

The Norwegian Twittersphere

Structure and Dynamics

Axel Bruns & Gunn Enli

Abstract

This article takes a new approach to the comprehensive study of an entire national Twittersphere. It identifies, to the extent that this is possible with the data made available through Twitter's Application Programming Interface (API), all accounts operated by Norwegian users and institutions, analyses patterns in their public profile information, and maps their follower/followee connections with each other. This provides new insights into the historical development of the Norwegian Twittersphere, its current network structure and the presence of diverse interests and issues amongst the nearly one million accounts within this community. Its findings also constitute important background information for future Twitter research that takes the familiar hashtag studies route: its observations enable such studies to filter their datasets for confirmed Norwegian accounts only, and to examine the presence of accounts with specific interest profiles, as determined by the present study, in their datasets.

Keywords: Norway, social media, Twitter, network mapping, echo chambers

Introduction

Twitter is by now well established in most developed nations as a key social media platform for public communication. Compared to larger rivals such as Facebook and Instagram, it is seen as especially relevant for relatively fast-paced, close to real-time interactions; this is due especially to its deliberately character-limited message format, which resembles oral communication more than longer-form writing. Additionally, a range of platform design choices – such as the ability to follow the updates posted by other users without, generally, needing those followees' permission, and the public visibility of almost all tweets, even to visitors to Twitter.com who have not registered an account of their own – have shaped the structure of Twitter's social network and the nature of user activities on the platform to a considerable degree. While in principle it is possible to set one's account to 'private', so that tweets are visible only to approved followers, only a very small percentage of Twitter users have chosen this option; the remainder are engaging, more or less actively, in a globally distributed public conversation.

This has meant that Twitter has been rightly recognised as highly influential, especially for its role in public information sharing and discussion, in the context of major

breaking news, key shared events in society, culture and sports, and the viral distribution of memes (Papacharissi & de Fatima Oliveira 2012; Highfield 2013, 2016; Vis 2013; Hermida 2014), which are often driven and loosely coordinated by one or more central hashtags. In such cases, Twitter users gather together rapidly to form *ad hoc* (Bruns & Burgess 2015), and sometimes affective (Papacharissi 2014), publics that compile, curate and share information, opinion, memes and other content related to their shared issue of interest; these publics persist for as long as the issue itself is alive, but often dissipate soon after. Indeed, a considerable proportion of Twitter research to date focusses on such temporary publics, and might be categorised under the overall banner of ‘hashtag studies’ (see Rambukkana 2015 for a recent collection of such work).

Such studies have produced valuable insights into the role of Twitter in the context of acute events (Burgess & Crawford 2011), but remain limited in their ability to contextualise their own observations against the broader range of activities taking place on Twitter at any one point. For instance, they might report a substantial absolute number of users participating in the discussion of a breaking news event, but are generally unable to express that number as a percentage relative to the total number of Twitter users active at the time; this means that it is impossible for them to say whether even an event as widely recognised as the Arab Spring (cf. Meraz & Papacharissi 2013) attracted a considerable proportion of Twitter users as followers of or participants in the discussion, or whether the majority of the global Twitter population simply continued their everyday activities without taking much notice of the dramatic events unfolding in Tunisia, Libya, Egypt and elsewhere. Further, the focus of many such studies on a small set of hashtags and keywords in their data gathering also means that they are likely to miss out on relevant tweets that (accidentally or deliberately) fail to include the keywords tracked; this especially includes follow-on @replies that might respond to a hashtagged tweet but do not themselves contain the hashtag again (Burgess & Bruns 2015). Finally, too, the very limited range of immediately available user profile metadata makes it difficult for these studies to target their findings to a set geographical area, even where this would be very valuable – for instance, to examine only those Arab Spring tweets originating in countries of the Middle East and Northern Africa, or to study only those tweets discussing the presidency of Donald Trump that were posted by users based in the United States. This means that the majority of current Twitter research is forced to make only global observations which do not necessarily offer much insight into the uses of Twitter in specific national and local contexts.

This article addresses some of these limitations by taking a new approach to the comprehensive study of an entire national Twittersphere. Its main purpose is to investigate the key characteristics of the historic and present-day Norwegian Twittersphere; in particular, in doing so it also examines the role of distinct national languages such as Norwegian in a globalised social medium where English dominates, and assesses the degree to which ‘echo chambers’ (Sunstein 2009) are affecting networked communication patterns. Norway’s relatively small population of 5 million, its advanced ICT infrastructure and its high usage of online and social media alongside a similarly high usage of public service media and mainstream news media (Syvertsen et al. 2014) position it as a useful case study, whose findings might subsequently be compared with other country-level Twittersphere studies in Europe and beyond. Previous studies of Twitter in Norway and Scandinavia have been concerned with relatively limited groups

of accounts and focussed on specific thematic clusters; most typically, they have studied Twitter activities around politicians and journalists (Larsson & Moe 2013; Skogerbø & Moe 2015).

In contrast, by identifying, to the extent that this is possible with the data made available through Twitter's Application Programming Interface (API), all accounts operated by Norwegian users and institutions, analysing patterns in their public profile information, and mapping their follower/followee connections with each other, we are able to provide new insights into the historical development of the Norwegian Twittersphere, its current network structure and the presence of diverse interests and issues amongst the nearly one million accounts within this community. Our findings are valuable in their own right – and represent only the second national Twittersphere to be mapped in this way, after Australia (Bruns et al. 2014, 2017) – but also constitute important background information for future Twitter research that takes the familiar hashtag studies route: our observations enable such studies to filter their datasets for confirmed Norwegian accounts only, and to examine the presence of accounts with specific interest profiles, as determined by our study, in their datasets.

The major contributions made by this article are twofold. First, and most importantly, we find that the Norwegian Twittersphere is composed of diverse but highly interconnected communities that emerged in earnest in 2009, but received an even more substantial boost in membership when Twitter, Inc. began to offer a Norwegian user interface for the platform in late 2011. Thematic interests in these communities range from business and politics to entertainment and interpersonal communication, with several specialist and minority topics also catered to. The analysis of the overall structure of the Norwegian Twittersphere which we perform in the present article provides a basis for further research into the uses of Twitter in Norway: drawing on our findings, researchers will be able to follow, analyse and compare the activities of selected Norwegian accounts and groups (as selected from the network clusters we have detected, or based on the metrics we have calculated) against other account populations, for example, and thereby to study activity on Twitter for a purposive sample of Norwegian accounts with defined characteristics. This enables research approaches that are considerably different from the still-predominant approach of gathering tweets (and, by extension, accounts) simply because those tweets contain specific keywords or hashtags: instead, they are able to capture all public tweets by these accounts, and thereby to produce a more comprehensive picture of the full range of Twitter activities within a specific community of accounts. We aim to build on the work presented here by pursuing such opportunities, and we encourage other Nordic Twitter scholars to join us in this effort.

