

Chapter Title: Dutch political Instagram Junk news, follower ecologies and artificial amplification

Chapter Author(s): Gabriele Colombo and Carlo De Gaetano

Book Title: The Politics of Social Media Manipulation

Book Editor(s): Richard Rogers, Sabine Niederer

Published by: Amsterdam University Press. (2020)

Stable URL: <https://www.jstor.org/stable/j.ctv1b0fvs5.7>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



This book is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0). To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>.



Amsterdam University Press is collaborating with JSTOR to digitize, preserve and extend access to *The Politics of Social Media Manipulation*

5 Dutch political Instagram

Junk news, follower ecologies and artificial amplification

*Gabriele Colombo and Carlo De Gaetano*¹

Abstract

The research examines junk news, followers of problematic sources as well as artificial amplification on Instagram during the 2019 Dutch provincial and European parliamentary elections. First, this study looks at the circulation of junk content in high-engagement political spaces on Instagram. Second, it takes up the question of the mainstreaming of Dutch junk news providers by looking at the intersection between the followers of Dutch political entities and those of junk news sources. Third, it looks at the presence of artificial engagement tactics (specifically fake followers) employed by Dutch political entities and news sources on Instagram. In all it was found that Dutch political Instagram is a relatively healthy space, but not for all issues or political entities.

Keywords: Instagram, artificial engagement, junk sources, fake followers, digital methods

Introduction: Fake followers, computational propaganda and their detection on Instagram

Though Facebook has been labelled the ‘hyperpartisan media machine’ (Herrman, 2016) and Twitter studied as a matter of routine, owing to the availability of datasets, Instagram, when scrutinized, has been found to perform well as an outlet for junk or hyperpartisan news circulation, artificially amplified engagement and other types of problematic content and users.

¹ The research was undertaken together with Rama Adityadarma, Joris van Breugel and Vic Krens.

The relationship between Instagram and different types of 'problematic information' (Jack, 2017) has been studied in connection with the Russian effort to influence the 2016 American elections. A study by New Knowledge (2018) found that Russian propaganda tactics played well on Instagram. The report analyses data from a variety of platforms, in order to detect efforts by the Internet Research Agency (IRA) to spread disinformation and divisive content. The study found that Instagram, with '187 million engagements' of '116,000 Instagram posts across 133 accounts' (New Knowledge, 2018: 7), to be at the forefront of an IRA operation, with better performing fake accounts and overall higher engagement than on Facebook.

The significance of Instagram, which 'outperformed Facebook' (New Knowledge, 2018: 8) as a battleground in the Russian disinformation enterprise, is linked, according to the report, to two possible causes. First, since it is a platform designed around sharing visual materials, Instagram may be well suited for the so-called 'image-centric memetic (meme) warfare' (2018: 8), that is, the weaponized use of image macros to stir conflict and foster division online. Second, the report states that the considerably high engagement of content from the IRA's accounts on Instagram may also be the result of click-farm activity, and some of the accounts in the dataset appeared indeed to be linked to 'a live engagement farm' (2018: 8).

With respect to the Dutch case, Russian influence has been studied mainly on Twitter, with the detection of trolling activities, especially in the aftermath of tragic and divisive events both in the Netherlands and in Belgium. For example, two journalistic studies found peaks in Russian trolling activity following the downing of MH17 in 2014 (Kist and Wassens, 2018; van der Noordaa and van de Ven, 2018a), while another study uncovered a (rather unsuccessful) organized Russian effort in spreading anti-Islam content on Twitter after the 2016 Brussels airport attacks (van der Noordaa and van de Ven, 2018b). Despite the lack of empirical research regarding Russian influence on Instagram, one study from the *NRC Handelsblad* (Kist and Wassens, 2018) suggests that a larger organized trolling activity may be found on other platforms beyond Twitter, including Instagram.

The use of computational means to amplify misinformation and hyperpartisan content on Instagram has not been linked exclusively to the Russian propaganda operation in the West, but it has also been described as a domestic tactic, adopted by national campaigners as well. A comparative, global study of social media manipulation in 48 countries (Bradshaw and Howard, 2018) describes different computational tactics for political influence online, including the use of fake accounts to attack other users, automated accounts generating artificial engagement, and human-curated

accounts that employ automation to be more efficient. With respect to the Netherlands, the study found such automated activity to be mainly linked to the boosting of Geert Wilders' hashtags on Twitter. While the study describes Twitter as the platform where automation thrives the most, signs of 'cyber troop activity' (Bradshaw and Howard, 2018: 13) are also to be found in other platforms, including Instagram, among 25% of the countries studied.

That fake following and artificial engagement flourish on Instagram may also be noted from reported cases in the news. As a case in point, in June 2017, the Russian journalist, Vasily Sonkin, posted an image of a vending machine, placed inside a shopping centre in Moscow, that lets users buy Instagram followers and likes. The news that for the (cheap) price of 50 Russian roubles (about EUR 0.70) one could buy 100 fake Instagram likes was reported by numerous tech or news media outlets (Matsakis, 2017; Feldman, 2017; Tan, 2017).

There have also been efforts by Instagram itself to counter artificially amplified activities on the platform. In December 2014, Instagram announced a crackdown on fake (or improperly obtained) profiles, in the so-called 'Instagram rapture' (Lorenz, 2014) that resulted in the deletion of hundreds of thousands of accounts. And later, in December 2018, a 'Christmas crackdown' (Lorenz, 2018) resulted in the shutdown of 500 meme accounts, some of which with millions of followers, suspected of using stolen or traded profiles. On the same note, in April 2019, Facebook filed a lawsuit against one company based in New Zealand, accusing it of providing 'fake likes, views and followers to Instagram users' (Romero, 2019). The lawsuit is presented as part of a larger effort by the platform to prevent 'inauthentic behaviour' on Instagram. While actions have been taken to cope with artificial engagement as well as fake or improperly obtained accounts, the platform has been apparently less active in limiting the spread of extremely coloured or hyperpartisan content, but rather has become the 'Alt-Right's new favourite haven' (Sommer, 2018), offering refuge to extreme right-wing personalities, after their accounts are deleted from Twitter.

The presence of a large automated engagement infrastructure on Instagram is also indicated by the deluge of fakeness-detection tools offered by commercial services online. The detection of inauthentic automated activity on Instagram may work by fingerprinting one account's follower base by nationality, and associate specific geographical locations, such as Brazil, Turkey or China, to suspected bot activity (Maheshwari, 2018). For example, among the available tools, *HypeAuditor*, a '100% AI-powered' service to expose 'fake followers and engagement' on Instagram, flags certain countries, such as Brazil, as geographical locations that may signal the presence of fake followers (Komok, 2018).

