



Analyzing user activity on Twitter during long-lasting crisis events: a case study of the Covid-19 crisis in Spain

Bernat Esquirol¹ · Luce Prignano^{1,2} · Albert Díaz-Guilera¹ · Emanuele Cozzo^{1,3}

Received: 21 September 2023 / Revised: 30 January 2024 / Accepted: 1 February 2024 / Published online: 29 March 2024
© The Author(s) 2024

Abstract

A pandemic crisis like the Covid-19 outbreak is a complex event, involving numerous aspects of the social life on multiple temporal scales. Focusing on the Spanish Twittersphere, we characterized users' activity behavior across the different phases of the Covid-19 first wave. Firstly, we analyzed a sample of timelines of different classes of users from the Spanish Twittersphere in terms of their propensity to produce new information or to amplify information produced by others. Secondly, by performing stepwise segmented regression analysis and Bayesian switchpoint analysis, we looked for a possible behavioral footprint of the crisis in the statistics of users' activity. We observed that generic Spanish Twitter users and journalists experienced an abrupt increment of their tweeting activity between March 9 and 14, in coincidence with control measures being announced by regional and state-level authorities. However, they displayed a stable proportion of retweets before and after the switching point. On the contrary, politicians represented an exception, being the only class of users not experimenting this abrupt change and following a completely endogenous dynamics determined by institutional agenda. On the one hand, they did not increment their overall activity, displaying instead a slight decrease. On the other hand, in times of crisis, politicians tended to strengthen their propensity to amplify information rather than produce it.

Keywords Online human behavior · Twitter · Covid-19 · Segmented regression · Bayesian switchpoint

1 Introduction

The COVID-19 outbreak is the first global pandemic of the information age. A large-scale public debate took place in a hybrid media ecosystem in which online social media (OSM) in general, and Twitter in particular, play a central role (Gallotti et al. 2020; Alshaabi et al. 2020; Tsao et al. 2021). The online debate addressed all the aspects of the pandemic induced crisis that the public opinion considered to be relevant: from the origin of the virus to the controversy surrounding vaccination campaigns; from the necessity (or lack of it) for lockdowns and other social distancing measures, to the effectiveness of the different national health systems. These are evidence that people used Twitter

to make sense of the pandemic crisis, i.e., they try to understand, through interaction with others, an ongoing event that is new, uncertain, and confusing (Heverin and Zach 2012). However, unlike other crisis events, as for instance earthquakes or terrorist attacks, a pandemic crisis is a complex phenomenon spanning different time scales. In particular, the first wave of a new infectious disease is an even more peculiar phenomenon, being completely unseen and unexpected. We can distinguish *a priori* three different phases for the unfolding of the first wave in a given region: an early phase in which the virus and information about it starts to spread, a phase in which the epidemic hits a particular region and the debate about control and mitigation measures arises, and a third phase in which these measures are finally implemented. Our hypotheses, building on some findings (Alshaabi et al. 2020; Dewhurst et al. 2020), are that online users will react differently to different phases and that different type of users will react in different moments.

In fact, it is well known that OSM react abruptly to external and internal perturbations, eventually resulting in deep transformations without warning (Britt 2015). This inherent volatility suggests that online behavior will change in

✉ Emanuele Cozzo
emanuele.cozzo@ub.edu

¹ Universitat de Barcelona, UBICS, Barcelona, Spain

² Heuristica Lab, Barcelona, Spain

³ Universitat Oberta de Catalunya, CNSC/IN3, Barcelona, Spain

response to crises and the associated interventions. In this sense, OSM can serve as an early warning system for any kind of emergency worldwide, from natural disasters to terrorist acts (Crooks et al. 2013; Fischer-Preßler et al. 2019).

The efforts made in the last few years to measure and characterize human and human-like activity in OSM mainly consist in the search for characteristic patterns in users activity data (Ferraz Costa et al. 2015). For instance, it has been shown that some aspects of human communication are described by universal statistics which may help distinguish among different types of human and human-like users (Gilani 2019; García-Gavilanes et al. 2014). Nonetheless, when an emergency occurs, all the typical patterns of activity are probably altered since individuals react to the unfolding of the events by changing their behavior. Thus, it is important to develop formal and computational methodologies to exploit digital traces in order to understand how actors at the micro scale adjust their behavior in response to macro scale phenomena, while collectively making sense of them Venturini et al. (2015). This is even more important if we take into account the misinformation processes that accompany an ongoing crisis event, named ‘infodemics’ in the case of a pandemic (Gallotti et al. 2020). In order to properly identify harmful activity in the context of a crisis, either abrupt or prolonged, we need to know how different classes of users behave in such circumstances.

In this work, we characterize online collective behavior throughout the different phases of the first wave of Covid-19, and, in particular, in response to the prevention and control measures adopted, using Spain Twittersphere as a case study. In particular, we will measure the daily posting frequency of a sample of Twitter users looking for abrupt changes in it. In this way, we will be able to exploit the OSM intrinsic volatility to analyze at the same time the effect of the different phases of the crisis on individual behaviors and whether a collective response emerges, thus estimating the long term effects of that on the typical patterns of activity for the duration of the crisis.

To this end, we consider the whole tweeting activity, regardless of the topic of the content, instead of only tracking tweets directly related to the Covid-19 outbreak. We are thus able to compare online behaviors before and during the crisis from an overall perspective. We partition the Spain Twittersphere in three different classes of users: politicians, journalists, and generic users, since we consider that they provide a good representation of the division of roles in a networked public sphere (Ausserhofer and Maireder 2013). Beyond the case study at hand, the proposed methodology can be applied in general to track and analyze the response of a well-defined OSM-sphere to exogenous and endogenous perturbations.

2 Methods

2.1 Data collection

We collected Twitter users’ timelines, *i.e.*, the time stamped list of tweets (including retweets) published by a user, of three different types of Spain-based profiles: journalists, politicians, and random users. In the next subsection, we describe the methods used for the construction of the three set of users.

Journalists’ accounts We observed that, usually, the news media follow on Twitter the journalists working for them and journalists usually have some variation of the word ‘journalist’ in their bio. Leveraging on these two elements, in order to sample Spanish journalists’ Twitter profiles, we searched the Twitter accounts of the top news media in Spain (list in appendix A) and got their friends (users being followed by the media account). Among them, we retained those who have in their bio some version of the word ‘journalist’. In this way, we obtained 1722 journalists’ profiles.

Politicians’ accounts We picked the members of the *Congreso XIV Legislatura* public list¹ that contains all the members in the Spanish congress during the February–June 2020 period. There are 350 congresspersons in Spain, but not all of them have a Twitter account: altogether, we retrieved 261 politicians’ profiles.

Random users: followers and friends Getting unbiased random samples of Twitter users for a given location is not an easy task given that geolocated tweets are sparse. We developed a method with the aim to reduce as much as possible the bias toward highly active users.

Our method works as follows:

1. We got the 100 most recent tweets from users located in Spain through the Twitter Standard Search API
2. We filtered those 100 tweets so that the creator of the tweet (u_0) had more than a given threshold number of followers, friends, and actions (55, 95, and 1000, respectively).
3. We got one follower and one friend of u_0 at random.

We repeated the process until we retrieved around 8.000 profiles in each random set, *i.e.*, followers and friends.

Step 2 is to filter out bots or extremely inactive users. Step 3 is to avoid a bias toward high-frequency users that we would have by taking user u_0 directly. In this way, we get two types of random users: random-friends are those that we get sampling the friends of u_0 , and random-followers those obtained by sampling the followers. The first group

¹ [Public Twitter list](#) containing Spanish congress persons.

