Hilary Sun
Professor Persily
LAW 806C
Fall 2017
12/10/2017

The Role of Open Source Technology in the Battle Against Fake News

Introduction

Social media has dramatically changed how people receive and disseminate news. While most social media platforms are not necessarily "free," since signing up for accounts means users give the platform permission to collect and sell their data, these services are still made accessible to a wide user base: essentially anyone with an Internet connection. Platforms like YouTube and Twitter continue to grow their users at an astounding pace; just this June, Facebook reported reaching two billion monthly active users, doubling their user count in less than five years (Constine). According to a Pew Research study, about four-in-ten Americans now get their news online, whether it be through searching for content themselves or clicking on links shared by their online networks on these social media platforms. Two-thirds of Facebook users read their news on their timelines, which is a over half the general population (Matsa and Lu). As these numbers continue to rise, so do the variety of actors that utilize these platforms to spread mis- and disinformation for commercial or political gain, including those who spread fake news.

Fake news that go viral on social media have an especially dangerous impact on society. Social media platforms are engineered to enhance viral sharing, such as showing the top Twitter hashtags or the most trending Facebook topics, so it is no surprise that viral content with higher visibility is more likely to influence a larger network of people. This sort of design "attribute[s] legitimacy to popularity," meaning that users are exposed to a lot of popular and unverified content that could include fake news (Deb, et al. 8). They could then be more inclined to believe

the most popular content even if it is not true. Virality also brings into question how the Internet challenges democratic ideals. While the Internet has provided access to news and sources from a variety of different viewpoints, malicious actors that disseminate fake news take advantage of how visibility is controlled on social media platforms. They act against democracy by having a high volume of their fake messages drown out other voices in society to advance their own personal goals (Persily 71-72). In response to these detrimental effects, developers and researchers have begun creating open source technologies that the public can use.

In order to make technologies to combat fake news, the developers must first of all define what "fake news" is: a difficult task in itself. Propaganda, satire, bias, and conspiracy theories can all fall under fake news, in addition to news that is simply false (Persily 68). The intentions behind a fake news story may also vary; while some forms of fake news like satire are used for parody or humor, others, such as propaganda, may have a more malicious design. Some ideologically-biased news sources pick and choose material to include and exclude in their reporting, and governments even use online smear campaigns to attack journalists and activists that speak up against them (Woolley 13).

What should be categorized as fake news is also highly subjective. Different users' own beliefs and political orientations may cause their definition of "fake news" to vary from others'. In this period of intense political polarization in America, negative views of the opposing side are growing, which means that those who read news not aligned with their viewpoints are more likely to ignore it or label it as fake ("Political Polarization"). An article published by a news source that is seen as more liberal can be labeled as false by a more conservative reader even if the content itself is sound, and vice versa, meaning that developers need to ensure their own technologies are not influenced by political bias.

Making technologies for the fake news problem is also difficult when the exact extent of the problem is still unknown. Even now, statistics about the extent of the Russian interference, which included malicious actors spreading fake news about the candidates during the 2016 election, are still coming out. More than 131,000 messages were published on Twitter, more than 1,000 videos were uploaded on YouTube, and 126 million users on Facebook were reached because of this interference with the election. These are numbers that were far beyond the extent that was reported originally (Isaac and Wakabayashi). It is difficult to guess the number of people who were affected by any fake news spread by Russian actors or others. About a third of Americans say that they often see fake political news stories online and a fifth have reported that they have shared fake news, but it is difficult to trust self-reported statistics. Although consumers say that they are relatively confident that they are able to tell a fake news story from a real one, it has been shown that humans are incredibly bad at determining what is fake or not (Barthel, et al.). In fact, the most important factors that help users decide whether a story is accurate is their own perceptions of the news source and the person who posted the link ("How people decide what news to trust"). So if a user's trusted network is posting links that lead to fake news, the user is more likely to believe the posted headlines without fact-checking against other news sources. The uncertainty of the extent of fake news' impact complicates the problem.

Since technology was part of the problem, making technologies that users can trust can also be difficult, which makes the transparency of open source technologies especially important. As a response to pressure from the public and the government to "fix" its fake news problem, Facebook edited its newsfeed in an attempt to remove instances of fake news. It has partnered with widely-accepted fact-checking websites to help fact-check its content. Facebook describes its methodology as taking the most-reported news articles, running it through some algorithms,

and then sending it to the fact-checkers ("Clamping down on viral fake news"). But its lack of transparency makes it unclear how effective this strategy is (Simonite). Additionally, because of the spread of fake news on social media platforms, users are now much less likely to trust news that come from these platforms (Mitchell, et al.). Not making the news feed methodology clear makes the improvement hard for users to accept when their trust levels with the platforms are already lower. Recent accusations from various conservative and liberal critics include claims that Facebook is left-leaning. This can also make any of their strategies seem politically biased, and when its top executives admit their own political views to the public, it is hard to not associate these views with the product itself (Herrman and Isaac).

