

# Identifying the Events That Connect Social Media Users: Charting Follower Accession on Twitter

**Axel Bruns**

**Darryl Woodford**

Queensland University of Technology

## Contributor biographies

Dr Axel Bruns is an Associate Professor in the Creative Industries Faculty at Queensland University of Technology in Brisbane, Australia, and a Chief Investigator in the ARC Centre of Excellence for Creative Industries and Innovation (<http://cci.edu.au/>). He is the author of *Blogs, Wikipedia, Second Life and Beyond: From Production to Producership* (2008) and *Gatewatching: Collaborative Online News Production* (2005), and a co-editor of *Twitter and Society* (2013), *A Companion to New Media Dynamics* (2012) and *Uses of Blogs* (2006). Bruns is an expert on the impact of user-led content creation, or producership, and his current work focusses on the study of user participation in social media spaces such as *Twitter*, especially in the context of acute events. His research blog is at <http://snurb.info/>, and he tweets at @snurb\_dot\_info. See <http://mappingonlinepublics.net/> for more details on his current social media research.

Darryl Woodford is a Research Fellow in the ARC Centre of Excellence for Creative Industries & Innovation, based at Queensland University of Technology. Prior to this, he completed a BEng at University of Essex in the United Kingdom in Computer Games & Internet Technology, and completed an MSc at IT University of Copenhagen, with a thesis focusing on the experience of players in games through a study of avatars and agency. He has taught into 'New Media' courses at QUT, as well as game studies courses at IT University of Copenhagen. Recent publications include work comparing regulatory structures of the gaming and gambling industries, and methodologies of Virtual World Research, while his current research concentrates on social media strategy and analytics, online video and the global games industry.

## Relevant Disciplines

Media and communications

## **Academic levels**

Advanced Undergraduate, Postgraduate

## **Methods Used**

Twitter API, statistical analysis

## **Keywords**

social media, Twitter, big data, social network analysis, public communication, follower accession

## **Abstract**

Twitter is the focus of much research attention, both in traditional academic circles and in commercial market and media research, as analytics give increasing insight into the performance of the platform in areas as diverse as political communication, crisis management, television audiencing and other industries. While methods for tracking Twitter keywords and hashtags have developed apace and are well documented, the make-up of the Twitter user base and its evolution over time has been less understood to date.

Recent research efforts have taken advantage of functionality provided by Twitter's Application Programming Interface to develop methodologies to extract information that allows us to understand the growth of Twitter, its geographic spread, and the processes by which particular Twitter users have attracted followers. From politicians to sporting teams, and from *YouTube* personalities to reality television stars, this technique enables us to gain an understanding of what prompts users to follow others on Twitter.

This article outlines how we came upon this approach, describes the method we adopted to produce accession graphs, and discusses their use in Twitter research. It also addresses the wider ethical implications of social network analytics, particularly in the context of a detailed study of the Twitter user base.

### **Learning Outcomes**

By the end of the case you should:

- Understand how Twitter data can be used to conduct different types of research.
- Understand the methodological challenges of analysing Twitter data, particularly relating to large samples and API restrictions.
- Be aware of the importance of understanding users' follower networks when discussing the role they play on a particular social network.
- Understand the evolution of research methodologies and the transition between quantitative and qualitative approaches.

### **Introduction**

Research into the uses of social media platforms has grown considerably in breadth and sophistication over the past few years, along with the uses of social media in various public communication contexts themselves. Twitter has proven a particularly productive space for researchers because the platform has traditionally provided a comparatively open and powerful Application Programming Interface (API) which offers structured and standardised access to a variety of data on real-time communication processes taking place on the

platform, as well as on background details such as users' profiles and follower/followee connections. With the growing commercialisation of Twitter data, such access has become increasingly restricted, as Jean Burgess and Axel Bruns have highlighted, but the data available to researchers and developers through the free Twitter API remain useful for a wide range of research applications.

### **Evolution of Research Approaches**

Over the past 5 years, a significant number of international researchers have developed methodologies to investigate specific communicative events on Twitter, building usually on tweet datasets which are defined by common keywords or hashtags. This research provided an opportunity to examine how Twitter users discuss specific issues and events, from political scandals through natural disasters to major television experiences. In the Further Readings below, there are a number of references for this type of 'traditional' Twitter analysis. The piece by Axel Bruns and Jean Burgess, in particular, highlights a range of methodological approaches for such research which vary from data-driven approaches such as the one we outline here to content and sentiment analysis of tweets (which might be used, for example, to measure the audience response to a television show or the dissemination of emergency information following a natural disaster). By deploying a range of standardised methods and tools for the study of such datasets, comparative analysis – comparing for example the magnitude of such communicative phenomena or the breadth of Twitter user participation in them – also became possible.