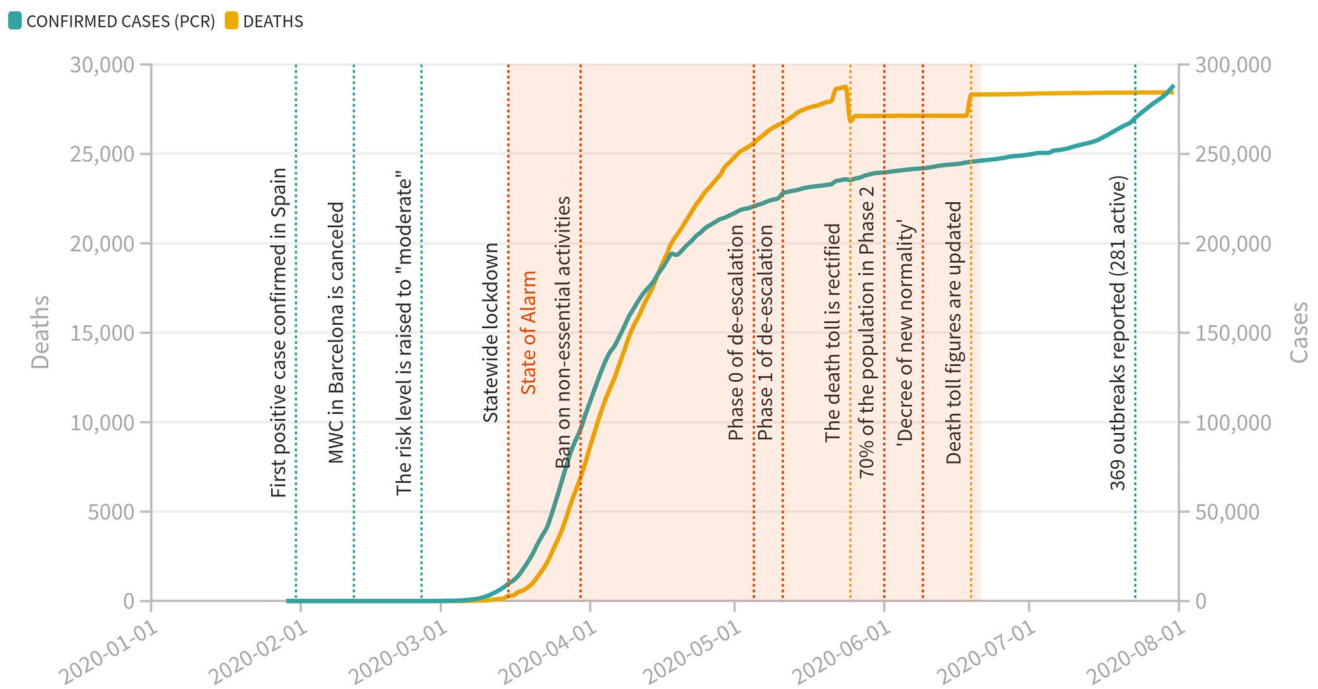


Fig. 1 Timeline of Covid-19 pandemic-related events in Spain: restriction measures, confirmed cases, and deaths during the first wave. A interactive version of the infographic is available at <https://public.flourish.studio/visualisation/4432020/>

is biased toward users following many other accounts, and the other is biased toward users having many followers. We used both sets separately and compared the results of the analyses performed on them at each stage. The underlying idea is that, whenever we identify a trait common to both random-followers and random-friends, such a trait can be regarded as a property of random users in general.

The Twitter Search API allowed us to retrieve 3200 tweets for each timeline, starting from the most recent one. Hence, depending on the user’s activity, we were able to cover a more or less large time period. The first data retrieval was conducted on April 28, then we periodically collected new tweets from the sampled users via the Twitter Search API.

2.2 Analysis of the activity

We characterized the activity of an individual user both from an overall (time-aggregated) perspective and analyzing its evolution through time, during the different phases of Covid-19 first wave (see Fig. 1).

First of all, we characterized individual profiles based on two variables. On the one hand, we compared their tendency to publish their own contents and to share other people’s tweets; on the other hand, we analyzed how distributed or concentrated their retweets are, in terms of the accounts that receive them (sources). In this way, we could highlight differences between the four samples of users under study, thus tracing a

preliminary behavioral profiling of the corresponding classes of accounts.

Secondly, we looked for possible abrupt jumps that the activity of a user may have experienced in some particular days, and we quantify them by means of the relative change in the mean activity. To identify such jumps, two different types of analysis have been performed and the results compared.

2.2.1 Production versus amplification of information

In this first preliminary time-aggregated analysis, we characterized each individual users according to their propensity to produce new information, *i.e.*, tweeting, or to amplify information produced by others, *i.e.*, retweeting. In addition, we quantify the diversity of the sources of the amplified information. With this aim, we introduce for each user *u* two parameters: the fraction of replicated content ρ_u and the mean entropy of the sources h_u , a measure of how generalist a given user is in terms of her information/content sources. Given the number of total tweets including retweets M_u emitted by a users and the number of retweets R_u ,

$$\rho_u = \frac{R_u}{M_u}. \tag{1}$$

The fraction of retweets of a given user to the same source *s* is $p_u^s = \frac{r_u^s}{R_u}$, where r_u^s is the number of retweets to the source

s. Thus, the normalized entropy of the sources for the user u is defined as

$$h_u = \frac{-\sum_s p_u^s \log p_u^s}{\log R_u}, \tag{2}$$

where the denominator is the entropy of the sources of a user with the same number of retweets, but being each one to a different source, that is, the maximal entropy attainable with that number of retweets.

In this way, we could compare each set of users with the others by visual inspection of their distribution in the $h - \rho$ plane. Quantitatively, we use the Kullback–Leibler (KL) divergence to measure the disparity between different set of users of the distributions of ρ and h separately. The KL divergence, given by the expression

$$D_{\text{KL}}(P \parallel Q) = \sum_i p(x_i) \log \left(\frac{p(x_i)}{q(x_i)} \right), \tag{3}$$

is one of the most used disparity measures between probability distributions. It is a measure of the information lost when Q is used to approximate P , being P and Q on the same finite support.

2.2.2 Stepwise segmented regression analysis

The first approach that we adopted for studying possible changes in the activity of different types of Twitter users was a stepwise segmented regression (SSR) analysis (Britt 2015) of the timeline of each profile. Given the list of timestamps of the tweets of a user, we optimized a piecewise linear fit that divided the timeline into segments. In this way, for each timeline we obtained a number of breakpoints—the end points of each segment—marking the moments of abrupt change in the tweeting frequency of the user. The optimal breakpoints were estimated by minimizing the residual sum of squares (Bai and Perron 2003), while their number was fixed by minimizing the Bayesian Information Criterion of the model.

In this way, we could associate to each user u the set of breakpoints $\tau_1^u, \tau_2^u, \dots, \tau_n^u$, *i.e.*, the days in which their activity changed abruptly, if any. Finally, we characterized users' breakpoints by the relative jump

$$J_i^u(\tau_i) = \frac{L_i^u - L_{i-1}^u}{m^u} = \frac{\Delta L_i^u}{m^u} \tag{4}$$

where L_i^u is the tweeting frequency of user u in the interval $[\tau_i^u, \tau_{i+1}^u)$ and m^u is the mean frequency on the whole user u retrieved timeline.

We could hence describe the overall activity by summing separately positive and negative jumps:

$$J^\pm(\tau) = \sum_u J_i^u(\tau_i) \Theta(\pm J_i^u(\tau_i)), \delta_\tau^{\tau_i^u} \tag{5}$$

where Θ is the Heaviside function, selecting for positive and negative jumps, respectively, and δ is the Kronecker delta, selecting jumps occurring at $\tau_i = \tau$.

2.2.3 Bayesian switchpoint analysis

The second approach was based on Bayesian inference. We performed a Bayesian switchpoint (BS) analysis, which looks for a single moment of change in the rate at which events occurred. The model assumes, for each user, Poisson-distributed tweeting events with constant but potentially different rates, before and after a switchpoint.

$$\lambda_1^{(0)} : \text{Exponential (rate} = \alpha) \tag{6}$$

$$\lambda_2^{(0)} : \text{Exponential (rate} = \alpha) \tag{7}$$

$$\tau : \text{Uniform [0, 1)} \tag{8}$$

For $i = 1 \dots n$:

$$\lambda_i = \begin{cases} \lambda_1^{(0)} & \text{if } \tau \leq i/n \\ \lambda_2^{(0)} & \text{if } \tau > i/n \end{cases} \tag{9}$$

$$m_i : \text{Poisson (rate} = \lambda_i) \tag{10}$$

Our prior for $\lambda_{(1,2)}$ was an exponential distribution with rate α , while our prior for τ was a uniform distribution over the whole time span of the timeline. α is the only hyperparameter in this Bayesian graph. From the expected values of the distribution we have:

$$E(m_i) = E(\text{Poisson}(\lambda_i) \mid \lambda_i) = E(\text{Exponential}(\alpha) \mid \alpha) = \frac{1}{\alpha}. \tag{11}$$

By fixing $E(m_i) = m$, *i.e.*, equating the expected values of the messages per day to the observed mean m , we have $\alpha = \frac{1}{m}$. In this way, we converted our prior in an informative prior. Another option would have been to consider two α_i , one for every Λ_i , but the values of those α_i would not be so clear and unbiased.

An Hamiltonian Monte Carlo method (Neal 2011) as implemented in TensorFlow (Abadi et al. 2016) was used for sampling from the relevant posterior distributions of $\lambda_{(1,2)}$ and τ .

