Noname manuscript No.
(will be inserted by the editor)

Evolving activity cascades on socio-technological networks

Javier Borge-Holthoefer · Pablo Piedrahita · Alex Arenas

Received: date / Accepted: date

Abstract Networks are the substrate on which social contagion propagates, from the growth of social movements to the adoption of innovations. In the complex networks community, it took some time to realize the difference between simple propagation –e.g. the spread of disease–, in which a single active node is sufficient to trigger the activation of its neighbors, and complex contagion, in which node activation requires simultaneous exposure to multiple active neighbors. Rooted in the social science literature, complex contagion has settled as the driving mechanism for behavior cascades on social networks. However, our access to digital traces of social interaction reveals, besides and beyond complex contagion, bursty activity patterns, repeated agent activation, and occasionally a form of synchronization under the form of trending topics and hypes. Thus, the threshold model –the paramount example in the tradition of complex contagion– needs to shift from a standpoint in which agents become irreversibly active ("one-off" events), to another in which agents continuously change their state and whose activity shows oscillating patterns. Here we review a mechanistic model that, within the logic of complex contagion, accounts as well for the temporal evolution of behavior cascades. In it, agents follow the dynamics of integrate-and-fire oscillators. The affordances of

J. Borge-Holthoefer
Internet Interdisciplinary Institute (IN3), Universitat Oberta de Catalunya, Barcelona
Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, Spain
E-mail: jborgeh@uoc.edu

P. Piedrahita
Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, Spain
E-mail: ppiedrahita@gmail.com

A. Arenas
Departament d'Informàtica i Matemàtiques, Universitat Rovira i Virgili, Tarragona
IPHES, Institut Catala de Paleoecologia Humana i Evolució Social, Tarragona
E-mail: alexandre.arenas@urv.cat

the model –and of some recent variations on it– will open a discussion and outlook for future developments.

Keywords Complex networks · Recurrent activity · Threshold models

1 Introduction

In the last fifteen years, complex networks researchers have heavily focused on the question of information and behavior propagation on social systems. It is a natural hot topic in the area, since it affects a wide range of phenomena with social, commercial, political and communicational implications, among others. Along this rich tradition, the distinction between *simple* and *complex* contagion represents a milestone. In short, the spreading mechanism of simple contagion implies that a node may switch its state to "active" solely from the influence of a neighboring active peer. On the contrary, complex contagion implies that, for a node to become active, it needs to be exposed to multiple active neighbors. While the former clearly suffices for many diffusion dynamics, e.g. epidemic spreading [1] and perhaps the propagation of low-value, low-risk information [2], it is the latter that better explains how agents adopt and spread further, for example, risky or costly behaviors: these demand reinforcement from multiple sources [3].

This is not to say that complex contagion came along just some years ago: collective effects of social influence have been largely studied through threshold models [4–6] in sociology; however, their development under the light of complex networks [7–9] is relatively recent. In the networked context, the group of reference is now determined by connectivity in the network and it changes from actor to actor –as opposed to activation of individual thresholds responding to global information, which lies at the core of the original formulations. Subtle as it may appear, the adoption of the idea of networked complex contagion to explain social propagation actually meant a reformulation of the problem, frustrating the possibility of solely mapping concepts, formalisms and results from network epidemics (e.g. [10]) to the social context: in the end, simply switching "susceptible" to "ignorant", or "infected" to "spreader", did not suffice.

From the moment it became clear that the threshold model was the standard tool to capture cascading processes in the social medium, many developments followed, specially focused on the role of topology in a contagion spread: the importance of the seed size, i.e. the fraction of initially active nodes in the dynamics, the presence of correlations and/or modularity in the structure, or the temporal characterization of the cascading processes [11–14], see [15] for an overview. All variations of the threshold model share, however, two important elements: first, that activation is modelled as a step function that goes from "inactive" to "active" (or 0 to 1) when thresholds are reached; and second, that thresholds can only be reached once, that is, activation is assumed to be a "one-off" event.