Despite being understudied, specifically in the Dutch context, Instagram appears to be a platform prone to the presence of various instantiations of junk and fakeness. There is the presence of content that can be described as false as well as merely hyperpartisan and divisive, but deliberately pushed online in order to stir conflict in a political space, both from outside the country and from within. It is also a platform prone to various computational tactics (such as bot work, fake likes and fake following) employed as a means to artificially amplify that same content.

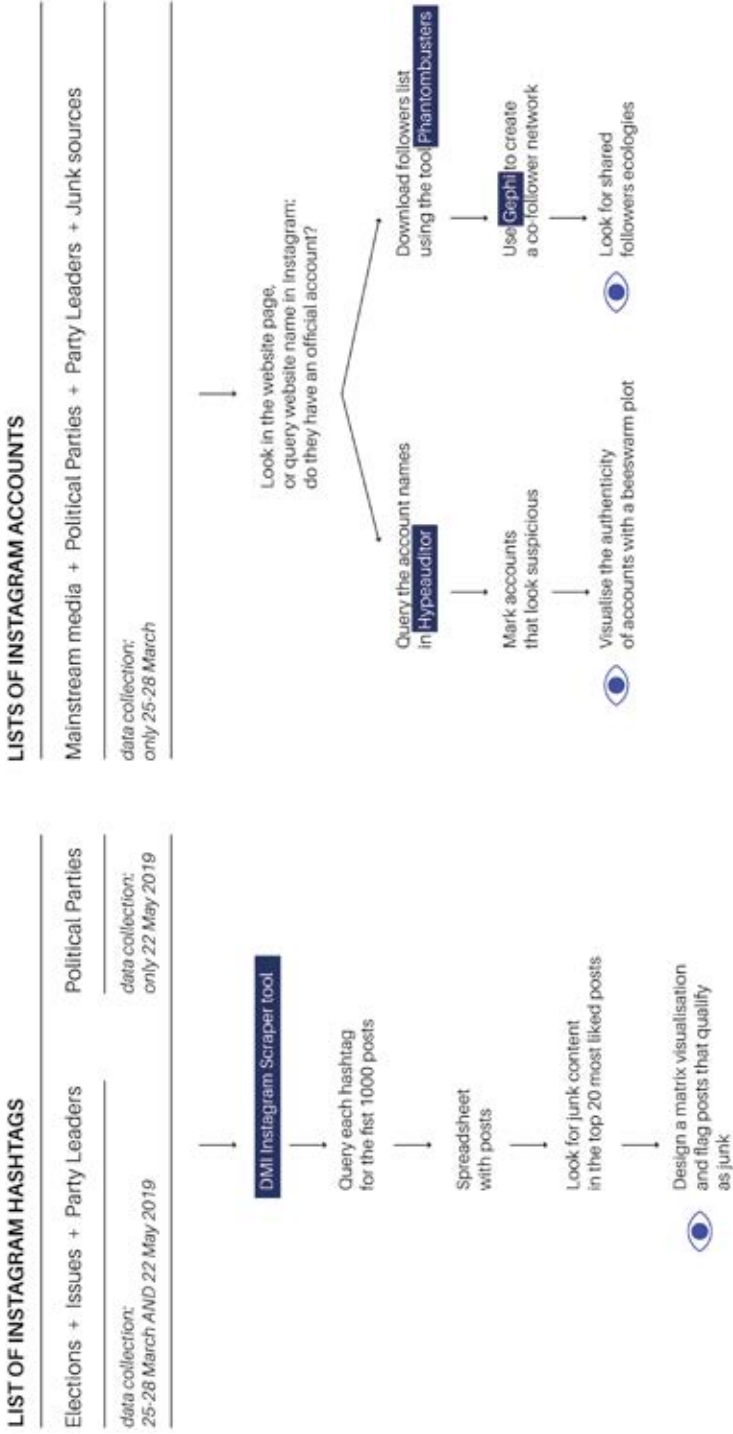
Junk content and artificial amplification in the Dutch political space on Instagram

As mentioned, a few studies have described Instagram as fertile ground for the distribution of inflammatory content in the form of memes, but also as a well-performing infrastructure for the artificial amplification of engagement. In this empirical research project, we devised three complementary approaches for the assessment of Dutch political Instagram in order to test these premises. They study the amount of junk content shared on the platform, the dubiousness of the most relevant information sources within the space, and the inauthenticity of followers that may generate artificial engagement (see Figure 5.1).

In the first part of this study, we search for levels of junk content shared on the platform, by asking to what extent the most liked content in a demarcated Dutch political space on Instagram can be defined as junk (i.e., disinformation, conspiracy, clickbait or hyperpartisan). Second, as the estimation of junk also can be made through 'online source criticism' (Rogers and Niederer, this volume), we expand the work by detecting it on Instagram at a source level. Here, we study the mainstreaming of junk sources by exploring the affinity of the follower bases of Dutch political entities with those of junk news providers (flagged as such by experts). We ask, to what extent do Dutch political entities share an audience with junk news sources on Instagram? Thirdly, in order to study the tactics of artificial engagement that political parties and news sources may employ on Instagram to boost their content, we search for signs of inauthentic activity in the follower bases in the Dutch political space and inquire into the extent of any efforts at artificially boosting (by means of fake followers) present around divisive topics on the Dutch Instagram.

In this research we employ a 'digital methods' approach (Rogers, 2013), repurposing Instagram-specific features to gather data from the platform. In particular, we compile a list of hashtags and profiles in order to demarcate the Dutch political space on Instagram. Within this space, we collect and

Figure 5.1 Diagram of the research protocol, showing the type of hashtags and accounts used for querying Instagram, and the tools used to collect, visualize and analyze the data



analyze most liked posts (i.e., posts that receive a certain number of likes) to study junk in the shared content. Furthermore, we collect followers of the accounts of Dutch political entities, mainstream news sources, and junk news sources, in order to study the intersection between their audiences, and more generally to assess the degree of junk in the Dutch political space at the level of the sources. Finally, we rely on profile features (such as the 150 characters bio in one's profile, or post captions), to evaluate the number of fake followers of the Dutch political space.

Detecting junk in the most liked content

The aim of the following analysis is to identify engaging content in the Dutch Instagram political space and observe the extent to which it contains junk – in the sense of how much engagement is generated by content that is either disinformation, conspiracy, clickbait or hyperpartisan.