Further, the present article documents the methodological steps taken in identifying the accounts that comprise the Norwegian Twittersphere on a large scale, and in analysing their activity and interconnection patterns. Adjusting for variations in national and cultural contexts, we hope that this outline of our methods will enable other research teams to replicate the study presented here for different national or regional Twitterspheres, and perhaps even to apply these approaches to the analysis of other social media platforms.

Identifying Norwegian Twitter accounts

Our study largely follows the data gathering and filtering process outlined in Bruns et al. (2017). At the time of our data gathering in early 2016, Twitter user IDs were arranged sequentially in order of account creation, counting up from 0 to close to 5,000,000,000. We used the Twitter API to systematically retrieve the profile information associated with each ID, until in February 2016 we found accounts with IDs above 4,900,000,000 that had been created only minutes before we gathered their profile information. This meant that we had exhausted the then current Twitter ID numberspace and reached the leading edge of continuing account creation. It should be noted here, of course, that not all of these IDs are associated with an account: the earliest profile information we were able to retrieve was for Twitter co-founder and CEO Jack Dorsey (@Jack), with ID 12, and we also noted considerable blocks of ID ranges that did not have any users associated with them. In addition to natural churn – where Twitter accounts are created at one point and subsequently suspended or deleted, leaving IDs unassigned – we assume that some such ID ranges were left unassigned for testing purposes at various stages of the development of the Twitter platform. In total, as of late February 2016, our approach gathered profile information for nearly 1.38 billion accounts.

To identify Norwegian accounts within this global userbase, we utilised a number of the available datapoints for account profiles. Contrary to the approach taken by Bruns et al. (2014, 2017), which in part drew on the display timezones users had set for their Twitter profiles, this field proved useless for our purposes: in contrast to the distinct Australian timezones that the previous study was able to utilise, Twitter does not offer a dedicated Norwegian timezone setting, and users are instead forced to use one of the other options representing Central European Time (e.g. '(GMT+01:00) Copenhagen'). At the same time, however, we were able to filter our global accounts dataset by the interface language that users had chosen for their accounts (this was impossible for the earlier Australian study, due to the prevalence of English as the primary language in that country).

We further constructed a list of the 100 largest cities in Norway, as well as of the twenty Norwegian counties, and filtered our global data for accounts that mentioned these locations in the free-form 'Location' and/or 'Description' fields of their profiles. This list was further supplemented with a number of more generic terms, including 'Norge', 'Norsk', 'Norway' and 'Norwegian'. The first results of this filtering process were manually reviewed in order to identify and eliminate false positives resulting from these filter terms occurring in other contexts; for instance, this was used to distinguish Bergen in the Norwegian county of Hordaland from Bergen op Zoom in the Netherlands, or Oslo, Norway from Oslo, Minnesota. At this stage, we chose to exclude Norwegian cities Alta and Ski from our filter altogether: 'Alta' (meaning 'high') is a very common term in several Romance languages, and 'ski' referred most frequently to the sport rather than the town; in both cases the false positives overwhelmed the true positives. While this could mean that accounts from these locations may be underrepresented in our final dataset, we note that a significant number of such accounts would also have exhibited secondary traits of Norwegianness that we will have identified, including mentions of their counties or of Norway itself (e.g. by using 'Alta, Finnmark, Norge' as a location text) or a Norwegian interface language setting.

In total, by using these filters we identified some 621,000 accounts with a language setting of 'no' (for the generic Norwegian setting), 10,000 accounts with 'nb' (for

Bokmål) and four accounts with ‘nn’ (for Nynorsk).¹ We also identified some 392,000 accounts with matching ‘Location’ information, and 100,000 accounts with matching ‘Description’ information. As many accounts matched two or three of our selection criteria, this resulted in a total dataset of some 988,000 accounts, and the remainder of this article discusses our analysis of this Norwegian userbase dataset. We note here that our approach will necessarily have missed Norwegian accounts that were not inherently identifiable as Norwegian, using these criteria; to be classified as false negatives by our filters, such accounts would have to use Twitter with an interface language other than Norwegian and not include any of the locations or other keywords from our list in their ‘Location’ and ‘Description’ information. This may especially apply to the most recent accounts at our time of gathering data, if their users had not yet made the effort to fully set up their profiles. The exact number of such false negatives is inherently unknowable; however, while such false negatives are excluded from the analysis we present in the following sections, it is unlikely that their absence would significantly skew our overall results.

Further, of course, our dataset constitutes a snapshot of the identifiable Norwegian Twittersphere as of February 2016: many more Norwegian accounts may have been created in the meantime, but equally some of the accounts we did identify at that point may have been deleted or suspended since then, and many other accounts might have been created in the past but already been deleted before we gathered our data and would therefore leave no trace in our dataset. This is an unavoidable aspect of any research that examines as dynamic a space as Twitter or the other major social media platforms.

The history of Twitter in Norway

We begin our analysis of this dataset by examining the historical development of Twitter use in Norway. Based on the account creation dates available for each Twitter account, Figure 1 shows the total number of new Norwegian accounts created per month since the launch of Twitter in March 2006, and still in existence by the time we gathered our dataset in early 2016. The overall pattern of growth in the early years of the platform mirrors similar dynamics at a global level, and as reported in Bruns et al. (2014) for the Australian Twittersphere: Twitter largely remained a niche platform until it emerged to greater media and popular attention in early 2009. We see at this point a sudden increase in new account creations to up to 20,000 new accounts per month (and recognising the possibility of subsequent account deletions, the real number at the time would have been even higher). After this initial spike in sign-ups, the Norwegian Twittersphere continued to grow steadily if slowly for another two years.

Early in 2012, however, Twitter experienced another, even more substantial and rapid influx of new Norwegian users. It is notable here that – unlike the earlier cohort – the vast majority of accounts created from this point onward are today using Twitter with a Norwegian-language Web interface; it is almost certain, therefore, that this new influx is related to Twitter’s launch of a Norwegian interface on 21 December 2011 (@twitter 2011). (The small percentage of older accounts that are now using a Norwegian interface language setting would have switched to Norwegian between December 2011 and our data gathering in 2016.) This pattern demonstrates that even for a country such as Norway, with a population widely able to read and write English,

the availability of Twitter in the native language still makes a considerable difference for the platform's popularity; additionally, of course, it is also likely that such adoption would have been boosted by contemporary Norwegian media reports about the new availability of Twitter in Norwegian. From 2012 onwards, therefore, the total number of new Norwegian accounts created each month rose substantially, and remained at a comparatively high level until 2015; the number of new Norwegian accounts using Twitter in English, by contrast, dropped substantially from its pre-2012 levels.




