Get Out of the Nest! The Infodemic of the #TwitterMigration

Keywords: Mastodon, Twitter, migration, social influence, DOSNs

Introduction. After years of steady growth, many popular social media are experiencing shifts in user engagement. In Twitter, such a shift has been especially abrupt after business magnate Elon Musk purchased the platform on October 26^{th} 2022. The acquisition itself, as well as several controversial management decisions taken by Musk shortly after (including massive layoffs, the suspension of some journalists' accounts, and the discontinuation of free API access) threw the platform at the center of a media storm and motivated many Twitter users to seek substitute services to migrate to.

Decentralized Online Social Networks (DOSNs), which have taken off sharply in the last few years [4], have been perceived as a captivating alternative. Among them, Mastodon has emerged as the most popular open-source analogue of Twitter [6]. Mastodon allows communities to manage their own *instances* (i.e., servers) and to connect them with others in a federated fashion thanks to a common protocol — similar to email services. On Twitter, users advocating for a transition to Mastodon promoted the #TwitterMigration movement, which led to the proliferation of new Mastodon instances and to spikes in registration requests. This phenomenon arguably represents one of the largest online social migrations in the history of the Web and, as such, it is worth investigating. Despite being triggered by exogenous circumstances, it is plausible that the migration unfolded organically through mechanisms of social influence, with users exerting pressure on peers by signaling their decision to migrate.

In this context, we study the #TwitterMigration phenomenon for the first time, find that it is compatible with an information contagion process driven by people signaling their commitment to change, and that such a process is best defined at the level of individual social communities on the Twitter follower graph. Our work sheds light on a process of behavioural change whose nature and scale are rarely observed and measured, which will hopefully help furthering our general understanding of collective change.

Methodology. Through Twitter's Academic API, we collected tweets related to the #TwitterMigration posted from October 26^{th} 2022 (date of Musk's takeover) to January 19^{th} 2023. To ensure a good recall, we considered tweets that (i) contain one of the 13 hashtags that we found through a snowball expansion of trending hashtags right after the acquisition, or (ii) mention Mastodon's Twitter account (i.e., @joinmastodon), or (iii) contain the keyword "mastodon". Also, we excluded retweets. We ended up with 1.3M+ English tweets by $\sim$0.5M users.

In the attempt of recreating their existing social network on Mastodon, Twitter users advertised their Mastodon *handles* to their followers. To find the Mastodon profile corresponding to a Twitter user, we used regex-based heuristics to identify Mastodon handles in the username, description, and tweets of the user, thus obtaining 75k unique Mastodon handles which we could match with existing Mastodon accounts. To validate the associations, we searched the retrieved Mastodon user profiles for mentions of Twitter handles and found that, when available, the handles matched the original Twitter accounts in 98% of the cases. As a result, we created a directed social graph where vertices are Twitter users that have a corresponding Mastodon account, and links are Twitter follower relationships. All nodes are annotated with a timestamp of migration, which corresponds to the date of creation of the Mastodon profile.

To verify whether the temporal pattern of migration is compatible with an information diffusion phenomenon, we adopted the hypothesis of *infodemic spreading* [8], and fit two *compartmental epidemiological models* [1] (SIR and SIRS) to our data. These models assume that information spreads like an infection through social connections and they are defined by an infection rate β, a recovery rate γ and, in the case of SIRS, a re-infection rate ε. From this set of parameters, we can extract the basic reproduction number $R_0 = \beta/\gamma$, i.e., a proxy of the typical number of infections generated by an infected individual and an indicator of the speed of contagion. In our data, the set of infected (I) includes Twitter users who, at time t had created a Mastodon account and are tweeting about the #TwitterMigration; the set of susceptible (S) includes users with no Mastodon account at t; and the remaining set of recovered (R) is determined by the models as a function of β. We estimated the parameters of both models using least square estimates [7]. To gain insight into whether the diffusion process had different properties in different regions of the network, we replicated the fit on the largest network communities identified by the *Louvain* community detection method.