To outline the Dutch political space in Instagram, we compiled a list of hashtags (see Table 5.1) that are used on the social network to identify the leaders of Dutch political parties (e.g., #markrutte), the 2019 Dutch provincial elections (#PS2019) as well as politically charged issues such as climate change (#klimaatverandering). We used the *Instagram Scraper* tool,² offered by the Digital Methods Initiative, to collect the 1,000 most recent posts per hashtag (data collected between the 25th and 28th of March 2019), together with their metadata (date of the post, media URL, caption, number of comments and number of likes). For each hashtag we selected only the 20 most liked posts, manually filtering out posts that are not relevant to the search criteria, or identical posts that prevent more diverse results from reaching the top 20.³

In this curated list of most liked posts, we conducted a close reading by looking at post captions and embedded media (images and videos) to understand how political party leaders and politically charged topics are discussed within the limits of the Instagram Dutch political space, and specifically to flag the presence of junk content.

As a result of this evaluation (see Figure 5.2), we found that out of the 400 most liked posts within our dataset there are (only) 45 posts that can

2 The tool is available at this link: <https://wiki.digitalmethods.net/Dmi/ToolInstagramScraper>

3 For example, we filter out posts about the Slovak professional footballer Marek Hamšík, who plays with the number 17 and is referred to in Instagram with the same hashtag of the Malaysia Airlines Flight 17 (#MH17). We also do not include in the dataset the identical posts of condolence messages for the Utrecht attack posted by Dutch national football team players with the hashtag #Utrecht.

be flagged as junk, 4 satirical posts, and 351 posts that do not appear to be junk. Looking at the engagement generated by these posts, junk content was liked 79,466 times, satirical content 37,532 times, and non-junk content 838,794 times.

In Figure 5.3, the 400 most liked posts are divided in hashtag-dedicated columns, in which they are also ranked from the most liked post in the first row to the least liked one in the last. Junk content is flagged using three different colours: light blue for hyperpartisan content, magenta for conspiracy, and blue for click-bait. Satirical posts are color-coded in dark blue. Finally, columns are ordered from left to right according to the amount of junk content, calculated on the total number of likes for each hashtag.

The analysis shows that the #zwartepiet, #geertwilders and #tunahankuzu hashtags represent the most divisive political spaces, with respectively 56.1%, 42.8% and 42.7% of their total amount of likes directed to junk content. Moreover, we find that the majority of the posts flagged as such can be considered hyperpartisan, mostly supporting and/or opposing particular ideology or figures, while only one post can be considered as clickbait, and one conspiracy. Generally, we did not find any trace of disinformation linked to the content that receives the most likes. The findings suggest that certain issues or political leaders, such as the Zwarte Piet debate and the leader of Denk political party, Tunahan Kuzu, draw more divisive content than others. Of the 20 most liked posts, however, we found no strong presence of junk.

In general, we found a relative scarcity of junk content in this high-engagement political space. In the top results for the Dutch provincial elections, #PS2019, we found only positive content, either celebrating preliminary poll results or encouraging people to exercise their right to vote. The hashtags, #24oktoberplein and #utrecht, returned mainly condolence posts and the news that the attacker was spotted and arrested. Almost all of the content we considered as junk is hyperpartisan. We found no presence of disinformation in the most liked results within the demarcated political space.

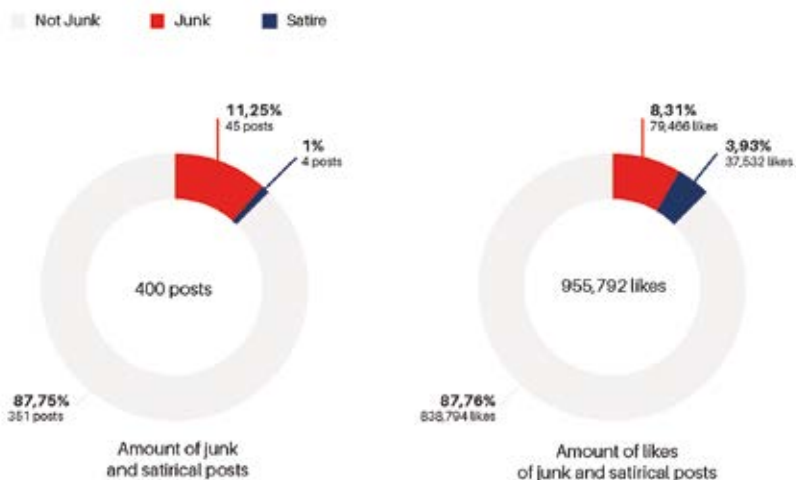
In order to ascertain the presence of junk content on Instagram surrounding the 2019 European Parliamentary elections in the Netherlands, we conducted a second hashtag analysis concerning content posted in the months before the election day (23 May). With the goal of demarcating the Dutch political space around the 2019 European elections, we compiled a new list of hashtags (see Table 5.2) used to identify Dutch political parties (e.g. #fvd) and their leaders (e.g. #thierrybaudet), the European elections

Table 5.1 Lists of hashtags pertaining to political leaders and politically charged discussions used to demarcate the Dutch political space on Instagram around the 2019 provincial elections

<i>Hashtags related to Dutch political party leaders</i>	<i>Hashtags related to politically charged discussions</i>
#markrutte, #rutte, #geertwilders, #wilders, #thierrybaudet, #baudet, #jetten, #tunahankuzu, #jesseklaver, #lodewijkasscher, #alexanderpechtold, #gertjansegers, #sybrand-buma, #mariannethieme	#PS2019, #klimaatverandering, #immigranten, #utrecht, #zwartepiet, #MH17, #24oktoberplein

Figure 5.2 Proportions of most liked content shared around the 2019 Dutch provincial elections, categorized as junk, satire, and not junk

Top 20 most liked posts per hashtag (400 posts in total)