However, what was largely missing from these approaches and analyses was a consideration of the way in which Twitter users were connected (what we may term 'networked interaction') on the platform. While in theory any tweet posted from a publicly visible account on Twitter could be observed by any other Twitter user who searched for one

of the hashtags or keywords occurring in the tweet, or by visitors to the sender's public profile on the Twitter Website, in practice most tweets are primarily visible to the followers of the originating account. Thus, Twitter datasets which include all the tweets containing a specific keyword or hashtag will combine tweets which enjoyed a potentially very large audience (because they were sent from an account with a substantial number of followers) alongside tweets which were almost certainly seen only by a small group of users (because they were sent from accounts with very few followers). Until now, little has been done to separate these types of tweets, that is, to understand the viewing audience of tweets as we might do with other forms of media (e.g. by considering newspaper circulation figures or television ratings). In order to develop a more comprehensive account of the uses of Twitter in public communication, it is therefore important to incorporate aspects of network analysis into the research process.

Additionally, like the real-time communication between Twitter users itself, the shape of the Twitter network (that is, which users follow other users and vice versa) can also be highly changeable. The sudden notoriety of specific Twitter users, as a result of their being involved in significant public events or through being retweeted or @mentioned by already prominent Twitter users, can lead to a sudden and rapid influx of new followers. For example, anyone @mentioned or retweeted by @BarackObama is likely to see their content exposed to a much larger audience than it had been previously, and will almost certainly gain a substantial number of new followers. Similarly, major controversies and other events, inside or outside of Twitter, may also lead to a substantial loss of followers over a short time period. Indeed, an observation of these changes in the follower base of any one specific Twitter account can point to the key online or offline events that the user was involved in at the time.

However, the datasets necessary to examine these questions are extremely large. The Twitter user base is estimated to be over 500 million accounts, spread over several billion

user IDs. Even on a smaller scale, it is not uncommon for an individual popular Twitter user to have more than one million users following them. Until recently, obtaining these data from the Twitter API was extremely time-consuming; however, improved and optimised data collection techniques have enabled us to gather more comprehensive datasets and thus to answer a number of the questions which had been beyond the scope of our research so far. Here, we provide more detail on our approach to investigating changes to the follower base of individual Twitter accounts, while also considering the potential research applications of the larger datasets on follower networks amongst the overall Twitter user base which we are now able to gather, and the ethical and technical obstacles which need to be addressed in conducting such research

## **Research Method**

As we began to collect data to conduct a large-scale study on the Twitter user base, we reviewed the literature that discussed how other researchers had been considering such data until now. During this literature review, we came across the approaches developed by a group led by Brendan Meeder at Microsoft Research, as well as the work of Tony Hirst at Open University. Both of these approaches had, seemingly independently, described the research opportunities emerging from the study of what Tony Hirst described as Twitter follower accession. Their research and our research initiatives utilise the Twitter API to retrieve a full list of all the followers for a given account. For the more technically minded, this is a call to the Twitter API's '*followers/ids*' function, which is documented on the Twitter Developers site: <https://dev.twitter.com/docs/api/1.1/get/followers/ids>.

Our approaches are based on the fact that, at least at present, the API returns this list of followers in reverse chronological order; those users who most recently *became* followers of

the target account are listed first. Further calls to the Twitter API (this time using *'users/show'*: <https://dev.twitter.com/docs/api/1.1/get/users/show>) enable us to identify for each of these follower accounts exactly when each account was created – information which becomes crucially important for our further research approach. For those without access to large-scale data collection tools, Tony Hirst's approach uses the creation dates of a random sample of followers, rather than retrieving all followers' account creation dates as we do. This simplified approach should not significantly affect the quality of the analysis for target accounts with a substantial number of followers.

