



UvA-DARE (Digital Academic Repository)

Digitally networked grassroots

Social media and the development of the movement for black lives and immigrant rights movement in the United States

van Haperen, S.P.F.

Publication date

2019

Document Version

Other version

License

Other

[Link to publication](#)

Citation for published version (APA):

van Haperen, S. P. F. (2019). *Digitally networked grassroots: Social media and the development of the movement for black lives and immigrant rights movement in the United States*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 10

Dynamics of Digitally Networked Leadership in #blacklivesmatter¹⁸⁸

Sander van Haperen, Justus Uitermark

188 This chapter is under review for publication at the time of writing, August 26, 2019.

Abstract

Black Lives Matter activists have given new impetus to aspirations of ‘leaderful’ organizing. Disavowing the centralization of leadership and rethinking the history of the struggle for civil rights, many have sought to develop and revamp models of decentralized leadership. Digital networking features prominently in these models, but the literature suggests its role is ambiguous: while social media provide affordances for decentralized leadership, they may also induce the centralization of leadership. To examine such leadership dynamics, we adopt a computational approach. Drawing on a dataset of roughly 18 million tweets collected over a 30 month period, we examine two aspects of leadership dynamics: the concentration of leadership (i.e. inequality in terms of centrality within networks) and the consolidation of leadership (i.e. the degree to which central figures retain prominence over time). Engaging with the measurement of network inequality by differentiating temporal aspects of concentration and consolidation, we show that leadership dynamics fluctuate strongly over time. Aside from a select few, there is little consolidation of leadership: new voices can consistently rise to prominence. Contributing insight in today’s preeminent social movement, we show how computational analyses can be adopted in ways that do justice to the dynamic and diverse nature of movements.

Introduction

While progressive movements have long adopted participatory democracy for both strategic and ideological reasons (Polletta, 2002), new communication technologies have revamped and reinvigorated ideals of ‘ad hoc, leaderless, participatory, and horizontalist’ decision-making (Tufekci, 2014b, p. 13). Scholars have argued that the use of social media allows for digital networking among large and diverse groups of activists without suppressing specific identities and concerns (Bennett and Segerberg, 2012). It is within these decentralized networks that community-based leaders can rise to prominence (see also Castells, 2012; Juris, 2005b; Tufekci, 2017). In this context, Black Lives Matter activists have invoked the concept of a ‘leaderful movement’. Moving beyond the notion that leadership is inherently problematic, they argue that leadership should not be delegated to a single leader but distributed over a diversity of movement participants (cf. Raelin, 2003, 2005; Wood and Ladkin, 2008). In the words of the originators of the #blacklivesmatter hashtag, Opal Tometi, Alicia Garza, and Patrisse Cullors-Brignac:

There are important implications for the possibilities that this new layer of leadership can offer the movement as a whole. We create much more room for collaboration, for expansion, for building power when we nurture movements that are full of leaders, and allow for all of our identities to inform our work and how we organize. This then allows for leadership to emerge from our intersecting identities, rather than to be organized around one notion of Blackness. Because of this, we resist the urge to consolidate our power and efforts behind one charismatic leader. When we center the leadership of the many who exist at the margins, we learn new things about the ways in which state sanctioned violence impacts us all (Tometi et al., 2015)

Although activists warn against the centralization of power, they do not reject leadership as such. On the contrary, it is vitally important for historically marginalized communities to bring forth leaders who can inspire, organize, and coordinate. This is especially true for groups that were pushed to the background as charismatic leaders took central stage, including women, LGBTQI communities, and other marginalized groups (Davis, 1983). While leadership is vital to the articulation of demands and the coordination of protest, leaders should not claim to speak for the movement as a whole and should not block others from rising to

prominence and having their say. Ideally, then, leadership is shared, transitory, and distributed. However, there are reasons to believe that social media impede egalitarian participation and contribute to the centralization of power. As we elaborate below, researchers have stressed that the network logics of social media foster highly unequal relationships within social movements (Gerbaudo, 2012; Nunes, 2014; Swann and Husted, 2017; Tufekci, 2017).

