



# Conductance and influence-capital: modeling online social influence

Rohit Ram<sup>1\*</sup> and Marian-Andrei RizoIU<sup>1,2\*</sup>

Handling Editor: Daniel Romero

\*Correspondence:

[rohit.ram@student.uts.edu.au](mailto:rohit.ram@student.uts.edu.au);  
[marian-andrei.rizoIU@uts.edu.au](mailto:marian-andrei.rizoIU@uts.edu.au)

<sup>1</sup>University of Technology Sydney,  
Sydney, Australia

<sup>2</sup>Data61, Sydney, Australia

## Abstract

Human interactions are mediated by social influence. During crises like the COVID-19 pandemic, social influence can be critical in determining whether life-saving information is adopted, public health measures are observed, or immunization campaigns meet their targets. The literature on online social influence presents notable limitations across disciplines. Psychosocial approaches effectively characterize the nature of influence by measuring how social factors impact these phenomena, but they lack computational modeling capabilities and rely on slow, non-scalable measurement methods that struggle to generalize to current issues. Conversely, computational approaches, while data-driven and based on network and event analysis, often fail to incorporate the critical social factors that underlie these phenomena. Our work bridges this gap through two main contributions that integrate the strengths of both approaches. First, we present a data-driven Generalized Influence Model that incorporates two novel psychosocial-inspired mechanisms: the conductance of the diffusion network and the influence-capital distribution. This model not only outperforms existing state-of-the-art approaches but also corrects the inherent biases introduced by the widely used follower count metric. Second, we empirically test long-held sociological hypotheses regarding influence, social class, and expertise in the online domain by applying our influence model to discussions around COVID-19. We quantify the influence and content veracity for more than 21.5 million X/Twitter users in relation to their professions. Our model suggests that executives, media, and military figures exert greater influence than pandemic-related experts such as life scientists and healthcare professionals. Worryingly, by leveraging existing COVID-19 misinformation datasets, we show that some of the most influential occupations also spread the most misinformation. These findings raise questions about the effectiveness of information dissemination by experts in situations of crisis.

**Keywords:** Social influence; Online social networks; Hawkes processes; COVID-19

## 1 Introduction

Understanding public opinion formation remains an enigma with ample impact on our increasingly digitized society. The polarization of public opinion signals breakdowns in social trust [1], extreme opinions form pathways to radicalization [2], and in recent years, opinions opposing expert medical advice have led to loss of life by undermining immuniza-

© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, 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 you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. 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-nc-nd/4.0/>.

tion campaigns and stoking vaccine hesitancy. While we are beginning to understand the factors that drive opinion formation, we lack scalable modeling and quantification tools. These tools are essential when trying to mobilize our peers to enact change—whether to accept the vaccine or alter their consumption patterns to avert catastrophic climate change. One of the foundational factors in opinion formation and in enacting change is *influence* [3]. This force governs interpersonal relationships and contributes to establishing societal institutions and reforms.

Social media ubiquity has expanded the role of influence by providing a fertile ground for influence mechanisms to unfold and for a minority of users to exert disproportionate control. There exists ample evidence that online dynamics have offline repercussions in terms of collective movement [4] and radicalization [5]. Although there is a rich trove of literature on social influence, stretching from the psychosocial to computational domains, a quantitative approach to measuring online influence remains elusive [6, 7]. The computational approaches focus heavily on measuring message propagation propensities and exposure to messages through one's social network [8–10]; however, social influence is more accurately defined as the ability to change the behaviors, opinions, or beliefs of others [11]. Many psychosocial complexities of influence (like the social attributes of the source of influence) are lost when influence is narrowly interpreted as merely a reaction to exposure. Therefore, we ask **how can online influence be quantified, such that it relates to the theoretical psychosocial influence phenomenon?**

Furthermore, sociologists have long espoused that influence is most directly related to concepts such as social class and expertise [12]. Social class, a multifaceted construct encompassing education, wealth, occupation, and subculture, strongly aligns with established determinants of influence such as authority and likability [13]. While historically linked to political power, in contemporary egalitarian societies with social mobility, class is predominantly indicated by occupation. Expertise, on the other hand, is related to a specific context and typically derives influence through demonstrated outcomes. Despite the compelling appeal of these sociological theories of influence, they remain largely untested on a large scale in online domains. Evaluating these theories is crucial, as adherence to expert advice reflects optimal societal behavior, and the extent to which social class either hinders or reinforces such advice can have critical implications. For instance, in an ideal scenario, populations would follow expert advice (e.g., from epidemiologists, biomedical scientists, healthcare professionals) during pandemics. However, the emergence of the anti-vaxxer movement and the proliferation of vaccine misinformation illustrates a contrasting reality. Therefore, we ask **how does influence correlate with its proposed determinants in the online domain?**

In this paper, we model online influence on Twitter/X. We achieve this in a sequence of two parts. In *the first part*, we build the Generalized Influence Model (GIM) – a novel contribution that bridges the divide between psychology and the quantitative approaches [14]. In Sect. 2, we review an essentially quantitative model [15] for online influence based on exposure. In Sect. 3, we build GIM by augmenting this model with two factors known in psychosocial literature to modulate the exertion of social influence: *conductance* [13, 16] and *social attribution* [13, 17–19]. We endogenize these factors by incorporating a social network conductance and a influence-capital distribution mechanism and propose several flexible implementations that maintain scalability. In Sect. 4, we show GIM to outperform several baselines, including the current state-of-the-art influence quantification [15], and

to correct the biases introduced by the widely used follower count metric. In *the second part*, in Sect. 5, we empirically test the hypothesis that social class and expertise are the primary influence determinants. Following prior work, we assume that social class is primarily associated with one's occupation [20] and expertise is associated with the quality of the sources they share. On a Twitter/X dataset containing discussions about the COVID-19 pandemic, we utilize GIM to identify the occupational groups that yield the highest influence. We determine user occupation using the O\*NET taxonomy [21], and numerically quantify the veracity of information users spread using a dataset of COVID-19 related misinformation [22]. The analysis reveals experts (i.e., epidemiologists, biomedical scientists, and healthcare professionals) are not always the ones that shape the discussions. So who does exert the greatest influence?

Successfully modeling online influence identifies the actors who shape societal views. It serves as the first step toward addressing societal issues—such as spreading misinformation amid a pandemic.

**The main contributions of this work include:**

- The **General Influence Model** which introduces two psychosocial inspired mechanisms: *conductance* and *influence-capital distribution*, to move toward psychosocial influence.
- An **evaluation analysis** of GIM, where we find the optimal hyperparameters, show that GIM outperforms the state-of-the-art, and illustrate how GIM is an unbiased estimator.
- An **analysis of sociological determinants of influence**; investigating the relationship between social class, expertise, and GIM social influence.

To aid the reader, we briefly define the key terms used throughout this paper. *Conductance* quantifies how readily influence flows between two users: analogous to physical materials conducting electricity or heat, social relationships conduct influence to varying degrees, and higher conductance makes influence transmission more likely. *Influence-capital* captures the latent capacity of an individual to exert influence, accumulated through the endorsements they receive in information cascades. The *branching structure* is the latent tree of who-influenced-whom within a retweet cascade. *Ground truth influence* refers to a peer-perceived influence ranking obtained via crowdsourced pairwise comparisons, used to calibrate and evaluate our model.

## 2 Preliminaries & related works

Quantitative influence approaches exploit readily available metadata, such as timing and network features, which has historically been underutilized in psychosocial literature. These approaches, and the axioms from which they are built, form a strong foundation for an inquiry into influence. Influence is traditionally difficult to observe, as we cannot easily observe the minds of the influenced. We build upon prior works that take a longitudinal view and hypothesize that the endorsement patterns over extended periods reveal the influence structures in an observed cohort. On Twitter/X, this endorsement is signaled via retweeting, which is widely accepted as a form of endorsement of an emitter (the person being retweeted) by a receiver (the retweeter) [23]. However, platforms typically do not expose the structure of these endorsements. Nearly all social media platforms have a 'resharing' behavior (with varying degrees of API accessibility). The technique of reconstructing influence structures relies on acquiring the timing of reshare events. Rizoïu

et al. [15] infer these structures by assuming retweets arrive following a Hawkes point process. In this section, we review approaches to influence quantification and articulate the Hawkes-modeled influence [15] from which GIM is built.

**Prior Influence Quantification.** Common approaches to quantitative models of influence include: using centrality on social graphs as an influence proxy [24, 25], approximating solutions to the influence maximization task [8, 9], and explicitly modeling popularity through approaches such as Hawkes Processes [26]. However, these approaches suffer from the conflation of social influence with other concepts such as popularity or activity. Several studies have highlighted the discordance between measures of popularity and influence [27–29]. Recent work has proposed several novel approaches to influence measurement. Nickel and Le [30] efficiently model entity interactions via a Multivariate Hawkes Process (though this approach is entirely based on propagation propensity rather than grounded influence), and Smith et al. [29] estimate the impact by removing confounders through a causal inference network framework (though this approach must effectively reconstruct the network, makes restrictive diffusion assumptions, and focuses on influence operator classification). Unlike these prior methods, our work directly accounts for psychosocial factors that mediate influence propagation: influence-capital and network conductance.