Thus, we could associate to each user a probability distribution for the switchpoint and a probability distribution for the two different tweeting rates, before and after the switch, respectively. To carry out an overall analysis of the behavior of the Spanish Twittersphere, as for the SSR analysis, we

needed to characterize the switchpoints. By approximating the distributions of the $\lambda_{(1,2)}$ with a delta centered in their mode, we obtained the following expression for the relative jump:

$$J^u(\tau_i) = P(\tau^u = \tau_i) \frac{\tilde{\lambda}_2^u - \tilde{\lambda}_1^u}{m^u} = \frac{\Delta \lambda^u}{m^u} \quad (12)$$

where $\tilde{\lambda}_{(1,2)}$ is the mode of the respective posterior distribution, and $P(\tau^u = \tau_i)$ is the posterior probability of the jump occurring at day i that acts as a weight. We could then characterize the overall activity by summing separately positive and negative jumps of each user weighted by the jump probability:

$$J^\pm(\tau) = \sum_u J^u(\tau_i) \Theta(\pm J^u(\tau_i)) \delta_\tau^{\tau_i} \quad (13)$$

3 Results

3.1 Overall activity of different classes of users

Comparing the different sets of users with respect to their tendency to produce new information rather than amplify already existing contents (ρ) and how they distributed their retweets among different sources (h), we observed that the two sets of random profiles—followers and friends—were practically indistinguishable (Fig. 2, upper panels). They both showed a nearly flat distribution of ρ (vertical axe), covering equitably all the spectrum from pure replicator ($\rho = 1$) to pure content producer ($\rho = 0$). For what concerns h (horizontal axe), both groups displayed a markedly right skewed distribution, a clear indication that most of them tended to distribute their retweets almost uniformly among different sources. Journalists (Fig. 2, lower right panel) appeared to be quite similar to random users, but avoiding extreme values of ρ , in line with the expectation that a journalist will always tend to retweet her media and colleagues as well as produce new contents. Finally, politicians (Fig. 2, lower left panel) displayed the most peculiar activity patterns. They are more similar to journalists in ρ but showing a completely different distribution of h (symmetric, with the mean and the mode falling slightly above 0.5).

Such differences and similarities are quantitatively described by the KL divergence for the distributions of ρ and h , represented as heatmap matrices in Fig. 3.

3.2 Time evolution

Once we clarified that random users differed from politicians and, although to a lesser extent, journalists in terms of their overall (re)tweeting activity, the next step was to compare

the evolution of their activity rate during the first wave of Covid-19.

For a preliminary general inspection, we considered all the users together, regardless their class. From both the SSR and BS analyses (Fig. 4), we obtained a clear indication that in the week between March 9 and 14, 2020 the majority of Spanish twitter users abruptly changed their tweeting behavior, starting to tweet with higher frequency. These days coincided with the first control measures being introduced at regional level on Monday 9 and the declaration of the state of alarm by the prime minister of Spain on Friday 13, which implied enforcement of a state level lockdown.

The BS analysis also detected a noticeable decrease in the tweeting activity in those days, but much less intense than the opposite behavior.

As a final note, the SSR and BS analyses were discordant near the boundary of the considered time window. However, this fact can be easily explained considering that these two models are biased in opposite ways in these regions. The analysis had hence to be regarded as inconclusive for what concerns both October 2019 and May 2020.

We repeated the analyses for the four sets of profiles separately (Fig. 5). In this case, we could observe that the politicians were the only group that did not display any change in their activity in the week of the lockdown declaration. Contrarily, both analyses indicate a positive jump in the activity near January 7, the day on which the two parties forming the coalition government found an agreement, as well as a pronounced negative jump near February 12, day in which the composition of the government was announced by the prime minister. Finally, journalists also showed an interesting behavior. Relative positive jumps concentrate in the aforementioned week, but they also presented a secondary peak on March 2, the day before the first Covid-19 death was officially confirmed in Spain.

Finally, we compared the distribution of h and ρ for the four classes of users before and after March 9 and starting on the first of January (Fig. 6). Contrary to what we observed for the activity, the only type of users that change their behavior are politicians. In particular, we observed a steep increase in their propensity to amplify information produced by others instead of publishing their own content.

4 Conclusions

This study highlights the dynamic relationship between crisis events and individual behavior changes on social media, particularly focusing on the Spanish Twittersphere during the COVID-19 pandemic. Our analysis reveals that political discourse surrounding crisis management significantly influences collective and individual online behaviors, more so than the crisis events themselves.

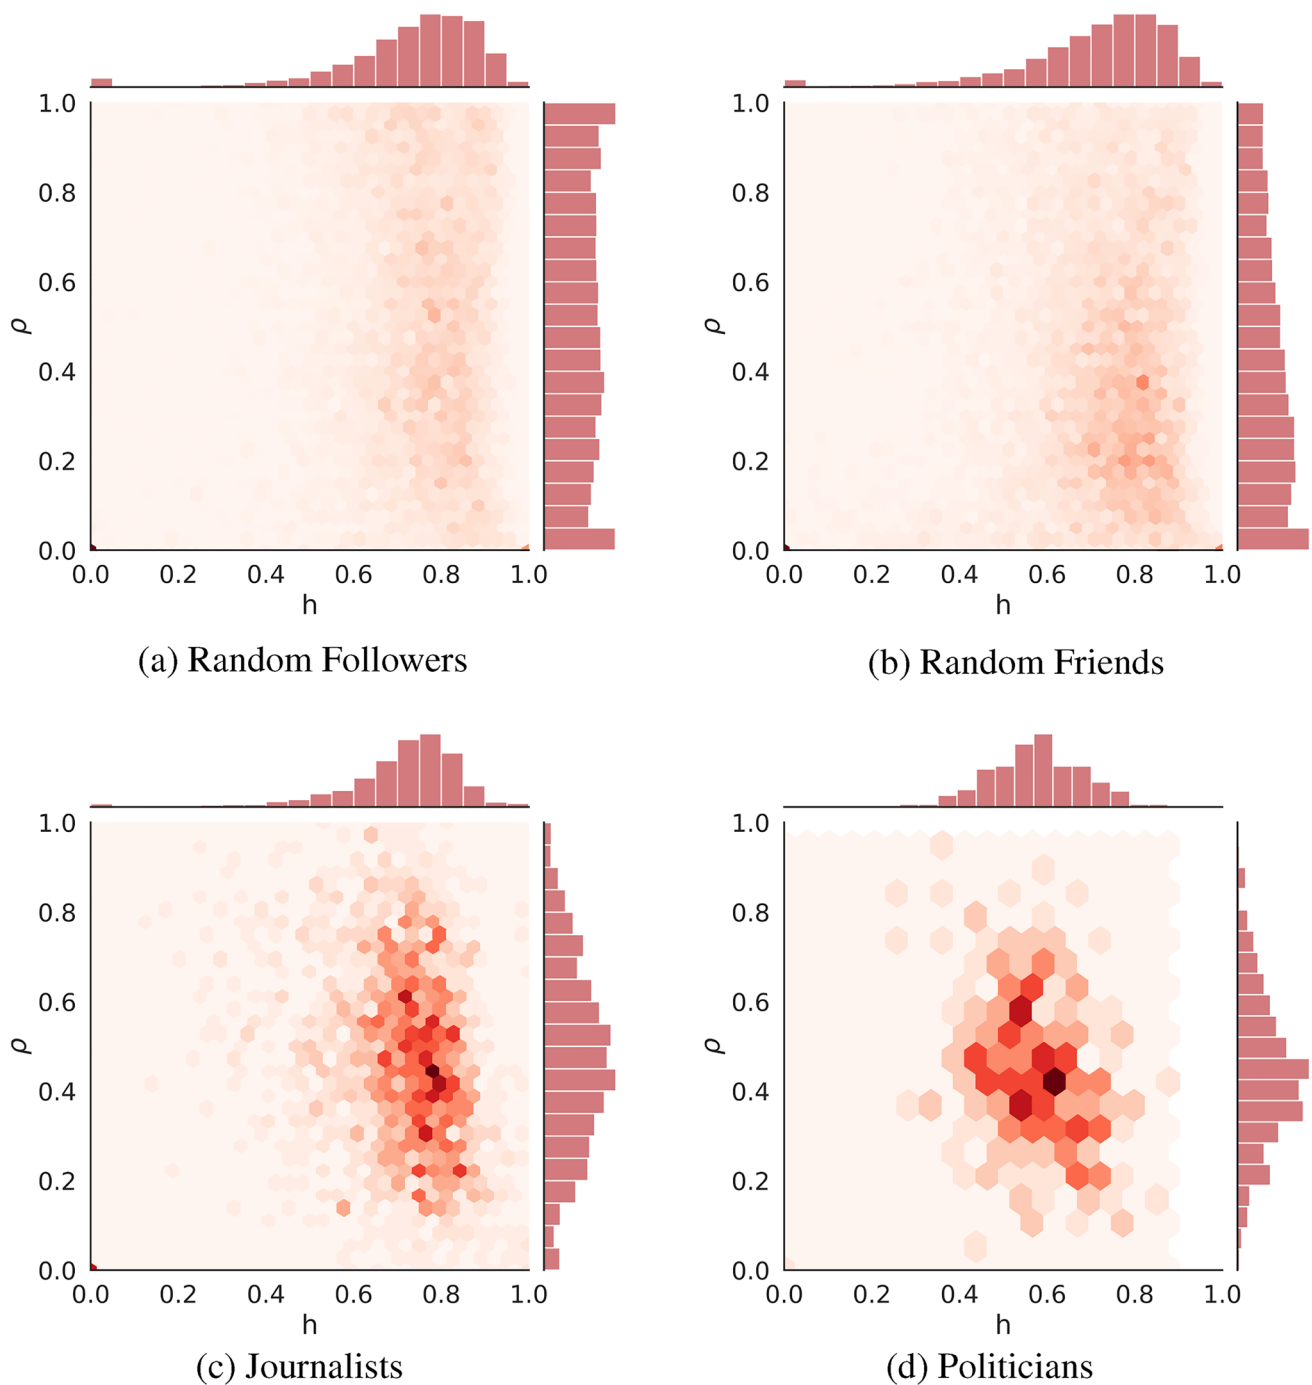


Fig. 2 Distribution of the fraction of retweets ρ and normalized entropy h for each one of the considered classes of users

Our analysis detected a notable shift in individual user activity on Spanish Twitter that coincided with the announcement of COVID-19 control measures. This shift in behavior is primarily attributed to the political response to the crisis rather than the epidemic outbreak itself. This focus on individual-level behavior change aligns with findings from the U.S.A. at the aggregate level, where an increase in OSM usage was noted following the declaration of a

national emergency (Valdez et al. 2020). Additionally, our results correspond with patterns of language use on Twitter, where there was a marked increase in discussions about control measures from early March, while general COVID-19-related terms began appearing in January (Dewhurst et al. 2020).