A conundrum of online organizing is, in short, that the very social media affordances that ostensibly facilitate egalitarian participation also potentially induce inequality. In this light, we aim to examine dynamics of leadership empirically. How does online leadership develop in #blacklivesmatter? While qualitative studies have begun to examine how competing tendencies towards centralization and decentralization play out (Bennett and Segerberg, 2012; Gerbaudo, 2012; Swann and Husted, 2017; Tufekci, 2017; Uitermark, 2017), we advance computational analysis to assess this dynamic. We adopt a relational approach to analyze how communities form online and to examine dynamics of leadership within them.

In the following, we first expand the discussion on leadership dynamics and develop a relational conceptualization of leadership within digitally networked communities. We examine leadership by distinguishing analytically between concentration and consolidation. The methodological section then provides details on data collection, the identification of communities, and measures of both leadership concentration and leadership consolidation. After presenting our findings, we conclude with a discussion of implications for theory development and activist practice.

Digitally networked leadership

Our understanding of digitally networked leadership extends efforts in social movement research to examine contention relationally. One recurrent theme in the literature is the disconnect between theories that stress the relational nature of contention and methodologies that rely on categorical data (Diani and McAdam, 2003; White, 2008). Although scholars of social movements have suggested to focus on the structure of relations among players in contentious politics (Duyvendak and Jasper, 2015a), there is a ‘persistent tendency to treat social movements as aggregates of discrete elements (be they individual or organizational actors, or events), rather than systems of relations’ (Diani, 2013, p. 145). However, the

advent of social media and the increasing availability of network data have given new impetus to relational approaches to social movements. A blooming branch of research examines empirically how interactions among social media users generate topological structures (Barberá et al. 2015; Centola 2010; González-Bailón et al. 2013; González-Bailón and Wang 2015) and facilitate information diffusion (Givan et al., 2010; Goel et al., 2012; S. González-Bailón et al., 2013; Rane and Salem, 2012). We build on and contribute to this expanding body of work by furthering conceptualizations and methods to study digital leadership dynamics relationally.

Leadership: concentration and consolidation

Traditional views of leadership seek to explain prominence in terms of individual attributes. A range of personal properties and skills such as charisma or eloquence allow leaders to inspire action by articulating what a movement stands for (Gusfield, 1966; Weber, 1946). The rise of New Social Movements in the 1960s and 1970s, which innovated with decentralized organization, prompted stronger emphasis on the context of communities to which leaders relate. For instance, researchers stressed that leadership requires knowledge of ‘local idioms, values, and practices’ to unite diverse groups into a mass base (Barker et al., 2001; Veltmeyer and Petras, 2002). The advent of network analysis in social movement studies has further reinforced the focus on leadership as fundamentally relational (Diani and McAdam, 2003; Ganz, 2010). Understood relationally, leadership is not a personal attribute but a quality of a group whose members refer and defer to leaders. While the analysis necessarily identifies leadership with particular persons, we need to attempt to ‘see through the personalities, to dissolve them into the network of processes which have brought them to our attention’ (Collins, 1998, p. 4). We build on this relational conceptualization by considering leadership as an emergent phenomenon that develops in ongoing interactions. In the online environment we study—Twitter—such interactions take the form of follower relations, mentions, or retweets. While in bureaucratic organizations leadership is associated with formal authority, on social media platforms like Twitter leadership is first and foremost associated with centrality in chains of interaction.¹⁸⁹

189 This is not to say that only centrality matters. For instance, Gerbaudo (2012) documents how activists who manage central accounts can use their administrative privileges to filter or push messages, a role similar to the connective leadership on Facebook described by (Poell et al., 2016), while Uitermark (2017) shows how hackers used their computer skills as well as secret chat channels to steer Anonymous activism. A limitation of our approach is that we only examine public expressions on a particular social media platform.

As mentioned in the introduction, the patterns emerging from these interactions are far from trivial, as tendencies towards the centralization of leadership may interfere with philosophies of egalitarian democracy. We know that self-organizing networks are conducive to the concentration of power in the form of unequally distributed social ties (Barabasi and Albert, 1999; Watts, 1999). Degree distributions of natural, technical and social networks are often heavy-tailed, forming networks with few high-degree nodes and a majority of small-degree nodes (Broido and Clauset, 2019). On social media, too, network connectivity is generally highly uneven, also in the case of nominally egalitarian movements (Uitermark, 2017). A likely generating mechanism in the case of prominence on social media is preferential attachment (Barabasi and Albert, 1999; Dorogovtsev and Mendes, 2002; Khamis et al., 2017). Activists who operate the most central accounts—those with the most followers, subscribers, or friends—hold disproportionate power in setting the movement’s agenda and choreographing protest (Gerbaudo, 2012; Swann and Husted, 2017; Zhu and Lerman, 2016). As Zeynep Tufekci notes, ‘ostensible leaderlessness does not stop de facto leadership from springing up, and de facto leadership is often composed of those with the most time, tenacity, extroversion, pre-existing social status, and even plain aggressiveness’ (cf. Freeman, 1973; Tufekci, 2017, p. 79). In her study of the Chilean student movement, Von Bülow (2018) found that social media use reinforced inequalities within the movement in spite of efforts by some groups of activists to mitigate such inequalities.