Figure 1. *New Norwegian Twitter accounts created per month, with interface language settings current as of February 2016.*

This pattern was interrupted, however, by a precipitous drop in new account creation between November 2013 and February 2014. Surprisingly, this decline only affected the Norwegian interface language component of the total Norwegian Twitter userbase; new sign-ups of accounts that used Twitter with an English-language interface remained steady throughout this period. We assume, therefore, that the Norwegian interface language option might have disappeared – by accident or by design – from the Twitter Website during this time, and that this served to put off prospective new users who might otherwise have joined during this time. It is notable here that the creation of accounts selecting Bokmål as their interface language commenced around the same time – we speculate, therefore, that the Twitter interface language selection process might have changed to offer dedicated ‘Bokmål’ and ‘Nynorsk’ options, but that – perhaps due to non-Norwegian developers’ confusion over these language choices – the generic ‘Norwegian’ language option was temporarily removed. In turn, the creation of new accounts using a Bokmål interface ceased in July 2015, suggesting another change to the Twitter language options; at the time of writing, the Twitter settings now again offer only a generic ‘Norsk – Norwegian’ option.

Other, smaller fluctuations in new account sign-ups were most likely related to domestic and international events that might have drawn public attention to the utility

of Twitter as a platform. Space available in this article does not permit us to examine these in detail; we do note, however, a substantial increase (from below 7,500 new accounts per month to more than 10,000 accounts) in July 2011, and suggest that this indicates a response to the role played by Twitter as a medium for crisis communication during the terrorist attack in Oslo and Utøya that month (Perng et al. 2013; Kaufmann 2015). Such events can have a lasting impact on public attitudes towards social media platforms, especially if the tenor of mainstream media coverage of such platforms also changes as a result.

Mapping the Norwegian Twittersphere

Network structure

In addition to compiling this general overview of the dynamics of Twitter adoption in Norway, we subsequently gathered the follower/followee connections for the 988,000 accounts we had identified: in total, the 988,000 accounts followed some 98.7 million Twitter accounts in all parts of the world. We further filtered these connections for links between Norwegian accounts only, thus excluding international connections. Partially excluded from this dataset are also the 39,000 accounts (or 4 per cent) set to 'private': for these it is impossible to determine what other Twitter accounts they themselves are following, though we are able to detect when other Norwegian accounts are following them. Of course, accounts which neither follow nor are followed by other Norwegian accounts also disappear from our analysis and mapping of the Norwegian Twitter network; this is particularly likely to eliminate remaining false positives that were misrecognised as 'Norwegian', since such accounts are most probably not connected to genuine Norwegian accounts. It will also remove inactive accounts that were abandoned soon after their creation and thus never developed a significant network of follower relationships. In total, then, this leaves some 716,000 of the total number of Norwegian accounts we had detected (or 72 per cent) that have at least one follower/followee connection to another Norwegian account; between these accounts, there are some 32.5 million distinct connections.

We imported this dataset into the open source network analytics software Gephi (Bastian et al. 2009), using the Force Atlas 2 algorithm (Jacomy et al. 2014) to visualise the network structure. Force Atlas 2 uses the common force-directed approach to positioning individual network nodes (our Twitter accounts) in relation to each other: each edge (our follower/followee connections) between two accounts draws these accounts closer together, while accounts without such connections repel each other. Aggregated over the entire dataset, this determines the positioning of all nodes in relation to each other, and usually leads to the placement of tightly connected groups of nodes in dense clusters, which in turn are at some distance from other clusters with which they share few connections. This provides a visual indication of distinct regions in the network, with the implication that such cluster formation is determined by shared interests or identities amongst the nodes located in each cluster. Figure 2 shows the overall structure of the network; the accounts with the largest number of Norwegian followers are shown in red here.




Figure 2. *Norwegian Twittersphere network. Edges not shown. Nodes positioned using Force Atlas 2 algorithm in Gephi (Jacomy et al. 2014); node colour gradient from blue to red represents increasing number of followers.*

We further utilised the Louvain community detection algorithm (Blondel et al. 2008) as implemented in Gephi in order to establish an independent computational perspective on the existence of such clusters. This algorithm takes a different approach to the detection of groups of unusually densely connected nodes in the network; in particular, it accepts a resolution parameter that determines how coarse or fine the individual communities it detects will be. Testing a number of different resolution settings, we found that a resolution parameter of 0.5 in Gephi produced a set of communities that largely matched the visual distinctions identified by the network visualisation.

In a final qualitative step in the analysis, we then examined the composition of each of the thirty largest clusters in order to infer a shared interest or identity that had guided the formation of the cluster; together, these thirty clusters contain some 97 per cent of the 716,000 accounts in the overall network. To do so, we focussed on the twenty accounts in each cluster that had the largest number of connections in the network, and reviewed their available profile information; these central, leading users tend to offer a good indication of key themes, topics, interests and identities within the larger cluster community. This labelling of clusters through a close reading of profile information is necessarily an interpretive exercise that relies on the researcher's familiarity with the Norwegian Twitter population, but does provide an important tool for the explanation of our subsequent quantitative findings.

The overall structure of the Norwegian Twittersphere that emerges from this analysis is of a broad division between news and politics, and other matters in public debate, to the right of the network, and more popular, everyday, and interpersonal uses to the left. Additionally, there is a substantial cluster of accounts focussing largely on football (and, by extension, on other sports) at the top of the map. Notably, strong geographic divisions

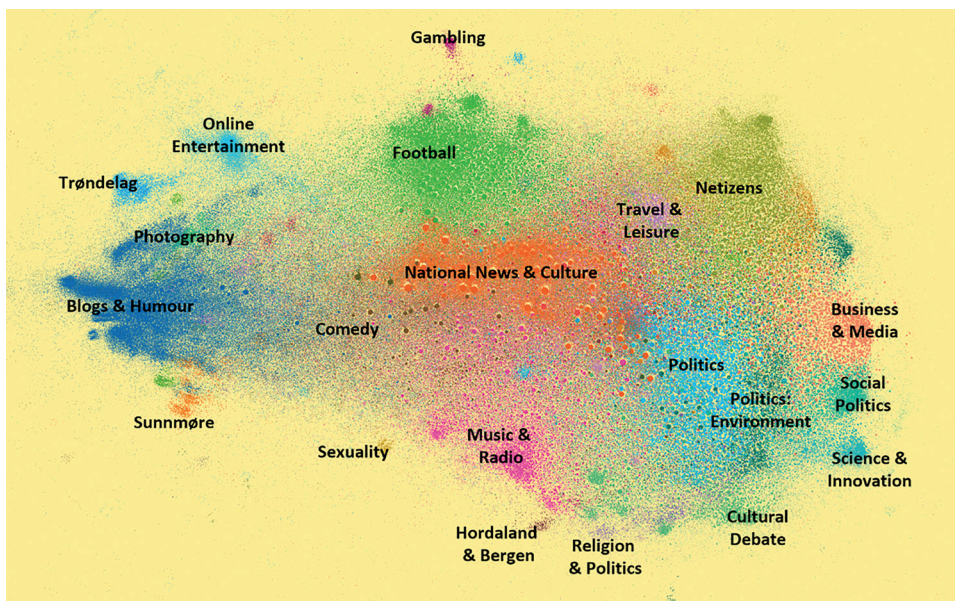


Figure 3. Norwegian Twittersphere network. Node colour based on Louvain community detection algorithm (modularity resolution 0.5; Blondel et al. 2008), implemented in Gephi. Clusters labelled following qualitative review of leading accounts in each community.