Social media platforms open to the public have become the breeding ground for fake news, and yet a wide range of other open technologies have appeared to combat the fake news problem. It is important to look at these technologies that are unassociated with the social media platforms, analyze how they interact with the public and try win its trust, and see how their techniques can be applied to foster the growth of neutral and publicly-accessible technological solutions.

The Definition and Importance of Open Source Technology

A lot of different products and tools exist on the web to combat fake profiles and fake news. Some are free and can be used by anyone, but many are private and require compensation for their use. Some are extremely open about how the technology itself works, but some keep their algorithms hidden from the public. One sort of technology that reveals its source code and methodology that is free for anyone on the Internet to use and access is open source technology.

The open source technology that I will examine will rely on the definition of openness that is based on the "condition of access" (Pénin 135). In his research on the idea of openness, Julien Pénin considers two levels of openness: a strong one, which means the resource is free and open to the public, or a weak one, which means the user has to ask the owner for permission to use it or may have to pay a reasonable price (135). I will focus on the former--technology that is free-of-charge for the public to use and also makes its algorithms and source code public. This distinction is especially important since the platforms that proliferated fake news are free for anyone who has Internet access and accepts the terms and conditions to use. My case studies will focus on technologies that are available for use by any user who currently uses Facebook, Twitter, or any other social media site.

The success and popularity of open source is uncontested. For example, Linux and Apache, two open source systems that are extremely complex and well-maintained, are used and improved upon by a wide audience of companies and independent developers (Weber 2). Although some of the technologies that I will examine are not as well-developed as these examples, open source technology will play an important role in the future battle against fake news, especially as a response to closed-off "solutions" by corporations or the government that people may see as politically biased, such as Facebook's own attempt to suppress fake news on their platform. Open source is arguably a lot more innovative and up-to-date than secret software from corporations because the inner-workings of the technology are exposed to the public. Anyone can spot bugs and vulnerabilities, suggest fixes and additional features, and improve upon it by branching off on the original project (Woolsey and Fox). Christopher Kelty describes open source software as "self-determining, collective, [and] politically independent": its very existence is a check on institutions of power and a form of democracy (xi, 7). The same can be

said about any form of open source technology, whether it was built with code or not. By creating open source technologies, individuals can challenge rules or restrictions imposed by those in power. For example, B.S. Detector, a tool that I will discuss in my later case studies, was developed directly in response to "Mark Zuckerberg's dubious claims that Facebook is unable to substantively address the proliferation of fake news on its platform" (B.S. Detector). In this battle for democracy, the role of open source is even more important. A technology that has the potential to be viewed by any member of the public is more likely to be trusted and accepted by the public. A technology that has the ability to be critiqued and improved by the public, regardless of their experience or political affiliation, is more likely to be seen as bipartisan and politically independent, as we will see in our later case studies.

While the open source technologies that I will examine all expose their algorithms and source code, they do not necessarily all advocate for the same measure of open innovation to all users, especially users who are not as familiar with computer science. We will see later that these decisions affect how effective a technology is.

Types of Existing Open Source Technology

My case studies focus on a broad definition of technology. These open source technologies fall into three case studies: open source databases, open source tools, and open source algorithms. The audience of these open source technologies vary. Some are based upon the belief that information should be directly disseminated to the public through a machine. Others believe that technology should be used to help speed up the fact-checking process, and that human fact-checkers should be making the final decision on the validity of a news source. Some face the fake news problem head-on by releasing a ready-to-use solution to the public.

Others recognize the limitations of artificial intelligence and believe that smaller algorithmic problems need to be solved first.

This paper is not meant to be a thorough analysis of all existing technologies. While other types of technology exist to combat fake news, I found these three categories to be the most relevant and impactful on the Internet.

Case Study 1: Open Source Databases

Although databases are not exactly a type of software, they are still a valuable resource. A lot of software rely on well-curated databases to run and test their algorithms, so having public databases is extremely important to open source development. Reliable and clean data is especially hard to find, since the process of aggregating this data requires a lot of time and resources.