**Retweet Cascades** consist of an original tweet and subsequent retweets. We denote a marked cascade up to time  $T$  as  $\mathcal{H}_\zeta(T) = \{v_1, v_2, \dots\}$ , where  $v_i = (t_i, \zeta_i)$  denotes the  $i$ th tweet,  $t_i$  is the event time relative to the original tweet ( $t_1 = 0$ ), and  $\zeta_i \in \mathbb{R}$ —dubbed the *mark*—is the event meta-data. We denote a possible realization of the endorsement structure as  $\mathcal{G}$ , and the set of all potential endorsement structures as  $\Upsilon$ .

**The Branching Structure** of retweet cascades is a latent graph  $(G, E)$  where  $G$  are the tweets, and  $E$  contains direct retweet relations. Given a cascade of  $n$  events, we define the set of all valid branching structures as  $\Upsilon = \{G | (v_i, v_j) \in G \text{ so that } t_i < t_j\}$ . Rizoiu et al. [15] show that  $|\Upsilon| = (n - 1)!$ . Figure 1 shows an example retweet cascade, its most likely branching structure in solid lines and other potential branching structures in dashed lines. Next, we compute the probability mass function over the branching structure set  $\Upsilon$ .


**Hawkes Processes** [31] have emerged as a powerful tool for modeling social media events and information diffusion, providing a mathematical framework for capturing the self-exciting nature of online user activities [32, 33]. The self-exciting property—the arrival of an event increases the likelihood of future events—is applicable here due to the property of social affirmation (i.e., past social actions encourage incoming actions).

In Hawkes processes, events arrive following the conditional intensity

$$\lambda(t | \mathcal{H}_\zeta(t)) = \mu(t) + \sum_{v_i \in \mathcal{H}_\zeta(t): t_i < t} \zeta_i^b \phi(t - t_i) ,$$

where  $\mu(t)$  is the baseline intensity; each event  $v_i$  increases the overall intensity by  $\zeta_i^b \phi(t - t_i)$ ;  $\zeta_i^b$  is one way to model marks where  $b$  mediates the marks' effect, and the kernel  $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$  controls the event intensity decay. The exponential  $\phi_{exp}(t) = e^{-rt}$  and the power-law  $\phi_{pow}(t) = (t + c)^{-(1+r)}$  are common parametric forms.

Hawkes and Oakes [34] propose the branching representation of Hawkes processes, where each event  $v_i$  generates offspring following a non-homogenous Poisson process of intensity  $\zeta_i^b \phi(t - t_i)$ —illustrated by the magenta areas in Fig. 1. Lewis and Mohler [35] use



the branching representation to estimate the probability that  $v_j$  is a direct offspring of  $v_i$  as

$$p_{ij} = \frac{\zeta_i^b \phi(t_j - t_i)}{\mu(t) + \sum_{t_k < t_j} \zeta_k^b \phi(t_j - t_k)}, t_i < t_j. \tag{1}$$

Intuitively,  $p_{ij}$  is the proportion of intensity that  $v_i$  contributed to the total intensity at time  $t_j$ . Note that for retweet cascades  $\mu(t) = 0$ , and  $p_{ij} = 0$ , when  $t_i > t_j$ . Finally, the probability of a valid branching structure is  $\prod_{(v_i, v_j) \in E(G)} p_{ij}$ .

**Hawkes-modeled Influence.** Rizoiu et al. [15] measure influence as the expected number of offspring of a tweet across all valid branching structures. They formalize influence as  $\varphi(v_i) = \mathbb{E}_{\mathcal{G} \in \Upsilon} \left[ \sum_{t_j > t_i} \mathbb{1}(\mathcal{G}_{i \rightarrow j}) \right]$ , where  $\mathbb{1}(\mathcal{G}_{i \rightarrow j})$  indicates a path between  $v_i$  and  $v_j$  in the branching structure  $\mathcal{G}$ . Using the *independent cascades* assumption [8] (that generating a tweet at  $t_j$  is independent of the diffusion structure up to  $t_j$ ), Rizoiu et al. [15] devise an efficient iterative procedure for computing  $\varphi(v_i)$ . They introduce  $m_{ik}$  the pairwise influence exerted by  $v_i$  on  $v_k$ , either directly when  $v_k$  is a direct offspring of  $v_i$  or indirectly when  $v_k$  lies on the same diffusion path as  $v_i$ . Formally,  $m_{ik} = \sum_{j=i}^{k-1} m_{ij} p_{jk}$ ,  $i < k$ ,  $m_{ik} = 1$  when  $i = k$ , and  $m_{ik} = 0$  when  $i > k$ . Intuitively,  $m_{ik}$  is the sum of the probabilities of all valid paths between  $i$  and  $k$ . Consider an example with three tweets,  $v_1, v_2, v_3$ , the influence  $m_{13}$  is computed as  $m_{13} = m_{11} p_{13} + m_{12} p_{23} = p_{13} + p_{12} p_{23}$  representing all possible paths between  $v_1$  and  $v_3$ . A tweet's influence is the total influence it exerts, i.e.  $\varphi(v_i) = \sum_{k=i}^n m_{ik}$ . The  $\{m_{ik}\}$  matrix is computed in  $n$  matrix multiplications, and the total time complexity is  $O(n^3)$ . Finally, a user's influence is the average influence of all their tweets.

**Towards Psychosocial.** Our work is inspired by improved performance, via theory-based social features, in tasks related to social influence such as the influence prediction task [36, 37] (classifying whether users perform a behavior after exposure), social graph embedding [38, 39], and recommender systems [40]. This suggests that incorporating psychosocial theory into computational models [41] can yield more accurate and nuanced representations of social influence processes. Our work aims to bridge the gap between computational efficiency and psychosocial theory by introducing mechanisms that capture the relational qualities of influence while maintaining computational scalability.

### 3 Methodology

Prior work estimates influence through endorsement timing and accessible user metadata; however, these approaches fail to consider the nature of relationships between people and the structure of these relationships. To address these limitations we propose the Generalized Influence Model (GIM) that quantifies influence based on observed information cascades. Our proposed GIM generalizes the Hawkes-modeled influence (see Sect. 2) by incorporating two mechanisms that model crucial social information, depicted in Fig. 1. These mechanisms – *conductance* ([13, 16], see Sect. 3.1) and *social attribution* ([13, 17–19], see Sect. 3.2) – are factors that psychosocial literature has identified as modulating the exertion of social influence.

We first formally introduce GIM, before detailing its mechanisms. Given a retweet cascade  $\mathcal{H}_t(T) = \{v_1, v_2, \dots\}$ , GIM quantifies the social influence of a tweet  $v_i$  as:

$$\varphi_\gamma(v_i) = \sum_{\mathcal{G} \in \Upsilon} \sum_{t_j > t_i} \underbrace{\mathbb{P}_\gamma(\mathcal{G}_{i \rightarrow j})}_{\text{Conductance}} \underbrace{\Psi(\mathcal{G}_{i \rightarrow j})}_{\text{Capital Distrib.}}, \tag{2}$$

where  $\Psi(\mathcal{G}_{i \rightarrow j})$  is the influence-capital allocated to  $v_i$  along the endorsement pathway from  $v_j$ , and  $\mathbb{P}_\gamma(\mathcal{G}_{i \rightarrow j})$  is the conductance-mediated probability that a path exists between  $v_i$  and  $v_j$  in  $\mathcal{G}$  (defined in Sect. 3.1).  $\Psi$  is defined at the level of edges, preserving the efficiency of the Hawkes-modeled influence computation (see Sect. 3.3). The influence of a user is the average influence of their tweets.

#### 3.1 Conductance

Conductance models how readily influence flows between two users. Just as physical materials conduct electricity or heat with varying efficiency, social relationships conduct influence to varying degrees: a user is more likely to be influenced by someone they are similar to or connected with than by a stranger. The conductance of an edge modulates the probability of influence transmitted along that edge (see Eq. (3)).

Quantitative models often assume that social ties are the primary influence channels, while other conductive channels (like homophily) are underexplored. Cognitive science [42] and social psychology [13, 16, 43] literatures suggest that some relationships are more influential than others. Notably, people in the same community (or who share similarities) are more influential to each other. Conductance assumes that different types of relations between users propagate influence more effectively (e.g., one might be more influenced by close family than by distant work colleagues); this makes some people more likely sources of influence than others. Intuitively, the social system propagates influence similarly to physical materials conducting electricity or heat. The conductance of a social connection

modulates the likelihood of adopting an opinion. The higher the conductance, the more likely the receiver will adopt the opinion of the emitter. The influence conductance encapsulates user relationships; by measuring which links are more conducive to influence, we can infer more realistic endorsement pathways.