are largely absent, even in spite of the considerable distances between Norwegian population centres outside of the Oslo-Bergen axis. A small number of regional communities do emerge, usually on the edges of the network, for areas such as Sunnmøre, Trøndelag, Østfold, and even Bergen, but for the most part these only constitute a subset of the total population of accounts associated with these locations, and may mainly represent the accounts of local public and media institutions that have made a concerted effort to follow and be followed by as many local accounts as possible. Overall, by contrast, the Norwegian Twittersphere is fairly well connected throughout, which should mean that information will be able to circulate throughout the network with relative speed.

Activity metrics

Table 1 presents a number of key statistics for each of the thirty largest clusters; the remaining smaller clusters, amounting to 3 per cent of the total account population, have been combined as ‘Others’. To begin with, this clearly shows the very substantial role of the two largest clusters, ‘Blogs & Humour’ and ‘Oslo/National News & Culture’. Comprising some 131,000 and 118,000 member accounts, respectively, these two clusters account for nearly 35 per cent of the entire Norwegian Twitter population; each is more than twice as large as the next biggest cluster, ‘Sports: Football’. There are also a number of very significant differences between these leading clusters, however. First, ‘Blogs & Humour’ contains a considerably larger number – 4,855 – of ‘protected’ accounts (that is, accounts whose owners have chosen to make their tweets visible only to approved followers), while ‘Oslo/National News & Culture’ includes fewer than 1,000

such accounts. This is in keeping with the distinction between the more everyday and interpersonal uses of Twitter in ‘Blogs & Humour’, where we might expect to see a greater number of ordinary, private users, as compared with the greater number of official, public accounts in ‘Oslo/National News & Culture’. As we have seen in Figure 2, ‘Oslo/National News & Culture’ and other clusters in the same region of the network also feature a greater proportion of accounts with large numbers of followers, further underlining this distinction.

At the same time, the two clusters have also produced very different volumes of tweets: while the members of ‘Blogs & Humour’ had generated more than 100 million tweets (or an average of 767 tweets per account) by the time of our data gathering in early 2016, the ‘Oslo/National News & Culture’ accounts had posted only some 7.7 million tweets during their time on Twitter (65 tweets per account). This, too, aligns well with the different thematic focus of the two clusters: the interpersonal, everyday topics addressed by ‘Blogs & Humour’ appear likely to lead to more activity by a broad range of accounts; in ‘Oslo/National News & Culture’, posting activity could be expected to be dominated by a handful of official media, political and government accounts, while others in this cluster might be more likely to engage in ‘listening’ activities (Crawford 2009) that do not themselves produce a substantial number of posts. We note, by contrast, that some of the smaller clusters relating to ‘Political’ as well as ‘Cultural Debate’, as well as to various areas of policy, have generally tended to produce a much more substantial volume and rate of tweeting activity.

This perspective is also borne out by the figures on the average and median numbers of tweets per day for accounts in each cluster. Here we make use of the fact that the Twitter API provides us with the exact date that each individual account joined Twitter; combined with the similarly known total number of tweets posted by each account, this enables us to calculate an average rate of tweets posted per day over each account’s lifetime. In Table 1, we thus present the average and median tweeting rates for the account population in each cluster; for the two leading clusters, this shows both that the average rate of tweeting is considerably greater in ‘Blogs & Humour’ (at 0.6840 compared to 0.0966 in ‘Oslo/National News & Culture’), and that the distinctions between comparatively active and inactive accounts are starker in ‘Oslo/National News & Culture’ (its median rate of only 0.0007 means that a large part of that cluster hardly tweets at all, and that the majority of tweets from that cluster originate from a small group of leading accounts). Notably, across all metrics, ‘Oslo/National News & Culture’ also performs significantly below the Norwegian Twittersphere as a whole, as a comparison with the ‘Total’ row of Table 1 shows.

The join dates for each account are also useful for retracing the dynamics of the adoption of Twitter as a social media platform in Norway. For each of the clusters we have identified, Table 1 also lists the median join date of accounts in the cluster – in other words, the date by which exactly half of the accounts that we found in each cluster at our time of data gathering in early 2016 had joined Twitter. The most significant outlier on this metric is the ‘Netizens’ cluster, representing software engineers, technologists, social media experts, geeks, nerds, and other users with a particularly strong technological affinity: perhaps unsurprisingly, Twitter recruited its earliest adopters from this demographic, and at least half of the ‘Netizens’ cluster was established as early as mid-2010 (and thus well before Twitter introduced a Norwegian-language user interface).