Furthermore, the initiation of the lockdown did not result in a further increase in online activity, which was a deviation

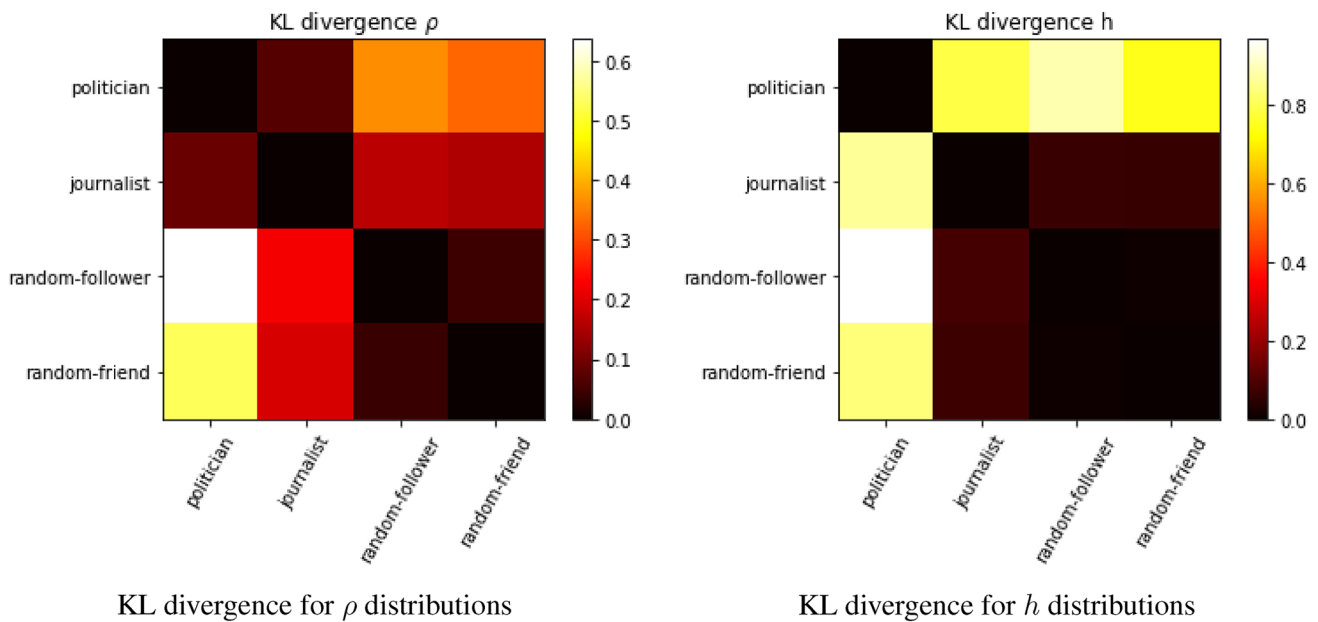


Fig. 3 Kullback–Leibler (KL) divergence of ρ (left panel) and h (right panel) distribution between classes of users

from initial expectations. Rather, the activity had already escalated during the discussions about impending measures, and the lockdown served to maintain this elevated level of engagement. This pattern, marked by sustained engagement during discussions about state-imposed measures, deviates from typical responses observed in political debates.

The work of Durazzi et al. (2021) complements our findings. In their study of the network structure of the English-typing COVID-19-related Twittersphere, they found that during the critical phase of the pandemic, users were increasingly focusing their attention on local authorities, political leaders, and key influencers. This observation underscores the crucial role of political responses and crisis-related discourse in shaping both individual and collective behavior on Twitter during a public health emergency. Our results directly relate these changes in collective behavior to synchronized abrupt changes at the individual level, further emphasizing the significant impact of political discourse during crises.

Previous studies (Boon-Itt and Skunkan 2020) have shown that users on Twitter discussed three main themes during the course of the first wave of Covid19: the emergency of the COVID-19 pandemic, control measures, and reports on the COVID-19 pandemic. The results of our analysis consistently point to the fact that it was the discussion on control measures that triggers abrupt behavior changes at individual levels in the form of increased activity. It is worth mentioning that the phase in which we observe major changes in users' activities roughly matches the transition between two different phases of the associated social epidemics (Strong 1990) as shown in Aiello et al. (2020). In

particular, extrapolating results from Aiello et al. (2020), jumps in activities should be located between the suspended reality phase, in which people express anger about the looming feeling that things were about to change, and the acceptance phase, started after the authorities imposed physical-distancing measures, in which people adjust to a "new normality." This observation calls for a deeper understanding of the relation between online activity and psychological states during a crisis.

Additionally, our findings reveal distinct dynamics in the activities of politicians and journalists. Politicians' activity patterns were mostly endogenous, responding primarily to the institutional agenda, as it is evident from their activity peaks during political campaigns. As for the type of activity—production or amplification of contents—, we observe a sharp concentration around middle levels of amplification. However, their behavior shows a concentrated distribution of retweets from select sources, displaying lower diversity compared to other user types. This aligns with Jungherr (2014), suggesting that politicians across various contexts use Twitter mainly for disseminating campaign activities and policy statements, and when interacting, they predominantly engage with other politicians or journalists. Unlike other user groups, politicians exhibited a change in their behavior, favoring amplification over content creation during certain periods. Policymakers tended to suppress their individual voices, likely in the attempt to provide a more coherent view on the topic, while reducing their involvement in other debates.

In contrast, ordinary citizens and journalists displayed more stable content production and consumption habits.

