AUDITING FACEBOOK ALGORITHMS
THE ELAPSED EFFECTS OF FACEBOOK NEWS FEED TO ENGAGEMENT WITH GUARDIAN ARTICLES

Naoise McNally
University College Dublin

Marco Bastos
University College Dublin

Extended abstract

The Facebook News Feed is the homepage feature of the social network that displays an endless scrolling of individual pieces of content such as status updates, photographs, videos, and web links/articles. The order to which content is arrayed is defined using multiple algorithms working simultaneously and concomitantly. These algorithms optimize the feed to prioritize content deemed most relevant to users (Mosseri, 2016), but the exact nature of this optimization is obfuscated by the proprietary nature of Facebook's algorithms and can only be inferred indirectly through changes enacted in the underlying algorithms. The auditing of such algorithms is therefore not a trivial task, as both the data and the models underpinning the algorithms are precluded from scrutiny by the public or the academic community.

This study addresses this gap by exploring the elapsed effects of changes in the Facebook News Feed on engagement with Guardian news articles. We tracked and documented public or otherwise known changes to the algorithms through Facebook public announcements, industry research, and information leaked to the press to parametrize a model that account for the variation in user engagement with Guardian news articles. To this end, we queried the Guardian API to collate a database of all Guardian articles published between 1st January 2010 and 31st December 2020 and subsequently queried the CrowdTangle API to retrieve Facebook engagement metrics for Guardian articles. Changes made to the algorithms were parameterized as high or low impact and were coded with a valence of positive, negative, or neutral positions towards trusted news sources.

The temporal series explored in this study present important variations. During this 10-year period, news organizations such as the Guardian struggled to transition from a print-and-advertisement to a digital-and-subscription business model. Substantial cuts and redundancies in editorial workforce followed, with quantifiable impact on content production reflected in the sharp decline Guardian articles published in print, towards the end of 2014, and then online around the end of 2016 (Figure 1a). Additionally, Facebook engagement data available via CrowdTangle has a more restrictive definition for public shares than Facebook does. CrowdTangle share counts are limited to articles shared within public pages, groups, or users with 'public' profiles, who usually are public figures like politicians or celebrities. CrowdTangle data can nonetheless be used as a benchmark, since the number of CrowdTangle public shares is a lower bound for the number of total public shares (Allen et al., 2021). Another important shortcoming of CrowdTangle is that the indexing of Guardian data is not uniform, with many articles published prior to 2014 not available on CrowdTangle (Figure 1b).
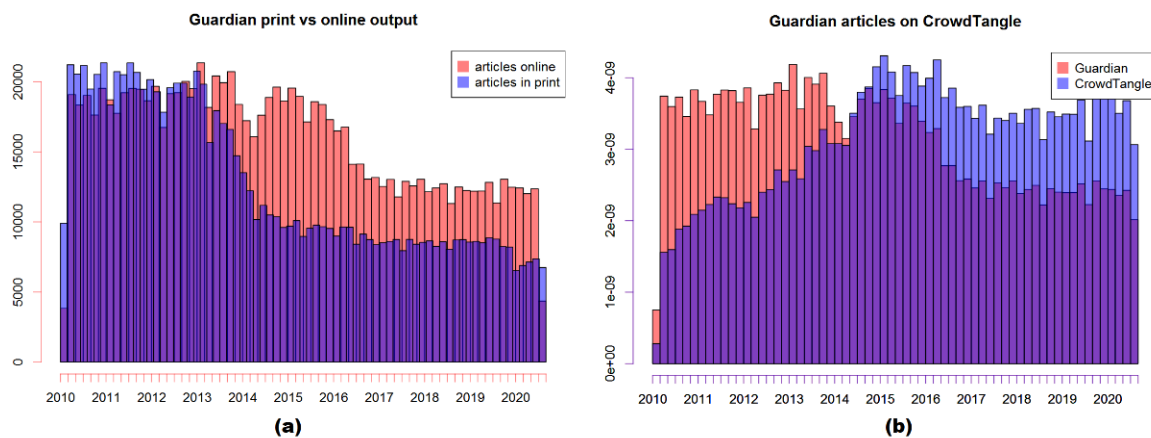

Figure 1(a): News articles published by the Guardian online and in print. Figure 1(b): Guardian articles with Facebook engagement metrics available on CrowdTangle (y-axis shows densities).