Remarkably, these decade-old developments coincide with the consolidation of the Web 2.0, the advent of the Internet of Things and a general adoption of mobile technologies. The availability of large amounts of those data have convinced researchers that theories can (and must) be mapped to real scenarios, and put into empirical test [16]. It is in this new scenario that scholars have learned, if only by observation, that many real cascades in online social networks (OSNs) present recurrent activation (i.e. agents that transition from 0 to 1 and viceversa repeatedly) that develop in a set of cascades with intriguing scaling properties [17]. While it remains true that some information propagation phenomena falls under the category of single-shot events, e.g. e-mail chain letters [18], threshold models fall short when it comes to reproduce a large fraction of what happens in OSNs: much activity is triggered not as mere content relaying (URL forwarding [19,20], "retweets" [21], "likes" [22], etc.), but rather as a reaction to activity itself. In short, threshold models cannot account for *unfolding events*, including the use of some OSNs as an instant messaging system, in which non-identical pieces of information around a topic may be circulating (typically over short time spans) in many-to-many interactions, along direct or indirect information pathways [23].

This important distinction calls for minimal mathematical definitions. As said, information-based cascades revolve exclusively around "content chains": the basic criterion to include a node i in a diffusion tree starting at j is to guarantee that (i) i and j became friends at t_1 (the notion of "friend" changes from SNS to SNS, and must be understood broadly here); (ii) i received a piece of information from j, who had previously sent it out, at time t_2; and finally (iii) the node i sends out the same piece of information at time t_3. Note that no strict time restriction exists besides the fact that $t_1 < t_2 < t_3$, the emphasis is placed on whether the *same* content is flowing. On the other side, the definition of behavior-based cascades relaxes the content copy condition in (iii), but tightens up the temporal restrictions: not only $t_1 < t_2 < t_3$, but also $t_3 - t_2 \leq \Delta\tau$, i.e. the time between reception and delivery of information must observe a certain sequentiality, separated at most by $\Delta\tau$, where τ is an arbitrary time lapse.

The definition of a behavior-based cascade is useful to measure the type of diffusion less focused on content and more on behavior [17,24,25]: it uncovers how –and how often– users get involved in sequential message exchange, for which the strict repetition of the same content is not necessary. In this sense, behavior-based cascades conform a set that includes information-based avalanches, but not viceversa. Along this line, our objective in this work is to make the case for, and present the main features of, a family of models that generalizes the threshold model –retaining the fundamental idea of complex contagion, and extending it to the temporal dimension.

2 A model for unfolding and recurrent cascading phenomena

Here we take the challenge to propose a mechanistic model for unfolding and recurrent cascading phenomena in OSNs. First, we revise the acclaimed Watts' threshold model [8] from which we will borrow the computation of the condition for large cascades to occur in networks. Second, we review a recent proposal of a mechanistic model that can be used as the core of a new set of models that enlarge the capabilities of threshold models to account for complex contagion phenomena, and adds the recurrence of cascades as a natural outcome of a continuous dynamical system. After we will analyze its possible parameterizations and discuss the applicability to real scenarios.

2.1 Watts' threshold model

The seminal work on thresholds model proposed by Granovetter [4], and its extension to networks proposed by Watts [8], are extraordinarily stylized abstractions of social dynamics, where the essentials on behavior adoption is based on the pressure of neighbors in such a way that, the more friends adopt that same behavior, the more likely is that the agent adopts it herself. Briefly, Watts' threshold model assigns a fixed threshold τ, drawn from a distribution $0 \leq g(\tau) \leq 1$, to each node (individual) in a complex network of size N and an arbitrary degree distribution p_k, being k the degree (number of acquaintances). Each node is initially marked as *inactive* except a set of seeding active nodes, usually this fraction of nodes is denoted $\Phi_0 = 1/N$. Being a_i the number of active neighbors, a node i with degree k_i updates its state becoming active whenever the fraction of active neighbors $a_i/k_i > \tau_i$. The simulation of this mechanistic process evolves following this rule until an stationary state is reached, i.e., no more updates can occur. Given this setup, the *cascade condition* in degree-uncorrelated networks can be derived from the growth of the initial fraction of active nodes, who on their turn might induce the one-step-to-activation (vulnerable) nodes. Therefore, large cascades can only occur if the average cluster size of vulnerable nodes diverges. This condition is met at [8,11]

$$F = \sum_k k(k-1)\rho_k p_k = \langle k \rangle \tag{1}$$

where ρ_k is the density of nodes with degree k close to their activation threshold, p_k is the fraction of nodes of degree k and $\langle k \rangle$ is the average degree [8].