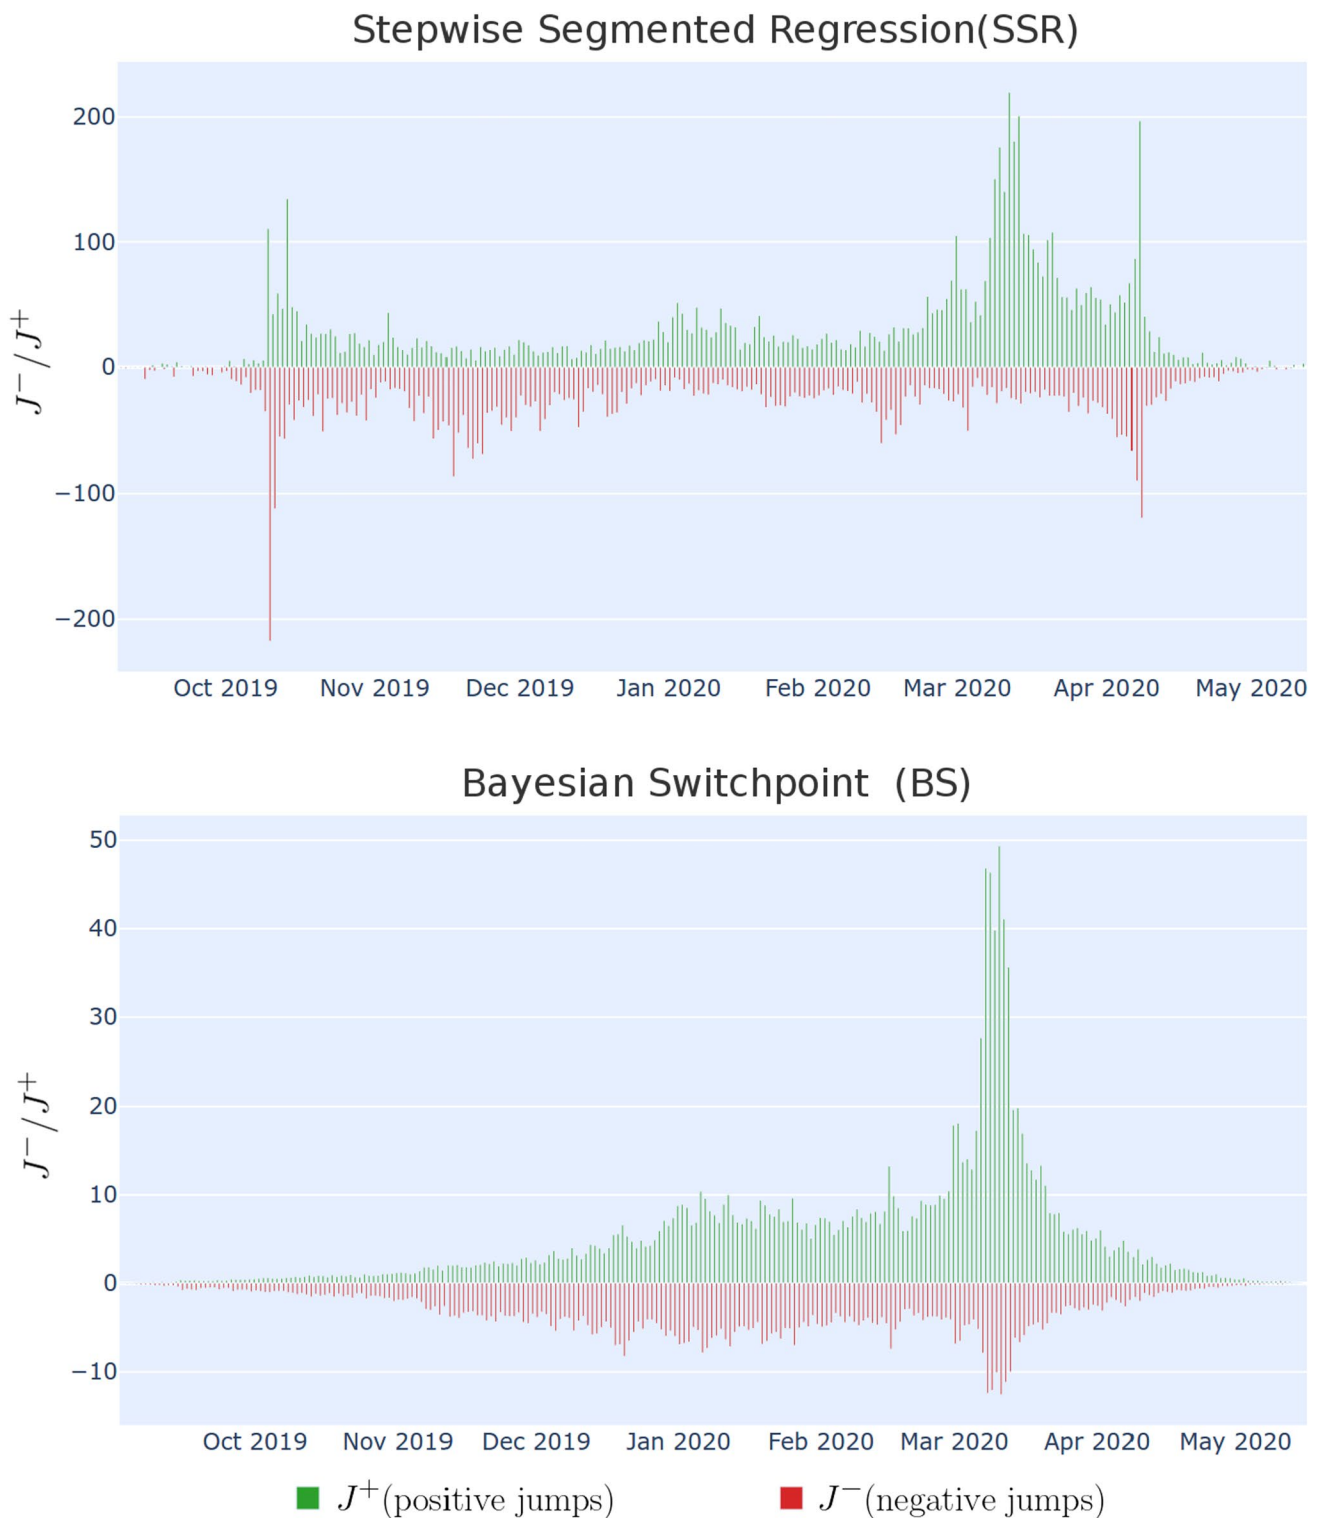


Fig. 4 Total daily positive activity jumps (green) and total daily negative activity jumps (red) calculated by performing stepwise segmented regression (SSR, upper panel) and Bayesian switchpoint analysis (BS, lower panel) for the whole set of users (Color figure online)

Citizens seemed to integrate the pandemic as a persistent theme in their usual online activities without significant changes in their engagement or source diversity. Journalists,

while mirroring ordinary citizens in their content habits, also played a leadership role by anticipating trends, as indicated by the existence of a secondary peak on March 2, although

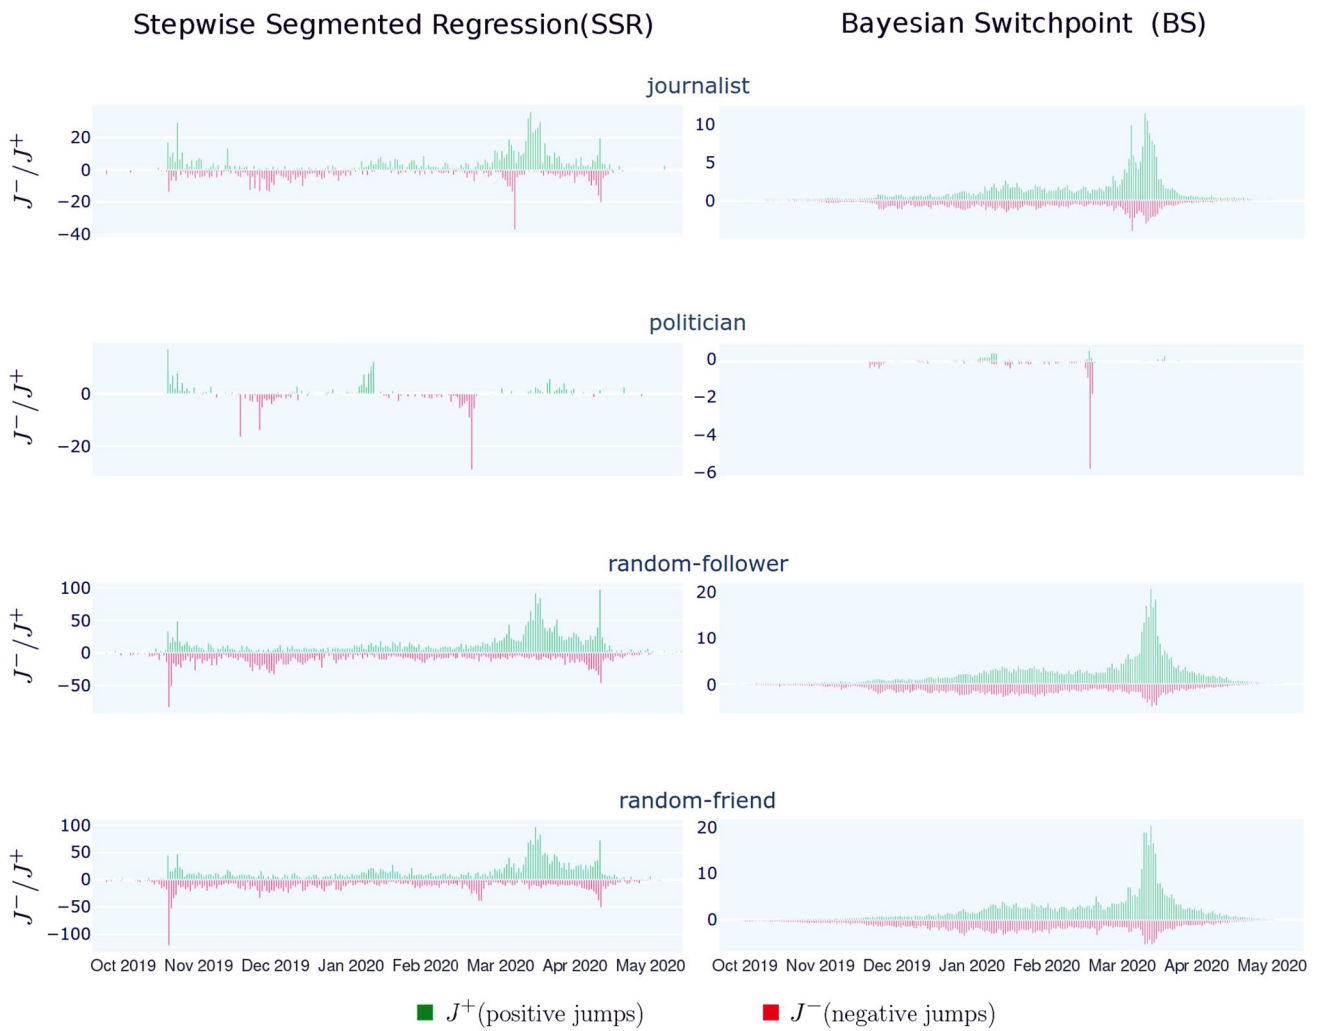


Fig. 5 Total daily positive activity jumps (green) and total daily negative activity jumps (red) calculated by performing stepwise segmented regression (SSR, upper panels) and Bayesian switchpoint analysis (BS, lower panels) for each individual class of users (Color figure online)