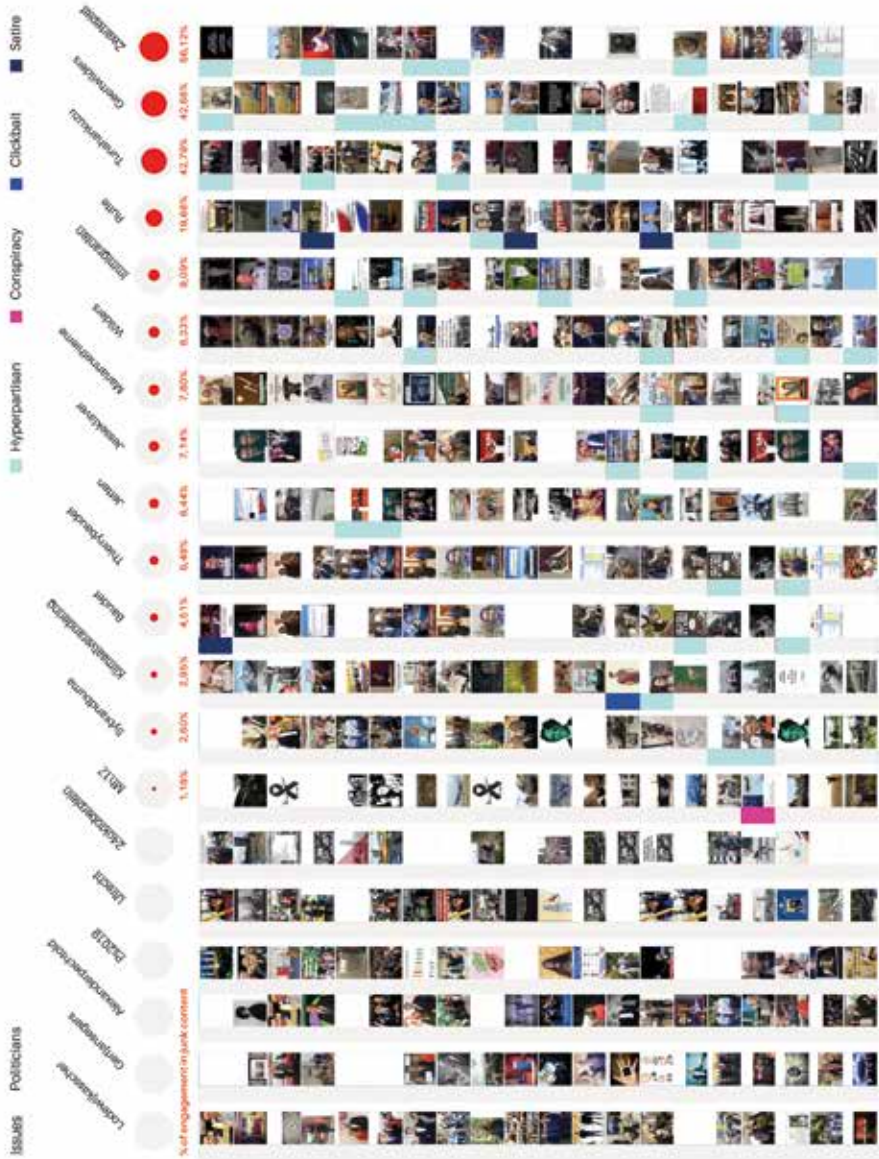


Data source: Instagram Scraper; data collection: 25-28 March 2019; pie charts

(#EUverkiezingen2019, #EUverkiezingen), and various politically charged issues such as immigration (#immigratie, #migratie, #immigranten) and climate change (#klimaat, #klimaatverandering).

With the Instagram Scraper tool, we collected the 1,000 most recent posts per hashtag (data collected on the 22nd of May) and their metadata. For each hashtag we only retained posts shared after the 28th of March, in order to focus on the detection of junk in the period prior to the European elections, but after that of the Dutch provincial elections. For each hashtag we selected the 20 most-liked posts, excluding those included in

Figure 5.3 20 most-liked posts per hashtag shared around the 2019 Dutch provincial elections, sorted from right (most junk) to left (least junk)



Data source: Instagram Scraper; data collection: 25-28 March; image wall

Figure 5.4 Examples of the posts flagged as hyperpartisan or satire

Data source: Instagram Scraper; data collection: 25-28 March; image wall

the dataset but not relevant to the search criteria.⁴ Subsequently, in order to ascertain the amount of junk in the dataset, we looked at embedded media and textual captions and flagged each post as junk or not (making the additional distinctions between disinformation, conspiracy, clickbait and hyperpartisan content).

The analysis (see Figure 5.5) confirmed the relative lack of junk content in the Dutch political space, also around the 2019 European parliamentary elections: out of 452 most liked posts, we found only 41 that can be considered junk (specifically hyperpartisan), counting for less than 10% of the total amount of posts. Moreover, hyperpartisan posts score low even in terms of engagement, generating only 4.66% of likes out of the total amount.

In Figure 5.6, most liked posts are organized in hashtag-dedicated columns. Columns are grouped by type of hashtag and sorted from right to left according to the number of likes generated by hyperpartisan content. The

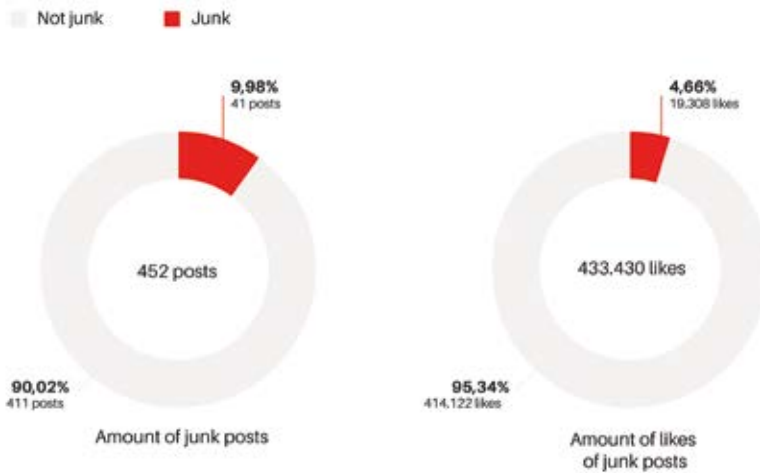
4 The query for some of the less popular hashtags returned less than 20 posts in the specified date range.