A number of contributions stress that there is not only power concentration at any particular moment in time but such uneven patterns persist over time due to preferential attachment or other mechanisms of cumulative causation (Michels, 1915; Rijt, Moon, Restivo, and Patil, 2014). Nunes concisely describes such tendencies when he writes that networks generally ‘do not appear to self-organise their way *out of* power laws’ (Nunes 2014: 32, emphasis in original). For at least two reasons, it is important to make an analytical distinction between the concentration of leadership at any moment in time and the consolidation of leadership over time. The first reason is that many activists themselves do not reject leadership as such, accepting or even endorsing leadership as a benign and necessary aspect of organizing for progressive change. Many do, however, attempt to mitigate the consolidation of leadership since this would fortify inequalities and foreclose the opportunity for new voices to gain prominence. Patisse Khan-Cullors puts it as follows: ‘The consequence of focusing on a leader is that you develop a necessity

for that leader to be the one who's the spokesperson and the organizer, who tells the masses where to go, rather than the masses understanding that we can catalyze a movement in our own community' (cited in Cobb, 2016). A second reason for the distinction is that social media activity is generally volatile: discussions come and go at a pace rapid beyond 24-hour news cycles (Bruns and Burgess, 2012; Stieglitz and Dang-Xuan, 2013; Yang and Counts, 2009). It seems *prima facie* plausible that this would then also be the case for the prominence of individuals within those discussions. For these reasons, we suggest two distinct dimensions of leadership centralization. We refer to uneven network connectivity as leadership *concentration*; a network has a high degree of leadership concentration when the bulk of connections is concentrated in a small portion of its nodes. In addition, we examine leadership *consolidation*. Consolidation occurs when prominent people retain their position over the course of time, transforming from ad hoc leaders into an oligarchic elite.

Materials and methods

Data collection: Twitter

We use data collected from Twitter to analyze concentration and consolidation of leadership in networks. We obtained data from the social media platform Twitter by monitoring the streaming API for posts containing #blacklivesmatter, between June 15th 2015 and December 15th 2017. In collecting and reporting data that users have posted on Twitter, we observe the Twitter terms of service which stipulate how data and user information can and must be represented in research. We note that these terms, as well as many recent studies, provide mostly legal considerations about consent but little concern about risk of harm (Moreno et al. 2013). Given the contentious nature of movements such as Black Lives Matter and the potential vulnerability of activists (Townsend and Wallace, 2016), we adopt the framework proposed by Williams et al (2017) who show that, despite agreeing to legal terms, users tend to be unaware that their online behavior might be published in research papers. Accordingly, we report only aggregate patterns. We do not mention individuals and we do not present information through which individual users can be identified, such as the textual or visual content of posts.

In our network analyses, every unique user represents a node, each mention and retweet a tie. There are 23,991,665 ties among 4,632,350 nodes. Because

we conceptualize leadership relationally in terms of user prominence, we focus primarily on indegree, measured as the sum of the number of mentions received and the frequency of being retweeted.

Identifying supporters and opponents

While many studies equate the activity around specific hashtags with social movement activity generally, our methods acknowledge that many users who use the #blacklivesmatter are not supporters. Because we are interested in dynamics of leadership within the movement, it is vital that we distinguish between supporters on the one hand and opponents on the other hand. For these reasons, we used the following procedure.

In a first step, we build a training sample from 18,501,785 unique tweets. We build the sample progressively, sampling 0.1% of each day, to make sure posts from throughout the period are selected rather than only the busiest days. These 18,501 tweets are our training set. We then manually code tweets from the training set as positive, negative, or unclear. Any uncertainty results in a tweet coded as 'unclear', so that only tweets unambiguously in support or opposition of Black Lives Matter are coded as such. Using the coded tweets, we derive two lexicons from the training set: one with positive content and one with negative content, with content including both phrases and URLs. These lexicons then inform algorithmic coding of the full dataset to determine whether users are supporters or opponents. Users with at least three coded tweets, of which at least two thirds are coded as positive, are considered as supporters, while users with more than two thirds negative tweets are considered opponents.