OpenSources is an example of a database made publicly available for use by anyone on the Internet. The project was started by Melissa Zimdars of Merrimack College and her research team, but it has since stalled, with its last update in April 2017. Zimdars has stated on Gitter in October 2017 that "there isn't an update planned for the near future as other groups have already taken this data set and expanded it substantially...while others are attempting to automate it." The slowness of updates shows how inefficient humans can be in compiling these databases compared to the fast pace of fake news creators on the web. However, analyzing its limitations and how it was used can give us valuable insight into how data about fake news should be disseminated to the public.

OpenSources contains a list of credible and not so credible news sources. Its mission, as stated on the website, is "to empower people to find reliable information online" (*OpenSources*).

The database list contains the website URL along with a tag that categorizes it in one of the following categories: fake news, satire, extreme bias, conspiracy theory, rumor mill, state news, junk science, hate news, clickbait, process with caution, political, and credible. The database does not categorize individual news articles; rather, it uses a set of clearly-stated criteria on its website to evaluate the news source website as a whole. It does not have any specific source code to make available, but it does make its methodology public on its website. Their technique to categorize news sources includes analyzing whether it has had questionable content in the past as well as looking at overall site aesthetic, source use, and writing style (*OpenSources*).

The value of providing databases of fake news sites to the public has been recognized by some fact-checking websites. For example, PolitiFact, a fact-checking website that recently teamed up with Facebook to help analyze possible fake news sources, published its own Fake News Almanac of unreliable websites ("A new database of fake news sites"). The list itself is still being updated, and as of November 2017, it has a total of 330 news sources listed. While the data is downloadable, it is not in an easy-to-use format for software developers to use since the user would have to create their own database from the data and manually update it. A user using OpenSources could update their database with a simple command. PolitiFact is also not as open and transparent as OpenSources is, since it takes no feedback directly from its users and is not as open about how it works.

Like OpenSources, the list categorizes websites into parody or joke sites, imposter sites, fake news sites, and sites that are duped by other fake news sites. But its exact methodology is not stated on the website; rather, it relies on a general description of each category to describe its criteria (Gillin). While useful, it seems more like an effort to make transparent the progress that Facebook is attempting to make. Additionally, this list is limited to unreliable news sites that

PolitiFact finds through its work with Facebook. Its effectiveness in combatting the fake news problem is limited by its engagement with fake news sources that have passed through Facebook's algorithms. Additionally, while it is open for anyone to use, it does not provide a way for users to give immediate feedback on the list itself and does not directly rely on public contributors. The extent of its user base is unclear, while OpenSources has clearly paved the way for a lot more complex databases and tools based on its level of engagement with its users.

OpenSources is different because it has been open to receiving more feedback from its users. Its website contains a survey to receive feedback from its users, including questions on how users use the database and how a more rigorous analysis of its data would affect its usage. The survey contains a section where survey respondents can ask for a follow-up to their feedback. Additionally, the team provides a chatroom on Gitter, which allows its users to ask questions on the tagging of news sources, request advice on other untagged news sources, and crowdsource volunteers and users for various projects that use OpenSources.

But despite these successes, OpenSources also has limitations. As of April 2017, it has only analyzed 834 news sites: nowhere close to the millions of news sources on the web that people read. The maintainers of the database could only focus on certain sites, so there are also not a lot of sites marked as "reliable." The team seemed to specifically seek out websites they saw as unreliable or fake and only added any reliable sources that they came across or were suggested by their users. While the methodology of how they categorize news sites is clearly stated, the specific reasoning for individual websites are not as concrete or descriptive. Some websites listed have a short reason listed beside it for why it was tagged the way it was, but not all of them do. Users can ask in Gitter why a source was marked a certain way, but they then have to wait for a response from the team.

Case Study 2: Open Source Tools

OpenSources and other databases were used by many developers to create their own open source tools. These sorts of tools include Chrome extensions, web pages, and other software that can be installed on the user's computer. While a database is simply an aggregation of data, tools have an interface that the user can directly interact with. Because of their ease of use, open source browser extensions have been a popular choice for developers. One example of such a tool is B.S. Detector.

B.S. Detector is an extension that was originally created to help address fake news on Facebook. B.S. Detector utilized OpenSources by pulling data about different websites from its database. When the extension was installed, it looked up any website that the user was currently browsing in the OpenSources database; if the website was labeled as unreliable, B.S. Detector displays a banner at the top of the page warning the user of the possible unreliability of the website. The banner also shows which OpenSources category the site is labeled as. For example, The Onion is labeled as satire. Like OpenSources, B.S. Detector is no longer under development, but it was also one of the most popular browser extensions out there. As of November 2017, the Chrome extension had around 24,000 users.