Table 5.2 Lists of hashtags pertaining to political leaders and politically charged discussions used to demarcate the Dutch political space on Instagram during the months before the 2019 European elections

<i>Hashtags related to Dutch political party leaders</i>	<i>Hashtags related to Dutch political parties</i>	<i>Hashtags related to politically charged discussions</i>
#markrutte, #rutte, #geertwilders, #wilders, #thierrybaudet, #baudet, #jesseklaver, #jjetten, #mariannethieme, #tunahankuzu	#cdavandaag, #pvv, #socialistischepartij, #pvda, #christenunie, #partijvoordedieren, #50pluspartij, #groenlinks, #fvd, #stempiraat, #volt nederland, #d66, #degroenen	#duurzaamheid, #klimaat, #klimaatverandering, #immigratie, #migratie, #immigranten, #mh17, #zwarte piet

Figure 5.5 Proportions of most-liked content shared around the 2019 European elections, categorized as junk and not junk

Top 20 most liked posts per hashtag (452 posts in total)



Data source: Instagram Scraper; data collection: 22 May 2019; pie charts

analysis shows that the hashtags related to political parties attracting more divisive content are #pvda and #christenunie, with respectively 27.6% and 25.1% of likes directed to hyperpartisan content. Compared to the dataset around the provincial elections, Geert Wilders (#geertwilders, #wilders) remains the political leader receiving the highest percentage of likes directed to hyperpartisan content (28.4%), followed by Jesse Klaver (#jesseklaver) with 13.4%, who instead scored low in terms of junk content in the previous analysis. We did not find traces of hyperpartisan content in the most-liked

posts around other political leaders. As was the case with the hashtags used to refer to the Dutch provincial elections, #EUverkiezingen2019 and #EUverkiezingen are related mainly to invitations to exercise the right to vote. Among the issues under study, #zwartepiet remains the most divisive one, with 22.9% of likes directed to hyperpartisan content.

In general, we did not find evident signs of dubiousness in the most-liked content around the 2019 European elections, except for a few hyperpartisan posts. The finding is aligned with that of the hashtag analysis conducted around the 2019 Dutch provincial elections.

Follower ecologies and the relevance of junk sources

In order to detect the relevance of junk news sources within the Dutch political space on Instagram, and to assess whether and how much junk news sources are becoming mainstream, we studied the overlap between followers of Dutch political entities, mainstream news sites and Dutch-language junk news sites.⁵ Specifically, we asked, to what extent are followers of junk news providers shared with those of Dutch political entities?

First, we demarcated the Dutch political space on Instagram, by compiling three lists of profiles: a list of Dutch political parties and their leaders, a list of Dutch mainstream media outlets, and the profiles of Dutch information sources flagged as junk in the expert list (see Appendix 6.2 in Hagen and Jokubauskaite, this volume). We then used the API Instagram Follower Collector by Phantombuster⁶ to collect the follower list of each Instagram account, and then, by creating a co-follower network, we looked at the amounts of shared followers between the political entities and the dubious Dutch information sources from the expert list.

In mapping the follower network of the Dutch political space, we found three distinct follower ecologies (see Figure 5.7). First, an ecosystem of followers of mostly established mainstream news organizations, such as the Dutch public broadcasting station, NOS. The follower bases of these news organizations are the largest in the network, which suggests that the Dutch mainstream news providers are still more relevant than those flagged

5 We use the list of sites flagged by the *Hoax-Wijzer* (www.hoax-wijzer.be), which was edited and enhanced by University of Amsterdam researchers, and is dubbed the 'expert list' (see Appendix 6.2 in Hagen and Jokubauskaite, this volume).

6 Phantombuster is an API store that 'provides ready-made cloud APIs to collect data from various social networks and improve marketing strategies' (phantombuster.com).

as junk, at least in terms of follower count. Few sites from the expert list are close to (or part of) the cluster of mainstream news organizations, due to a relatively high number of shared followers. Shared followers among mainstream news organizations and junk news sites may indeed suggest a special affinity among them, or rather be the signal of the mainstreaming of junk news providers.

A second ecosystem is made up of political parties and their youth organizations. The distribution of parties is laid out from left-wing to right-wing parties, whilst still being tightly clustered together. This may suggest that most followers either follow multiple parties on the same side of the political spectrum or follow all political parties regardless of political leaning. What can also be observed is the relative distance of the cluster of political parties to that of news organizations, suggesting that followers of political entities are mostly not shared with those of news organizations.

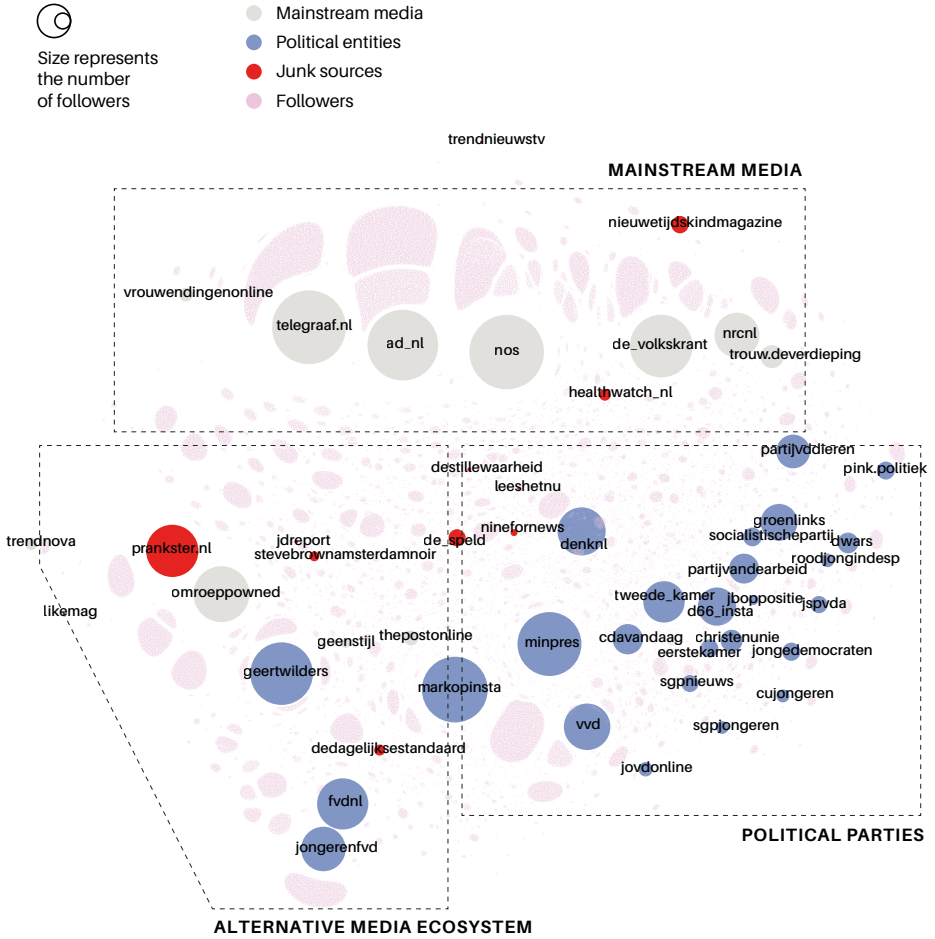
A third cluster is made up of right-wing political entities, which are far from other political entities, closer to few hyperpartisan or clickbait sites and to few, less established, mainstream news providers. Within this cluster, the account of PVV leader Geert Wilders is surrounded by *GeenStijl*, a tendentious 'shock blog' and PowNed, the public broadcasting station that is an offshoot of *GeenStijl*. The official profile of FvD (Forum for Democracy) and the youth organization of the same party are even more distant and isolated from other parties: they are surrounded by individual political commentators and share a high number of followers with the hyperpartisan news site, *De Dagelijkse Standaard*. This topology may suggest that although these parties and personalities share some followers with those from other sides of the political spectrum, they are mostly on their own and produce content consumed by a unique audience.