The conductance mechanism modulates the likelihood of cascade pathways using user lexical and following similarity, and relationship ties. We denote as  $\gamma_{i,j}$  the conductance of an edge  $(v_i, v_j)$ , and define the updated probability that  $v_j$  is a direct offspring of  $v_i$  as

$$p'_{ij} = \frac{\zeta_i^b \phi(t_j - t_i) \gamma_{i,j}}{\mu(t) + \sum_{t_k < t_j} \zeta_i^b \phi(t_j - t_k) \gamma_{k,j}}, t_i < t_j. \tag{3}$$

We consider two choices for conductance: *topological* (users' social network) and *homophilic* (users' similarity with others). We further operationalize homophilic conductance, using two lenses: *following* and *lexical*.

**Topological conductance** assumes influence flows between users who are connected in the social graph (i.e., the follower relationship). Each valid edge  $(v_i, v_j)$ , with  $t_i < t_j$ , has a baseline conductance  $\beta_{top}$ , regardless of whether  $u_i$  is connected to  $u_j$ , accounting for alternative influence conduits (such as news feeds or users following topics and hashtags). Formally, we define the topological conductance  $\gamma_{ij}^{top} \in [0, 1]$  as  $\gamma_{ij}^{top} := \beta_{top} + (1 - \beta_{top})a_{ij}$ , where  $a_{ij} = 1$  if  $u_j$  follows  $u_i$  and 0 otherwise.

**Homophilic conductance** of an edge, models the connection between similarity and influence, i.e., people similar to us influence us more [44, 45]. For each user  $u_i$  we first build a user representation  $h_i \in \mathbb{R}^n$ . Next, we quantify  $\gamma_{ij}^{hom}$  the homophilic conductance between two users using the cosine similarity between their user representations plus the baseline conductance  $\beta_{hom}$ . Formally,  $\gamma_{ij}^{hom} := \beta_{hom} + (1 - \beta_{hom})\text{cosine}(h_i, h_j)$ .

The *following lens* leverages the observation that similar people consume similar content. On social media, following popular users is akin to consuming content. Accordingly, we measure the similarity between two users based on whether they follow the same people. For the following lens, we identify the 1000 most followed users, and collect the followers of all the users in our dataset (i.e., users they follow). We represent a user  $u_i$  as  $h_i \in \mathbb{R}^{1000}$ , where  $h_i[j] = 1$  if  $u_i$  follows the  $j$ th most followed user ( $h_i[j] = 0$  otherwise).

The *lexical lens* exploits the insight that similar people use similar language. The vocabulary and language style of users can be a strong indicator of their community, and we measure the similarity of users based on their choice of language. For the lexical lens, we construct user documents by concatenating the user tweets; we represent them using TF-IDF (Term Frequency-Inverse Document Frequency); and a feature hashing dimensionality reduction technique [46]. Finally, we represent each user  $u_i$  as  $h_i \in \mathbb{R}^{1,048,576}$ . Note that the homophilic conductance with the lexical lens does not require knowledge of the user following graph, which can be prohibitively expensive to obtain for Twitter. Related to our work, Centola [43] characterize the relationship between homophilic ties and the influence phenomena, however their network is artificial and they illustrate only one homophilic lens. We note that other attributes—such as partisanship, gender, or race—may serve as stronger indicators of homophily than the lexical and following lenses used here. We opt for lenses that are readily computable at scale without requiring sensitive demographic data, and leave the incorporation of richer sociological features into the conductance mechanism to future work.

### 3.2 Influence-capital distribution

The endorsement pattern is not a sufficient explanation of the ability of individuals to exert influence. Sociologists have long pointed to the importance of weak-ties [47] and the accumulation of social capital by *bonders* [48]. Consider the example in Fig. 1. Alice influences Bobbie, who influences several people; should Alice (the initiator) or Bobbie (the connector) be considered more influential?

Psychologists have recognized that particular characteristics in individuals are correlated with influence, namely authority [17], attractiveness [18], likeability and others [13]. We define *influence-capital* as the congruence of characteristics that enable the exertion of influence. We propose that users possess latent influence-capital, such that an emitter with higher influence-capital exerts more influence. We posit that a distribution mechanism that explains endorsement pathways, must measure this influence-capital. Concretely, we model influence-capital as a quantity that accumulates through the endorsements a user receives: when a user is retweeted, a portion of the retweeter’s capital flows upstream to them. Users who trigger large or far-reaching cascades accumulate more capital, and thus are estimated to be more influential. The question remains, how do we allocate influence-capital to explain the exerted influence along an inferred endorsement pathway?

We propose a *influence-capital distribution* that leads to the accumulation of influence-capital and offers a post-hoc explanation for the inferred endorsement patterns. The mechanism transfers a proportion of a node’s (tweet’s) influence-capital to its parent in the information cascade, allowing upstream and highly connected nodes to accumulate influence-capital (which translates to influence). The distribution mechanism that we propose here aims to explain the endorsement pathways inferred from the data, and it is related to the concept of value-allocation schemes [19, 49]. Several studies [19, 50, 51] have utilized allocation schemes to recognize the role of *bonders* by assuming benefits, generated by a node, decay with distance within a social graph. We migrate this intuition to a directed diffusion scenario. Intuitively, the allocation rewards users responsible for bonding (i.e. connectors), and those who reach distant communities (i.e. initiators).

We construct the influence-capital distribution as follows. Whenever the user  $u_i$  directly influences  $u_j$  (i.e.,  $v_j$  is a direct offspring of  $v_i$  in an endorsement pathway),  $u_j$  transfers a portion of their influence-capital to  $u_i$  (denoted as  $\pi_{ij}$ ). Each tweet is endowed with 1 influence-capital for participation; it pays a proportion  $\alpha \in (0, 1)$  of all its capital to its parent (if they exist) and keeps  $(1 - \alpha)$ . Formally:

$$\pi_{ij} = \begin{cases} \alpha, & 1 \leq i < j, \quad j \text{ transfers } \alpha\% \text{ capital to } i, \\ (1 - \alpha), & i = j \neq 1, \quad j \text{ keeps } (1 - \alpha)\% \text{ capital,} \\ 1, & i = j = 1, \quad \text{The initiator does not transfer.} \end{cases}$$

Figure 1 illustrates the mechanism. Charlie passes Bobbie  $\alpha$  of her endowed capital (keeping  $1 - \alpha$ ), and Bobbie receives capital from her endowment (1), her three direct influencees ( $3\alpha$ ), and one indirect influencee ( $\alpha^2$ ). She keeps  $(1 - \alpha)$  of this sum. The capital distribution of a path  $\mathcal{G}_{i \rightarrow j}$  is  $\Psi(\mathcal{G}_{i \rightarrow j}) = \prod_{(k,l) \in E(\mathcal{G}_{i \rightarrow j})} \pi_{kl}$ , and a user’s social influence is proportional to the total influence-capital they accumulate via the capital distribution mechanism. The scheme naturally conserves total value, which is equal to the number of participants.

### 3.3 Iterative computation


GIM can be computed efficiently by extending the pairwise influence  $m_{ik}$  (introduced in Sect. 2) to incorporate the concepts of conductance and influence-capital distribution. Formally,  $m_{ik} = \sum_{j=i}^{k-1} m_{ij} p'_{jk} \pi_{jk}$ ,  $i < k$ , where  $m_{ik} = p'_{ik} \pi_{ik}$  when  $i = k$ , and  $m_{ik} = 0$  when  $i > k$ . Consequently (and similar to the Hawkes-modeled influence), we obtain  $\varphi_{\gamma}(v_i) = \sum_{k=i}^n m_{ik}$ . The full derivation of the latter, from Eq. (2) via Eq. (3) and Sect. 3.3, is shown in the online appendix [52]. GIM recursively generates all possible paths in the same way as the Hawkes-modeled Influence [15], which allows an efficient iterative algorithm of temporal complexity  $O(n^3)$ . Visibly, the Hawkes-modeled influence [15] is a special case of GIM, with  $\pi_{jk} = \gamma_{jk} = 1, \forall j, k$ .

## 4 GIM evaluation

Having developed the theoretical foundation for our Generalized Influence Model, we now turn to its empirical evaluation. This section addresses three key questions: (1) what combination of conductance and distribution mechanisms yields optimal performance? (2) how does GIM compare to existing influence estimation methods? and, (3) can GIM overcome the known biases of traditional influence metrics?