Table 1. Key activity metrics for the 30 largest clusters

Cluster Name	Accounts in group	Protected	Verified	Median join date	Total tweets	Tweets per account	Average tweets per day	Median tweets per day
Blogs & Humour	131,311	4,855	3	23.07.2012	100,736,580	767.16	0.6840	0.0258
Oslo/National News & Culture	118,483	980	5	05.08.2013	7,721,647	65.17	0.0966	0.0007
Sports: Football	55,107	1,047	61	01.05.2012	32,649,497	592.47	0.4078	0.0146
Music & Radio	38,094	1,106	72	07.12.2011	17,978,740	471.96	0.2811	0.0118
Online Entertainment	36,969	618	6	12.06.2013	15,610,716	422.27	0.5028	0.0199
Comedy	34,845	1,297	2	17.01.2012	12,155,711	348.85	0.2022	0.0071
Netizens	34,033	1,652	10	29.05.2010	17,344,820	509.65	0.3689	0.0129
Others	22,126	910	6	26.08.2012	8,240,985	372.46	0.3333	0.0056
Travel & Leisure	21,715	881	12	28.04.2012	13,018,976	599.54	0.4179	0.0051
Development & Environment Policy	19,256	752	18	23.08.2012	7,017,327	364.42	0.2932	0.0100
Troms/Nordland News	18,285	691	1	16.05.2012	4,246,332	232.23	0.1527	0.0031
Western Country News	17,478	630		08.04.2012	3,664,045	209.64	0.1487	0.0036
Business & Innovation	17,299	616	4	07.07.2012	3,535,449	204.37	0.1665	0.0036
Political Debate	16,763	450	6	27.01.2012	9,751,181	581.71	0.3220	0.0167
Cultural Debate	16,711	614	5	30.01.2012	6,656,536	398.33	0.2273	0.0103
Science, Innovation & Education	14,309	610		11.04.2012	5,140,675	359.26	0.2200	0.0089
Health Policy	13,742	507	5	20.09.2012	3,780,448	275.10	0.1906	0.0052
Music & Entertainment	11,242	178		06.10.2012	927,497	82.50	0.0667	0.0009
Religion & Politics	10,076	397	1	29.02.2012	2,840,294	281.89	0.1665	0.0068
Social Elites & Satire	9,854	277	2	05.06.2012	1,017,314	103.24	0.0841	0.0007
Business & Media	9,240	757	1	23.04.2012	27,420,052	2,967.54	1.9127	0.0890
Trøndelag News	7,271	468		17.04.2012	3,454,163	475.06	0.3217	0.0190
Food & Drink	6,957	259		14.05.2012	1,993,346	286.52	0.2221	0.0048
Sunnmøre News & Culture	6,567	316		30.03.2012	2,787,846	424.52	0.2706	0.0158

Table 1. Cont.

Cluster Name	Accounts in group	Protected	Verified	Median join date	Total tweets	Tweets per account	Average tweets per day	Median tweets per day
Photography	6,055	343	1	01.03.2012	2,075,153	342.72	0.2139	0.0099
Irony and Curiosa	5,278	194		04.04.2012	1,579,698	299.30	0.2087	0.0086
Hordaland/Bergen News	4,493	165		05.06.2012	1,097,836	244.34	0.1451	0.0035
Østfold News	3,550	185		24.03.2012	1,078,426	303.78	0.1702	0.0029
Oslo Region	3,165	267		17.06.2012	3,058,934	966.49	1.0260	0.0314
Gambling	2,976	116		20.08.2012	670,343	225.25	0.1566	0.0042
Sexuality	2,826	132		19.11.2012	1,448,780	512.66	0.6536	0.0066
Total	716,076	22,270	221	08.07.2012	320,699,347	447.86	0.3544	0.0066

Comments: List ordered by number of accounts in each cluster. 'Others' is a merged category of remaining, small clusters.

A second, and somewhat less obvious early cluster in the Norwegian Twittersphere is ‘Music & Radio’; here, an interest by Norwegian users in following the Twitter activities of international artists as well as the need for Norwegian bands to use Twitter to market themselves to an international audience might be assumed as key drivers of this comparatively early adoption of the platform, and at least half of the cluster was already present in December 2011 when Twitter launched its Norwegian interface.

For the vast majority of clusters in the Norwegian Twittersphere, on the other hand, median join dates range through 2012, and largely reflect the considerable increase in new account sign-ups that Figure 1 has also shown; a substantial number of the accounts in these clusters were created only after Twitter became available in the Norwegian language. Somewhat surprisingly, ‘Oslo/National News & Culture’ (alongside the roughly contemporary ‘Online Entertainment’) is the most recent of these clusters to be fully established: it took until August 2013 for the first half of its members to join Twitter. This suggests a different dynamic of Twitter adoption, and may serve to explain at least in part the significantly diverging activity metrics we have already identified for this cluster: many of its members joined only well after Twitter’s breakthrough as a major social media platform in Norway, entering into an already much more settled Twittersphere in which it might have seemed more acceptable simply to follow other accounts and consume their information rather than to contribute actively by tweeting themselves.

Cluster interconnectivity

In a final step in our analysis, we examine the interconnections between the individual clusters we have identified here. This draws especially on the E-I Index, proposed by Krackhardt and Stern (1988), as a measure of the interconnectivity between clusters. First, for each of the clusters we have identified, we are able to count the number of follower/followee connections that members of the cluster direct to other members of the same cluster; we designate these as internal links. Similarly, we are able to count the number of connections directed by members of a cluster to accounts *outside* of that cluster; those are external links. We note here that all such connections are directional: the links we count here mean that account *a* follows account *b*, and by extension that an account in cluster *A* follows an account in cluster *B*, but this does not imply that *b* necessarily also returns the attention and follows *a*. Aggregated for each cluster, this exercise thus results in a total count of internal as well as external links from accounts in the cluster; these counts are listed in Table 2.

Second, the E-I Index converts these raw counts into a normalised index on a scale from -1 to +1, using the formula

$$E-I\ Index = \frac{\# External\ Links - \# Internal\ Links}{\# External\ Links + \# Internal\ Links}$$

Using this calculation, a cluster whose members exclusively follow other members of the same cluster (that is, whose number of external links is zero) would receive an E-I Index of -1; a cluster with an equal number of external and internal links would receive an E-I Index of 0; and a cluster whose members do not follow each other, and instead only connect to accounts outside of the cluster, would receive +1. Since the very concept of network clusters presupposes a community of nodes that are more densely connected

to each other than to the rest of the network, clusters with an E-I Index of +1 are unlikely to occur; nonetheless, the relative differences in E-I Index values between clusters offer a useful measure of how inward- or outward-looking the members of each cluster are.

The analysis of such interconnectivity between clusters is important especially in the context of present debates about ‘echo chambers’ (Sunstein 2009) and ‘filter bubbles’ (Pariser 2012). Unfortunately, the ‘echo chamber’ and ‘filter bubble’ concepts remain poorly defined in current literature: while Sunstein’s books *Republic.com 2.0* (2009) and *#Republic* (2017) have done much to popularise the term, for instance, they do not provide a clear framing of what exactly an ‘echo chamber’ is, or what criteria should be used to detect it. In their work on echo chambers amongst Democrat- and Republican-leaning Twitter users in the United States, Colleoni et al. (2014: 318) suggest that ‘the mechanism through which this fragmentation of political discourse operates is homophily, defined as the tendency of similar individuals to form ties with each other’, and we follow this perspective by seeing echo chambers as related predominantly to the follower/followee connections made by Twitter users; these are most crucial in determining the potential reach of any actual tweets posted or retweeted by these users (in keeping with the metaphor, they circumscribe the distance that the echo can travel without additional amplification).