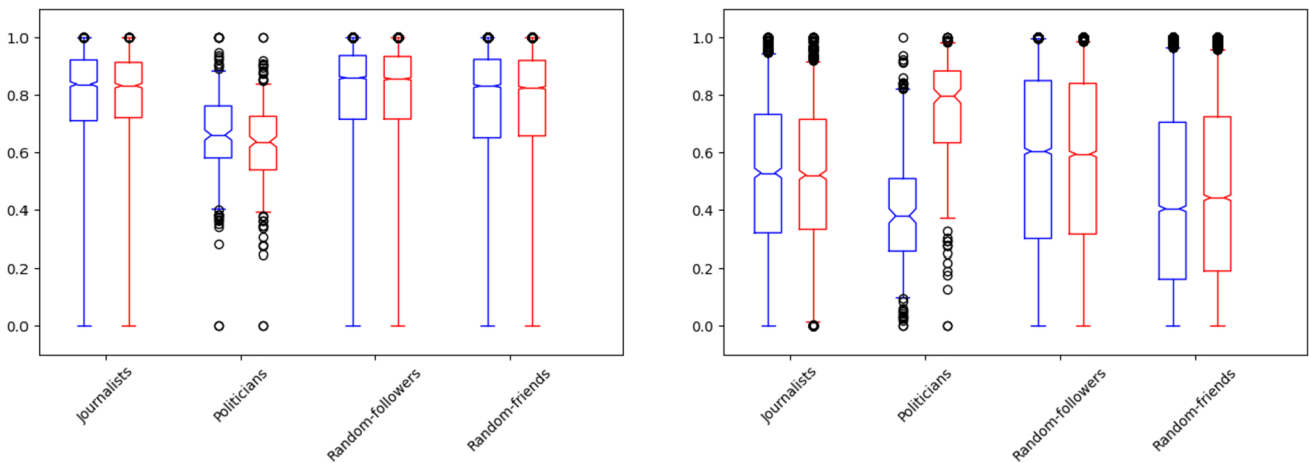


Fig. 6 h (left panel) and ρ (right panel) statistics before (blue) and after (red) March 9 of, from left to right, journalists, politicians, random followers, and random friends. Whiskers are set at the 5th and

95th percentiles. Dots represent outliers. Notches represent 95% CI of the medians. Median values obtained by bootstrapping data 1000 times (Color figure online)

their anticipation was relatively short-term compared to the evolving pandemic.

Methodologically, this study highlights the importance of dissecting individual behaviors to comprehend observed aggregated behaviors on social media. The increase in total volume of daily posts could result from various microscopic mechanisms, such as new users joining the platform, a gradual increase in existing users' activity, or, as in our case, synchronized abrupt jumps in individual activity levels. These different mechanisms signify distinct social phenomena and have varied implications for interventions. Particularly considering the infodemic associated with an ongoing crisis, the fact that the rise in overall volume stems from an increase in the activity of a large portion of users, without changes in their consumption habits (i.e., amplification vs content creation), has significant implications for assessing misinformation amplification levels. This is also relevant for automated bot detection if post frequency is a classification feature. We have proposed a statistical method that allows for distinguishing the microscopic mechanisms responsible for observed collective behavior and the associated time-scales during a prolonged crisis. This method, as demonstrated through our case study, also enables insights into the potential causes of individual behavioral shifts and the timing of these shifts. Further analysis of interaction structures and content, which was not performed in this study, would complete the picture and provide a more comprehensive understanding of these dynamics.

In conclusion, our study provides a nuanced understanding of the interplay between political discourse and individual behavior on social media during a public health crisis. It highlights the complex effects of crisis management

strategies on individual online engagement, essential for effectively addressing and mitigating the impacts of future crises.

Appendix A List of news media Twitter handlers

el_pais, JotDownSpain, eldiarioes, elespanolcom, revistamongolia, la_ser, _infoLibre, EFEnoticias, elmundoes, elconfidencial, indpcom, ctxt_es, publico_es, ondacero_es, cuatro, LaVanguardia, europapress, laSextaTV, rtve

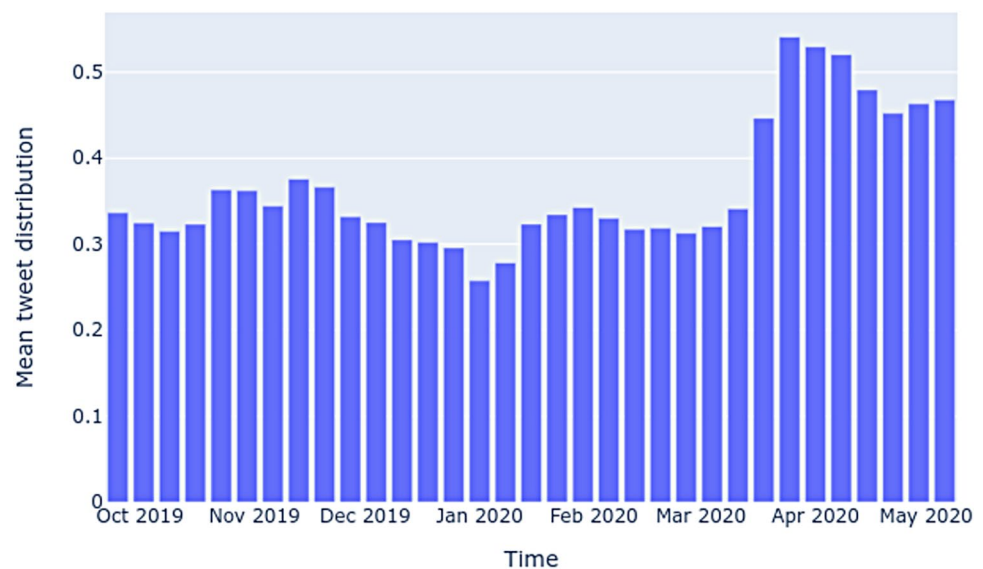
Appendix B Overall activity

In Fig. 7, we represent the mean number of tweets per day for each week. We observe that the overall activity increases in the week after the one in which we observe the majority of individual changes.

Appendix C Distributions of the number of breakpoints in the stepwise segmented regression analysis

In Fig. 8, we show the fraction of the number of optimal breakpoints for each user type. We can see that the distribution is quite similar among different user types. We can observe that the maximum of the distribution is always centered at 2 breakpoints, and that few users reach the 5 breakpoints. This serves as an *a posteriori* validation of the maximum number of breakpoints we set in our model.