**Ground Truth Influence.** Social influence is inherently difficult to quantify, particularly in online settings where direct behavioral change is challenging to observe. By *ground truth influence* we mean empirically measured social influence recovered from crowdsourced pairwise comparisons of Twitter/X profiles. We utilize our previously developed empirical influence ground truth [53] which employs a human-in-the-loop active learning approach to measure peer-perceived influence. The method leverages crowdsourcing workers who perform pairwise comparisons between users, determining which is more influential. For each comparison, workers are shown the two users' profiles, recent tweets, and links to their public pages, and are asked to judge which user is more influential. The limitations of this approach are explored in detail in Ram and RizoIU [53]. To optimize the required number of comparisons, the method implements a Quicksort-based active learning algorithm that selects the most informative pairs to compare, reducing the quadratic complexity of naive approaches to a loglinear one. The comparisons are then used to fit a Bradley-Terry model, which produces a complete ranking of social influence scores. By constructing rankings from pairwise comparisons, this approach avoids the ordinality distortions that arise when psychometric constructs—which lack standardized units of measurement—are coded as numerical scores [54]. The resulting dataset consists of influence rankings for 500 Twitter users selected from the #ArsonEmergency dataset [55]—containing Twitter discussions about the Australian bushfires collected between November 2019 and January 2020. This empirical influence measure has been validated by showing strong correlation with the Big Two of social cognition (agency and communion), theoretical determinants of social influence. We utilize this empirical influence ranking to calibrate and evaluate GIM.

**Evaluation metrics.** We evaluate GIM ranking against the ground truth (described above) using two measures: NDCG-AUC (see next) and MAPE. The information retrieval literature uses the *Normalised Discounted Cumulative Gain* (NDCG) to measure the overlap between two rankings. It privileges the correct ranking of the top-ranked positions and discounts errors in the lower rankings. Applied to influence, NDCG@k aims to order the  $k$  most influential users correctly. We compute the Area Under Curve for NDCG@k

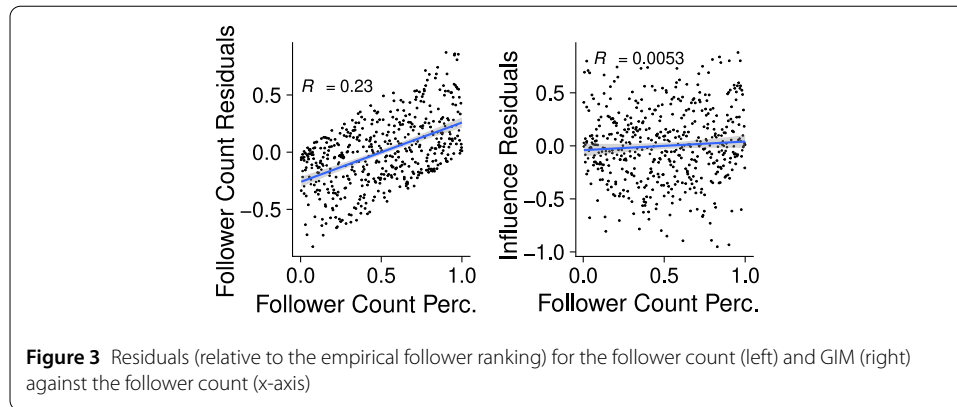


(NDCG-AUC) by varying  $k$  and producing a single metric value. We also compute the Mean Absolute Percentage Error (MAPE) of the difference in the ranking percentiles for each target. Having established our evaluation methodology, we next introduce the baseline methods against which we compare GIM.

**Influence ranking baselines.** We compare GIM to four baselines. Two baselines are widely used heuristics; *PageRank* [56, 57] assumes influence flows via random-walks on constructed social graphs (here the follower network) and *retweet influence* [58], counts the retweets of a user’s authored tweets. They are centrality- and feature-based approaches, respectively. The other two baselines are purpose-built state-of-the-art influence estimators: *Hawkes-modeled influence* [15] and *ProfileRank* [25] (a PageRank variant). Next, we conduct a comprehensive search to identify the optimal configuration of our proposed model.

**GIM search.** For each combination of conductance (topological, lexical, and following) and distribution mechanism (influence-capital, and none), we perform a grid search over the hyper-parameters  $\beta$  (conductance) and  $\alpha$  (distribution). At each grid point, we compare the influence scores obtained by GIM against the ground truth empirical influence ranking. Figure 2 shows the baselines and GIM with several conductance-distribution combinations in the space of the performance measures: negative MAPE (x-axis) and NDCG-AUC (y-axis) (the top-right corner optimizes both measures).

Analysis of the experimental results reveals three key observations. First, GIM consistently Pareto-dominates (i.e., outperforms) all baselines for almost every hyperparameter combination, showing that our psychosocial-inspired mechanisms render automatic influence quantification closer to the human judgment. Among baselines, the next best performing is the Hawkes-modeled influence, followed by PageRank, retweet influence, and the purpose-built ProfileRank. Second, we observe that the homophilic conductances



**Figure 3** Residuals (relative to the empirical follower ranking) for the follower count (left) and GIM (right) against the follower count (x-axis)

(i.e., lexical and following) typically outperforms the topological conductance, and between homophilic conductances lexical outperforms the following conductance. Third, only two models are not Pareto-dominated: the topological-none (best NDCG-AUC) and the lexical-influence-capital (best neg. MAPE). Notice, however, that the topological conductance requires recovering the follower network. In practice, this is prohibitive for large datasets (such as the #COVID-19 dataset) due to rate limitations of the Twitter/X API. Based on these findings, in our analysis in Sect. 5 we use the homophilic lexical conductance ( $\beta = 0.18$ ) with the influence-capital distribution ( $\alpha = 0.02$ ). The 18% baseline conductance indicates that the accounted channels do not explain a relatively large proportion of conductance. Note that while passing 2% of the influence-capital to the parent might not seem much, this adds up for nodes with high degrees (particularly given the long-tail distribution of follower count [59]). Having established GIM's superior performance against existing models, we now examine its ability to address a key limitation of traditional influence metrics.

**Debiasing the follower count.** The follower count is widely used as a proxy for influence [58, 60, 61]; however, it has been repeatedly shown to be biased [27–29]. We examine whether GIM can overcome this known limitation. Figure 3(left) shows that the follower count residuals—the difference between the follower count percentile and the empirical score percentile—are positively correlated with the follower count percentile ( $R^2 = 0.48$ ). In other words, the follower count overestimates the influence of the highly followed users and underestimates the lowly followed users. In contrast, Fig. 3(right) shows that GIM residuals are not correlated with the follower count percentile ( $R^2 = 0.073$ ). That is, GIM is an unbiased influence estimator with respect to the follower count. Having demonstrated GIM's effectiveness at quantifying influence in an unbiased manner, we next proceed to apply it to real-world data to investigate the relationship between social class, expertise, and influence in online discourse.

## 5 Social class, expertise, and influence during COVID

Having established GIM as an effective influence quantification model that outperforms state-of-the-art approaches and corrects inherent biases, we now apply this tool to examine fundamental sociological hypotheses. Sociologists have long espoused that power and influence are primarily related to social class and expertise [12]. The influence of experts, and trust in their advice, is important for the optimal functioning of society. Failure to adhere to advice can lead to unnecessary death (e.g., vaccine hesitancy) and existen-

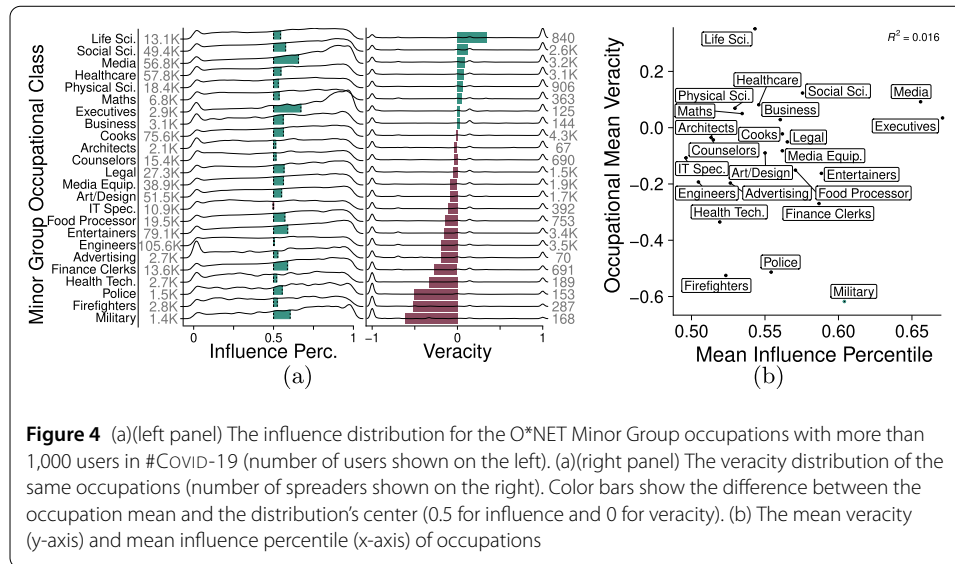
tial threats (e.g., climate change). Social class is often nebulously defined but is related to wealth, occupation, and subculture. In our increasingly egalitarian societies, occupation provides a suitable proxy. We choose occupation as our primary lens not because it is the strongest correlate of homophily—attributes such as partisanship, gender, or race may exhibit stronger homophilic ties—but because of its theorised relationship to influence through authority, prestige, and expertise [12, 13]. Moreover, occupation is among the most accessible and verifiable attributes at scale, as users frequently self-report it in their social media profiles. The COVID-19 pandemic presents an ideal context for investigating these relationships, as it constitutes a crisis where expert communication should be paramount for public safety. Ideally, influence, expertise, and social class are aligned; however, this does not seem to have been the case as Bailo et al. [62] observed that far-right accounts, which had been peripheral during previous crisis events, managed to assume more central positions in online COVID-19 discussions.