By contrast, although – in the absence of data on Norwegian accounts’ day-to-day tweeting activities and content – we do not address this concept in detail here, we would then define filter bubbles as related to what those accounts tweet (and in particular, what existing information they share and retweet), as ‘filtering’ implies a conscious case-by-case decision about which incoming information is passed on and which is not; a filter bubble would then emerge when a group of Twitter users (within the same cluster or across multiple clusters) all make the same decisions about what information they pass on to their followers, and what information they ignore. The two concepts are clearly interrelated, of course: what information is available for filtering to a given user depends in part (but not fully) on whom they follow in the network, and how far the information that passes through a user’s filter travels is determined in turn in large part by the structure of their network that follows them. However, in spite of a frequent conflation of the two terms in scholarly and popular literature, we suggest that a useful distinction can be made here: extending the underlying metaphor, the filtering provides (or re-broadcasts) the sound source for the echo, and the echo chamber determines how far that sound may travel.

Proponents of these concepts suggest that the structure and affordances of social media platforms encourage homophilous networking patterns and synchronised filtering practices, and therefore mean that users will increasingly encounter only that subset of all available views and information that is especially likely to agree with their own pre-existing interests and preferences – endlessly repeating the same topics, ideas, beliefs, values, and ideologies. At face value, the existence of the thematic network clusters that we have shown for the Norwegian Twittersphere might be seen as supporting this thesis: users who follow accounts from only one of these clusters will see in their Twitter feed only the tweets that circulate within that cluster (but we note that this already assumes that such users do not also use Twitter’s search functions, or follow the hashtags and trending topics that the Twitter platform may also recommend to them). However, if members of the clusters we have identified also follow a substantial number of accounts

from outside, the ‘echo chamber’ thesis is undermined: such members would be exposed to a wider range of views and topics than might be available in the cluster itself. Further, if they act on that information by retweeting it to their own followers, this would also challenge the ‘filter bubble’ concept as we have defined it above.

We therefore operationalise the E-I Index as a useful measure of the extent to which such ‘echo chambers’ do exist in observable reality. Table 2 presents both the raw counts of external and internal follower/followee links originating from accounts in each cluster, as well as the E-I Index calculated from these counts using the formula above; as with all other findings presented in this article, these figures cover only links within the Norwegian Twittersphere, and do not include any further follower/followee relationships to non-Norwegian Twitter accounts. It is immediately obvious that, with one exception, the E-I Indices for all clusters are considerably above zero; this shows that, even in spite of the predominant thematic interests expressed in their cluster membership, most users also follow a diverse range of other Norwegian Twitter accounts. The major exception from this pattern is the large ‘Blogs & Humour’ cluster, whose E-I Index remains close to zero: even this indicates only an even balance between inward and outward focus in following patterns, however, and cannot be used as evidence that this cluster would exhibit pronounced ‘echo chamber’ characteristics.

The surprisingly high E-I Index values observed for most of these clusters, however, could be seen as challenging the understanding of these communities of accounts as clusters in the first place: if these groups of accounts are largely following others not in their group, it is reasonable to ask whether they can still be meaningfully understood as belonging to the same group. To address this question, finally, Table 2 also provides an indication of the percentage of accounts in each cluster that are being followed by other members of the same cluster. These percentages indicate that, in all clusters, a substantial subset of all cluster members is followed by their peers in the same community; this is true especially for many of the smaller clusters. For instance, 91 per cent of the 6,055 members of the ‘Photography’ cluster are being followed by other photography enthusiasts, and it is this dense and mutual interconnection amongst members that designates this community as a cluster that can be detected by the Louvain algorithm we have used in our analysis, even if members of this community *also* follow over 25,000 other Norwegian accounts outside of their own cluster. While these external links are diverse and directed at other clusters across the entire Norwegian Twittersphere, the internal links are considerably more focussed and help the community stand out as a distinct cluster.

Overall, then, the picture that emerges from this analysis is of a truly national Norwegian Twittersphere, without exclusive barriers around specific clusters, in which important information has a reasonable chance to be seen well across the network. This is notably different from past observations of highly polarised network structures, for instance in social media discussions of U.S. politics (e.g. Conover et al. 2011). In addition to the fact that many such studies are built on datasets based on self-selecting hashtag practices, rather than the comprehensive follower/followee network data that we are working with here, this may also be explained by the comparatively much smaller size of the Norwegian Twittersphere, which is likely to work against such polarising tendencies: the fewer communicative partners are available for connection, the less potential there is for cliques to form within the overall network. (In spite of its internal connectivity, the Norwegian Twittersphere as such could still form a kind of national echo chamber if it

Table 2. Key interconnectivity metrics for the 30 largest clusters

Cluster Name	External connections (within Norway)	Internal connections	E-I index	Accounts receiving in-cluster edges (%)
Blogs & Humour	105,580	108,279	-0.01	84
Oslo/National News & Culture	137,996	38,945	0.56	34
Sports: Football	148,587	42,748	0.55	77
Music & Radio	131,715	28,498	0.64	74
Online Entertainment	56,123	23,501	0.41	63
Comedy	103,527	24,287	0.62	65
Netizens	128,251	30,409	0.62	87
Others	68,682	14,714	0.65	67
Travel & Leisure	91,932	16,338	0.70	76
Development & Environment Policy	107,269	16,478	0.73	90
Troms/Nordland News	67,720	15,938	0.62	86
Western Country News	60,573	14,160	0.62	82
Business & Innovation	93,165	13,412	0.75	77
Political Debate	117,442	13,749	0.79	81
Cultural Debate	93,687	14,193	0.74	85
Science, Innovation & Education	93,286	12,394	0.77	88
Health Policy	82,485	11,490	0.76	87
Music & Entertainment	14,267	3,985	0.56	37
Religion & Politics	59,466	8,831	0.74	84
Social Elites & Satire	28,997	3,086	0.81	31
Business & Media	137,708	8,397	0.89	88
Trøndelag News	22,950	6,380	0.56	86
Food & Drink	53,193	5,400	0.82	79
Sunnmøre News & Culture	27,212	5,583	0.66	88
Photography	25,621	5,217	0.66	91
Irony and Curiosa	14,470	4,584	0.52	84
Hordaland/Bergen News	33,295	3,294	0.82	72
Østfold News	30,769	3,076	0.82	86
Oslo Region	7,067	2,549	0.47	81
Gambling	22,391	2,463	0.80	84
Sexuality	15,278	1,765	0.79	61

Comments: List ordered by number of accounts in each cluster (see Table 1).

is poorly connected with the global Twitter network, perhaps; as some two-thirds of the total number of 98.7 million accounts followed by the Norwegian Twitter population are non-Norwegian accounts, however, this does not seem to be the case.)