The code for the tool is also available publicly on GitHub, a popular website for collaborating on software projects. Developers could branch off on the code to develop their own projects based on B.S. Detector, or they could request to integrate their own features into the original code. For users who are not used to reading code, not a lot of information about how the extension works was provided. But B.S. Detector is very open in its interactions with its users. It

offers an open chat on its website for users to share their own tools, ask questions, and provide feedback. Users could also file issues in the GitHub asking for bug fixes or feature requests.

Despite its ease of use and wide accessibility, B.S. Detector still was not entirely effective or convincing for certain users. There was still a lot of concern from users that the extension influenced what types of news that they were reading and how they perceived certain types of news. Users complained about perceived liberal or conservative bias in the flagging. Even in the Chrome extension reviews, users argue about sites that are marked and whether or not they are correctly labeled. One user argues that the tool "seems to target websites that are pro-life and pro-natural family," while another claims that it actually flags more left-leaning sites as clickbait ("B.S. Detector").

Other users complained that the tool was too simplistic and does not give enough convincing information for why a site was marked or not. As the user CobaltBW mentions in his Chrome review: "Anyone who is internet-savvy is already going to understand what kind of site they're looking at without the help of the extension, and anyone who actually believes the stuff on the site is not going to be convinced by the ratings from this extension" ("B.S. Detector"). They argue that providing more evidence for why a site was marked the way it was would help convince more users of a site's unreliability.

FiB is another extension specifically focused on the Facebook news feed. It was developed by a group of students for the HackPrinceton hackathon in 2016. Instead of reading from a database, its underlying structure is an algorithm written by its developers that uses image recognition, keyword extraction, source verification, and Twitter searches to verify a Facebook post (*FiB*). When installed, it goes through the Facebook news feed and verifies the posts, marking them as verified or not verified by putting a small tag on the top right of the post. When

users post or share content, the extension also checks the new posts to warn users if they are spreading any unverified information. It is even more simplistic in its design than B.S. Detector: while B.S. Detector gives the user a brief explanation of why the news site was marked, FiB does not give any reasoning to its labels. Because of this, FiB has also gotten many complaints on its algorithm. One user says that "clearly something is wrong" with the algorithm since it is marking sites such as The Onion as verified ("Project Fib").

FiB's source code is available to the public through GitHub, and various contributors have contributed to it, but its last update was in early 2017. As of December 2017, the Chrome extension only had around 723 users. Even though the developers briefly describe the algorithm used to classify verified and unverified posts, most of it is hidden within the code. Users who are not familiar with code would not be able to understand how it works. The developers also do not seem to make an attempt to clarify their algorithm. A user who tried to contact the developers to better understand the algorithm was also not satisfied; "all [he] got was the turnaround" ("Project Fib"). The development team does not seem to have won the public's trust; its low user base and unsatisfied reviews clearly demonstrate this.

Although FiB and B.S. Detector are similar in that they both work to label content, their methodology is very different. FiB is an artificial-intelligence-centered solution that purely relies on the machine to classify something as verified or not, while B.S. Detector is simply reading in data from a human-maintained database and presenting it in a more usable way. B.S. Detector is limited to warning users of when websites have had specific content that was found to be unreliable, while FiB works on marking specific content.

rbutr, another extension, takes a different approach to combatting fake news. Instead of labeling, its goal is to "improve critical thinking skills and foster a culture of critical thinking in

all internet users" by leaving the more subjective decision making to the user (*rbutr*). Instead of directly telling the user whether a news site or article is fake or not, it provides a list of other articles that have a rebuttal to the article's content. This list includes either articles that directly argue against the content or others that provide contradictory evidence. To do this, it relies on users crowdsourcing and submitting rebutting articles. When a user is on a website, if there exists rebuttals, a small pop up appears in the top right of the website. The user can then click on the Chrome extension to view a list of articles that provide rebuttals. It also makes its rebuttals database public for third-party use and general viewing, allowing open innovation.

Through this strategy, rbutr avoids accusations of bias that B.S. Detector and other tools like it face. Because disputed news sources have lists of actual articles that show a conflicting viewpoint, users can go to these links and decide for themselves if they believe the original article or the ones that are provided as rebuttals. Although the ease of use that was present in B.S. Detector was taken away, humans are given the ultimate decision-making power in marking what is false or not.

However, this tool also has many limitations. Because the tool is crowdsourced, a majority of news articles will not have any rebuttals simply because it takes more effort on the part of the users to find rebutting articles and put them in. Fake news articles that have just come out might not necessarily be marked immediately and could still have a negative impact on its readership. As of November 2017, around 7,000 users use the extension, which is still not as large of a user base as B.S. Detector. This sort of tool has great potential with a large number of users who are willing to contribute, but without a large user base, it is not as powerful.