Following our procedure, we identify 89,203 users as supporters (2.1%) and 144,630 users as opponents (3.4% of all 4,632,350 unique users). Not all of those supporters receive mentions: 20,849 or 0.5% of all unique users, appear in the tweets of others in our dataset. The proportion of users classified as supporters or opponents is fairly low because of our decision to filter out users who tweet with #blacklivesmatter incidentally (fewer than 3 times) or whose tweets are neutral or ambiguous. Because we are interested in the development of online leadership, we opt for accuracy over quantity and focus on unambiguous supporters and their relations in our further analyses.

Measuring concentration of leadership

We measure concentration of leadership as the distribution of prominence. More precisely, we use the method proposed by Clauset et al (Clauset et al., 2009) to determine the power law exponent over the indegree distributions, with a higher exponent indicating a more unevenly connected network. The exponent α captures the probability density $p(x)$ of picking at random a fraction of nodes from the empirical network on day t with a degree lower than k .

Practically, we fit power law models above an estimated threshold using the R *igraph* package (Csardi and Nepusz, 2006). To do so, we quantify the distance between the empirical observations and modeled power law distributions. By varying the threshold incrementally, the optimal x_{min} is determined by estimating a Kolmogorov-Smirnov statistic. In practice, this means that all users with a degree lower than the estimated threshold are excluded from the fitted power law model.

Measuring consolidation of leadership

We examine consolidation as the sustained prominence of leaders over time. We do so in two ways. First, we measure network consolidation using Spearman's rank. More precisely, we evaluate the correlations between the indegree of users at different time intervals. This is an adapted application of a method used to examine fame trends in traditional media (van de Rijt, Shor, Ward, and Skiena, 2013). The correlations indicate how strongly and in what direction the prominence of users in a particular period is associated with the preceding period. We reason that consolidation occurs when prominent users continue to be prominent, diminishing the potential attention space for newcomers. Or, more formally: if there is no relationship between early prominence and prominence in later periods, user ranks would be distributed randomly between two periods. To test this expectation we calculate Spearman's rank correlations for indegree. After determining the indegree of all users for each month in the observation, we calculate Spearman's rank correlations for indegree at various time intervals for the same user, taking the first month the user appears in our dataset as a starting point, t_1 . This is paired with the degree of that same user in the following month, t_2 . Then, t_1 and t_3 are paired, et cetera (van de Rijt et al., 2013). The increase or decrease in the correlation between these pairs provides insight in prominence over time: A user who sustains a similar degree rank position over time contributes to higher overall correlation for the population as a whole. Users are grouped in cohorts according to thresholds of average indegree. By reporting increases or decreases

for different cohorts of prominence, we can compare tendencies towards more or less consolidation for both the leadership and rank-and-file. While this measure usefully indicates correlations of prominence ranks, it comes with limitations: baselines are arbitrary, outliers are excluded, and it doesn't track individuals over longer periods. Baselines for prominence are determined by comparing the first appearance of a user in our dataset at later intervals. While users are aggregated to mean degree thresholds and correlations are calculated iteratively (and can theoretically be negative), there are many exterior reasons that may influence rank at t_0 . Moreover, statistically non-significant values and outliers warrant further examination, exactly because outliers include the most prominent leaders. Finally, the measure accurately measures overall decay of consolidation at various ranks, but we are also interested in how specific leaders consolidate their prominence over longer periods.

Second, we calculate a complementary measure of consistency over time, which addresses the limitations of the Spearman's ranking discussed. Do the same people remain prominent or do new prominent voices arise as time progresses? To examine this, we compile a list of the 100 most prominent leaders across the entire period, and compare this to the top 10 most prominent leaders daily. For every day we determine a top 10 of users receiving the most mentions and retweets. On that day, these user was prominent in terms of receiving attention (indegree). We then determine the number of days that users are among the ten most prominent recipients of mentions. We select the 100 users with the most number of days they are among the ten most prominent leaders. We refer to this as the 'established leadership'. We calculate which percentage of the daily top 10 leaders are part of the established leadership, with a higher percentage indicating more consolidation.

