to understand *both* the effects of disinformation as well as the properties of non-credible news on facebook.
- (iii) Most work on understanding fake news looks at production — how these articles are written. We analyse the consumption — how people react to these articles. For instance, Gu, Kropotov, and Yarochkin (2017) show that fake news are written in a way to elicit extreme reactions. One may ask: *do they elicit such reactions?* We show that they do.
- (iv) Our results show that although some of the fake-news properties are similar to those in twitter datasets, there are differences too. We hypothesise that such differences are not only a characteristic of the platform (facebook) but of the country (Finland) and the societal trust in media.

Dataset

Our data-collection process starts by gathering a list of news websites that typically spread disinformation in Finland. The list was compiled by collecting material from different independent fact-checkers in Finland, including (i) the Finnish National Broadcasting Company (YLE) (Vehkoo, 2016; Tyry, 2016), independent fact-checking websites, like, Huhumylly,⁵ and (ii) different independent journalists (Tyry, 2016). The final list contains the 20 most popular news domains that spread disinformation.⁶

Next, we obtained all the articles published by these 20 websites. The crawling process yielded a total of 64 577 distinct news articles. For these articles, using the facebook Graph Search API, we collected the following *public* information: (i) users who shared the article and time of share, (ii) facebook reactions to the article (*like, haha, angry, love, sad, and wow*) and the timestamps of the reactions, (iii) users who shared the original facebook post. We were able to obtain this information for 15 295 articles (the remaining articles were either not shared publicly or did not receive much attention). We refer to this set of articles as DISINFO.

For comparison purposes, we also obtained a similar dataset of credible news stories from six popular credible

⁵<https://huhumylly.info/>

⁶The list of credible media and disinformation websites can be found here: <https://pastebin.com/PGyRYJDH>

⁴<http://bit.do/eFiBg>

Table 1: Number of articles and users in our dataset.

Dataset	# Articles	# Users	# Shares	# Reactions
DISINFO	15 295	52 889	33 175	261 805
CREDIBLE	15 281	121 127	22 887	204 353




Figure 1: Articles ranked by the number of shares (left) and reactions (right).

news media in Finland, such as Helsingin Sanomat.⁶ We randomly sampled roughly the same number of articles (15 000) and obtained facebook data for these articles. We refer to this set as CREDIBLE. Table 1 lists the statistics of the dataset.

Analysis

We start by observing from Table 1 that while the two datasets, DISINFO and CREDIBLE, contain roughly the same number of articles, the size of the user base in the credible news dataset is approximately 2.3 times larger than in DISINFO. This indicates that the credible news stories have a wider spread in terms of number of users reached compared to the biased news in Finland. However, even though the size of the user base in the biased news dataset is smaller, biased news stories have 30% more social interactions in terms of shares and reactions than the credible news stories, which suggests that the average biased article tends to get more shares and reactions than the average credible article.

To further understand this behavior, we examine the distributions of the reactions and shares for the two datasets. Figure 1 illustrates the differences between the datasets by ranking the articles by their share and reaction counts. The distribution is heavy tailed, with most of the articles having a small number of shares and reactions. However, the 100 most shared DISINFO articles have been shared 2.4 times more than the 100 most popular credible articles. Similarly, the top 100 articles in DISINFO in terms of number of reactions have 1.7 times more reactions than the top 100 articles in CREDIBLE. This results are in line with the findings of Fletcher et al. (2018); Silverman (2016), who report that a small number of fake news articles have a higher reach than credible news articles.

Next, we look into the type of reactions elicited by the articles. Figure 2 shows the percentage of each reaction type for the two datasets. The default option, “like” dominates the reaction types in both datasets. DISINFO stories have 18% of “angry” reactions while credible stories have only 6%. Likewise, biased news articles acquired double the number of “haha” reactions compared to the credible




Figure 2: Reaction type distribution of the biased news (left) and credible news (right).




Figure 3: Users ranked by the number of shares (left) and reactions (right).

news. These observations indicate that articles in DISINFO elicit strong reactions from the users. Again this result is in line with previous work: previous studies have shown that fake news and misleading stories usually consist of negative words and strong emotions, which could trigger strong reactions from the users (Janze and Risius, 2017; Gu, Kropotov, and Yarochkin, 2017; Conroy, Rubin, and Chen, 2015; Weeks, 2015). Our results provide evidence that users indeed respond to such articles with extreme emotions.

Having compared the share and reaction distributions of the articles, we look at similar distributions of the users in Figure 3. We observe that the 250 users who share the most articles in DISINFO, share 2.3 times more articles than the top 250 users in CREDIBLE (Figure 3 (left)) and have 4.3 times more reactions.

Next, we look at rate of spread of articles. We plot the cumulative distributions of the share time of the articles in Figure 4. Credible news stories get approximately 70% of their shares within the first hour on average while biased stories get only 54% of their shares. Similarly, while the average credible news story reaches 90% of its shares within the first 20 hours after its first share, it takes 42 hours for the average biased story. In other words, even though biased news stories get more shares on average, they spread slower than credible articles. This finding is in contrast to the results reported in a recent study on the spread of lies, where lies are found to spread faster than the truth (Vosoughi, Roy, and Aral, 2018).