Case Study 3: Open Source Algorithms

While rbutr and B.S. Detector sought to directly put a solution into the public sphere, other developers chose to focus on researching and developing algorithms to improve the fact-checking process. These technologies break up the fake news problem into a small subset of problems to solve rather than trying to solve the entire problem at once. While not usable in the short-term or not understandable by the general public, these technologies are very important to ensuring that we have strong algorithms that are being used in the tools that we develop.

The most well-known example is the Fake News Challenge, created and funded by Dean Pomerleau from Carnegie Mellon University and Delip Rao of Joostware. The goal is "to explore how artificial intelligence technologies, particularly machine learning and natural language processing, might be leveraged to combat the fake news problem," particularly in automating parts of the human fact-checking process (*Fake News Challenge*).

Since its creation, it has released and completed the first fake news challenge, which was to improve stance detection. A team from Talos Intelligence ultimately won with their algorithm of combining a gradient-boosted decision tree that uses deep learning and a deep convolutional neural network, different machine learning techniques where the machine determines patterns within the data by itself. One common stage of the fact-checking pipeline is comparing an article to other articles that are on the same topic. The challenge was therefore to take an input of a headline and body text either from the same news article or from two different ones and output whether the headline and the body agree, disagree, discusses, or are unrelated to each other. An algorithm that automates stance detection is a technology that can be used by any fact-checking or news organization, regardless of their political affiliations. Because of this, the Fake News Challenge avoids any accusations of bias. Additionally, any team is welcome to participate, and participants are judged based on how well their algorithm scores. Participants were required to

make their code publicly available on GitHub, where anyone would be free to parse through it. With the open source code, anyone on the Internet with a knowledge of how to use machine learning algorithms can innovate on it or use it for their own products. Although Fake News Challenge has only finished one stage of their competition with the stance detection challenge, a lot of excitement has already arisen from the results, and their publicly available Slack has been buzzing with anticipation about the next Fake News Challenge.

The creators of Fake News Challenge state that instead of taking on the task of labeling a story as true or false, stance detection is a good problem to tackle first because of how difficult the process of assigning a label to articles is. In their discussions with journalists and fact-checkers, they found that these professionals would "rather have reliable semi-automated tool [*sic*] to help them in do [*sic*] their job better rather than fully-automated system whose performance will inevitably fall far short of 100% accuracy" (*Fake News Challenge*). Part of this is also because of how limited software is in understanding the nuance of language (Simonite). Artificial intelligence still is far from fully being capable of human judgment, which is especially needed in a problem where the definition of fake news is not even agreed on by humans. These limitations are demonstrated even with the winning team's score. Each team's algorithm was evaluated by testing it against a test set of headlines and body texts that it was not previously exposed to. The winning team's algorithm correctly categorized 82.02% of these test examples, meaning that even though they got a majority of the stance detections correct, they were nowhere near categorizing 100% of the known examples given in the dataset.

Additionally, the algorithms developed from the Fake News Challenge are not fully understandable by the general public unless they have a thorough understanding of artificial intelligence technologies. These algorithms would have to be combined into tools that would

then be eased into the fact-checking pipeline. Since the Fake News Challenges only solve parts of the problem, the solutions are not useful unless they are integrated into newsrooms and fact-checking organizations, which we are currently not sure if they are. Even if they are integrated, these organizations would need to have their own teams of software engineers to continuously improve upon the algorithm created by the participants of the Fake News Challenge. Again, putting these solutions into the fact-checking pipeline is not something that can be done immediately, unlike how open source tools are directly available for use by the public.

Discussion

None of these technologies are perfect or fully developed enough to be considered a solution to the fake news problem. Developers are struggling to figure out how much to prioritize ease of use and how much to prioritize giving the decision-making to the human. Some technologies were released to the public and marketed to the public as a plausible solution even though they are not yet complete.

OpenSources, an open source database, puts well-curated data into the public sphere for developers to use to build tools. However, it is also limited in scope because no human team can update and maintain a database fast enough to keep up with the pace that fake news proliferates in the Internet. This is especially true for OpenSources, which was maintained by a handful of researchers. B.S. Detector was one of the open source tools developed that used the data from OpenSources to flag news sites as reliable or not. However, because it was limited only to the data available in the OpenSources database, some of its users believed that it was politically-biased or gave too little information about why a site was flagged. These same criticisms were made about FiB. rbutr, another tool that aggregated rebuttals to specific news articles, gave

humans the decision-making power by allowing them to critically think about the content they were reading. However, because the rebuttals are collected through crowdsourcing, its user base is too small for the tool to be very useful. Additionally, when there are rebuttals, the user has to read through the articles, requiring more of the user's time. Another open source technology is open source algorithms, such as the ones developed through the Fake News Challenge. The creators of the challenge believe that time needs to be spent on developing useful algorithms that help speed up the human fact-checking process instead of having a machine tell the user whether news is fake or not. But the algorithms are not useful to the general public unless they are integrated into newsrooms and fact-checking companies that also have a dedicated software engineering team to maintain and improve upon the technology.