In this section, we leverage our validated GIM framework to empirically examine the interplay between social influence, expertise, and occupational class during the COVID-19 pandemic. We first introduce the COVID dataset and use GIM to compute the influence of all users in the #COVID-19 dataset. Next, we develop methods to quantify two key variables: users' occupations as proxies for social class, and the veracity of information they spread as an indicator of expertise. We extract users' occupations and the veracity of the information they spread, and we tabulate their influence and veracity against their occupation. Finally, we analyze the relationships between these variables to assess whether influence patterns align with theoretical expectations about expertise and social class.

**Dataset.** The dataset was collected from Twitter/X in the context of the COVID-19 pandemic. The #COVID-19 dataset was constructed using the keyword *covid19* during August 2020 and contains 143,356,591 tweets by 21,527,913 users. Note, this dataset is different from the #ArsonEmergency dataset, we used to tune GIM.

**Determine the occupations of Twitter/X users.** To establish occupational classifications as proxies for social class, we implement a systematic approach for identifying users' professions [63]. We match user occupation against the Minor Group Occupational Classes of the O\*NET occupational taxonomy [21] using textual fuzzy matching [64]. O\*NET described users better than other taxonomies investigated. We search each user's Twitter/X description and select the first matched occupation (following Sloan et al. [20])—assuming people list their actual occupation first, before hobbies and other information. We validated the classifier on 100 labeled users, where ground-truth labels were derived by two annotators, with disagreement resolved by discussion. The classifier has a mean macro-F1 of 0.54, comparable to the classification performance reported by literature [65].

**Quantifying Expertise.** We assume that expertise is correlated with the veracity (i.e., quality) of the content users share. To operationalize expertise in the online domain, we develop a method for measuring the quality of information shared by users. We use the links and tweets that users share to quantify the veracity of the information they spread. We follow prior research [66] and compute a veracity score for the domain names of the URLs. First, we extract from the CoAID dataset [22] all the links and tweets associated with *true* and *fake* information related to COVID-19. The dataset contains full URLs, and some URL domains appear many times. For each URL domain within CoAID, we count the true (#*R*) and fake (#*F*) entries recorded in the dataset. Finally, we generate for each



**Figure 4** (a)(left panel) The influence distribution for the O\*NET Minor Group occupations with more than 1,000 users in #COVID-19 (number of users shown on the left). (a)(right panel) The veracity distribution of the same occupations (number of spreaders shown on the right). Color bars show the difference between the occupation mean and the distribution's center (0.5 for influence and 0 for veracity). (b) The mean veracity (y-axis) and mean influence percentile (x-axis) of occupations

domain a normalized score as  $\frac{\#R-\#F}{\#R+\#F}$ . The score is between  $-1$  (domain is fully unreliable and spreads misinformation) and  $1$  (fully reliable). Furthermore and similar to [66], we curate a set of *high-quality health sources* (HQHS) from prominent health websites (e.g., CDC, WHO, and Mayo Clinic), and medical journal websites (e.g., The Lancet, and Nature). We compute the veracity score of a post with a link as follows: (1)  $-1$  or  $1$  if the full link appears in CoAID as fake or true, respectively; (2)  $1$  if the link domain appears in the HQHS set; (3) the domain's veracity  $\in [-1, 1]$ , if neither (1) nor (2) applies. A user's veracity score is the mean veracity of the posts they share.

**The influence and veracity of occupations.** Having quantified influence, occupation, and information veracity, we now analyze how these variables interact across different professional groups during the pandemic. Figure 4a shows the distribution of user influence and user veracity scores, for occupations with more than a thousand users in the #COVID-19 dataset. Our model estimates that *Executives*, the *Media*, *Entertainers* and the *Military* are among the most influential occupations. This result is hardly surprising for the former three. Notably, *Media* and *Entertainers* are not only influential but have prominent online presences, perhaps on account of their attention-related business models. The latter (*Military*) is an occupation with high bipartisan support and respect in the US – home of most English-speaking Twitter/X users. Examining the veracity distributions reveals important patterns in information quality across occupational groups. For all occupations, Fig. 4a shows that the veracity score has a bimodal distribution. One mode is around  $-1$  (the users who spread mainly misinformation) and another at  $1$  made of users who spread high-quality information. Users typically do not engage with both types of information, probably due to homophily and online polarization. As all occupational subpopulations contain both types of users, the mean occupation veracity (colored horizontal bars in the right panel of Fig. 4a) represents the ratio between misinformation and high-quality information spreaders within an occupation. Contrary to theoretical expectations about expertise and information quality, our analysis suggests that several occupations traditionally associated with authority spread substantial misinformation. Surprisingly, we observe that a significant proportion of *Military*, *Firefighters* and *Police* users spread misinformation. Closer investigation shows they frequently share from controversial publishers, such

as *foxnews.com*, *zerohedge.com*, and *breitbart.com*. To better understand the relationship between influence and information quality, we examine their correlation across occupational groups. Figure 4b shows the scatterplot of the occupations, in the space of mean influence (x-axis) and mean veracity (y-axis). First, we observe that the two quantities are uncorrelated ( $R^2 = 0.016$ ). This suggests a tension with theoretical expectations that expertise (measured as information quality) would align with influence. We also see that *Life Scientists* (including epidemiologists) and *Social Scientists* typically spread high-quality information; however, our model estimates limited influence for these groups. *Healthcare* workers are moderately influential and have a good proportion of high-quality information spreaders; however, *Healthcare Technicians* including paramedics and medical technologists, have a significant proportion of misinformation spreaders. We note that these findings reflect users active in the tracked COVID-19 discussions on Twitter/X, and that keyword-based data collection may bias the observed user sample.

## 6 Discussion

Existing methods for online influence measurement rely on heuristic definitions and fail to model complex phenomena. In this work, we have developed an online social influence framework. Our conductance mechanism departs from traditional approaches in the literature which frequently considers the social network (i.e. the topological lens) as the primary channel of influence [8, 9]. As online interactions are increasingly mediated by recommender systems, web search, and cross-platform engagement, modeling the similarity between users (i.e. the homophilic lens) provides a simple and effective method for modeling influence channels. Our findings demonstrate that lexical features are readily attainable and perform influence modeling well (see Fig. 2). The conductance mechanism we have introduced extends beyond the current application and could be used to mediate diffusion models or as a more general conception of the social graph, for social tasks (i.e. recommendation) or social network analysis.

A foundational challenge in network diffusion research is disentangling peer influence from homophily: observed behavioural contagion may reflect genuine peer-to-peer influence or may arise because similar individuals independently adopt similar behaviours [67]. Prior work, such as Aral et al. [67], treats influence and homophily as independent, competing causes and decomposes the observed effect additively. GIM takes a fundamentally different position: rather than treating homophily as a rival explanation to be separated from influence, the conductance mechanism models homophily as the *channel through which* influence propagates. Higher similarity between users increases conductance, making influence transmission more likely along homophilic ties. We acknowledge that fully disentangling causal peer influence from homophily-driven correlation requires experimental or quasi-experimental designs beyond the scope of this observational study.

A growing body of evidence suggests that the spread of controversial and political content on social media follows complex contagion dynamics, where adoption requires reinforcement from multiple exposures rather than a single contact [68–71]. While GIM employs the independent cascades assumption [8] for computational tractability, the underlying Hawkes process captures a form of social reinforcement consistent with complex contagion: the conditional intensity at any time is the sum of contributions from *all* prior events in the cascade, so the likelihood of a new event reflects cumulative exposure rather than a single contact (see Sect. A). However, GIM does not model the non-linear threshold

effects that are also associated with complex contagion, such as wide bridges and complex paths [72]. Incorporating such non-linear effects would significantly impact the scalability of the model, and we leave this extension to future work.

The influence-capital distribution mechanism in our GIM framework provides an intuitive approach for allocating influence-capital to explain influence (as perceived by peers). This approach significantly improves upon the arbitrary allocations of prior influence models and integrates naturally with the scalable procedure for estimating influence over the stochastic diffusion graph.

