FeedReflect: A Tool for Nudging Users to Assess News Credibility on Twitter

Md Momen Bhuiyan
Virginia Tech
Blacksburg, VA 24061
momen@vt.edu

Kexin Zhang
Georgia Tech
Atlanta, GA 30332
kexin@gatech.edu

Kelsey Vick
Virginia Tech
Blacksburg, VA 24061
kvick31@vt.edu

Michael A. Horning
Virginia Tech
Blacksburg, VA 24061
mhorning@vt.edu

Tanushree Mitra
Virginia Tech
Blacksburg, VA 24061
tmitra@vt.edu

CSCW '18 Companion, November 3–7, 2018, Jersey City, NJ, USA
ACM 978-1-4503-6018-0/18/11.
https://doi.org/10.1145/3272973.3274056

Abstract
In recent years, the emergence of fake news outlets has drawn out the importance of news literacy. This is particularly critical in social media where the flood of information makes it difficult for people to assess the veracity of the false stories from such deceitful sources. Therefore, people oftentimes fail to look skeptically at these stories. We explore a way to circumvent this problem by nudging users into making conscious assessments of what online contents are credible. For this purpose, we developed *FeedReflect*, a browser extension. The extension nudges users to pay more attention and uses reflective questions to engage in news credibility assessment on Twitter. We recruited a small number of university students to use this tool on Twitter. Both qualitative and quantitative analysis of the study suggests the extension helped people accurately assess the credibility of news. This implies *FeedReflect* can be used for the broader audience to improve online news literacy.

Author Keywords
Social Media; News Credibility; User engagement; Reflection; FeedReflect

CCS Concepts
•**Human-centered computing** → **User centered design;**

Mainstream source
ABC News
BBC
CBS News
NPR
The Economist
Wall Street Journal

Table 1: A sample of sources in the mainstream category

Non-Mainstream source
Breitbart News
Drudge Report
Info Wars
Project Veritas
RT
The Onion

Table 2: A sample of sources in the non-mainstream category

Figure 1: Content highligted due to question in the comment

[1] http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/

[2] http://www.journalism.org/2016/12/15/many-americans-believe-fake-news-is-sowing-confusion/

Introduction

"*Twitter is honestly my main source of news ... but I feel like more recently I am not actually reading the stories as much. Because I don't know. It gets a little repetitive ... So sometimes I read headline and I am like that's enough information.*" – Anonymous Twitter user's response when asked about her news reading habits on the platform.

Social media platforms, such as, Twitter and Facebook play a vital role in information dissemination. Nearly 60% of Twitter's 330 million active users get news on the site[1]. The increasing popularity of Twitter as a news source has also led to the emergence of a vast number of outlets disseminating information ranging from carefully vetted news to completely fabricated rumors. For example, ABCNews.com.co, with its striking similarity to the traditional professional media ABCnews.com, publishes hoax stories, which often go viral on social media sites. Repeated exposure to false stories originating from these sources could create an illusion of truth[2]. This, in essence, challenges the fundamental value of a democratic society: a well-informed citizenry. How can we enable users to be conscious consumers of news on Twitter while preserving their freedom of choice?

In recent years, "nudges," such as, warnings, reminders, and recommendations, have been used to steer user behavior in particular directions, without sacrificing the freedom of choice [3]. Scholars have shown that a nudge is more effective when it requires conscious or deliberate thinking instead of subconscious processing [4]. Research on information credibility also suggests that warning a user about potential misinformation is preferable than correcting the misinformation. Drawing from these studies, we designed a system that facilitates conscious news consumption by applying a nudge-by-warning cue on tweets and requiring reflecting on the information's credibility.

To nudge users about news tweets that require accurate credibility assessment, we use visual cues such as highlighting and obscuring tweets. These visual cues act both as an element to draw attention and nudge users about the veracity of the source while they are reading news on Twitter. The nudge is followed by engaging them in conscious assessment of the credibility of information. Designs that facilitate reflection have widely been used in the educational domain to engage users in any activity [1]. For example, one study achieved higher engagement in journaling by asking users to reflect on past journal entries. Drawing from this research, we use a reflective element, i.e., questions to engage users in news credibility assessment.