Using such comparatively simple information – the place of each user in the numbered list of followers for the target account, and the creation date of each follower account – we are then able to begin graphing the follower accession curve for a given account. To illustrate this, we used the case of then-Australian Prime Minister Kevin Rudd (@KRuddMP), whose follower base we examined in late June 2013, in the following discussion. Figure 1 (Caption: Preliminary follower accession curve for @KRuddMP, in late June 2013) shows a preliminary follower accession curve for Rudd's account, which was generated by plotting the account creation date for each follower (on the horizontal axis) against the position of the follower in Rudd's list of then around 1.2 million followers (on the vertical axis):

To understand the graph, take a point on the vertical axis – for example, around the 600K mark. If you then follow a horizontal line from that point across the graph, each of the blue points which are placed on that horizontal line represents one individual follower of @KRuddMP who began following that account at a time when Rudd had around 600,000 followers. Moving from left to right through the graph, we find that some users who began following Kevin Rudd at that time had already been active on Twitter for some years (their own accounts had been created in 2006 or 2007, as the horizontal axis at the bottom of the graph indicates), but a greater number of followers had joined Twitter only some time after

the start of 2009, like the vast majority of Twitter accounts existing today. Following the same horizontal line, eventually you reach a sharp edge on the right side of the graph: this represents the most recently created accounts which followed @KRuddMP around this time.

A clear leading edge is obvious for the entire graph in Figure. 1. This is due to the simple fact that Twitter users could only have become followers of the target account *after* creating their own accounts; no account following @KRuddMP before, say, December 2009 could itself have been created after December 2009. For Twitter users with a reasonably substantial number of followers, however, there will always be new followers who become followers very shortly after joining Twitter themselves – as a result, the bottom edge of a follower accession graph provides a good approximation of the overall follower growth curve for the target account.

Put simply, if the Twitter account of @KRuddMP's 600,000th follower was created on 1 December 2009, then it cannot have become a follower before that date; its earliest possible accession date to the follower base is 1 December 2009. Due to the fact that the follower list is chronological, this also means that followers number 600,001, 600,002, and so on must have become followers after 1 December, since they became followers only after 600,000 joined. By the same token, if the account of follower 600,500 was created on 10 Dec. 2009, then any accounts who followed @KRuddMP after it did must have become followers after that date. From this, we can begin to assign earliest possible accession dates to each of @KRuddMP's followers: 1 Dec. 2009 for accounts, 600,000-600,499; 10 Dec. 2009 for accounts 600,500 onwards, and so on. (In reality, these assignments are often considerably more fine-grained.) What results from this, then, is a fairly precise approximation of the follower growth curve for the target account: using the same example, we can assume that @KRuddMP passed 600,000 followers on or shortly after 1 December 2009, 600,500 followers ten days later, and so on.



One significant *caveat* applies to this approach, however: using the Twitter API, we can only retrieve information on the users who are *currently* following our target account – it does not (and for privacy reasons, should not) provide us with any information about users who did once, but now no longer follow the target account. It would be perfectly possible (if unlikely), for example, that by 2009, @KRuddMP had amassed more than a million followers, but that 400,000 of those followers have since disappeared again. If we see in Figure 1 that by 1 December 2009, @KRuddMP appears to have accumulated some 600,000 followers, then, what we are really seeing is that by 1 December 2009, some 600,000 of those accounts *who are still following Rudd today* had already begun to follow him.

This is an unavoidable limitation, and one which needs to be remembered in analysing Twitter follower accession data. It does not invalidate the analytical uses which our method enables; we are still able to examine, for example, what events in the career of a Twitter account have led to an especially strong influx of long-term followers (i.e., of followers who have stuck around). What we cannot examine through this method is the level of day-to-day fluctuation or churn in follower numbers, or any sudden mass unfollowing events (for example as a result of negative occurrences in the career of the target account). Such unfollowing events could be studied, by regularly (e.g. weekly or monthly) gathering the current follower lists for a given account, and cross-checking these lists to determine the number of users who have followed or unfollowed the account from one point to the next. Just as we argue that Twitter should not present this information for privacy reasons, however, as researchers we must weigh the merits of such methods against the ethical challenges of presenting details of users' past behaviours, particularly where information is revealed about users who are significantly less 'public' than the target accounts considered here.

## Analytical Opportunities

Collecting and plotting the data from the Twitter API is only the first step in drawing conclusions about the Twitter user base. Examining the plotted data in turn provides a number of subsequent research questions, both generic and specific to a particular dataset. In the first place, we will usually be interested in key spikes in the follower accession, and in investigating the likely causes of such events. At a theoretical level, there are a number of possible explanations, including genuine events in the career of the target account, but also an influx of 'spam' followers or other platform-dependent events. In the case of @KRuddMP, Figure 1 shows a rapid increase in the number of followers between late June 2009 and late January 2010 – a time during which more than 600,000 new accounts join as followers. This prolonged spike in the accession rate of new followers is especially unusual because it begins and ends quite abruptly; at 200-250 new followers per day, the accession rates before and after this period are comparable to each other, but considerably lower than the 2,000-6,000 new followers per day that Rudd received during the period itself.