For $F < \langle k \rangle$ all the clusters of vulnerable nodes are small, and the initial seed can not spread beyond isolated groups of early adopters; on the contrary, if $F > \langle k \rangle$ then small fraction of disseminators may unleash –with finite probability– large cascades. The cascade condition has been analytically determined for different initial conditions [11] as well as for modular and correlated networks [12,13], placing the threshold model in the more general context of critical phenomena and percolation theory [26].

The model, as described, has some scope limitations since it can account only for one-shot events, for instance the diffusion of a single rumor or the adoption of an innovation. Also, this framework leaves no room for spontaneous initiative: even low-threshold nodes –those with higher propensity to participate in a cascade– will not be recruited unless their neighbors act upon them. Empirical evidence suggests, instead, that once an agent becomes active that behavior will be sustained, and reinforced, over time [27]. This creates a form of sustained activation that will be affected and affect other agents over time in a recursive way. Indeed, events evolve in time –and so do the cascades elicited therein [24], as a consequence of dynamical changes in the states of agents as dynamics evolve. Cascades are then events that brew over time in a system that holds some memory of past interactions. Moreover, the propensity to be active in the propagation of information sometimes depends on other factors than raw social influence, e.g., mood, personal implication, opinion, etc. Eventually, the solution to some of the previous requirements would consist in to adapt the Watts' model by injecting new seeds after avalanches stop, and giving more realism to the parameterization of the model in terms of: memory [9], directionality of links [28], weights [29], etc. However, none of the previous analyze a situation in which the thresholds themselves vary in time, i.e. $\tau_i(t)$.

2.2 A minimal model for recurrent activation (MRA) on social systems

Here we revise a model with self-sustained activity, where system-wide events emerge as microscopical conditions become increasingly correlated [30]. The main idea is to assign to every agent a continuous dynamical function parallel to a phase response curve (a curve describing the relationship between a stimulus and a response). To build up the MRA, the auhtors in [30] capitalize on the classical integrate-and-fire oscillator (IFO) model by Mirollo and Strogatz [31]. In this model, each node in a network of size N is characterized by a voltage-like state $m \in [0, 1]$ of an oscillator, which monotonically increases with phase ϕ until it reaches $m = 1$, and then it fires or *activates* (emits information to its coupled neighbors), and immediately deactivates (resets its state to $m = 0$). The pulsatile dynamics, makes that each time a node becomes active, the state of its k neighbors is increased by ϵ. More precisely, m is uniformly distributed at $t = 0$ and evolves such that

$$m = f(\phi) = \frac{1}{w} \ln(1 + [e^w - 1]\phi) \qquad (2)$$

parametrized by $w > 0$ to guarantee that f is concave down. Whenever $m_i = 1$ then *instantaneously* $m_j = \min(1, m_j + \epsilon)$, if the edge (i, j) exists. Thus, a node may reach activation either by itself (spontaneously, as its phase comes to an end) or because of its neighbors' action (which may pull it up to action).

In neuroscience, integrate-and-fire models have been extremely useful to assess the bursty behavior and the emergence of cascades in neural populations represented as lattices [32] or complex networks [33,34]. Inspired in these facts,