In short, Spearman's rank indicates how strongly prominence is correlated at various intervals for various cohorts overall, but doesn't allow for tracking the most prominent individual leaders over longer periods of time. To address those limitations we also examine how consistently the same leaders remain prominent.

Results

The concentration of leadership

Social media potentially enable many leaders to rise to prominence, but digital networking could also induce unevenness by concentrating most attention among

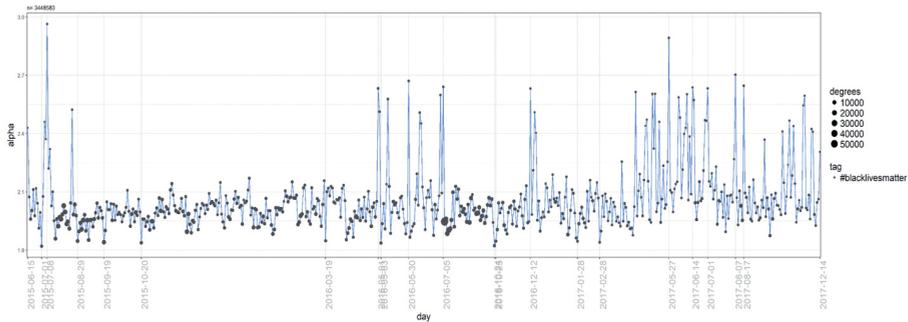
a select few. People engaging in online activity with #blacklivesmatter have together generated and sustained attention for the movement over the years. To what degree are Twitter networks of supporters of Black Lives Matter uneven? And can we observe a trend in the direction of more or less concentration of prominence? As explained above, we fit a power law on the indegree distribution and then use the power law's exponent as a measure of how unevenly mentions are distributed among users at various times (Figure 10.1).

Our findings do not suggest a marked trend in the direction of more or less concentration. A line that curves towards increased concentration in the later period is the best possible option but a very poor fit. The development of concentration over time is too erratic to capture with a regression line, particularly so during the later period under observation. If we consider the network of supporters as a whole and calculate the exponent for each day, we find a mean exponent of 2.08. Digital activity on other platforms exhibits similar distributions in combination with rapidly decaying prominence of anything but the long tail. Such pattern has been found for, for instance, Wikipedia links (Zlatić, Božičević, Štefančić, and Domazet, 2006, p. 3), online videos (Avramova, Wittevrongel, Bruneel, and De Vleeschauwer, 2009; Szabo and Huberman, 2010), Instagram popularity (Ferrara et al., 2014), and music popularity (Hu and Han, 2008; Sinha and Pan, 2006). Compared to similar networks, the network of supporters is slightly more evenly distributed but the differences are rather minimal. Whereas we find a mean exponent of 2.08, Ratkiewicz et al (2010:105) and Barabasi and Albert (1999) respectively found an alpha of 2.6 and 2.1 in their studies of online popularity.

While this suggests that online prominence among Black Lives Matter supporters is more evenly distributed than comparable networks, the difference is marginal, especially when we compare to other kinds of networks. For instance, Rijt et al. (2013, p. 275) found an alpha of 1.1 for traditional news media. In short, looking at the network of supporters as a whole, we find that it by and large conforms to what we would expect of social media networks. What further stands out is that concentration varies greatly on a daily basis. The exponent ranges between 1.81 and 2.96 for supporters (compared to 1.99 and 3.16 among all users). During particularly volatile periods, the exponent changes by factor of 10 or even 100 from day to day (note that the exponent's scale is logarithmic). In our estimations of the model fit, the range of the lower threshold falls between 1 and 5, with a mean of 2.34 and standard deviation 0.97, corresponding to 700,956

users (71.9% of all people receiving mentions, including both supporters and opponents). We also note that the power law models tend to increasingly diverge from the observed network in the upper ranges of the distributions. All in all, we find that (1) prominence is concentrated but (2) slightly less than in comparable networks and that (3) there is no clear trend toward more or less concentration over time.