We note again in this context also that our analysis of the connectivity patterns within the Norwegian Twittersphere extends well beyond the focus on *political* echo chambers that is common to many other studies. We take instead a broader approach to the ‘echo chamber’ concept, looking for comparatively isolated communities on any topic – and finding few that match this description. As we use the term ‘echo chamber’ in the present discussion, then, we use it to refer to communities of shared topical interests, rather than necessarily of shared sociocultural beliefs or political ideologies. At the same time, we

do note that those clusters which address overly political interests are also internally well-connected, with more than 80 per cent of accounts in any of these political or policy clusters being followed by other cluster members: we therefore do not find any strong evidence of internal polarisation *within* these clusters either. (Such observations could be extended through further analysis of the specific network structures within each of these clusters, which the available space for the present article does not permit us to present here.)

Such percentage metrics, then, also offer an insight into the internal network structure of each cluster. High connectivity percentages within a cluster indicate a comparatively flat internal network structure, where most members have a reasonable chance of being followed by one of their peers; lower percentages point more strongly to a hub-and-spoke structure in which a smaller group of leading accounts have attracted a disproportionate share of their peers' attention. This is the case especially for 'Music & Entertainment', 'Oslo/National News & Culture' and 'Social Elites & Satire' (the latter of which contains, for example, both official and parody accounts of the Norwegian royal family), where in each case only about one third of cluster members are being followed by their peers. This more centralised structure is not surprising given the themes of these clusters: in each of them, the dominance of a small number of nationally prominent celebrity, institutional and news accounts reflects established patterns of public attention in the Norwegian public sphere well beyond Twitter.

Conclusion and outlook

This article has presented the first ever comprehensive analysis of follower/followee network structures in the Norwegian Twittersphere. We have identified a wide range of network clusters that largely represent thematic interests (rather than geographic or demographic distinctions), and outlined the divergent dynamics that have led to the establishment of these clusters and continue to drive their further development. In particular, we have shown how a single key alteration to the Twitter platform itself – the introduction of a Norwegian-language user interface in December 2011 – triggered a substantial influx of new Norwegian accounts to the platform, considerably altering the structure of the network and affecting the balance of personal and professional, private and public interests represented in the Norwegian Twittersphere.

Our overview of the activity metrics for these diverse network clusters has pointed to further significant distinctions: while some clusters in the network are highly active and support the networking and communication between most of their members, others are comparatively quieter and focus attention on a smaller number of highly active, leading accounts, most likely operated by key individuals and institutions. This points to a number of different approaches to using Twitter amongst its Norwegian userbase: while some are using the platform to actively post their own updates and to engage with others, others are there predominantly to follow key information sources and to 'listen' to their news (Crawford 2009).

Finally, at least for a broader perspective that focusses on the intersections between overall thematic interests rather than specific sociocultural beliefs or political ideologies, we have also shown that even in spite of current concerns about the supposedly deleterious effects of social media on public debate – expressed in discussions about

‘echo chambers’ and ‘filter bubbles’, and moral panics about ‘fake news’ – there is very little evidence for the existence of such ‘echo chambers’ in the observable structure of follower/followee connections in the Norwegian Twittersphere. Rather, although we can readily detect densely connected regions of the overall network and describe these as communities of accounts that address shared themes and topics of interest, Norwegian accounts generally also choose to follow a large number and diverse range of *other* participants in the national Twittersphere, avoiding the trap of settling into purely homophilous networks.

We note here, however, that our study is based solely on an analysis of follower/followee connections, and does not take into account the information – the tweets – actually transferred through such connections on a day-to-day basis. We have therefore focussed only on the concept of ‘echo chambers’, which we see as more closely related to such underlying network structures. By contrast, the presence of ‘filter bubbles’ – a concept that we understand as speaking more directly to what users choose to *share* rather than whom they choose to *follow* – could be tested by observing (for each of the clusters we have identified, or for a selection of especially relevant clusters) the content of the tweets that accounts post on any given day. Such an observation could examine, for instance, whether different clusters draw their information from different outside sources (by analysing the URLs embedded in tweets) or to what extent cluster members are prepared to retweet content originating from outside of their own cluster. If a given cluster were to draw on a unique set of outside sources, or to retweet only its own members’ tweets, then such a cluster could be said to exhibit ‘filter bubble’ traits. The continuous tracking of the tweets posted by the 988,000 Norwegian Twitter accounts that we have identified here (or at least by a substantial subset thereof) that would be necessary for such an analysis is well beyond the scope of the present article, however, and must for now be left to further research.

The structural analysis of the Norwegian Twittersphere that we have presented here is valuable in its own right – it offers a comprehensive overview of the themes and interests represented in Norwegian Twitter, and may be usefully compared against further such studies that examine other national Twitterspheres, or even against the global Twitter network. It also constitutes a snapshot of the platform’s take-up in Norway at a particular point in time (in early 2016), and a future iteration of the same study should shed useful light on further changes in the Twitter userbase as the platform evolves and as it continues to compete with other social media services. It is just as important, however, to understand the data presented in this article as a crucial background benchmark for further research results: for future studies of Twitter activities in Norway, our observations make it possible not only to provide a general overview of tweeting dynamics (for instance, to indicate that a particular event or topic attracted a given number of unique participants who together generated a set number of tweets), but also to pinpoint which clusters in the network engaged especially strongly with these issues. A political controversy whose discussion does not spread beyond the nearly 17,000 members of the ‘Political Debate’ cluster, for example, might be important, but has failed to attract the interest of anyone but political insiders. By contrast, if that same controversy became the subject of discussion in thematically non-cognate clusters, then it should be understood as more significant: it has connected far more widely and engaged user communities not normally concerned with tracking political developments.

The lack of background data on the underlying structures of specific national Twitterspheres has traditionally made such assessments of the reach of specific issues difficult if not impossible. By comprehensively mapping the structure and historical dynamics of the Norwegian Twittersphere, the research presented here enables such deeper analysis for the first time.