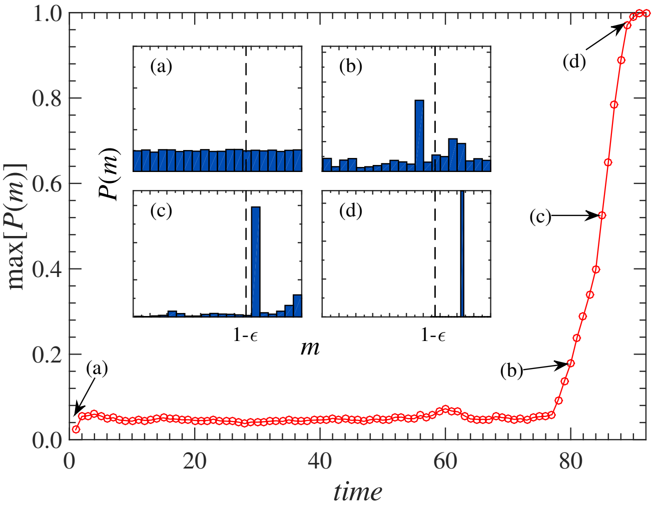

Fig. 1 Inset (a)-(d): Motivation m probability distributions of four different representative times along the synchronization window. Each snapshot depicts the m-state histogram of the N oscillators. The dynamics begins with a random uniform distribution of m-states –inset (a)– and it progressively narrows during the transition to synchrony –inset (d). Main: largest fraction of synchronized nodes across time. The path to synchronization evolves steadily at a low level, and eventually suffers an abrupt transition. Adapted from [30].

it has been proposed a MRA to mimic social activity as a complex network of IFOs representing the time evolving activation of individuals. Within this framework, spontaneous propensity –activation regardless exogenous factors– is guaranteed; while contagion is genuinely complex, i.e. the number of necessary external influences (if any) to show activity varies in time. Remarkably, activity is purely periodical only if the oscillators are either isolated (disconnected), dynamically uncoupled ($\epsilon = 0$), or in other very particular conditions of full synchrony. These three scenarios are irrelevant –in terms of social modeling. On the other hand, some traces of periodicity in users' activity in socio-technical platforms –like Twitter– have been widely studied at the aggregate level (see for instance [35]), and they exist also at the individual level. Therefore, the proposed MRA is willing to represent time-evolving cascades of activity, whose distribution will depend on the parameters of the system.

The model comprises two free parameters, w and ϵ, which are closely related. The parameter w may be interpreted as the willingness or *intrinsic propensity* of agents to participate in a certain diffusion event: the larger w, the shorter it takes for a node to enter the tip-over interval $1 - \epsilon \leq m < 1$. The other parameter, ϵ quantifies the amount of influence an agent exerts onto its neighbors when it shows some activity. Larger ϵ's forces agents to go more rapidly into the tip-over region. Both quantities affect the level of *motivation* m of a given agent. Note that ϵ in the current framework evokes τ in the classical threshold model, in the sense that both determine the width of the tip-over region. Finally, the phase is translated into time steps, and then prescribed as $\phi(t) = t$.

3 Numerical insights: regimes of activity

The preliminary numerical analysis of the proposed MRA already reveals several regimes of activity, and provides some intuitions regarding the evolving dynamics of these systems, see Figure 1. Typically, the initial conditions $\{m_i\}$ are randomly set at $t = 0$ (inset (a)), out of a uniform distribution. Then the system evolves, until global cascades appear. In more detail, Figure 2 shows the typical activity observed in an Erdős-Rényi network, with $N = 10^4$ and $\langle k \rangle = 4$. Note that at the beginning of the simulation ($t \lesssim 30$), only isolated nodes, or small groups of them, fire occasionally. Then, for $t \gtrsim 30$ a cluster of synchronized oscillators emerges abruptly, until it reaches about 80% of the total nodes and continues increasing its size –slowly– for $t \gtrsim 50$ (see [32] for another example of this activity on two-dimensional square lattices).

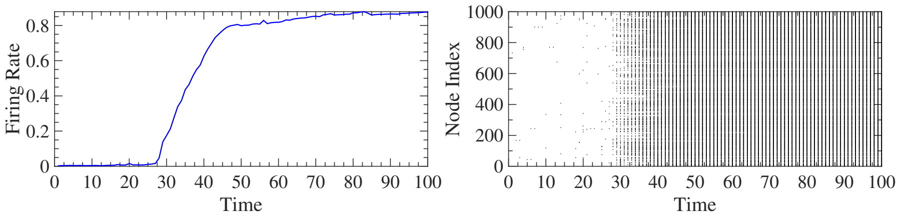

Fig. 2 Firing rate (left) and raster plot (right) for an Erdős-Rényi network with $N = 10^4$, $\langle k \rangle = 4$, $k_{min} = 2$, $\omega = 10$ and $\epsilon = 10^{-3}$ (threshold $\theta = 1$, always). Note that, although the network size is 10^4, the raster plot in this case only displays the activity of 10% of the population.