Fake followers and artificial engagement

In order to profile the follower base of the previously demarcated Dutch political space, we feed each account⁷ (of political entities, but also of mainstream media, and of those from the expert list) in the HypeAuditor tool to check the authenticity of the accounts and look for signs of artificial boosting and fake followers. With HypeAuditor one can profile an Instagram account to determine the authenticity of its follower base. To assess the

⁷ *HypeAuditor* analyzes only accounts with more than 1,000 followers. For this reason, we limited the detection of fakeness to accounts with more than 1,000 followers.

Figure 5.7 Follower ecologies in the Dutch political space, visualized as a co-follower network and manually annotated. In the network, accounts with higher amounts of shared followers (pink) are placed closer to each other.

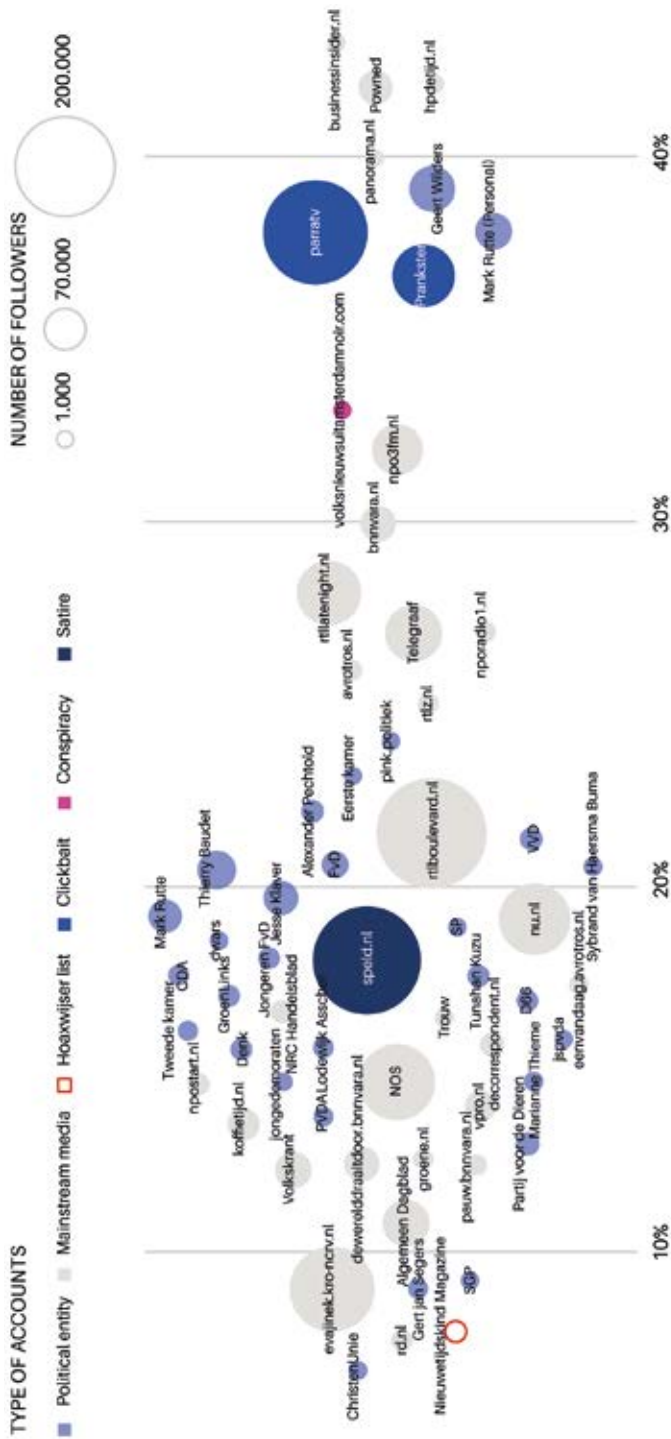


Data source: Phantombuster; data collection: 25-28 March; network graph

extent to which Dutch political accounts are employing artificial engagement tactics, we use reports from HypeAuditor, regarding the percentage of real followers, and their geographical origin.⁸ The percentage of fake followers

8 According to HypeAuditor, the geographical origin of one follower base is detected by analysing profiles biographies and place names in post captions (twitter.com/hypeauditor/status/1077143110432538624).

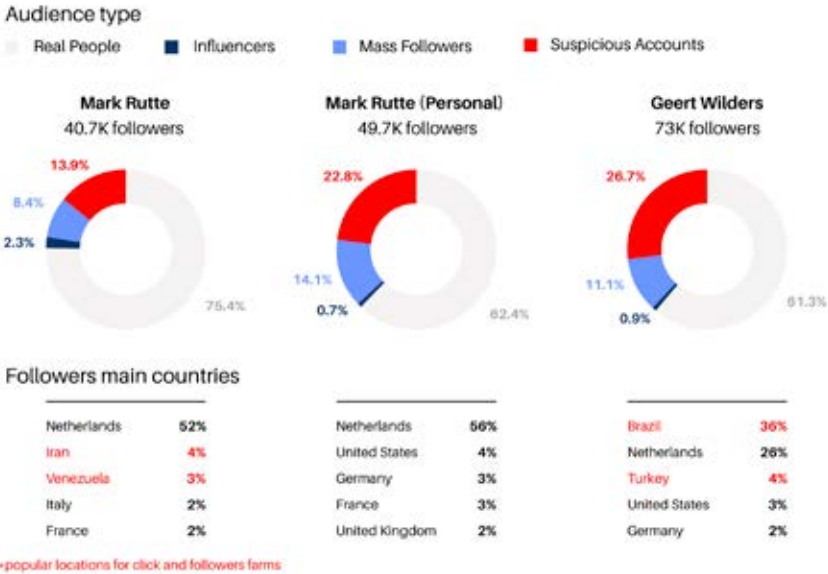
Figure 5.8 Degree of account fakeness according to report by the HypeAuditor tool. Accounts on the further right have more suspected 'fake followers' than accounts on the left side of the graphs.