The literature constructed the theoretical relationship between social influence and social class, with occupation serving as a key indicator of social class via educational attainment, income, and prestige. Through our analysis, we have estimated occupational influence in discussions around the COVID-19 pandemic (see Fig. 4b). Our results reveal that certain aspects of social class can explain our results; for example, the influence of *Executives* might be explained by their income and occupational prestige. Interestingly, our model suggests that individuals with high educational attainment are not influential in this context. Additionally, our findings suggest a significant determinant of influence, with both *Media* and *Entertainers*, having high influence. More broadly, the modeling of social influence can have broader implications for understanding how extreme opinions infiltrate mainstream discussions [73].

We acknowledge several limitations of our occupational analysis. While sociological theory posits a strong link between social class and influence, influence may also be related to other sociographic attributes beyond occupation. We chose occupation for its theorised relationship to influence determinants (authority, prestige, expertise) and its accessibility and verifiability at scale. Further empirical studies examining influence through additional sociographic lenses are warranted and left to future work. GIM is designed to quantify the perceived influence of individual users; the occupational analysis aggregates these individual estimates into cohort-level averages and is therefore a downstream application. The core model is rigorously evaluated against a human-derived ground truth, compared to four baselines across multiple mechanism combinations via grid search, and shown to debias the follower count metric (see Sect. 4), characterising its expected performance on new data. The occupational conclusions additionally depend on the observed user sample, which is shaped by keyword-based data collection and the heavy-tailed distribution of influence. Future work could further strengthen these findings by empirically measuring the influence of occupational groups and comparing them against the model estimates.

In conclusion, we have presented an online influence model that could be used to; provide insights into downstream modeling such as opinion dynamics, understand the spread of misinformation, or identify influential individuals in political/social campaigns. Our model suggests that *Media*, *Executives*, *Entertainers*, and the *Military* are among the most influential occupations; however the latter contain a large proportion of misinformation spreaders. In contrast, pandemic experts (i.e. *Life Scientists*) are estimated to have limited influence. These findings raise questions about the effectiveness of information dissemination by experts in situations of crisis.

## Appendix A: Complete derivation of GIM

This section shows the complete derivation from Eq. (2) to Sect. 3.3, both shown in the main text.

We define the influence of a tweet given the diffusion scenario  $\mathcal{G}$ , as:

$$\varphi(v_i|\mathcal{G}) = \sum_{v_j \in V(\mathcal{G})} \theta(i, j|\mathcal{G}) \tag{A.1}$$

where  $V(\mathcal{G})$  is the set of nodes in  $\mathcal{G}$ , and we denote as  $\theta(i, j|\mathcal{G}) := \Psi(\mathcal{G}_{i \rightarrow j})$  the influence-capital distribution along a path. We start from the definition of GIM given a retweet cascade (main text Eq. (2)):

$$\begin{aligned} \varphi_\gamma(v_i) &= \sum_{\mathcal{G} \in \Upsilon} \sum_{t_j > t_i} \underbrace{\mathbb{P}_\gamma(\mathcal{G}_{i \rightarrow j})}_{\text{Conductance}} \underbrace{\Psi(\mathcal{G}_{i \rightarrow j})}_{\text{Capital Distrib.}} \\ &= \sum_{\mathcal{G} \in \Upsilon} \mathbb{P}_\gamma(\mathcal{G}) \sum_{v_j \in V(\mathcal{G})} \theta(i, j|\mathcal{G}) = \sum_{\mathcal{G} \in \Upsilon} \mathbb{P}_\gamma(\mathcal{G}) \varphi(v_i|\mathcal{G}) \end{aligned} \tag{A.2}$$

Notably, due to the factorial number of diffusion scenarios in  $\Upsilon$ , computing the influence for each graph is intractable.

**Incremental construction of diffusion scenarios.** We leverage the *independent cascades assumption* (see Sect. 2) to construct an efficient influence computation that overcomes intractability. The key observation is that each tweet  $v_k$  is added simultaneously at time  $t_k$  to all diffusion scenarios constructed at time  $t_{k-1}$ .  $v_k$  contributes only once to the tweet influence of every tweet found on the path to which  $v_k$  is attached. The tweet influence is computed incrementally by updating  $\varphi(v_i), i < k$  at each time  $t_k$ . We denote by  $\varphi^k(v_i)$  the value of tweet influence of  $v_i$  after adding node  $v_k$ . As a result, we only track how the tweet influence increases over time steps, and we do not construct all valid diffusion scenarios.

GIM assumes that a user’s tweet is influenced by one of the precedent tweets, chosen stochastically from a discrete distribution over the valid edges. Alternatively, we can interpret that all previous tweets influence the new tweet proportionally to the same discrete distribution (a view inline with recent findings about influence and complex contagion). Let  $\Upsilon_{1:k-1}$  be the set of all possible diffusion scenarios at time  $t_{k-1}$ , and  $\mathcal{G}^- \in \Upsilon_{1:k-1}$  be one such diffusion scenario, with the set of nodes  $V^- = \{v_1, v_2, \dots, v_{k-1}\}$ . When  $v_k$  arrives, it can attach to any node in  $V^-$ , generating  $k - 1$  new diffusion scenarios  $\mathcal{G}_j^+$ , with  $V_j^+ = V^- \cup v_k$  and  $E_j^+ = E^- \cup (v_j, v_k)$ . We can write the set of scenarios at time  $t_k$  as:

$$\Upsilon_{1:k} = \left\{ \mathcal{G}_j^+ \mid \forall j < k, \forall \mathcal{G}^- \in \Upsilon_{1:k-1} \right\} \tag{A.3}$$

We write the tweet influence of  $v_i$  at time  $k$  as:

$$\varphi^k(v_i) = \sum_{\mathcal{G}^+ \in \Upsilon_{1:k}} \mathbb{P}_\gamma(\mathcal{G}^+) \varphi(v_i|\mathcal{G}^+) \stackrel{cf. (A.3)}{=} \sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \sum_{j=1}^{k-1} \mathbb{P}_\gamma(\mathcal{G}_j^+) \varphi(v_i|\mathcal{G}_j^+) \tag{A.4}$$

**Attach a new node  $v_k$ .** We concentrate on the right-most factor in Eq. (A.4) – the tweet influence in scenario  $\mathcal{G}_j^+$ . We observe that the terms in Eq. (A.1) divide into two: the paths from  $v_i$  to all other nodes except  $v_k$  (i.e. the old nodes) and the path from  $v_i$  to  $v_k$ . We

obtain:

$$\varphi(v_i|\mathcal{G}_j^+) = \sum_{\substack{v_l \in \mathcal{G}_j^+ \\ l > i, l \neq k}} \theta(i, l|\mathcal{G}_j^+) + \theta(i, k|\mathcal{G}_j^+)$$

Note that a path that does not involve  $v_k$  has the same influence-capital contribution in  $\mathcal{G}_j^+$  and in its parent scenario  $\mathcal{G}^-$ , i.e.  $\theta(i, l|\mathcal{G}_j^+) = \theta(i, l|\mathcal{G}^-)$ , for  $l > i$ , and  $l \neq k$ . We obtain

$$\varphi(v_i|\mathcal{G}_j^+) = \sum_{\substack{v_l \in \mathcal{G}^- \\ l > i}} \theta(i, l|\mathcal{G}^-) + \theta(i, k|\mathcal{G}_j^+) \stackrel{cf. (A.1)}{=} \varphi(v_i|\mathcal{G}^-) + \theta(i, k|\mathcal{G}_j^+) \tag{A.5}$$

Combining Eq. (A.4) and (A.5), we obtain:

$$\begin{aligned} \varphi^k(v_i) &= \sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \sum_{j=1}^{k-1} \mathbb{P}_\gamma(\mathcal{G}_j^+) \left[ \varphi(v_i|\mathcal{G}^-) + \theta(i, k|\mathcal{G}_j^+) \right] \\ &= \underbrace{\sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \varphi(v_i|\mathcal{G}^-) \sum_{j=1}^{k-1} \mathbb{P}_\gamma(\mathcal{G}_j^+)}_A + \underbrace{\sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \sum_{j=1}^{k-1} \mathbb{P}_\gamma(\mathcal{G}_j^+) \theta(i, k|\mathcal{G}_j^+)}_{m_{ik}} \end{aligned} \tag{A.6}$$

**Tweet influence at previous time step  $t_{k-1}$ .** Given the definition of  $\mathcal{G}_j^+$  in Eq. (A.3) and the independant cascades assumption, we obtain that  $\mathbb{P}_\gamma(\mathcal{G}_j^+) = \mathbb{P}_\gamma(\mathcal{G}^-)p'_{jk}$ . Consequently, part A in Eq. (A.6) can be written as:

$$\begin{aligned} A &= \sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \varphi(v_i|\mathcal{G}^-) \sum_{j=1}^{k-1} \mathbb{P}_\gamma(\mathcal{G}^-)p'_{jk} \\ &= \sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \varphi(v_i|\mathcal{G}^-) \mathbb{P}_\gamma(\mathcal{G}^-) \sum_{j=1}^{k-1} p'_{jk} \stackrel{cf. (A.2)}{=} \varphi^{k-1}(v_i) \end{aligned} \tag{A.7}$$

A is the tweet influence of  $v_i$  at the previous time step  $t_{k-1}$ . Note that  $\sum_{j=1}^{k-1} p'_{jk} = 1$  because  $v_k$  is necessarily the direct retweet of a previous nodes  $v_j, j < k$  of the retweet cascade.