Through discussions on our research Website at *Mapping Online Publics* (<http://mappingonlinepublics.net/>) and research on the history of the platform, we discovered that this curious rise in Rudd's follower accession rate could be explained by the fact that from late June 2009 onwards, Twitter changed its new user sign-up procedures: in particular, it introduced the first iteration of its Suggested User List, featuring a handful of important and notable Twitter users whom new users were all but forced to follow as part of the sign-on process. Rudd was one of the very few Australians on that (global) list, and obviously benefitted significantly. Subsequent changes to the Suggested User List made it easier for users to bypass the 'follow prominent users' stage of the sign-on process, leading to a

slowing of Rudd's accession rate later in 2009, and by early 2010 the new user sign-on process was remodelled considerably once again, returning Rudd's accession rate to a more 'normal' pace. As an aside, it is interesting that over 600,000 of the followers who joined him during that time did remain with him, even if their initial following was perhaps coerced to some extent by the Twitter sign-up process; although whether this represents genuine loyalty and interest or merely means that these users have not been provoked to unfollow remains unclear from our analysis to date.

In other cases, such sudden spikes in new followers may have more sinister reasons. During the 2013 Australian election campaign, for example, a sudden influx of more than 60,000 new followers was observed for then-Australian Opposition Leader Tony Abbott. As one can see in Figure 2 (Caption: Follower accession curve for @TonyAbbottMHR, 12 August 2013. Suspicious followers are shown in red.) below, this spike takes on a very different shape to that observed for Kevin Rudd. In this case, the vast majority of these new follower accounts had been created at almost exactly the same time during the past 60 days, and thus our subsequent analysis points much more strongly towards a conclusion that these new followers are 'fake accounts': accounts created with the sole purpose of boosting their target accounts' follower numbers. This does not mean that Abbott or his supporters orchestrated this influx, however. Given the negative publicity which resulted from this case, it is just as likely that one of the Opposition Leader's political enemies could have bought these fake followers for his account in order to embarrass him. (Subsequently, following mainstream media coverage, the fake accounts were quickly deleted by Twitter.)

Given that significant attention is paid to users' follower numbers as a means of assessing their standing and importance within the Twitter network, the approach we have outlined here allows us to begin to distinguish genuinely popular users from accounts whose follower numbers have been artificially boosted by 'fake' followers and gamed by 'guaranteed

followback' schemes. Such schemes will often appear as unusual patterns in the accession data, and their effects can usually be distinguished clearly from a more 'organic' growth in follower numbers. As a result, we are better able to estimate the realistic 'reach' of a particular account and its tweets; 'fake' followers should not be considered part of its audience, given that these followers are highly unlikely to be reading the tweets.

Beyond such individual analyses, however, a study of follower accession patterns also provides an opportunity to compare the follower growth of different accounts and to correlate these patterns with the events which affected them. Figure 3 (Caption: Follower accession curves (first 50,000 followers) for selected Australian political leaders) provides a comparison for key Australian political leaders between 2008 and mid-2013, for example, focussing on the first 50,000 followers gained by each account.

Through such comparison, a number of significant patterns are revealed, which may then become the starting point for an analysis across other groups of users. In Figure 3, it can be seen that politicians' accounts gain a substantial number of followers rapidly when they are elected to key positions (as with the election of Tony Abbott as Opposition Leader on 1 Dec. 2009, or the election of Julia Gillard as Prime Minister on 24 June 2010), but their opponents also gain followers during such times (this is the case for outgoing Opposition Leader Malcolm Turnbull when he is defeated by Abbott, and for Abbott himself, as well as Turnbull, when Julia Gillard becomes PM). It appears, in other words, as if major political events generally increase public attention for politics, and encourage Twitter users to follow their political leaders.