Amount of fake followers (in %, estimated by the HypeAuditor tool)

Data source: HypeAuditor; data collection: 25-28 March 2019; bee swarm plot

Figure 5.9 Visualization of the follower base of Mark Rutte’s personal and work accounts and Geert Wilders’ account, based on results from the HypeAuditor tool. Each follower base is segmented based on ‘audience type’ and geographical provenance. Popular suspicious countries, that may suggest an inauthentic follower base, are coloured in red.



Data source: HypeAuditor; data collection: 25-28 March 2019; pie charts

returned by the tool is then used to rank each account from less fake to more fake (see Figure 5.8). Furthermore, we zoomed in on those accounts with a higher percentage of fake followers, to observe their geographical provenance (paying particular attention to suspicious countries), as well as the segmentation of the follower base provided by HypeAuditor, which breaks down followers in ‘real people’, ‘influencers’, ‘mass followers’ and ‘suspicious accounts’ (see Figure 5.9).

Generally, we found that the majority of profiles do not have a suspicious follower base, with most accounts scoring higher than 70% in the real follower metrics provided by the tool. There are some accounts, however, that are suspect of having a fake follower base. For instance, the media entity PowNed has 32.6% of suspicious followers. The clickbait site Prankster also scores relatively high in terms of fake following. Within the group of political entities, the personal account of Mark Rutte and the account of Geert Wilders have the highest number of suspicious followers. Strikingly, the ‘work’ account of the prime minister, Mark Rutte, has a lower percentage

of fake followers than that of his personal account. On the other hand, the account for the political party, Christenunie, has hardly any suspicious followers, just as the SGP (Reformed Political Party) and that of the minister Gert-Jan Segers.

When we look closer to the nationality of the follower bases, we found no suspicious results, with most of the accounts followed by users based in The Netherlands. For both of Mark Rutte's accounts, the followers are mostly based in the Netherlands. On the contrary, Geert Wilders account has 36% of his followers from Brazil. This raises some questions regarding the legitimacy of Geert Wilders' follower base, for Brazil is often mentioned as one location that can signal the presence of fake followers (Maheshwari, 2018).

In all the follower analysis does not show an organized effort of artificial boosting within the Dutch political Instagram sphere, and it indicates, with the exception of Geert Wilders, a rather authentic follower base.

Conclusions: Findings and limitations

The goal of the present research is to detect the scope of junk news and the degree of artificial amplification in the Dutch political Instagram sphere. More generally, it can be considered an attempt at applying to the Dutch context the argument in the New Knowledge report (2018) that Instagram performs well in terms of junk content circulation and artificial amplification strategies. It also takes up the invitation from the *NRC Handelsblad* study to inquire into other platforms than Facebook and Twitter for disinformation campaigning and computational propaganda.

The presence of dubious content (or lack of thereof) has been studied on three levels: at the story level (by looking at the circulation of junk content in high-engagement political spaces on Instagram); at the source level (by looking at the intersection between the follower bases of Dutch political entities and that of news sources flagged as junk); and through the detection of artificial engagement tactics, specifically fake followers, among the profiles of Dutch political entities as well as Dutch information sources.

In general, we found a rather healthy political space. Most liked content in the Dutch political space proved to be junk to a very small degree, although we found a small amount of hyperpartisan and polarizing content centred around more divisive figures and issues in the 2019 Dutch provincial and European elections. With respect to the alignment of the audience of Dutch political parties with that of (mainstream or junk) news providers, we found

mainstream news organizations to be still more relevant in this political space, somehow confirming the argument that in the Netherlands 'the vast majority of news consumption remains of the mainstream sources' (Rogers and Niederer 2019, this volume). Furthermore, the analysis of the follower base of Dutch political entities (and that of news sites, both mainstream and junk) revealed an apparent authentic audience with almost no signs of artificial engagement.

Within a relatively healthy political spectrum, it is at the extremes that junk sources and artificial amplification surface. With the current research we have pointed out a special affinity between right-wing political entities and some information sources that may be defined as junk (or at least hyperpartisan). Furthermore, the few indications of artificial engagement we have found are located at the far end of the political spectrum, with Geert Wilders' account being the most suspected of inauthentic activity.

In the co-follower analysis, we found that extreme political entities to have a unique follower base, not shared with other parties or mainstream news sites. Right-wing political entities are also relatively closer (in terms of shared followers) to suspicious sources (a few of them flagged by the expert list). Above all, Geert Wilders' account is the closest (according to shared followers) to hyperpartisan news sources. Relatedly, Geert Wilders' account is the only one of those under study that may reveal signs of artificial engagement, as suggested by a geographically dubious follower base. This finding resonates with the 2015 scandal about a suspicious increase of the follower count of Geert Wilders' Twitter profile. In addition, the already mentioned comparative study of social media manipulation strategies by the Oxford Internet Institute (Bradshaw and Howard, 2018) also refers to Geert Wilders as making use of various artificial boosting strategies in The Netherlands, reporting on an analysis by a social media analytics firm that in February 2016 found 26 fake accounts amplifying the #geertwilders hashtag on Twitter.

The determination of the relative absence of junk content, dubious sources and fake followers in the scope of the current research has a series of methodological limitations. First, in the search for junk news in the shared contents, we collected data based on a limited list of hashtags related to Dutch politicians and controversial topics. One could repeat the analysis to include other politically charged issues. Furthermore, we have considered only the top 20 most-liked posts per hashtag, whereas we could have also counted the number of comments per posts to analyze most engaged-with content. Moreover, we could have included in the analysis a larger set of

posts that do not necessarily make it to the top (because they receive fewer likes, or have fewer comments), in order to evaluate the presence of junk news in less engaged-with spaces. In addition, given that for data collection we made use of the DMI *Instagram Scraper*, which ‘scrapes Instagram to retrieve posts’ (Digital Methods Initiative, 2019), this research is dependent on the limits of such scraping, including Instagram’s rate limits which are not documented and unknown security challenges (Instaloader, 2019). It is also not a platform that invites research through scraping. As others have pointed out, social media platforms are designed to increase a platform’s commercial value, rather than to meet researchers’ needs (Borra and Rieder, 2014). To overcome the limitations, one could use additional tools for data collection and compile a richer data set.