Results. The SIR and SIRS models accurately reproduce the cumulative migration curve, with an estimated $R_0 = 4.57$ (Figure 1, left) which, being $\gg 1$, indicates a highly infectious process. The value of R_0 is the same for the SIRS model, which hints at the fact that multiple manifestations of commitment by the same user (i.e., reinfections) have a negligible role. When shifting the focus to the mesoscopic perspective of the Louvain communities, we find that (i) the value of R_0 is heavily community-dependent, ranging from 3.85 to 9.04, (ii) R_0 seems uncorrelated with community size, and (iii) the estimated R_0 can vary considerably between SIR and SIRS in some communities, unlike in the overall network (Figure 1, right). Identifying the community features that correlate with their reproduction number is key to learn possible indicators of successful behavioural change in the context of online interactions. In ongoing work, we aim at verifying orthogonal hypotheses of determinants of beahvioural change, ranging from the committed minority hypothesis [2] to psycho-linguistic intepretations of opinion change [5]. To do so we are characterizing the communities in terms of structural patterns of the network, activity of committed individuals over time, and linguistic markers of social interaction [3].

[1] Norman TJ Bailey et al. *The mathematical theory of infectious diseases and its applications*. Charles Griffin & Company, 1975.

[2] Damon Centola et al. “Experimental evidence for tipping points in social convention”. In: *Science* 360.6393 (2018).

[3] Minje Choi et al. “Ten Social Dimensions of Conversations and Relationships”. In: *The Web Conference 2020*. New York, NY, USA: ACM, 2020.

[4] Anwitaman Datta et al. “Decentralized Online Social Networks”. In: *Handbook of Social Network Technologies and Applications*. Ed. by Borko Furht. Springer, 2010.

[5] Jürgen Habermas. *Communication and the Evolution of Society*. Beacon Press, 1979.

[6] Lucio La Cava, Sergio Greco, and Andrea Tagarelli. “Understanding the growth of the Fediverse through the lens of Mastodon”. In: *Appl. Netw. Sci.* 6.1 (2021).

[7] Junling Ma. “Estimating epidemic exponential growth rate and basic reproduction number”. In: *Infectious Disease Modelling* 5 (2020). ISSN: 2468-0427.

[8] Jonathan Skaza and Brian Blais. “Modeling the infectiousness of Twitter hashtags”. In: *Physica A: Statistical Mechanics and its Applications* 465 (2017).

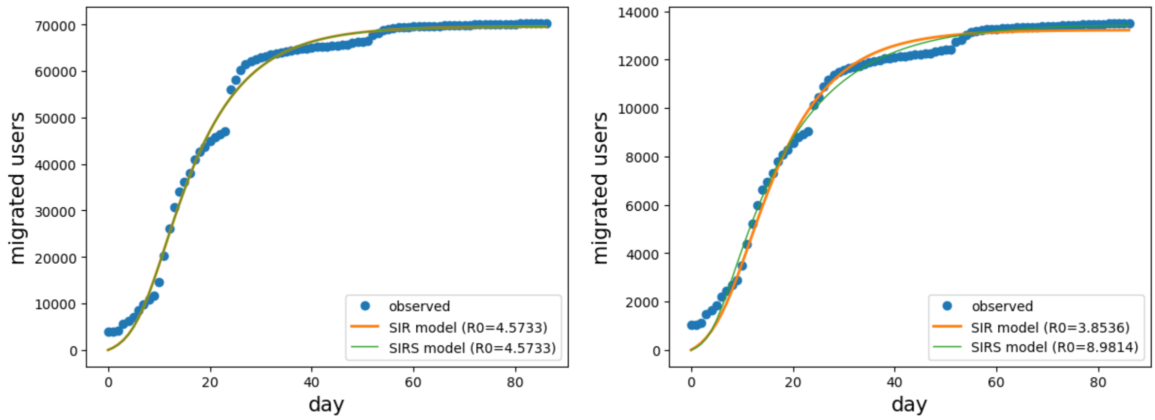

Figure 1: (Left): Cumulative number of Twitter users migrated to Mastodon over the course of 3 months since Elon Musk's acquisition of Twitter. Two fits estimated with compartmental epidemiological models (SIR and SIRS) are shown, along with their reproduction numbers R_0. (Right): The same cumulative curve and fits replicated on one single large Louvain community on the Twitter follower graph.