OpenSources and other databases available to the public and the tools such as B.S. Detector that build upon these databases are troubling because users are using them even though they are nowhere near complete. OpenSources and other databases are not large enough to encompass all sorts of fake news websites that exist; we can question whether or not such a database can ever exist given the fast proliferation of fake news sources. Even just identifying the most prevalent fake news sites would require massive resources and a large team of dedicated people. Any website not listed within OpenSources is simply not marked when a user of B.S. Detector visits the website. News that the FiB's algorithms cannot fully process are also not marked. This could lead users into thinking that pages that are not marked as unreliable by the extensions are more likely to be true. A recent Yale study evaluated Facebook's new technique of tagging fake articles with "disputed by 3rd party fact-checkers" to combat fake news, which is similar to the techniques used by the extensions. As mentioned above, Facebook has partnered with unbiased fact-checking organizations such as PolitiFact and Snopes to fact-

check news articles that have been flagged as possible fake news, or other news articles that are most salient. Since Facebook cannot possibly tag all fake news stories, Gordon Pennycook and David G. Rand find through their study that not only does tagging only modestly decrease an article's perceived accuracy, but something known as the "implied truth" effect makes untagged fake news stories seem more accurate, which could be even more dangerous than not tagging any fake news stories at all (1). The same idea can be applied here. Tagging a website does not necessarily mean a user will believe the tag, as we saw from the Chrome extension reviews for B.S. Detector. Additionally, it would be impossible for any database to tag every single fake news site. The concept could backfire, and users could be led to believe that untagged sites are more reliable than they actually are. The same "implied truth" effect could happen to rbutr users as well. If articles do not have rebuttals listed yet, the user could be led to believe that articles with no rebutting articles could be more accurate than those that do not. The effect may be less extreme in this case, however, since having rebuttals or no rebuttals does not necessarily mean that it is telling the user whether the article is true or not. It simply provides less information for the user to base their decisions upon. These open source databases are good for providing a sample dataset for developers to build their tools, but they should be clear to their users that the data is nowhere near complete. Developers such as those who created B.S. Detector should only use the data in OpenSources to build prototypes of what kinds of tools users could possibly use in the future when we have developed a more thorough way of identifying fake news; B.S. Detector, for example, does not ever warn the user that its underlying dataset is not complete and that users should proceed with caution when using the tool. It should have been clear that it was not putting forth a catch-all solution.

FiB's idea of utilizing a full artificial intelligence solution comes with other implications as well. The fake news problem is "ultimately about humans not machines," so fully relying on algorithms is not addressing the actual problem behind the fake news problem (Simonite). Similar to the mindset behind rbutr's approach to improve critical thinking, the creators of the Fake News Challenge realize that humans will have to make the ultimate decision on what is labeled as fake news or not. If people entirely rely on algorithms to tell them what is real and what is not, what happens when the people spreading fake news create a technology that can manipulate and overcome the fake news algorithm, making it so that their fake news is marked as verified? Because of the increased reliance on the algorithm, the public would be even more likely to believe these falsely verified news that they see. As one FiB Chrome extension reviewer puts it: "At the end of the day, we can't expect a software to think for us entirely" ("Project Fib"). Most of the public does not trust artificial intelligence either. A study found that Americans are more worried than excited about machines taking over human processes, believing that more harm would come from fully-automating solutions (Smith and Anderson). A machine-driven technology would have a hard time winning favor with the public, at least in the present.

The mindset behind rbutr, B.S. Detector, and FiB is that information should be provided directly in real-time to the user about the news they are reading. Fact-checking companies currently take a long time to verify certain news stories, meaning that the damage may have already been done before they are able to finish their verification. B.S. Detector and FiB work to just tell a user whether a news site is reliable or not on the spot, while rbutr puts the fact-checking process in the user's hands. However, many users do not have the time to parse through rebutting articles when scanning the news, while many other users cannot fully trust a tool like

B.S. Detector or FiB that may not be seen as a neutral technology. The Fake News Challenge seeks a different approach to increase the speed of fact-checking by automating some parts of the fact-checking pipeline for use by professionals who do have the time to parse through news articles and who have the credibility to gain the public's trust. Solutions to the fake news problem should be a combination of human and machine-based efforts such as those supported by Fake News Challenge to both speed up the process through automation but also retain the human aspect of fact-checking.

