

PhD Proposal: Detection of coordinated influence campaigns

Keywords: social media, influence campaigns, manipulation

Labs:

- LIX, Inria CEDAR team

Supervisors:

- Oana Goga (CNRS/LIX) – <https://www.lix.polytechnique.fr/~goga/>

Background and PhD topic description

Social media platforms have changed how users consume news and stay updated on current events, with nearly half of U.S. adults now turning to social media, especially Facebook, as their primary news source [12]. This reliance on Facebook for news brings both advantages and concerns. On the one hand, it enables effortless news dissemination, democratizes access to information, and allows users to exchange ideas and opinions with people. On the other hand, many organizations have raised concerns about the platform facilitating exposure to misinformation [1, 6]. One key enabling mechanism is the ease with which anyone can claim to be a news provider and share news-related content without verification. **Recent reports showed the emergence of organizations aiming to influence voters during elections by claiming to be local news providers** [2].

The goal of this PhD is to detect online coordinated campaigns aiming to influence citizens by masquerading as news providers.

Challenge 1: Automated detection of self-proclaimed news providers: Fostering a healthy news environment requires constant monitoring and auditing of content shared by both known and less-known self-proclaimed news providers. *Unfortunately, having a comprehensive view remains impossible, as Facebook does not disclose the list of self-proclaimed news providers on the platform.* In an attempt to audit the (mostly U.S.) news media ecosystem, known journalistic agencies, MediaBiasFactCheck and NewsGuard, have aggregated a list of 4k news media Facebook pages [10, 9]. As they are the only sources, many recent news-related studies have only considered established news providers listed by journalists [4, 13, 7, 8, 11]. However, we do not know to which extent these lists are comprehensive and, hence, to which extent relying studies provide an extensive view of the entire Facebook news ecosystem.

In this proposal, we propose an approach that relies on the assumption that Facebook pages claiming to be (and wanting to look like) news sources typically post news-related content. Therefore, our key idea is to perform a daily crawl that: (1) exploits the GNews API [5] to get a sample of news articles published by established news media in the past 24 hours and extract a set of corresponding *keywords*; (2) uses CrowdTangle [3], an API provided by Meta, to search for Facebook posts mentioning these keywords in the past 24 hours; and (3) filters only Facebook pages that self-identify as news media. Facebook pages that claim to be news providers usually list on their About page the corresponding news domain. This way, we will be able to have a list of both Facebook pages and domains that claim to be news providers. Our plan is to create the largest dataset Worldwide of self-proclaimed news providers.

Challenge 2: Clustering of news providers pertaining to the same entity: To be able to detect a coordinated influence campaign, we need to be able to group together all news providers that belong to the same entity. Previous works have exploited IP addresses to link together news domains [2]. However, with the emergence of public hosting infrastructure, this method is no longer effective. Our key idea is that news providers are usually displaying ads on their pages. When setting up the ad technology to be able to show ads, a website owner needs to create a file, called ads.txt, that will list its identifier in different ad networks. Our plan is to collect the ads.txt files of all the websites detected in the previous step and use clustering techniques to link together news websites that share the same identifiers across different ad networks.

Challenge 3: Identify influence campaigns: After we link together news sources that belong to the same entity, the next step is to distinguish between legitimate clusters (the same mother organization supports news sources in different regions of the country) and coordinated influence campaign clusters. For this, we will analyze the content of the posted articles using the NLP techniques, and we will measure biases in the way current issues are presented. We will also implement NLP techniques to detect the use of propaganda in the presented articles.

References

- [1] Michael Barthel, Amy Mitchell, and Jesse Holcomb. Many americans believe fake news is sowing confusion, 2016.
- [2] Priyanjana Bengani. Hundreds of ‘pink slime’ local news outlets are distributing algorithmic stories and conservative talking points. *Columbia Journalism Review*, 2019.
- [3] CrowdTangle. A tool from Meta to help follow, analyze, and report on what’s happening across social media, 2023.
- [4] Laura Edelson, Minh-Kha Nguyen, Ian Goldstein, Oana Goga, Damon McCoy, and Tobias Lauinger. Understanding engagement with u.s. (mis)information news sources on facebook. 2021.
- [5] GNews. A Python Package that searches Google News RSS Feed and returns a usable JSON response, 2023.
- [6] Ted Van Green. Few americans are confident in tech companies to prevent misuse of their platforms in the 2020 election, 2020.
- [7] Andrew M Guess, Brendan Nyhan, and Jason Reifler. Exposure to untrustworthy websites in the 2016 us election. *Nature human behaviour*, 2020.
- [8] Ro’ee Levy. Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 2021.
- [9] Media Bias Fact Check, 2023.
- [10] News Guard, 2023.
- [11] Michael Scharnow, Frank Mangold, Sebastian Stier, and Johannes Breuer. How social network sites and other online intermediaries increase exposure to news. *Proceedings of the National Academy of Sciences*, 2020.
- [12] Mason Walker and Katerina Eva Matsa. News consumption across social media in 2021, 2022.
- [13] Galen Weld, Maria Glenski, and Tim Althoff. Political bias and factualness in news sharing across more than 100,000 online communities. 2021.