Figure 10.1: The concentration of prominence among supporters of #blacklivesmatter fluctuates on a daily basis

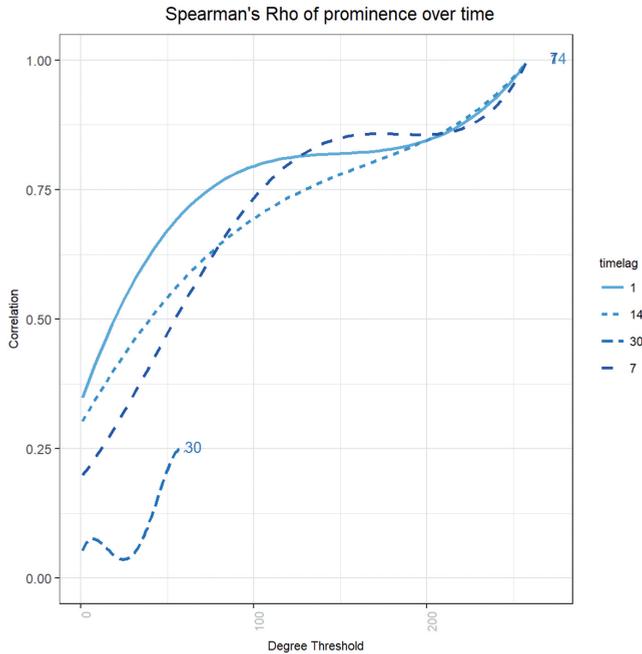


Note: Power law exponents by day, among supporters. Point size indicates the intensity of activity as the absolute cumulative indegree. Note that the graph is based on a calculation of the power law exponent for daily slices of the network. It takes into account minimum sample size, estimation of lower power law thresholds, and Kolmogorov-Smirnov p values for goodness-of-fit. Each point represents a cumulative degree distribution on a logarithmic scale: the distance between exponents of 1 and 2 is a factor of 10.

The consolidation of leadership

Social media potentially allow for the continual rise of new leaders to prominence, but digital networking could also induce inequality by consolidating most attention among the same people over time. Is there a trend among the leadership towards increasing or decreasing consolidation? First, we measure consolidation over time for different levels of prominence. Figure 10.2 shows the Spearman's rank correlation at various times, for users above a certain threshold of prominence. For example, the value 0.75 for users with at least 100 mentions indicates that out of all leaders with at least 100 mentions, 75 percent were already prominent a week before.

Figure 10.2. Only the most prominent leaders retain prominence



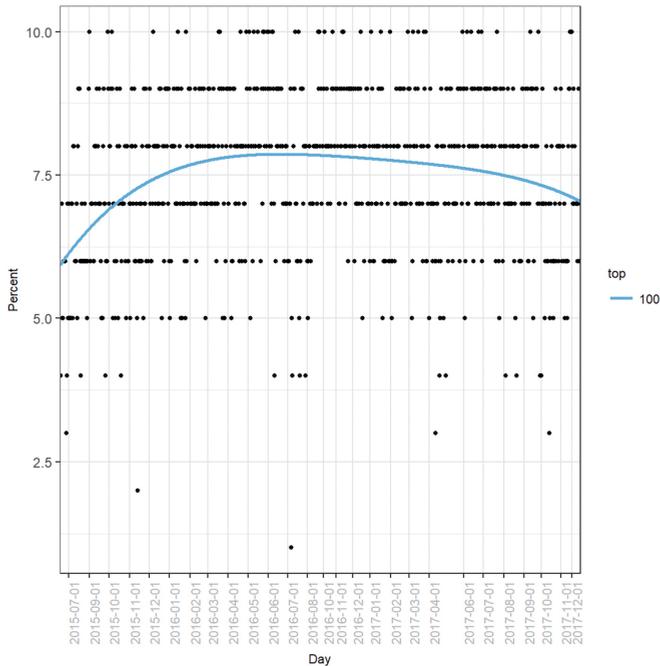
Note: Spearman's rank correlation (x axis) of prominence at different times, above varying threshold levels of prominence (y axis). Line type indicates the increase of days after initial measure of prominence. Note: only significant results are shown ($p \leq 0.05$).