A number of other patterns are also visible: Julia Gillard, for example, who was Australian Prime Minister from 2010 to 2013, had had a Twitter account since 26 Oct. 2009, but her follower numbers only rose (and rapidly so) from the moment she became PM. This may

indicate that her account had until then been run under a different name (as a private rather than public communication tool), or had even been set to ‘protected’, requiring would-be followers to be approved by the account operator. Only once she had been installed as Prime Minister, it seems, did Julia Gillard’s account accept followers. Each of these observations provides the opportunity for further quantitative or qualitative analysis (for example, it might be interesting to consider mentions of politicians’ Twitter accounts in the mainstream press, or to look in further detail at Julia Gillard’s tweet history to ascertain how she – or her social media team – developed her account). However, these observations also serve as a starting hypotheses for investigations of the follower base dynamics of Twitter accounts in other fields, such as television or sport.

Figure 3 thus provides an indication of the potential benefits of a follower accession analysis using the methods outlined here; specific research questions will depend on the nature of the target account or accounts to be examined, of course, and we do not seek to provide an exhaustive list of such possibilities here. Political, but also media and celebrity accounts, might be examined and compared to ascertain their respective popularity on Twitter, for example, and to track the growth or decline of their popularity as expressed in follower numbers over time. An analysis of more ‘ordinary’ users’ accounts may be valuable where such users suddenly find themselves thrust into the spotlight, perhaps because they suddenly become popular or notorious, or as a result of their activities in relation to unforeseen events such as breaking news or natural disasters. Alternatively, the analysis of follower growth patterns for a random sample of ordinary accounts, however constructed, may also shed light on how users’ Twitter networks develop over time.

Additionally, while we have outlined in detail here the methodology and a sample case for follower accession, almost identical methods could also be used to determine *followee* accession patterns for any given account – that is, the rate at which a specific account has

followed other users. For a celebrity account, such an analysis could determine, for example, whether the account simply automatically follows back every new follower (in this case, follower and followee accession graphs should be almost identical in shape), or whether the account follows back only a select number of its followers (and if so, whether such follow-backs occur at a steady pace or in more or less regular bursts as new followers are reviewed and followed back). Specific uses for such followee accession analyses are perhaps somewhat less obvious than for follower accession analysis, but should also continue to be explored.

### **Limitations**

Follower accession analysis as we have outlined it here depends crucially on the fact that the follower lists returned by the Twitter API are currently provided in chronological order, enabling us to chart a follower growth curve by comparing a user's position in the follower list with their own account creation date. Twitter has hinted that the format of follower lists may be subject to change in the future, which would render our method obsolete. For now, however, no such changes have been made, and the method produces reliable results.

Our method also generates substantially more precise results for accounts with larger numbers of followers and a comparatively regular rate of follower growth. Where accounts have as yet failed to amass significant follower numbers, the number of available data points remains too low to result in a reliable accession curve, and analytical results should be treated with some caution. At the same time, however, for accounts with a very large number of followers the access restrictions which Twitter imposes on its API generate a different set of problems: while here, the available data can generate very fine-grained accession curves, the retrieval of those data can take considerable time since Twitter has limited the number of

times its API may be accessed in each 15-minute period; generally, therefore, follower accession data cannot be retrieved speedily, nor can such analyses be repeated frequently.

Finally, as also noted throughout this article, the method outlined here cannot provide any information about users who may at one point have followed the target account, but have since unfollowed it again. What our follower accession data (and the analyses built upon such data) represent are the patterns of follower accession for those users who remain followers *at the time of data collection*. Only comparative longitudinal studies would be able to detect acts of unfollowing, by highlighting which users from a previously collected list of followers for a given account are no longer present on a subsequent list.

## **Conclusion**

We hope that the particular study of Twitter follower data we have highlighted here has inspired you to consider how such an approach could be used to examine the networks of different types of Twitter users, but also to explore more broadly how new research approaches could be brought to bear specifically upon Twitter as a platform, and across a variety of other social networks. We also invite you to consider how the mixed-methods approach discussed here (in which an initial question is answered by quantitative data, which in turn provides a number of subsequent research questions and hypotheses for qualitative analysis) may benefit your own research projects.

Finally, we hope that the case we have outlined here highlights the vast quantity of information that is made available through social networks, and the implications of that information for scholarly research. As researchers, we are in a position to synthesise the data made available through platforms and to make our findings available in a more accessible

format both to other academics and to the general public. When working with datasets of this nature, where the data normally exist in a fairly abstract format, the ethical responsibilities of the researcher in making research findings available to a wider audience should not be understated, and we invite you to consider the implications analysing user activity in social media spaces as you conduct projects that draw on ‘big data’ from such sources.