**Contribution of  $v_k$ .** With A being the influence of  $v_i$  at the previous time step, intuitively  $m_{ik}$  is the contribution of  $v_k$  to the influence of  $v_i$ . Knowing that:

$$\begin{aligned} \mathbb{P}_\gamma(\mathcal{G}_j^+) &= \mathbb{P}_\gamma(\mathcal{G}^-)p'_{jk} \text{ and} \\ \theta(i, k|\mathcal{G}_j^+) &= \Psi(\mathcal{G}_{i \rightarrow j})\pi_{jk} = \theta(i, j|\mathcal{G}_j^+)\pi_{jk} \stackrel{cf. 3.2}{=} \theta(i, j|\mathcal{G}^-)\pi_{jk} \end{aligned}$$

we write  $m_{ik}$  as:

$$\begin{aligned} m_{ik} &= \sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \sum_{j=1}^{k-1} \mathbb{P}_\gamma(\mathcal{G}^-)p'_{jk} \theta(i, j|\mathcal{G}^-)\pi_{jk} \\ &= \sum_{j=1}^{k-1} p'_{jk} \pi_{jk} \underbrace{\sum_{\mathcal{G}^- \in \Upsilon_{1:k-1}} \mathbb{P}_\gamma(\mathcal{G}^-) \theta(i, j|\mathcal{G}^-)}_{m_{ij}} = \sum_{j=1}^{k-1} p'_{jk} \pi_{jk} m_{ij} \stackrel{cf. 3.3}{=} \end{aligned}$$

## Appendix B: Efficient computation of GIM

We define two matrices. First, the transfer matrix  $T = [p'_{ij} * \pi_{ij}]$ , where the element  $p'_{ij}$  is the probability that tweet  $v_j$  is a direct retweet of tweet  $v_i$  (defined in Eq. (3)) and  $\pi_{ij}$  is the proportion of capital transferred from  $v_j$  to  $v_i$ ; Second, the influence accumulation matrix  $M = [m_{ij}]$ , with  $m_{ij}$  defined in Eq. (3.3) is the contribution of  $v_j$  to the influence of  $v_i$ . For each column  $j$  of  $M$ , we compute the first  $j - 1$  elements by multiplying the sub-matrix  $M_{[1..j-1, 1..j-1]}$  with the first  $j - 1$  elements on the  $j$ th column of matrix  $T$ , the  $j$ -th element is  $\pi_{jj}$ , and the remaining elements are 0. The computation of matrix  $M$  finishes after  $n$  steps, where  $n$  is cascade size.

## Appendix C: Value-allocation scheme

Value-allocation schemes over networks are an example of fair division games, concerned with how to allocate the value generated by a network of players among the players; for example, the allocation of advertisement revenue over a chain of marketers, or the allocation of utility revenue across a network of electricity infrastructure providers. Allocation scheme literature generally aims to understand the stability, efficiency, and fairness of network formation (usually in cooperative undirected networks). Our paper is concerned with the allocation of influence-capital in a non-cooperative setting over an ad-hoc directed-acyclic diffusion graph. The seminal work of Jackson and Wolinsky [50]'s *Connections* game defines the utility of a player as the sum of benefits to all other players decayed by the length of the path between them minus the cost of maintaining direct links for the player.

### Abbreviations

COVID-19, Coronavirus disease; GIM, Generalized Influence Model; O\*NET, Occupational Information Network; NDCG, Normalised Discounted Cumulative Gain; AUC, Area under the curve; MAPE, Mean Absolute Percentage Error; API, Application Programming Interface; URL, Universal Resource Locator; US, United States of America.

### Acknowledgements

The authors extend their gratitude to Dr. Manuel Cebrian who, during an informal discussion over coffee more than a decade prior to this publication, hypothesized the existence of a conductance property in social networks analogous to electrical conductance in physical systems.

### Author contributions

Both authors contributed substantially to the manuscript's conceptualization, design, data collection, analyses, interpretation, and drafting.

### Funding information

Open Access funding enabled and organized by CAUL and its Member Institutions. This research was supported by an Australian Government Research Training Program (RTP) Scholarship, the Advanced Strategic Capabilities Accelerator (ASCA), the Australian Department of Home Affairs, the Defence Science and Technology Group, the Defence Innovation Network, the National Science Centre, Poland (Project No. 2021/41/B/HS6/02798).

### Data availability

The datasets analyzed and code required to reproduce the results are available in the GitHub repository, <https://github.com/behavioral-ds/generalized-model-of-social-influence>. Auxiliary code and/or data is available from the corresponding author upon reasonable request.

## Declarations

### Competing interests

The authors declare no competing interests.

Received: 11 August 2025 Accepted: 19 March 2026 Published online: 17 April 2026

### References

1. Rapp C (2016) Moral opinion polarization and the erosion of trust. *Soc Sci Res* 58:34–45

2. McCauley C, Moskalenko S (2008) Mechanisms of political radicalization: pathways toward terrorism. *Terrorism Polit Violence* 20(3):415–433
3. Nowak A, Vallacher RR, Kus M, Urbaniak J (2005) The dynamics of societal transition: modeling nonlinear change in the Polish economic system. *Int J Sociol*
4. Greijdanus H, Matos Fernandes CA, Turner-Zwinkels F, Honari A, Roos CA, Rosenbusch H, Postmes T (2020) The psychology of online activism and social movements: relations between online and offline collective action. *Curr Opin Psychol*
5. Gunton K (2022) The impact of the Internet and social media platforms on radicalisation to terrorism and violent extremism. In: *Privacy, security and forensics in the Internet of Things (IoT)*
6. Peng S, Zhou Y, Cao L, Yu S, Niu J, Jia W (2018) Influence analysis in social networks: a survey. *J Netw Comput Appl*
7. Mason WA, Conrey FR, Smith ER (2007) Situating social influence processes: dynamic, multidirectional flows of influence within social networks. *Personal Soc Psychol Rev*
8. Kempe D, Kleinberg J, Tardos É (2003) Maximizing the spread of influence through a social network. In: *KDD*
9. Du N, Song L, Gomez-Rodriguez M, Zha H (2013) Scalable influence estimation in continuous-time diffusion networks. In: *NIPS*
10. Mishra S, Rizoiu M-A, Xie L (2016) Feature driven and point process approaches for popularity prediction. In: *CIKM*
11. Moussaïd M, Kämmer JE, Analytis PP, Neth H (2013) Social influence and the collective dynamics of opinion formation. *PLoS ONE* 8(11):78433
12. Kraus MW, Callaghan B, Ondish P (2019). *Social class as culture*
13. Cialdini RB (2001) *Influence: science and practice*. Pearson Education, Upper Saddle River
14. Asch SE (1961) Effects of group pressure upon the modification and distortion of judgments. In: *Documents of gestalt psychology*
15. Rizoiu M-A, Graham T, Zhang R, Zhang Y, Ackland R, Xie L (2018) #debatenight: the role and influence of socialbots on Twitter during the 1st 2016 us presidential debate. In: *ICWSM*
16. Newcomb TM (1953) An approach to the study of communicative acts. *Psychol Rev*
17. Milgram S, Godehus C (1978) *Obedience to authority*. Ziff-Davis Publishing Company
18. Horai J, Naccari N, Fatoullah E (1974) The effects of expertise and physical attractiveness upon opinion agreement and liking. *Sociometry*
19. Ghaffar F, Hurley N (2020) Structural hole centrality: evaluating social capital through strategic network formation. *Comput Soc Netw*
20. Sloan L, Morgan J, Burnap P, Williams M (2015) Who tweets? Deriving the demographic characteristics of age, occupation and social class from Twitter user meta-data. *PLoS ONE*
21. O\*NET Development, N.C. O\*NET OnLine. <https://www.onetonline.org/>. Accessed: 2021-06-30
22. Cui L, Lee D (2020) Coaid: Covid-19 healthcare misinformation dataset. arXiv preprint [arXiv:2006.00885](https://arxiv.org/abs/2006.00885)
23. Metaxas P, Mustafaraj E, Wong K, Zeng L, O’Keefe M, Finn S (2015) What do retweets indicate? Results from user survey and meta-review of research. In: *Proceedings of the international AAAI conference on web and social media*, vol 9, pp 658–661
24. Liu Q, Xiang B, Yuan NJ, Chen E, Xiong H, Zheng Y, Yang Y (2017) An influence propagation view of pagerank. *ACM Trans Knowl Discov Data*
25. Silva A, Guimarães S, Meira W Jr, Zaki M (2013) Profilerank: finding relevant content and influential users based on information diffusion. In: *SNKDD*
26. Mishra S, Rizoiu M-A, Xie L (2018) Modeling popularity in asynchronous social media streams with recurrent neural networks. In: *ICWSM*
27. Romero DM, Galuba W, Asur S, Huberman BA (2011) Influence and passivity in social media. In: *ECML PKDD*
28. Bakshy E, Hofman JM, Mason WA, Watts DJ (2011) Everyone’s an influencer: quantifying influence on Twitter. In: *WSDM*
29. Smith ST, Kao EK, Shah DC, Simek O, Rubin DB (2018) Influence estimation on social media networks using causal inference. In: *IEEE SSP*
30. Nickel M, Le M (2021) Modeling sparse information diffusion at scale via lazy multivariate Hawkes processes. In: *WWW*
31. Hawkes AG (1971) Spectra of some self-exciting and mutually exciting point processes. *Biometrika*
32. Rizoiu M-A, Lee Y, Mishra S, Xie L (2017) In: Chang S-F (ed) *Hawkes processes for events in social media*. Association for Computing Machinery and Morgan & Claypool, pp 191–218. <https://doi.org/10.1145/3122865.3122874>. <https://dl.acm.org/doi/10.1145/3122865.3122874>
33. Kong Q, Rizoiu M-A, Xie L (2020) Modeling information cascades with self-exciting processes via generalized epidemic models. In: *Proceedings of the 13th international conference on web search and data mining*. ACM, New York, pp 286–294. <https://doi.org/10.1145/3336191.3371821>. <https://arxiv.org/abs/1910.05451>. <https://dl.acm.org/doi/10.1145/3336191.3371821>
34. Hawkes AG, Oakes D (1974) A cluster process representation of a self-exciting process. *J Appl Probab*
35. Lewis E, Mohler G (2011) A nonparametric em algorithm for multiscale Hawkes processes. *J Nonparametr Stat*
36. Luceri L, Braun T, Giordano S (2019) Analyzing and inferring human real-life behavior through online social networks with social influence deep learning. *Appl Netw Sci*
37. Qiu J, Tang J, Ma H, Dong Y, Wang K, Tang J (2018) Deepinf: social influence prediction with deep learning. In: *KDD*
38. Gu Y, Sun Y, Li Y, Yang Y (2018) Rare: social rank regulated large-scale network embedding. In: *WWW*
39. Tsitsulin A, Mottin D, Karras P, Müller E (2018) Verse: versatile graph embeddings from similarity measures. In: *WWW*
40. Yu W, Qin Z (2019) Spectrum-enhanced pairwise learning to rank. In: *WWW*
41. Rizoiu M-A, Mishra S, Kong Q, Carman M, Xie L (2018) Sir-Hawkes: linking epidemic models and Hawkes processes to model diffusions in finite populations. In: *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW ’18*. ACM, New York, pp 419–428. <https://doi.org/10.1145/3178876.3186108>. <http://arxiv.org/abs/1711.01679> <http://dl.acm.org/citation.cfm?doid=3178876.3186108>
42. Lin LC, Qu Y, Telzer EH (2018) Intergroup social influence on emotion processing in the brain. *PNAS*
43. Centola D (2011) An experimental study of homophily in the adoption of health behavior. *Science* 334(6060):1269–1272