A Hypothetical Technological Solution

An ideal solution would be to continue developing algorithms to help automate parts of the fact-checking process, like Fake News Challenge is doing, and then work to integrate these algorithms into newsroom and independent companies. Future parts of the fact-checking process that could be automated include faster aggregation of news stories before they go viral, better ways to detect any bias in the news article, and quicker checking of an article's sources. These algorithms should then be maintained and continuously improved upon to ensure that their technology is up-to-date and ahead of the technology developed by malicious actors. Having an open source database that can be updated by all these news organizations and fact-checking companies could also build collaboration in the effort to fight fake news and aggregate information faster. This could be where the government steps in to help build collaboration between private news organizations and foster sharing of resources.

Developers, such as those working on B.S. Detector, FiB, and rbutr, should work on tools to make dissemination of fact-checking more readily available to the public. If such a tool would directly block news sites from the public, it could be seen as censorship. If this needed to happen,

the government should be involved to regulate and ensure that private parties are not censoring content themselves. These sorts of tools could be developed by a public entity or a private entity, but this team of developers should be an established neutral group that has creators from all parts of the political spectrum. Being marketed as a neutral team of creators could help avoid some of the accusations of bias that FiB and other tools face. Despite the premature nature of the open source tools investigated, they do have the benefit of being able to quickly pass on information to the public. Currently, a lot of fact-checking information is only available on websites, which takes more time for a reader to verify whether something they are reading is true or not. Tools should be built to help present information from professional fact-checkers in fast and effective way. More user research should be done on how people interact with different tools to get their news to ensure that ease of use is maintained while still giving the user enough information to trust the technology.

Not only would the team that works on developing these tools need to be marketed as unbiased, but the methodology itself would have to be presented as unbiased as well. This would enable people to trust the technology more. Clearly describing their methodology and how they define fake news would help with transparency and winning the public's trust. Additionally, educating the public about how the technology should be used is important as well. Some people who are not familiar with artificial intelligence believe or want to believe that a machine could be built to solve the fake news problem. However, the limitations demonstrated by FiB and the algorithms from the Fake News Challenge show that we are still far from a fully-automated process. Developers of these tools would need to ensure that users fully understand that the tool they are using is simply a tool and should not replace the user's own decision-making process.

Additional resources, such as a more cohesive and comprehensive database of fake news sources or rebutting articles such as those provided by rbutr, could be provided to the public. This would allow them to be involved in the problem-solving process. Putting a database maintained by a variety of news organizations and fact-checking organizations into the public sphere would also make it easier for more open source technologies to be created. Open source developers can then innovate in conjunction with private organizations, a collaboration that also needs to be strengthened in order to fully address the fake news problem.

References

Barthel, Michael, et al. "Many Americans Believe Fake News Is Sowing Confusion." *Pew Research Center*, 15 Dec. 2016, http://www.journalism.org/2016/12/15/many-americans-believe-fake-news-is-sowing-confusion/. Accessed 26 Oct. 2017.

"BigMcLargeHuge/opensources." *GitHub, Inc.*, https://github.com/BigMcLargeHuge/opensources. Accessed 5, Nov. 2017.

B.S. Detector. http://bsdetector.tech/. Accessed 26 Oct. 2017.

"B.S. Detector." *Chrome Web Store*, https://chrome.google.com/webstore/detail/bs-detector/dlcgkekjiopopabcifhebmphmfmdbjod/reviews?hl=en. Accessed 8 Dec. 2017.

Constine, Josh. "Facebook now has 2 billion monthly users...and responsibility." *TechCrunch*, 27 June 2017, https://techcrunch.com/2017/06/27/facebook-2-billion-users/. Accessed 22 Nov. 2017.

Deb, Anamitra, et al. "Is Social Media Jeopardizing Democracy?" *The Omidyar Group*, 14, Aug. 2017.

Fake News Challenge. http://www.fakenewschallenge.org/. Accessed 26 Oct. 2017.

"FiB." *Devpost*, https://devpost.com/software/fib. Accessed 8 Dec. 2017.

Gillin, Joshua. "PolitiFact's guide to fake news websites and what they peddle." *PunditFact*, 20 April 2017, http://www.politifact.com/punditfact/article/2017/apr/20/politifacts-guide-fake-news-websites-and-what-they/. Accessed 26 Nov. 2017.