Note

1. Norwegian has two written standards, Bokmål and Nynorsk.

References

- Bastian, Mathieu; Heymann, Sebastien & Jacomy, Mathieu (2009). Gephi: An Open Source Software for Exploring and Manipulating Networks, pp. 361–362 in *Proceedings of the Third International ICWSM Conference*. San José: AAAI Press. Available at <<http://www.aaai.org/ocs/index.php/ICWSM/09/paper/viewFile/154/1009/>> [Accessed 12 June 2017].
- Blondel, Vincent D.; Guillaume, Jean-Loup; Lambiotte, Renaud & Lefebvre, Etienne (2008). Fast Unfolding of Communities in Large Networks. *Journal of Statistical Mechanics: Theory and Experiment*, 10: P10008. doi:10.1088/1742-5468/2008/10/P10008.
- Bruns, Axel & Burgess, Jean (2015). Twitter Hashtags from Ad Hoc to Calculated Publics, pp. 13–28 in Rambukkana, Nathan (ed.) *Hashtag Publics: The Power and Politics of Discursive Networks*. New York: Peter Lang.
- Bruns, Axel; Burgess, Jean & Highfield, Tim (2014). A “Big Data” Approach to Mapping the Australian Twittersphere, pp. 113–29 in Arthur, Paul Longley & Bode, Katherine (eds.) *Advancing Digital Humanities: Research, Methods, Theories*. Houndmills, UK: Palgrave Macmillan.
- Bruns, Axel; Moon, Brenda; Münch, Felix & Sadkowsky, Troy (2017). The Australian Twitter-sphere in 2016: Mapping the Follower/Followee Network. *Social Media + Society*, 3(4): 1–15. doi:10.1177/2056305117748162.
- Burgess, Jean & Bruns, Axel (2015). Easy Data, Hard Data: The Politics and Pragmatics of Twitter Research after the Computational Turn, pp. 93–111 in Langlois, Ganaele; Redden, Joanna & Elmer, Greg (eds.) *Compromised Data: From Social Media to Big Data*. New York: Bloomsbury Academic.
- Burgess, Jean & Crawford, Kate (2011). Acute Events in Social Media. Paper presented at the *Association of Internet Researchers* conference, 10 October 2011, Seattle. Available at <<https://www.slideshare.net/jeangenie/acute-events-in-social-media>> [Accessed 12 June 2017].
- Colleoni, Eleanor; Rozza, Alessandro & Arvidsson, Adam (2014). Echo Chamber or Public Sphere? Predicting Political Orientation and Measuring Political Homophily in Twitter Using Big Data. *Journal of Communication*, 64(2): 317–332. doi:10.1111/jcom.12084.
- Conover, Michael D.; Ratkiewicz, Jacob; Francisco, Matthew; Goncalves, Bruno; Menczer, Filippo & Flammini, Alessandro (2011). Political Polarization on Twitter. Paper presented at the *Fifth International AAAI Conference on Weblogs and Social Media*. Available at <<http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2847>> [Accessed 12 June 2017].
- Crawford, Kate (2009). Following You: Disciplines of Listening in Social Media. *Continuum*, 23(4): 525–35. doi:10.1080/10304310903003270.
- Hermida, Alfred (2014). Twitter as an Ambient News Network, pp. 359–72 in Weller, Katrin; Bruns, Axel; Burgess, Jean; Mahrt, Merja & Puschmann, Cornelius (eds.) *Twitter and Society*. New York: Peter Lang.
- Highfield, Tim (2013). Following the Yellow Jersey: Tweeting the Tour de France, pp. 249–261 in Weller, Katrin; Bruns, Axel; Burgess, Jean; Mahrt, Merja & Puschmann, Cornelius (eds.) *Twitter and Society*. New York: Peter Lang.
- Highfield, Tim (2016). *Social Media and Everyday Politics*. Malden, Mass.: Polity.

Acknowledgements

This research was supported by the Research Council of Norway through the FRISAM project *The Impact of Social Media on Agenda-Setting in Election Campaigns*, and by the Australian Research Council through the Future Fellowship project *Understanding Intermedia Information Flows in the Australian Online Public Sphere* and the LIEF project *TrISMA: Tracking Infrastructure for Social Media Analysis*.

- Jacomy, Mathieu; Venturini, Tommaso; Heymann, Sebastien & Bastian, Mathieu (2014). ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLOS ONE*, 9(6): e98679. doi:10.1371/journal.pone.0098679.
- Kaufmann, Mareile (2015). Resilience 2.0: Social Media Use and (Self-) Care during the 2011 Norway Attacks. *Media, Culture & Society*, 37(7): 972–987.
- Krackhardt, David & Stern, Robert N. (1988). Informal Networks and Organizational Crises: An Experimental Simulation. *Social Psychology Quarterly*, 51(2): 123–40. doi:10.2307/2786835.
- Larsson, Anders Olof & Moe, Hallvard (2013). Twitter in Politics and Elections: Insights from Scandinavia, pp. 319–330 in Weller, Katrin; Bruns, Axel; Burgess, Jean; Mahrt, Merja & Puschmann, Cornelius (eds.), *Twitter and Society*. New York: Peter Lang.
- Meraz, Sharon & Papacharissi, Zizi (2013). Networked Gatekeeping and Networked Framing on #Egypt. *The International Journal of Press/Politics*, 18(2): 138–66. doi:10.1177/1940161212474472.
- Papacharissi, Zizi (2014). *Affective Publics: Sentiment, Technology, and Politics*. Oxford University Press.
- Papacharissi, Zizi & de Fatima Oliveira, Maria (2012). Affective News and Networked Publics: The Rhythms of News Storytelling on #Egypt. *Journal of Communication*, 62: 266–82. doi:10.1111/j.1460-2466.2012.01630.x.
- Pariser, Eli (2012). *The Filter Bubble: What the Internet Is Hiding from You*. London: Penguin.
- Perng, Sung-Yueh; Buscher, Monika; Wood, Lisa; Halvorsrud, Ragnhild; Stiso, Michael; Ramirez, Leonardo & Al-Akkad, Amro (2013). Peripheral Response: Microblogging during the 22/7/2011 Norway Attacks. *International Journal of Information Systems for Crisis Response and Management*, 5(1): 41–57.
- Rambukkana, Nathan (ed.) (2015). *Hashtag Publics: The Power and Politics of Discursive Networks*. New York: Peter Lang.
- Skogerbø, Eli & Moe, Hallvard (2015). Twitter på tvers – koblinger mellom journalister og politikere [Across Twitter – Connections between Journalists and Politicians]. *Norsk Medietidsskrift*, 22(3): 1–21.
- Sunstein, Cass R. (2009). *Republic.com 2.0*. Princeton, NJ: Princeton University Press.
- Sunstein, Cass R. (2017). *#Republic: Divided Democracy in the Age of Social Media*. Princeton, NJ: Princeton University Press.
- Syvertsen, Trine; Enli, Gunn; Mjøs, Ole J. & Moe, Hallvard (2014). *The Media Welfare State: Nordic Media in the Digital Era*. Ann Arbor: University of Michigan Press.
- @twitter (2011). Twitter is now available in four new languages: Danish, Finnish, Norwegian and Polish. [Twitter] 21 December 2011. Available at <<https://twitter.com/twitter/status/149299492917747712>> [Accessed 12 June 2017].
- Vis, Farida (2013). Twitter as a Reporting Tool for Breaking News: Journalists Tweeting the 2011 UK Riots. *Digital Journalism*, 1(1): 27–47. doi:10.1080/21670811.2012.741316.