Fig. 7 Activity per week for all our users



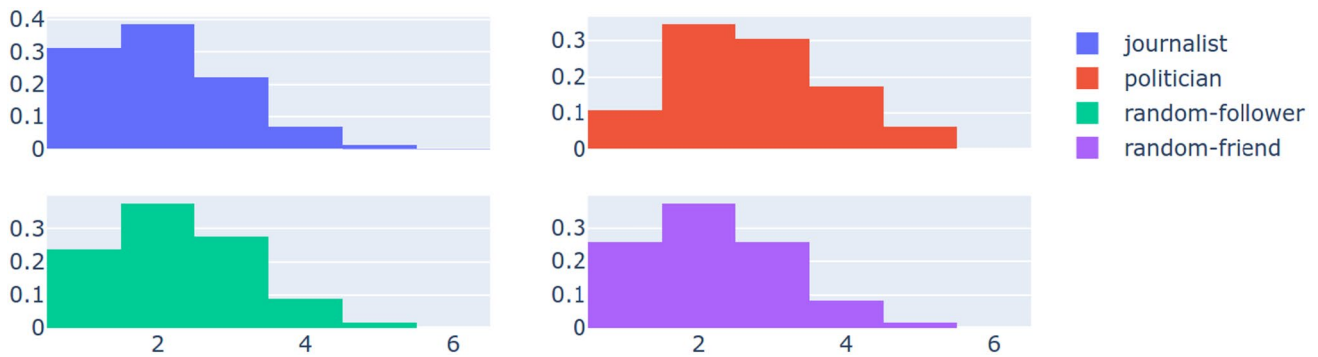


Fig. 8 Number of breakpoints for each user type

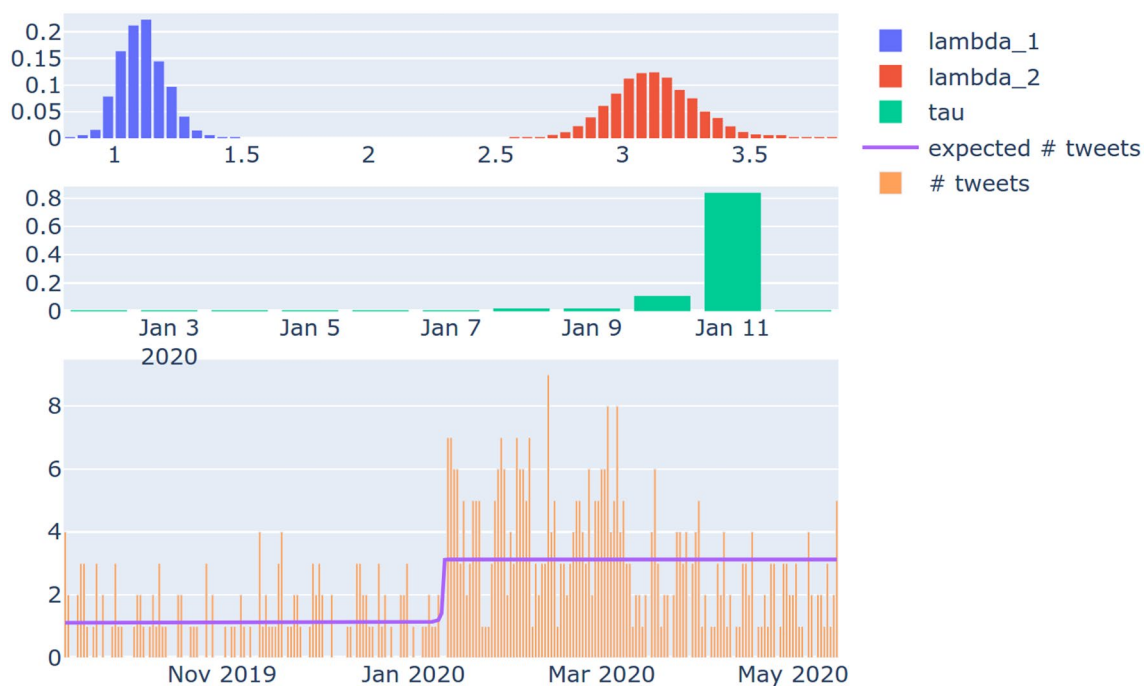


Fig. 9 Synthetic data with λ changing from 1 to 5 at the 2019-01-12 and from 5 to 2 at the 2020-03-02. For the model, we obtain a p -value = 0.46

Appendix D Bayesian switchpoint analysis in the presence of multiple switches

Users timeline may show more than one switch, while our model assumes there is just one. We show some examples of the results we get when synthetic timelines are generated with multiple switches (Figs. 9 and 10).

In both examples, we can see how in case of multiple switches the model will obviously get only one of them. We can conclude that we can use the p -value as a way of knowing if there are some existing breakpoints in the data that we missed. When the p -value is low, it does not

mean that the model has to be rejected, but rather that it is choosing one of the breakpoints present in the data, while other breakpoints probably exist.

Appendix E Bayesian switchpoint analysis with non-Poissonian models

For comparison, we also analyzed other models for a given example.

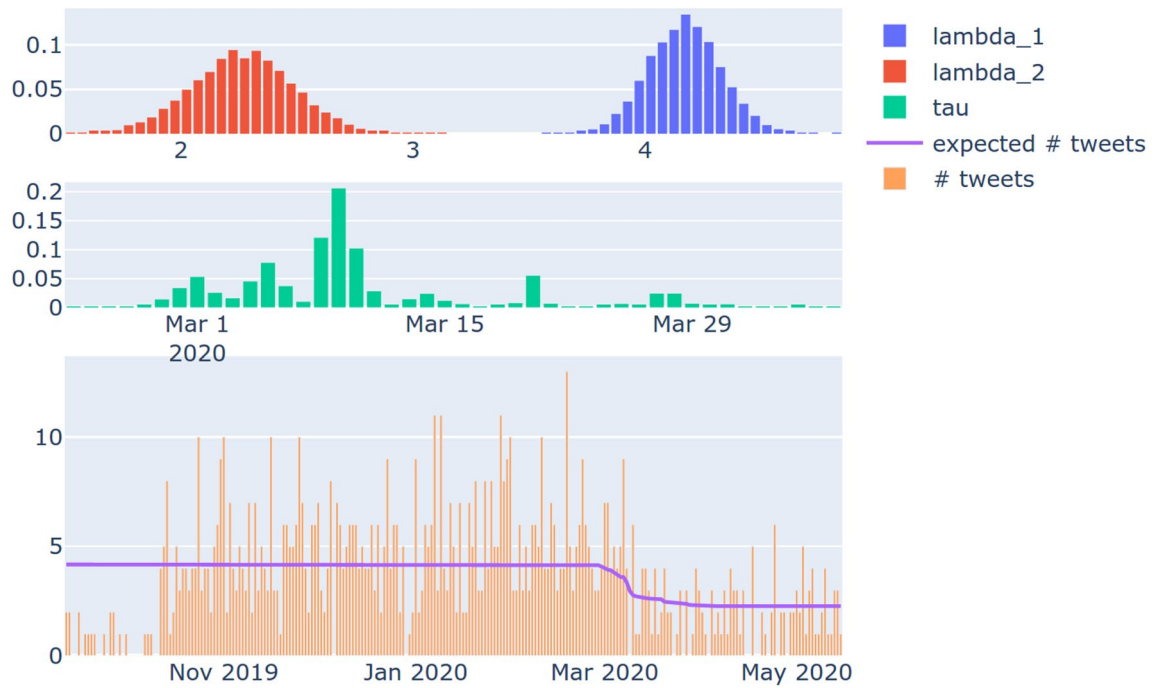


Fig. 10 Synthetic data with λ changing from 1 to 5 at the 2019-10-12 and from 5 to 2 at the 2020-03-02. For the model, we obtain a p -value = 0.50