Hahn, Robert W., ed. *Government Policy toward Open Source Software*. Brookings Institution Press and AEI, 2010, 1-11. Web. 26 Nov. 2017.

Herrman, John and Mike Isaac. "Conservative Accuse Facebook of Political Bias." *The New York Times*, 9 May 2016, https://www.nytimes.com/2016/05/10/technology/conservatives-accuse-facebook-of-political-bias.html?_r=0. Accessed 26 Nov. 2017.

"How people decide what news to trust on digital platforms and social media." *American Press Institute*, 17 April 2016, https://www.americanpressinstitute.org/publications/reports/survey-research/news-trust-digital-social-media/. Accessed 23 Nov. 2017.

Isaac, Mike and Daisuke Wakabayashi. "Russian Influence Reached 126 Million Through Facebook Alone." *The New York Times,* 30 Oct. 2017, https://www.nytimes.com/2017/10/30/technology/facebook-google-russia.html?_r=0. Accessed 22 Nov. 2017.

Kelty, Christopher M. *Two Bits: The Cultural Significance of Software*. Duke University Press, 2008. Web. 26 Nov. 2017.

Matsa, Katerina E. and Kristine Lu. "10 facts about the changing digital news landscape." *Pew Research Center*, 14 Sept. 2016, http://www.pewresearch.org/fact-tank/2016/09/14/facts-about-the-changing-digital-news-landscape/. Accessed 22 Nov. 2017.

MetaCert. https://metacert.com/. Accessed 26 Nov. 2017.

Mitchell, Amy, et al. "Trust and accuracy." *Pew Research Center*, 2016, July 7, http://www.journalism.org/2016/07/07/trust-and-accuracy/. Accessed 23 Nov. 2017.

OpenSources. http://www.opensources.co/. Accessed 5 Nov. 2017.

Owen, Laura H. "A new database of fake news sites details how much fakery has spread from Trump v. Clinton to local news." *NiemanLab*, 21 April 2017, http://www.niemanlab.org/2017/04/a-new-database-of-fake-news-sites-details-how-much-fakery-has-spread-from-trump-v-clinton-to-local-news/. Accessed 26 Nov. 2017.

Owen, Laura H. "Clamping down on viral fake news, Facebook partners with sites like Snopes and adds new user reporting." *NiemanLab*, 15 Dec. 2016, http://www.niemanlab.org/2016/12/clamping-down-on-viral-fake-news-facebook-partners-with-sites-like-snopes-and-adds-new-user-reporting/. Accessed 26 Nov. 2017.

Pénin, Julien. "Are You Open? An Investigation of the Concept of Openness for Knowledge and Innovation." *Revue Économique*, vol. 64, no. 1, 2013, pp. 133–148. *JSTOR*, JSTOR, www.jstor.org/stable/23485192.

Pennycook, Gordon and David G. Rand. "Assessing the Effect of 'Disputed' Warnings and Source Salience on Perceptions of Fake News Accuracy." Version 2.0. 15 Sept. 2016.

Persily, Nathaniel. "Can Democracy Survive the Internet?" *Journal of Democracy*, vol. 28, no. 2, 2017, pp. 63-76.

"Political Polarization in the American Public: How Increasing Ideological Uniformity and Partisan Antipathy Affect Politics, Compromise, and Everyday Life." *Pew Research Center*, 12 June 2014, http://www.people-press.org/2014/06/12/political-polarization-in-the-american-public/. Accessed 22 Nov. 2017.

"Project Fib." *Chrome Web Store*, https://chrome.google.com/webstore/detail/project-fib/njfkbbdphllgkbdomopoiibhdkkohnbf. Accessed 8 Dec. 2017.

rbutr. http://rbutr.com/. Accessed 26 Nov. 2017.

Simonite, Tom. "Humans Can't Expect AI To Just Fight Fake News for Them." *Wired*, 15 June, 2017, https://www.wired.com/story/fake-news-challenge-artificial-intelligence/. Accessed 26 Oct. 2017.

Smith, Aaron and Monica Anderson. "Americans' attitudes toward a future in which robots and computers can do many human jobs." *Pew Research Center,* 4 Oct, 2017, http://www.pewinternet.org/2017/10/04/americans-attitudes-toward-a-future-in-which-robots-and-computers-can-do-many-human-jobs/. Accessed 8 Dec. 2017.

Weber, Steven. *The Success of Open Source*. Harvard University Press, 2004.

Woolley, Samuel C. "Computational Propaganda and Political Bots: An Overview." *Can Public Diplomacy Survive the Internet? Bots, Echo Chambers, and Disinformation*. Ed. Shawn Powers and Markos Kounalakis.