44. Crandall D, Cosley D, Huttenlocher D, Kleinberg J, Suri S (2008) Feedback effects between similarity and social influence in online communities. In: KDD
45. Goel R, Soni S, Goyal N, Paparrizos J, Wallach H, Diaz F, Eisenstein J (2016) The social dynamics of language change in online networks. In: International conference on social informatics
46. Moody J (1989) Fast learning in multi-resolution hierarchies. In: NIPS
47. Granovetter M (1983) The strength of weak ties: a network theory revisited. *Social Theory*
48. Dekker P, Uslaner EM (2003) *Social capital and participation in everyday life*. Routledge, Oxfordshire
49. Shapley LS (1997) A value for n-person games. *Class Game Theory* 69
50. Jackson MO, Wolinsky A (2003) A strategic model of social and economic networks. In: *Networks and groups*. Springer, New York, pp 23–49
51. Subbian K, Sharma D, Wen Z, Srivastava J (2014) Finding influencers in networks using social capital. *Soc Netw Anal Min*
52. Appendix O (2021) Appendix: conductance and Influence-Capital: modeling Online Social Influence. [https://www.dropbox.com/s/wcaovmb5oeawnvq/supplementary\\_material.pdf?dl=0](https://www.dropbox.com/s/wcaovmb5oeawnvq/supplementary_material.pdf?dl=0)
53. Ram R, Rizioi M-A (2024) Empirically measuring online social influence. *EPJ Data Sci* 13:53. <https://doi.org/10.1140/epjds/s13688-024-00492-z>
54. Carpentras D, Quayle M (2023) The psychometric house-of-mirrors: the effect of measurement distortions on agent-based models' predictions. *Int J Soc Res Methodol* 26(2):215–231
55. Graham T, Keller T (2020) Bushfires, bots and arson claims: Australia flung in the global disinformation spotlight. *Conversation* 10
56. Page L, Brin S, Motwani R, Winograd T (1999) The pagerank citation ranking: bringing order to the web. Technical report, Stanford
57. Xiang B, Liu Q, Chen E, Xiong H, Zheng Y, Yang Y (2013) Pagerank with priors: an influence propagation perspective. In: IJCAI
58. Cha M, Haddadi H, Benevenuto F, Gummadi K (2010) Measuring user influence in Twitter: the million follower fallacy. In: ICWSM
59. Wu S, Rizioi M-A, Xie L (2019) Estimating attention flow in online video networks. In: *Proceedings of the ACM on human-computer interaction* 3, pp 1–25. <https://doi.org/10.1145/3359285>
60. Frantz TL, Cataldo M, Carley KM (2009) Robustness of centrality measures under uncertainty: examining the role of network topology. *Comput Math Organ Theory* 15(4):303–328
61. Riddell J, Brown A, Kovic I, Jauregui J (2017) Who are the most influential emergency physicians on Twitter? *West J Emerg Med*
62. Bailo F, Johns A, Rizioi M-A (2023) Riding information crises: the performance of far-right Twitter users in Australia during the 2019–2020 bushfires and the covid-19 pandemic. *Inf Commun Soc*, 1–19. <https://doi.org/10.1080/1369118X.2023.2205479>
63. Kern ML, McCarthy PX, Chakrabarty D, Rizioi M-A (2019) Social media-predicted personality traits and values can help match people to their ideal jobs. *Proc Natl Acad Sci USA* 116:26459–26464. <https://doi.org/10.1073/pnas.1917942116>
64. Turrell A, Speigner B, Djumalieva J, Copple D, Thurgood J (2018) Using job vacancies to understand the effects of labour market mismatch on uk output and productivity
65. Mukherjee S, Widmark D, DiMascio V, Oates T (2021) Determining standard occupational classification codes from job descriptions in immigration petitions. In: ICDMW
66. Singh L, Bode L, Budak C, Kawintiranon K, Padden C, Vraga E (2020) Understanding high-and low-quality url sharing on covid-19 Twitter streams. *J Comput Soc Sci*
67. Aral S, Muchnik L, Sundararajan A (2009) Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proc Natl Acad Sci USA* 106(51):21544–21549
68. Romero DM, Meeder B, Kleinberg J (2011) Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on Twitter. In: *Proceedings of the 20th International Conference on World Wide Web*. ACM, New York, pp 695–704. <https://doi.org/10.1145/1963405.1963503>
69. Bakshy E, Karrer B, Adamic LA (2009) Social influence and the diffusion of user-created content. In: *Proceedings of the 10th ACM conference on electronic commerce*. ACM, New York, pp 325–334. <https://doi.org/10.1145/1566374.1566421>
70. Mønsted BM, Sapiezzyński P, Ferrara E, Lehmann S (2017) Evidence of complex contagion of information in social media: an experiment using Twitter bots. *PLoS ONE* 12(9), Article ID 0184148. <https://doi.org/10.1371/journal.pone.0184148>
71. Guillebaud D, Becker J, Centola D (2018) Complex contagions: a decade in review. In: Lehmann S, Ahn Y-Y (eds) *Complex spreading phenomena in social systems*. Springer, Berlin, pp 3–25. [https://doi.org/10.1007/978-3-319-77332-2\\_1](https://doi.org/10.1007/978-3-319-77332-2_1)
72. Guillebaud D, Centola D (2021) Topological measures for identifying and predicting the spread of complex contagions. *Nat Commun* 12:4430. <https://doi.org/10.1038/s41467-021-24704-6>
73. Kong Q, Booth E, Bailo F, Johns A, Rizioi M-A (2022) Slipping to the extreme: a mixed method to explain how extreme opinions infiltrate online discussions. *Proc Int AAAI Conf Web Soc Media* 16:524–535. <https://doi.org/10.1609/icwsm.v16i1.19312>

## Publisher's note

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