Finally, we study the *network* of users sharing these news articles. We construct a user-user share network G , where an edge ($u \rightarrow v$) between users u and v indicates that u shares an article posted by v . Table 2 shows the properties of G . We can see that the graph for DISINFO is much denser (1.6 times) compared to the corresponding graph for CREDIBLE. The connectivity of these graphs can also be seen from the

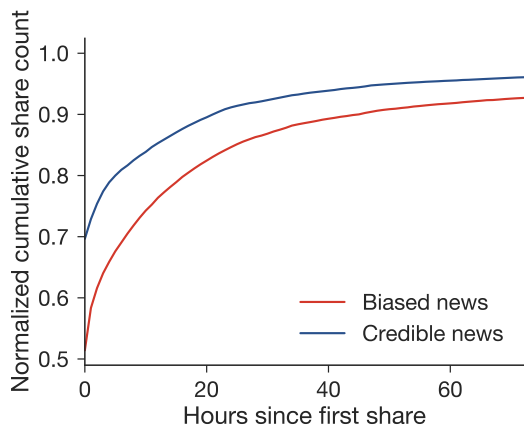


Figure 4: Average cumulative share time distribution of an article for the first 3 days since the first share.

Table 2: Size of the user-user share network. Size of the largest connected components in the brackets.

Properties	DISINFO	CREDIBLE
$ V $	5 142 (4 695)	4 653 (2 322)
$ E $	9 975 (9 637)	4 792 (3 222)

size of the connected component, which contains over 90% of the nodes for DISINFO while containing only 50% for CREDIBLE. These numbers indicate that users in the biased news dataset are well connected in a smaller community.

Looking at the degree distributions of the users in the graphs, we see that users in DISINFO have higher in-degree (Figure 5 (left)) and lower out-degree (Figure 5 (right)), indicating that in DISINFO a higher number of users reshare content. This observation reinforces the hypothesis that users in DISINFO form a tightly-knit community.

Discussion and Conclusions

In this study we collect and analyze a large dataset of credible and disinformation news articles on facebook from Finland. We found that disinformation is not wide spread in Finland, but there is a small minority of articles and users who create impact. Disinformation stories typically elicit highly emotional reactions, and hence acquire more social interactions. Disinformation also spreads *slowly* compared to credible stories. Users spreading biased news are more strongly connected compared to users who primarily share credible news stories. Our results both support and contradict existing studies from other countries (e.g., US and Italy) on other platforms (twitter).

This work is limited to Finland and the findings may not generalize to other countries. Despite of this, our results highlight the importance of such studies. We show that for Finland, where the trust in media is high, the impact of disinformation might not be as pronounced as it might be in other countries. Thus, we can not make claims on the effects of fake news based on a single platform or a single country and that the results may vary depending on the local context. To




Figure 5: Normalized indegree (left) and outdegree-rank (right) distribution of the user-user share network.

encourage further research in this area, we provide our code and data for other researchers to build upon.

References

- Conroy, N. J.; Rubin, V. L.; and Chen, Y. 2015. Automatic deception detection: Methods for finding fake news. *Proceedings of the Association for Information Science and Technology* 52(1):1–4.
- Ferrara, E. 2017. Disinformation and social bot operations in the run up to the 2017 french presidential election.
- Fletcher, R.; Cornia, A.; Graves, L.; and Nielsen, K. 2018. Measuring the reach of "fake news" and online disinformation in Europe. *Oxford: Reuters Institute for the Study of Journalism, University of Oxford*. Accessed 12 January 2019.
- Gottfried, J., and Shearer, E. 2016. News use across social media platforms 2016. *Pew Research Center*.
- Graham-Harrison, E., and Cadwalladr, C. 2018. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. *The Guardian*. Accessed 15 December 2018.
- Gu, L.; Kropotov, V.; and Yarochkin, F. 2017. The fake news machine: How propagandists abuse the internet and manipulate the public. *Trendlabs research paper, Trend Micro*.
- Guess, A.; Nyhan, B.; and Reifler, J. 2018. Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 us presidential campaign. *European Research Council*.
- Janze, C., and Risius, M. 2017. Automatic Detection of Fake News on Social Media Platforms. *PACIS*.
- Koponen, H. 2018. Finland remains resistant to fake news, disinformation. *International Press Institute*. Accessed 12 January 2019.
- Morgan, S. 2018. Fake news, disinformation, manipulation and online tactics to undermine democracy. *Journal of Cyber Policy* 3(1):39–43.
- Shao, C.; Ciampaglia, G. L.; Varol, O.; Flammini, A.; and Menczer, F. 2017. The spread of fake news by social bots. *CoRR* abs/1707.07592.
- Silverman, C. 2016. This Analysis Shows How Viral Fake Election News Stories Outperformed Real News On Facebook. *BuzzFeed*. Accessed 12 January 2019.

- Tyry, M. 2016. Disinformaatio, vihapuhe ja mediakasvatuk-
sen keinot. *Osuuskunta Mediakollektiivi and Koulukino*.
Accessed 26 December 2018.
- Vehkoo, J. 2016. Valheenpaljastaja: Varoituslista valeme-
dioista 1 luota nihin sivustoihin. *Yle*. Accessed 12 January
2019.
- Vosoughi, S.; Roy, D.; and Aral, S. 2018. The spread of true
and false news online. *Science* 359(6380):1146–1151.
- Weeks, B. E. 2015. Emotions, partisanship, and misper-
ceptions: How anger and anxiety moderate the effect of
partisan bias on susceptibility to political misinformation.
Journal of Communication 65(4):699–719.
- Zhao, Z.; Zhao, J.; Sano, Y.; Levy, O.; Takayasu, H.;
Takayasu, M.; Li, D.; and Havlin, S. 2018. Fake news
propagate differently from real news even at early stages
of spreading. *ArXiv e-prints*.