Secondly, we established the fakeness in the Dutch political follower base using the metrics provided by a single tool (HypeAuditor). We could have compared the results with those by other similar services (and audited the auditors, so to speak). Moreover, we searched for signs of inflated engagement in the Dutch political space only by looking at followers’ demographics, while we could have paid attention to other signals such as patterns of repetition in posts comments. For example, to account for other tactics of artificial engagement on Instagram, one could perform a co-hashtag analysis⁹ in a demarcated issue space, and detect signs of (semi-automatic) boosting, such as the use of long list of popular unrelated hashtags, deliberately added in the post captions to increase content visibility.¹⁰ Moreover, one could trace back the users involved in this activity and profile them in order to evaluate their authenticity.

9 In addition to the most recent lists of posts, the *Instagram Scraper* tool returns a network of hashtag co-occurrences, that is, a file that contains the hashtags used at least once together with the hashtag under study. For each pair of hashtags, the tool returns a numeric value representing the total number of posts in which the two hashtags appear together in the data set. A similar approach is largely used for empirical research on Twitter: with co-hashtag analysis one can gain a sense of the relationship between subtopics in a conversation (Borra and Rieder, 2014); or find additional and/or more ‘significant hashtags’ (Rogers, 2017) to be queried to expand a corpus of data; or spot hashtags practices aimed at enhancing the visibility of particular content (Wang et al. 2016), or overturning its original meaning through hashtag hijacking practices (Berg, 2017).

10 Unlike Twitter, which has a character limit of 280 characters, Instagram’s character limit is 2,200 characters, and users can include up to 30 hashtags in the caption and comment sections of the post. This results in certain users adding blocks of more or less related hashtags to the posts to enhance their visibility. Even if Instagram is applying countermeasures to block the use of certain hashtags (Drewe, 2016), there are several websites that provide lists of safe and popular hashtags that users can copy paste directly in their posts (for example, tagblender.net).

References

- Berg, Kati Tusinski (2017). 'Social Media, Hashtag Hijacking, and the Evolution of an Activist Group Strategy.' in *Social media and crisis communication*, edited by Lucinda Austin and Yan Jin. Routledge, pp. 141-156
- Bradshaw, Samantha and Phillip N. Howard (2018) 'Challenging Truth and Trust: A Global Inventory of Organized Social Media Manipulation', Computational Propaganda Data Memo, Oxford: Oxford Internet Institute.
- Borra, Erik and Bernard Rieder (2014). 'Programmed method: developing a toolset for capturing and analyzing tweets'. *Aslib Journal of Information Management*, 66(3), 262-278.
- Digital Methods Initiative (2015) Instagram scraper. Available at: <https://wiki.digitalmethods.net/Dmi/ToolInstagramScraper>. [Accessed 22 April 2019].
- Drewe, Nick (2016) 'The Hilarious List Of Hashtags Instagram Won't Let You Search', *The Data Pack*, 10 May. <http://thedatapack.com/banned-instagram-hashtags-update/>
- Feldman, Brian (2017) 'In Russia, You Can Buy Instagram Likes From a Vending Machine', *New York Magazine*, 8 June.
- Herrman, John (2016) 'Inside Facebook's (Totally Insane, Unintentionally Gigantic, Hyperpartisan) Political-Media Machine', *The New York Times*, 28 August.
- Instaloader (2019) Instaloader software, version 4.2.5, GitHub project, <https://instaloader.github.io/>.
- Jack, Caroline (2017) *Lexicon of Lies: Terms for Problematic Information* New York: Data & Society Research Institute.
- Kist, Reinier and Rik Wassens (2018) 'Russisch trollenleger ook actief in Nederland', *NRC Handelsblad*, 15 July.
- Komok, Anna (2018) 'How to Check Instagram Account for Fake Followers', *HypeAuditor*, 6 July. <https://hypeauditor.com/blog/how-to-check-instagram-account-for-fake-followers/>
- Lorenz, Taylor (2014) 'Instagram Rapture Claims Millions Of Celebrity Instagram Followers', *Business Insider*, 18 December.
- Lorenz, Taylor (2018) 'Instagram's Christmas Crackdown. No meme account is safe – not even @God', *The Atlantic*, 27 December.
- Maheshwari, Sapna (2018) 'Uncovering Instagram Bots With a New Kind of Detective Work', *The New York Times*, 12 March.
- Matsakis, Louise (2017) 'This Russian Vending Machine Will Sell You Fake Instagram Likes', *Motherboard Vice*, 7 June.
- New Knowledge (2018) 'The Tactics & Tropes of the Internet Research Agency', White Paper, Austin, TX: New Knowledge.

- Romero, Jessica (2019) 'Preventing Inauthentic Behavior on Instagram', *Facebook Newsroom*, 25 April. <https://newsroom.fb.com/news/2019/04/preventing-inauthentic-behavior-on-instagram/>
- Rogers, Richard (2013) *Digital Methods*. Cambridge, MA: MIT Press
- Rogers, Richard (2017) 'Foundations of digital methods: Query design' in *The Datafied Society. Studying Culture through Data*, edited by Mirko Tobias Schäfer and Karin Van Es. Amsterdam University Press, pp. 75-94
- Rogers, Richard and Sabine Niederer (2019) *The Politics of Social Media Manipulation*, The Hague: Ministry of Internal Affairs.
- Sommer, Will (2018) 'Instagram Is the Alt-Right's New Favorite Haven', *The Daily Beast*, 30 October.
- Tan, Yvette (2017) 'There's a vending machine selling fake Instagram likes, because this is what we've become', *Mashable*, 7 June.
- van der Noordaa, Robert and Coen van de Ven (2018a) 'Hoe Russische trollen inspelen op westerse angsten', *De Groene Amsterdammer*, 28 August.
- van der Noordaa, Robert and Coen van de Ven (2018b) '3 miljoen "trollentweets" onderzocht: hoe Rusland via sociale media ook in België verdeeldheid zaait', *Knack*, 27 November.
- Wang, Rong, Wenlin Liu, and Shuyang Gao (2016) 'Hashtags and information virality in networked social movement: Examining hashtag co-occurrence patterns.' *Online Information Review* 40(7): 850-866.

About the authors

GABRIELE COLOMBO is a design researcher. In 2018 he completed his Ph.D. research at Politecnico di Milano. His research and teaching activities revolve around the design of visual tools in support of digital social research, focusing on the design of novel strategies for the analysis of collections of images and videos.

CARLO DE GAETANO is a designer and researcher with the Visual Methodologies Collective, Amsterdam University of Applied Sciences. He focuses on data visualization for social and cultural research. Carlo is interested in speculative methods with machine learning, images as data and the mapping of social issues.