FeedReflect

We developed *FeedReflect*, a Google Chrome extension for Twitter, a popular social medium where people consume news. We intervene on users' Twitter feeds by highlighting content that require careful assessment of credibility. First, we emphasize tweets from mainstream news sources containing questions in the comments. Then we de-emphasize content from non-mainstream or dubious sources. News content with questions usually refer to less credible news events. We only highlight content from mainstream sources because it does not make sense to reflect on content from questionable sources. Instead, tweets from non-mainstream news sources are dimmed or grayed out to nudge users into thinking that content from those sources are potentially dubious. The literature on nudges also suggest that offering specific information pertaining to the nudge—such as reasoning behind the interventions—can enhance their effect [3]. Therefore, for each visual cue we add a tooltip that explains the reasoning behind the intervention. Finally, the system allows the user to reflect on those tweets by answering questions pertaining to the credibility of the news, thus named *FeedReflect*.

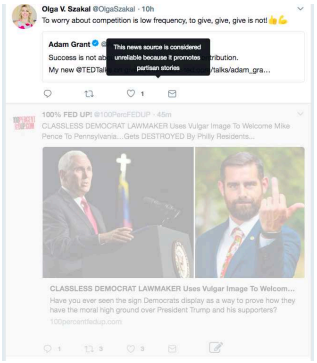

Figure 2: Content dimmed due to less credibility of the source

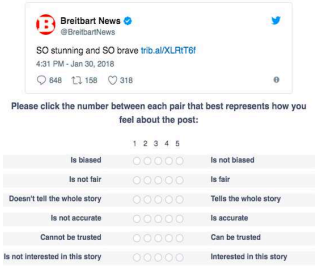

Figure 3: Survey for the purpose of reflection

[3] https://www.cision.com/us/2014/06/top-10-us-daily-newspapers/

[4] http://www.journalism.org/2014/10/21/political-polarization-media-habits/pj_14-10-21_mediapolarization-01/

[5] https://www.ap.org/about/news-values-and-principles/downloads/ap-news-values-and-principles.pdf

[6] opensources.co

To maximize the number of users who could use the system, we developed a tool that is independent of an operating system. A browser extension for Google Chrome can be used in any Web operating system. After the user installs the extension, it changes her Twitter feed by applying appropriate interventions.

Classifying News Sources

Mainstream News Source: There is no official definition of mainstream media in the literature. Thus, we used a heuristic approach designed by a journalism and communication-media expert—Michael Horning, a professor from the department of communication, who is also a co-author. Our expert compiled the list of mainstream news sources by referring to the most circulated[3] and the most trusted news sources from a Pew survey[4], followed by removing two of the sources, a news aggregator (Google News) and a local source (AMNewYork). Our final list comprise sources which generally follow ethical journalistic practices in their reporting [5]. Table 1 shows a sample.

Non-mainstream News Source: For non-mainstream news sources, we used a curated list created by a research team from Merrimack College [6]. This list organizes sources in a taxonomy which includes bias, clickbait, and satire. We removed the "reliable" category from this list as the sources in this category were neither well-known nor reported fabricated content. Table 2 shows a sample of these sources.

Applying the interventions

Visual Cues to Nudge User: We applied two visual cues on users' news feeds. The first visual cue highlights content from mainstream news sources that contain questions in the Twitter comments. Figure 1 shows how the highlighting changes a tweet. As a simplistic method to detect questions, we match the "?" mark in the comments to identify such tweets. While this method is robust, one limitation is that it neglects whether the question relates to the news. The nudge comes with a tooltip showing the first comment with a question. The goal is to drive the user's curiosity and encourage exploratory behavior. The second visual cue dims content from non-mainstream sources by reducing the opacity. Figure 2 shows the intervention in action. Along with dimming, a tooltip explains the reason behind this change. In this case, the tweet originates from an alternative media which doesn't follow standard or ethical journalistic practice. The purpose of the reduced opacity along with the explanatory message is to nudge users into believing that the source is dubious.

Questions to Make Users Reflect: The final component of our system provides questions for the users to reflect on the credibility of the information. Studies show that Gaziano's News Source Credibility Scale is effective in credibility assessment [2]. Figure 3 shows the questions we used from this scale. All tweets with visual cues contain a survey button at the bottom. When a user clicks on this button, the questions pop up. While answering them, the user need to actively reflect on the news.