There is a marked difference between rank-and-file and the most prominent leaders. This finding is consistent with studies of media celebrity using comparable methods that found that only the most prominent individuals consolidate over the long-term, while 'small fame quickly shrinks' (van de Rijt et al., 2013, p. 276). In our data, there is a very weak to weak correlation for retaining prominence more than a day among users being mentioned fewer than 10 times on average (20,035 users, about 96% of all 20,849 identified leaders, estimates ranging between 0.1 and 0.4, $\rho \approx 0.5$). For the middle cadre, users with 10 to 50 degrees (706, or 3.4% of leaders), results begin to diverge. Among this group there is a very weak to moderate correlation for being prominent the next day, but this effect quickly decays after more than a day (estimates ranging between 0.2 and 0.7, $\rho \approx 0.5$). The upper middle cadre, those between 50 and 100 mentions, correlations begin to range to strong, even after more than a day (20,035 users, about 96% of all 20,849 leaders identified, estimates ranging between 0.4 and 0.8, $\rho \approx 0.5$).

For the most prominent leaders, those receiving more than 100 mentions on average (31, or 0.1% of leaders), there is a moderate to strong correlation for remaining prominent even after two weeks (ranging between 0.5 and 1, $\rho \approx 0.5$). This correlation increases with rank. However, we find little evidence of retention after more than a month. There is a group of users with up to 50 mentions for whom is weakly correlated with prominence after a month. The graph does not show statistically non-significant results. Closer inspection of outliers suggests these are well-known leaders who do in fact benefit from prominence retention over longer periods.

Second, we examine such known leaders, to see whether they consolidate leadership over time. While the measure discussed above allows us to compare degree ranks, this additional measure allows us to examine consolidation over time of specific leaders. This is calculated as a percentage of the established leadership that is prominent on any given day (Figure 10.3).

Figure 10.3. The most prominent leaders are consistently prominent



Note: The percentage of the established leadership that is prominent on a daily basis. The vertical axis measures the percentage: how many of the 100 most prominent leaders overall reappear in the top 10 of any day (y-axis). A decline in the trend indicates less consolidation among the most prominent users.

Most of the prominent people on any given day are well-known leaders. In other words: on average, 7.5 of the 100 established leaders appear in the top 10 daily. This percentage ranges between 1 and 10, with a mean of 7.5 and standard deviation of 1.4. For instance, the value of 10 percent on September 1st 2015 indicates that 10 of the 100 established leaders appear in the top 10 of that day, or in other words, that all of the most prominent users that day are part of the established leadership. Although a poor fit, the trend suggests an early increase of consistency, when a larger share of the established leadership manages to consolidate their prominence. This share slowly declines after July 2016.

Discussion

The advent of social media has afforded activists with powerful new avenues for building decentralized movements. Such digitally networked movements would be less dependent on singular, charismatic leadership at the top of a hierarchy, promising instead a more diverse range of community-based leadership. Research examining digital networking often emphasizes these egalitarian aspirations (Bennett and Segerberg, 2012; Castells, 2000; Centola, 2015; Conover, Ferrara, et al., 2013; Goel et al., 2012; Rane and Salem, 2012; Tremayne, 2014) but also highlights the uneven nature of digital networks on social media platforms (Borge-Holthoefer et al., 2011; Shirky, 2003; van de Rijt et al., 2013). From our findings follow implications for further research of leadership in digitally networked movements.

First and foremost, we note the incredible number of people having engaged with the hashtag #blacklivesmatter over the years (over four and a half million unique Twitter users). Our findings suggest that distributions of prominence are highly uneven among the sizeable online leadership, in the sense that a small number of leaders receive a disproportionately large share of mentions and retweets. However, this inequality varies strongly from day to day, and only a select few most prominent leaders consolidate their position over time.

The methodological implication of our findings is that the use of power laws to describe relational structures in social movements warrants great caution. Our findings highlight at least two reasons for caution: variance over time and differences between concentration and consolidation. Power laws have been identified in a bewildering range distributions ranging from citation networks to earthquake

intensity (Broido and Clauset, 2018; Clauset et al., 2009) and social movement researchers have followed suit by observing power law distributions in digital networks (Borge-Holthoefer et al., 2011; González-Bailón et al., 2011). Contrary to the idea that power laws are inherent to digitally networked movements out of which they cannot self-organize (Nunes, 2013, p. 32), our findings suggest that power laws are less than invariant when analysed over time. While indications of an overall power law distribution may tell us something about the unevenness in what we might otherwise think of as strictly horizontal movements, aggregating all relations to a single timeframe obscures major day-to-day variance. By examining the patterns and causes of fluctuating degrees of power concentration, we get a better understanding of how network configurations change over time in social movements.