### **Exercises and Discussion Questions**

1. We detail how our research methodology evolved towards a focus on accession charts from a more general study of Twitter. What other methodologies that you are aware of might be useful when analysing Twitter data?
2. Within your particular field of study, how might a mixed-methods approach benefit your research? Do you prefer a strict methodology or an evolutionary approach?
3. What ethical concerns do you think apply to utilising social media datasets? How might these apply to ‘big data’ more generally?
4. What limitations applied to our work on Twitter accession charts? How might these limitations, or others imposed by the platform, impact on your own work?
5. Other than the presence of ‘fake’ followers, what other information do you think can be drawn from the accession charts discussed here?
6. To what extent do you think (a) internal events, such as the organic growth of the platform, and (b) external events, such as news coverage, are responsible for the growth in a user’s follower base?



## Further Readings

Hirst, T. (2013). What Happened Then? Using Approximated Twitter Follower Accession to Identify Political Events. *OUseful.info*, 4 Mar. 2013. Retrieved from <http://blog.ouseful.info/2013/03/04/what-happened-then-using-approximated-twitter-follower-accession-to-identify-political-events/>

Hirst, T. (2013). Estimated Follower Accession Charts for Twitter. *OUseful.info*, 5 Apr. 2013. Retrieved from <http://blog.ouseful.info/2013/04/05/estimated-follower-accession-charts-for-twitter/>

Bruns, A. (2013). Introducing Twitter Follower Accession Graphs. *Mapping Online Publics*, 8 July 2013. Retrieved from <http://mappingonlinepublics.net/2013/07/08/introducing-twitter-follower-accession-graphs/>

Bruns, A. (2013). Follower Accession: How Australian Politicians Gained Their Twitter Followers. *Mapping Online Publics*, 8 July 2013. Retrieved from <http://mappingonlinepublics.net/2013/07/08/follower-accession-how-australian-politicians-gained-their-twitter-followers/>

Bruns, A. (2013). Factcheck: Did Tony Abbott Buy Twitter Followers? *Mapping Online Publics*, 12 Aug. 2013. Retrieved from <http://mappingonlinepublics.net/2013/08/12/did-tony-abbott-buy-twitter-followers/>

## References

- Bruns, A. (2013). Factcheck: Did Tony Abbott Buy Twitter Followers? *Mapping Online Publics*, 12 Aug. 2013. Retrieved from <http://mappingonlinepublics.net/2013/08/12/did-tony-abbott-buy-twitter-followers/>
- Bruns, A., & Burgess, J. (2012). Researching News Discussion on Twitter: New Methodologies. *Journalism Studies*, 13(5-6), 801–14.  
doi:10.1080/1461670X.2012.664428
- Bruns, A., Burgess, J., & Highfield, T. (2014). A ‘Big Data’ Approach to Mapping the Australian Twittersphere. In P. Arthur and K. Bode, eds., *Repurposing the Digital Humanities*. Basingstoke, Hampshire: Palgrave Macmillan.
- Bruns, A., & Stieglitz, S. (2012). Quantitative Approaches to Comparing Communication Patterns on Twitter. *Journal of Technology in Human Services*, 30(3-4), 160–185.  
doi:10.1080/15228835.2012.744249
- Bruns, A., & Stieglitz, S. (2013). Towards More Systematic Twitter Analysis: Metrics for Tweeting Activities. *International Journal of Social Research Methodology*, 16(2), 91–108. doi:10.1080/13645579.2012.756095
- Burgess, J., & Bruns, A. (2012). Twitter Archives and the Challenges of “Big Social Data” for Media and Communication Research. *M/C Journal*, 15(5). Retrieved from <http://journal.media-culture.org.au/index.php/mcjournal/article/viewArticle/561/0>
- Hirst, T. (2013). Estimated Follower Accession Charts for Twitter. *OUseful.info*, 5 Apr. 2013. Retrieved from <http://blog.ouseful.info/2013/04/05/estimated-follower-accession-charts-for-twitter/>

Meeder, B., Karrer, B., Sayedi, A., Ravi, R., Borgs, C., & Chayes, J. (2011). We Know Who You Followed Last Summer: Inferring Social Link Creation Times in Twitter. In *WWW '11: Proceedings of the 20th International Conference on the World Wide Web*. New York: ACM. 517-526. doi:10.1145/1963405.1963479.