Result

To test *FeedReflect's* effectiveness in enabling credibility assessment, we recruited 16 university students, equally dividing them into treatment and control groups. The students participated in a three week study and used Twitter everyday as they would normally do. The extension activated the warning nudge (highlighting and obscuring) for the treatment group alone. Thus, the treatment users answered the credibility assessment questions for each highlighted and obscured item, while the control group answered questions without the nudge. The study resulted in 241 data points. For each tweet, we computed the mean of the credibility questions as the credibility score of that

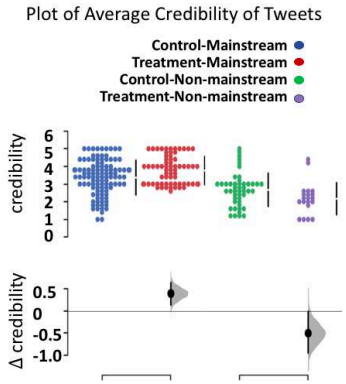

Figure 4: Estimation plot of the credibility score of mainstream and non-mainstream sources for treatment and control. Student's t-test (2-sided) shows significant difference, with $p < 0.001$.

"yeah, I actually clicked on the article, read through. I mean I do that for things that interest me. But even this time the things that didn't interest me, I clicked on it and read through."

"I'm probably more likely to look it up or do more research on it rather than just reading and like retweeting things that aren't true"

"I think it drew out the process"

"Actually I already had to do but it kind of like pushed the extent of that cuz it took a little bit more time and I was actively having to think about the questions that I was subconsciously doing. so it took a more active role I guess."

tweet. Figure 4 shows an estimation plot of this score. The Cronbach α of the score for both groups were above 90%, suggesting high coherence. We find that the treatment group scored the tweets from mainstream sources higher and non-mainstream sources lower compared to the control group, suggesting the effectiveness of our warning nudge.

After three weeks, we interviewed users in the treatment group to understand the effect of our nudge-by-warning intervention. A qualitative analysis of the interviews suggests that the extension can engage users in reflecting on the credibility of news in their social media feeds. Here we discuss a set of themes that emerged. We show quotes in margins, color coded corresponding to each theme.

(1) Make people reread and rethink the news. *FeedReflect* made people reread the news. Users would reread them even if the stories were not of interest to them. They were also more likely to second-guess the content.

(2) Make people use external resources for credibility assessment. One effect of *FeedReflect* was that users would use external resources to research the information.

(3) Make people actively participate in content credibility assessment. Users were more aware of the content they were viewing. Answering the survey questions made them actively participate in credibility assessment, rather than subconsciously making judgments.

Limitations

The interview pointed out several limitations. Most users complained about how news related to entertainment or opinion from editorials were highlighted—items which would not need credibility assessment. Another limitation was the choice of the system. Recent statistics suggests that 80% of Twitter users access the platform on mobile. Future research should aim for designing more sophisticated interventions that can work on mobile systems while automatically excluding entertainment and opinion news.

Contribution to HCI

FeedReflect is a new system to nudge users' credibility perception in social media. This is a third-party extension that (1) nudges users in content credibility assessment, (2) makes people reflect upon the social media content they see on their feeds, and (3) gives users control over the intervention as they are free to ignore it. We believe this system could improve people's news literacy on social media.

Conclusion

In this paper, we introduced a system—*FeedReflect*—to engage users in careful evaluation of news on Twitter. Deploying the intervention among a group of university students showed its effectiveness. When using *FeedReflect*, users perceived tweets from mainstream sources more credible, while those from non-mainstream sources as less credible.

REFERENCES

[1] Eric PS Baumer, Vera Khovanskaya, Mark Matthews, Lindsay Reynolds, Victoria Schwanda Sosik, and Geri Gay. 2014. Reviewing reflection: on the use of reflection in interactive system design. In *Proc. DIS'14*. ACM, 93–102.

[2] Cecilie Gaziano and Kristin McGrath. 1986. Measuring the concept of credibility. *Journalism quarterly* 63, 3 (1986), 451–462.

[3] David Halpern. 2016. *Inside the nudge unit: How small changes can make a big difference*. Random House.

[4] Cass R Sunstein. 2016. Do people like nudges? (2016).