Guardian data amounts to 1,020,163 news articles whereas the CrowdTangle data set includes 606,102 unique Guardian URLs. The database includes 76 variables, including Facebook expected and recorded engagement metrics, shares, comments, and Facebook Reactions ('Like,' 'Love,' 'Care,' 'Haha,' 'Wow,' 'Sad,' 'Care,' 'Thankful,' and 'Angry'). We aggregate share and engagement metrics by total and average (mean), which is the dependent variable in the temporal models fitted for this study. We performed robustness and Augmented Dickey-Fuller tests to ascertain that the series are non-stationary (Figure 2), and subsequently modelled the data using cross-correlation analysis, granger causality test, and anomaly detection. Cross-correlational analysis (CCF) not only identifies potential correlations between two series, but it also detects lagged influences between two time series. As such, CCF was crucial to characterize the temporal direction and the influence between the weighted temporal changes made to Facebook algorithms and engagement with Guardian articles over the 10-year period.
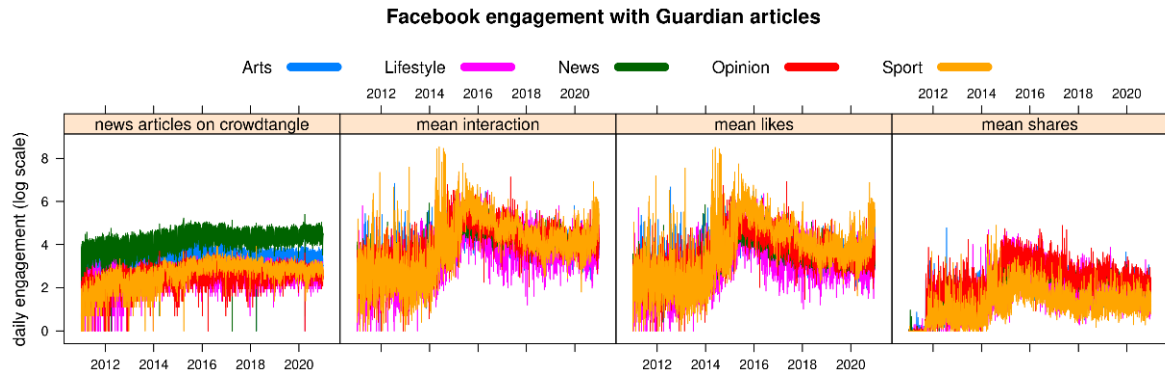
Figure 2: Time series of Facebook engagement with Guardian content 2010-2020

We performed cross-correlation analyses (CCF) to the series and tested for optimal lags that are significant at the 95% confidence interval level. As there is no significant pattern of correlations between opinion, lifestyle, and arts, we concluded that the series are co-occurring. News and sport, however, present significant relationships and we therefore concluded that the temporal series of algorithmic change does lead the temporal series of engagement with Guardian news and sport articles on Facebook. Figure 3 unpacks these findings, with upper and lower dashed lines showing that News and Sport are above the significance threshold at lags +21 and +24. In short, content of news value was found to be significantly more likely to be affected by changes made to Facebook News Feed algorithm within a window period of around 3 weeks.
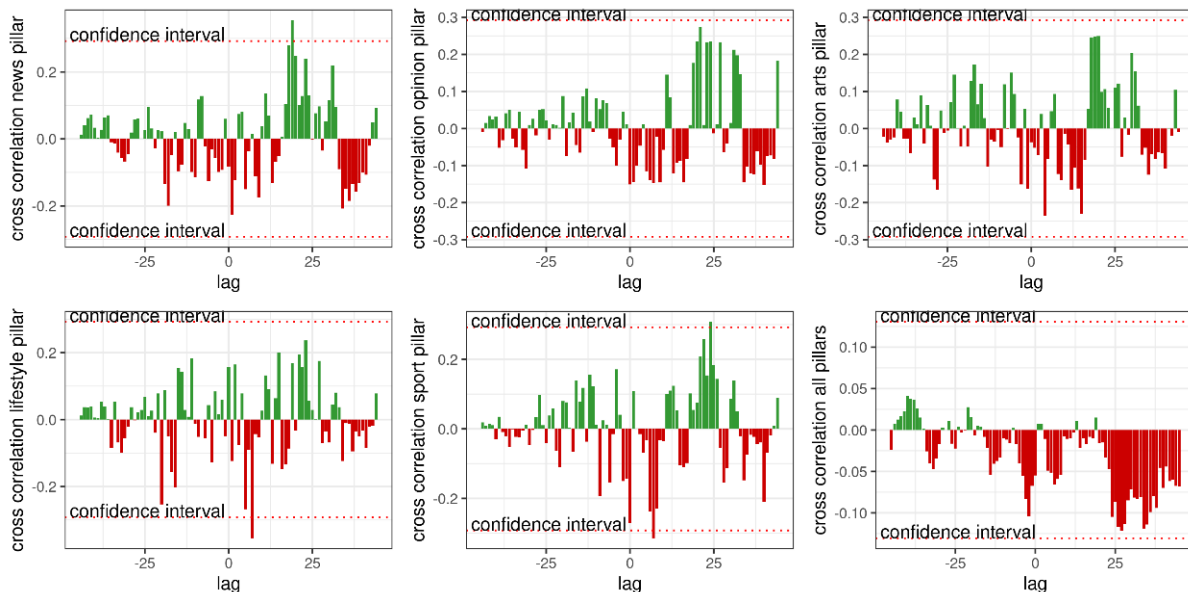

Figure 3: 3: Cross correlation between changes made to Facebook News Feed algorithms and Facebook engagement with Guardian content.