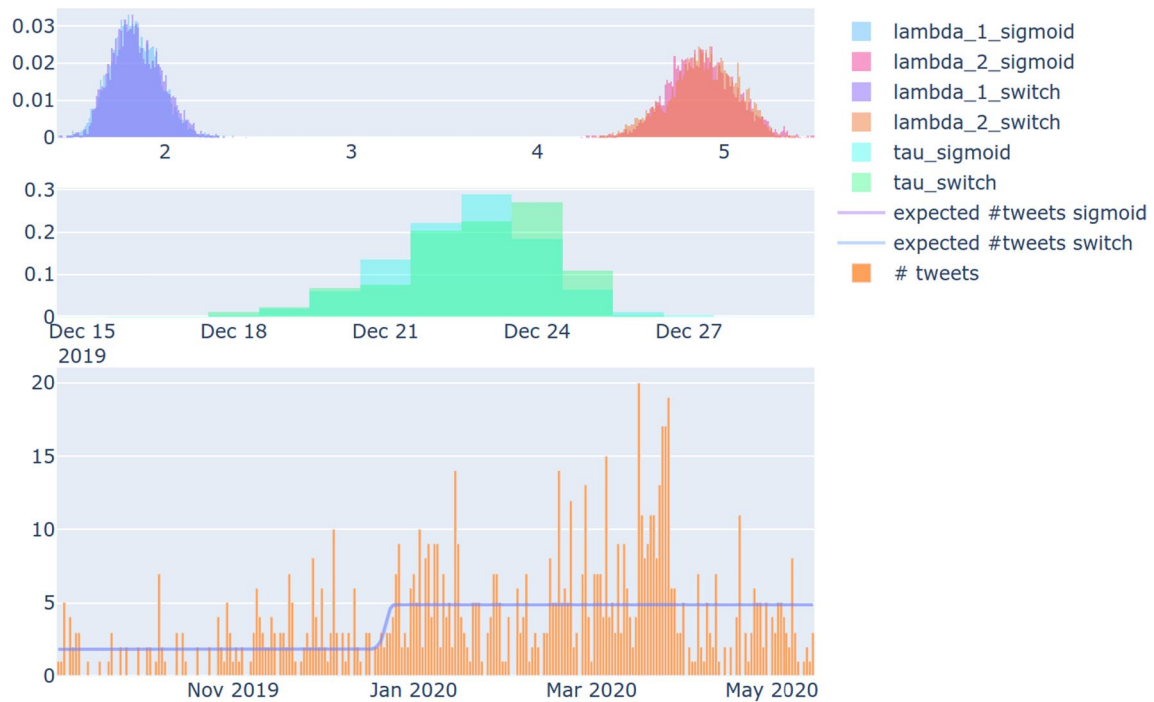


Fig. 11 Comparison between Poissonian and Sigmoidal model

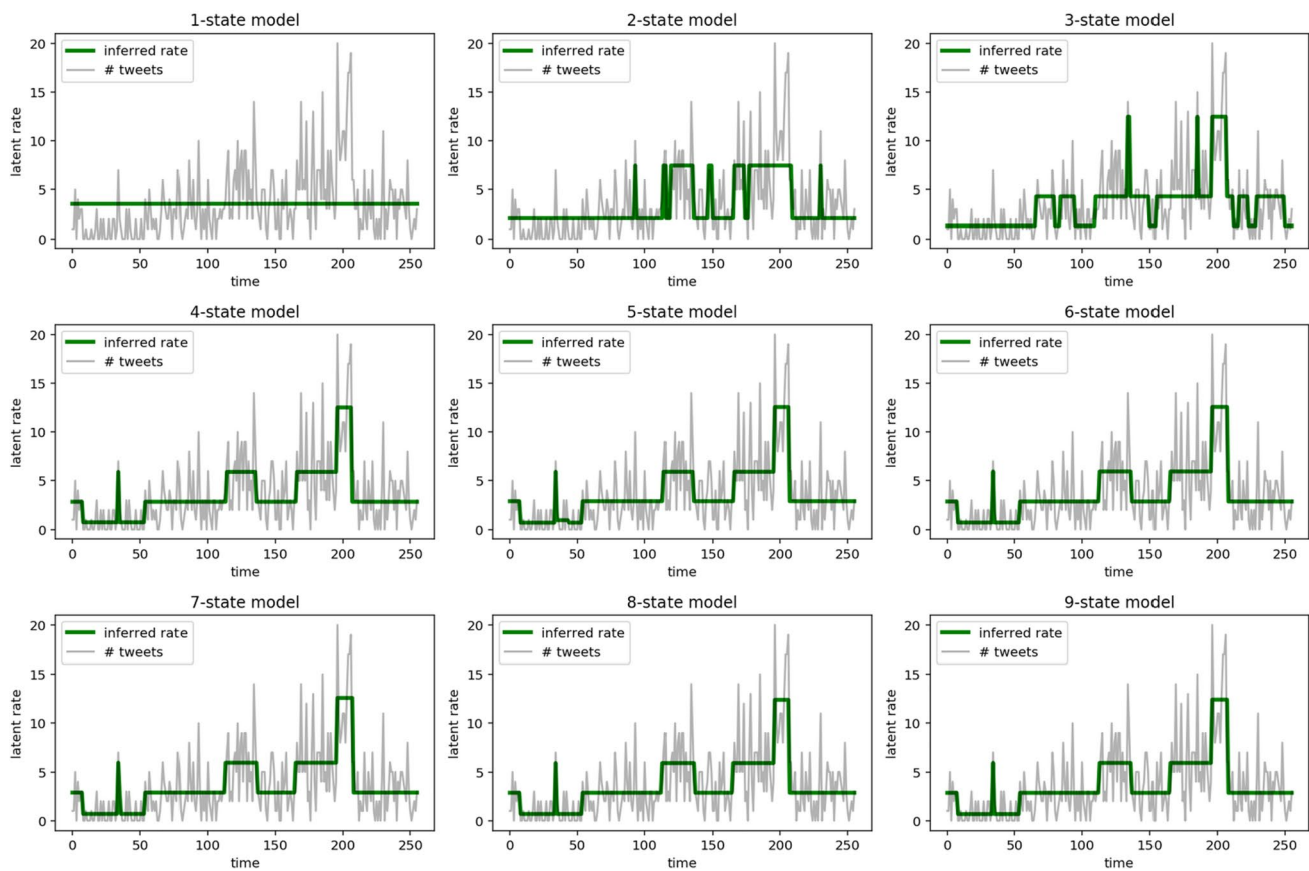


Fig. 12 Multistate model for different states

E.1 Sigmoidal

We analyze here the sigmoidal model. Instead of a sudden switch, we now have a smooth transition between the two λ s. While in this way we can reduce the uncertainty in τ , in general the two models give very similar results (Fig. 11).

E.2 Multistate model

As in the case of the breakpoints model in SSR, we can allow for multiple states or levels represented by more than two λ s. The maximum number of levels being (L_{MAX}).

We will have several realization of the model with different (L_{MAX}) and get the optimal model minimizing the Bayesian Information Criterion.

This model has the advantage of fitting the data more accurately, but the computation time and the simplicity of the result are worse than that of the linear breakpoints model. For example, we can see in Fig. 12 that the 4-state model has a small spike at a given day, this is useful to detect anomalies in the activity, but when we look for changes that persist through time, the linear breakpoints model introduces less noise.

Acknowledgements A.D.-G., E.C., and L.P. acknowledge support from the Spanish grants PGC2018-094754-B-C22 and PID2021-128005NB-C22, funded by MCIN/AEI/10.13039/501100011033 and "ERDF A way o making Europe"; and from Generalitat de Catalunya (2021SGR00856).

Author contributions BE, LP, and EC made substantial contributions to the conception and design of the work; BE worked on the acquisition of data. BE, LP, and EC performed analysis and interpretation of data. BE developed the software used in the work. All authors drafted the work and revised it critically for important intellectual content; all authors approved the version to be published and agree to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature.

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes

were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, Devin M, Ghemawat S, Irving G, Isard M et al (2016) Tensorflow: a system for large-scale machine learning. In: 12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16), p 265–283
- Aiello L, Quercia D, Zhou K, Constantinides M, Šćepanović S, Joglekar S (2020) How epidemic psychology works on social media: evolution of responses to the covid-19 pandemic. arXiv preprint [arXiv:2007.13169](https://arxiv.org/abs/2007.13169)
- Alshaabi T, Minot JR, Arnold MV, Adams JL, Dewhurst DR, Reagan AJ, Muhamad R, Danforth CM, Dodds PS (2020) How the world's collective attention is being paid to a pandemic: COVID-19 related n-gram time series for 24 languages on twitter. *PLoS One* 16(1):e0244476
- Ausserhofer J, Maireder Axel (2013) National politics on Twitter: structures and topics of a networked public sphere. *Inf Commun Soc* 16(3):291–314
- Bai J, Perron Pierre (2003) Computation and analysis of multiple structural change models. *J Appl Economet* 18(1):1–22
- Boon-Itt S, Skunkan Yukolpat et al (2020) Public perception of the COVID-19 pandemic on Twitter: sentiment analysis and topic modeling study. *JMIR Public Health Surveill* 6(4):e21978
- Britt B (2015) Stepwise segmented regression analysis: an iterative statistical algorithm to detect and quantify evolutionary and revolutionary transformations in longitudinal data. In: *Transparency in social media*, Springer, p 125–144
- Crooks A, Croitoru A, Stefanidis A, Radzikowski J (2013) #Earthquake: Twitter as a distributed sensor system. *Trans GIS* 17(1):124–147
- Dewhurst DR, Alshaabi T, Arnold MV, Minot JR, Danforth CM, Dodds PS (2020) Divergent modes of online collective attention to the COVID-19 pandemic are associated with future caseload variance. *CoRR*, [arXiv: 2004.03516](https://arxiv.org/abs/2004.03516)
- Durazzi F, Müller Martin, Salathé M, Remondini D (2021) Clusters of science and health related Twitter users become more isolated during the COVID-19 pandemic. *Sci Rep* 11(1):19655
- Ferraz Costa A, Yamaguchi Y, Juci Machado Traina A, Traina Jr C, Faloutsos C (2015) Rsc: mining and modeling temporal activity in social media. In: *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, p 269–278
- Fischer-Preßler D, Schwemmer Carsten, Fischbach K (2019) Collective sense-making in times of crisis: connecting terror management theory with twitter user reactions to the berlin terrorist attack. *Comput Hum Behav* 100:138–151
- Gallotti R, Valle F, Castaldo N, Sacco P, De Domenico M (2020) Assessing the risks of 'infodemics' in response to COVID-19 epidemics. *Nat Hum Behav* 4(12):1285–1293
- García-Gavilanes R, Kaltenbrunner A, Sáez-Trumper D, Baeza-Yates R, Aragón P, Laniado D (2014) Who are my audiences? A study of the evolution of target audiences in microblogs. Springer, Cham, p 561–572
- Gilani Z, Farahbakhsh R, Tyson G, Crowcroft J (2019) A large-scale behavioural analysis of bots and humans on twitter. *ACM Trans Web* 13(1):1–23
- Heverin T, Zach Lisl (2012) Use of microblogging for collective sense-making during violent crises: a study of three campus shootings. *J Am Soc Inf Sci Technol* 63(1):34–47
- Jungherr A (2014) Twitter in politics: a comprehensive literature review. Available at SSRN 2865150
- Neal RM (2011) Mcmc using hamiltonian dynamics. In: Brooks S, Gelman A, Jones G, Meng XL (eds) *Handbook of Markov chain monte carlo*
- Strong P (1990) Epidemic psychology: a model. *Social Health Illn* 12(3):249–259
- Tsao SF, Chen H, Tisseverasinghe T, Yang Y, Li L, Butt Z (2021) What social media told us in the time of COVID-19: a scoping review. *Lancet Digit Health* 3(3):E175–E194
- Valdez D, Ten Thij M, Bathina K, Rutter LA, Bollen Johan (2020) Social media insights into us mental health during the COVID-19 pandemic: longitudinal analysis of twitter data. *J Med Internet Res* 22(12):e21418
- Venturini T, Jensen Pablo, Latour B (2015) Fill in the gap: a new alliance for social and natural sciences. *J Artif Soc Soc Simul* 18(2):18–29

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.