Our findings further show that it is critical to distinguish conceptually between the concentration and the consolidation of leadership. Social media do not necessarily incubate decentralized and egalitarian movements. Prominence in the limited attention space of social media is not only distributed unevenly, it is highly ephemeral for all but those most prominent. We find strong correlations for the prominence of leaders between different intervals, an effect clearly tied to rank. The Spearman measure indicates that these correlations quickly wear off after a couple of days, for most. On the one hand, this finding suggests that social media allow rank-and-file users a rapid but temporary rise to ‘insta-fame’ (Khamis et al., 2017; Marwick and Boyd, 2011; Marwick, 2015; Tufekci, 2013). Indeed, in the volatile, competitive, and high-churn attention space of social media, prominence ranks are ephemeral for all but the most salient leaders. This ephemerality poses distinct challenges for social movements. Social media may have a democratizing effect in the sense that a large number of individuals can momentarily become prominent, potentially making it easier for voices to rise up, but that may come at the price of a potential lack of institutionalization (Tufekci, 2017). The high turnover of attention implies a lack of continuity: most leaders cannot build the recognition and trust necessary to guide protest beyond an initial phase of prominence.

On the other hand, while only a select few leaders consolidate their prominence for anything longer than two weeks, our additional measure indicates they do so consistently. This elite is made up of 0.1% (31) of all supporters. These well-known figures consistently receive a lot of attention from other people, starring frequently among the ten most prominent users on any given day (although not necessarily on

consecutive days, as the Spearman ranks indicates). These voices of *de facto* leaders are ascribed a disproportionate influence by others. Our data and analysis did not focus on the period of such paths to consolidated prominence, offering important avenues for further inquiry (Centola, 2015; Khan-Cullors and Bandele, 2018; Mckesson, 2018). If digitally networked movements are marked by ephemerality from all but a few, the question arises how fledgling online prominence becomes consolidated in such exceptional cases (Kang, 2015). Their prominence puts these leaders in a uniquely difficult position: they generate visibility for the cause, but due to that very visibility they bear the brunt of public scrutiny from both within and outside of the movement (Tufekci, 2017, p. 79). Even if dependency on a single charismatic leader decreases, the ephemerality of social media opens pathways for the rise of voices whose prominence may lead to an equally quick fall out of favour, reinforcing rather than abating oligarchic tendencies.

Conclusions

A key challenge for digitally networked movements is that social media both potentially facilitate participatory democracy, while also potentially inducing inequalities. For activists aspiring to leaderful movements, this means that a more diverse range of voices could rise up from communities in lieu of one or few figureheads. We conceptualize leadership in digitally networked movements in relational terms, emphasizing relations over individual attributes. While not necessarily apprehensive of other important aspects of leadership, this understanding reflects that leadership is bestowed by others, being attributed recognition of one's voice. In the limited attention space that develops under the unifying hashtag #blacklivesmatter, prominence thus critically depends on relational dynamics.

To study these dynamics of leadership empirically, we applied network analysis to a dataset of roughly 18 million tweets containing #blacklivesmatter, collected over 30 months. Our procedure allows for the discernment of supporters and opponents, and relational structures among them. The long period under observation allows for a more detailed view of network dynamics in digitally networked movements than is common in prior research.

We find, in short: a large number of leaders, but also unevenness among them and limited consolidation. There are concentrations of prominence, although these

vary from day-to-day. While occasionally attention is relatively evenly distributed among supporters of the movement, attention is commonly dominated by few individuals. These leaders form important chains in loosely connected digital networks, potentially introducing new vulnerabilities. In terms of consolidation, the leadership of #blacklivesmatter is consistently diverse, and new voices readily gain prominence. Many conversations are not led by the same people day after day. For the movement as a whole, 99% of the online leadership changes on a daily basis. Only the most select few leaders manage to consolidate prominence persistently. In sum, while we do find degrees of leadership concentration among supporters of #blacklivesmatter, there is very limited consolidation of such prominence in the highly volatile attention space of Twitter.

These findings pioneer computational methods that may complement qualitative accounts of leadership in digitally networked social movements. Power law distributions are the prevailing method of measuring inequality in networks. We advanced this method by differentiating concentration and consolidation, assessing with multiple measures how inequality develops over an extended period. In doing so, our findings provide insight in today's preeminent social movement, suggesting how computational methods can be employed in ways that take into account the dynamic and diverse nature of movements.