Next we performed Granger causality tests to determine the robustness of our findings. Figure 4 shows the results across the five pillars, with the News Feed algorithms predicting Facebook engagement with Guardian articles focused on news ($p < .0327$),

whereas changes made to the algorithm were not found to be significant with respect to arts, lifestyle, opinion, or sports. These results are consistent with the hypothesis that content of journalistic importance, such as news and politics, are more likely to be affected by algorithmic changes. Figure 4 also shows the anomalies detected by the S-H-ESD algorithm, largely concentrated in a period of intense changes made to the algorithm (2014-2016), with the yearly period of May-August accounting for more than half of the anomalies (99 out of 191). Two-thirds of the anomalies detected in the news pillar took place in 2016, with the vast majority concentrated in June 2016. The significant anomalies found in the news series before mid-2016 reflect critical changes introduced by Facebook, when posts from friends and family were further prioritized and ranked higher (Mosseri, 2016).
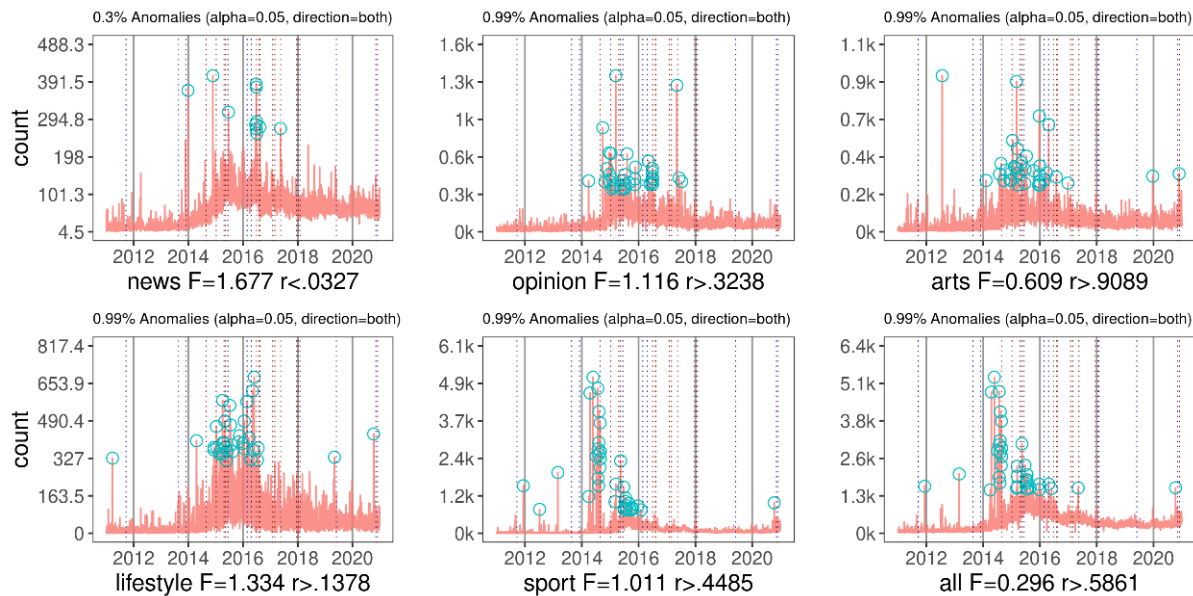

Figure 4: Granger causality tests with F-statistic and p values for each news pillar. Blue dots represents anomalies detected by the S-H-ESD algorithm.

In summary, this study presents a proof-of-concept study to audit the Facebook News Feed. Our results, however, speak only to the elapsed effect of these changes. We have no way to qualify the impact they may have had to the news industry and their revenue. We can only speculate that these changes were substantial. The asymmetric power exerted by social platforms on news organizations is such that news companies find themselves trying to meet the ever-changing demands of social platforms' algorithms. This period of persistent decline in revenue and trust in news organizations was accompanied by fundamental changes in the integrity of the information ecosystem, particularly with regard to how trusted information and news are obtained and consumed online. These issues have profound implications for the public discourse and the informed citizenry that sit at the center of democratic deliberation, a center ground now populated by digital intermediaries such as Facebook, Google, and Twitter whose algorithms define the integrity of information online.

References

Allen, J., Mobius, M., Rothschild, D. M., & Watts, D. J. (2021). Research note: Examining potential bias in large-scale censored data. Harvard Kennedy School Misinformation Review.

Mosseri, A. (2016, 2016, June 29). Building a Better News Feed for You. Facebook Newsroom. https://about.fb.com/news/2016/06/building-a-better-news-feed-for-you/