The dynamics in Figure 2 was set for a very small value of ϵ (10^{-3}). At this value of the coupling, attaining full synchronization in an ER network is still possible because there exists a large set of nodes that fire in-phase (for $t \gtrsim 50$) and its size increases over time. However, synchronizability is not guaranteed: given a particular topology and a value of ω, there is a critical value ϵ_c below which full synchronization (or even a cascade of size $\sim N$) can not occur.

Figure 3 explores precisely this, showing the typical activity observed in a scale-free network given two parametrizations of the dynamics. In the top panels, even for a system of agents with low intrinsic motivation ($\omega = 3$), a strong influence ϵ pushes the dynamics to complete synchronization early on. On the contrary, bottom panels show a system where agents are strongly motivated ($\omega = 48$), and even with initial conditions are set to favor full synchronization (notice large cascades observed in $t \lesssim 50$), there are values of

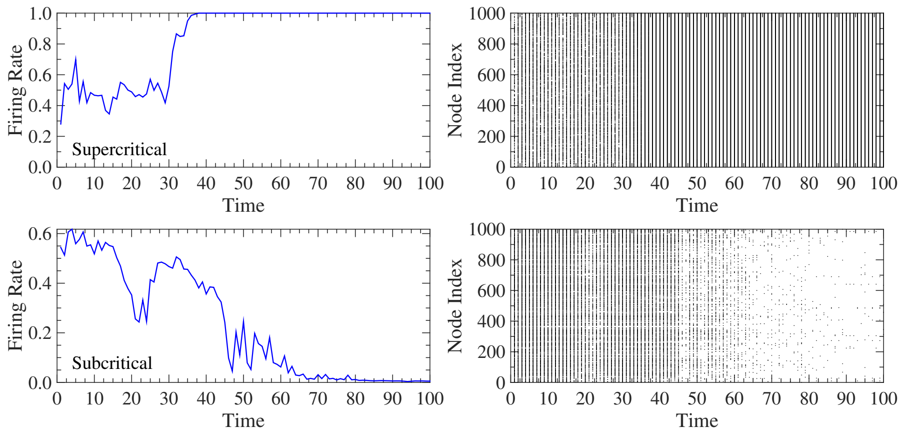

Fig. 3 Firing rates (left column) and raster plots (right column) for a SFN with $N = 10^4$, $\gamma = 3.0$, $k_{min} = 2$. In the top panels supercritical activity is shown for $\omega = 3$ and $\epsilon = 0.233$, and bottom panels display subcritical activity for $\omega = 48$ and $\epsilon = 10^{-3}$.

the coupling ϵ for which the population of oscillators not only does not sustain its cluster of synchronized agents (which can happen for values $\sim \epsilon_c$) but it tends to desynchronize rapidly the global activity.

These observations on the existence or not of global cascades (thus, the existence of not of an ϵ_c) can be extended to a larger parameter space. In Figure 4 the regions where macro-cascades ($O(N)$) are possible are color-coded for each cycle $0, 25, 75, 100, \ldots$ (see [30] for a precise definition of "cycle"). Dark blue is used in regions where cascades are not comparable to the system's size (labeled as N.C., "no cascades", in the figure). Note that if cascades are possible for a given cycle c, they will be possible also for any $c' \geq c$. This figure renders an interesting scenario: on the one hand, it again suggests the existence of critical ϵ_c, in coherence with the lower panel of Figure 3 (dark blue area in the phase diagram). On the other, it establishes how many cycles it takes for a particular $(\epsilon, \langle k \rangle)$ pair to attain macroscopical cascades (full synchronization) –which becomes an attractor thereafter, for undirected connected networks. Given the cumulative dynamics of MRA, in contrast with Watts' model, the region in which global cascades are possible grows with $\langle k \rangle$. Turning to the social sphere, these results open the door to predicting how long it takes for a given topology, and a certain level of inter-personal influence, to achieve system-wide events. Furthermore, the existence of a limiting ϵ_c determines whether such events can happen at all.

4 Analytical insights: homogeneous parametrization

The previous section illustrates the suggestive properties that MRA displays from a dynamical point of view –but also as a proxy to recurrent complex

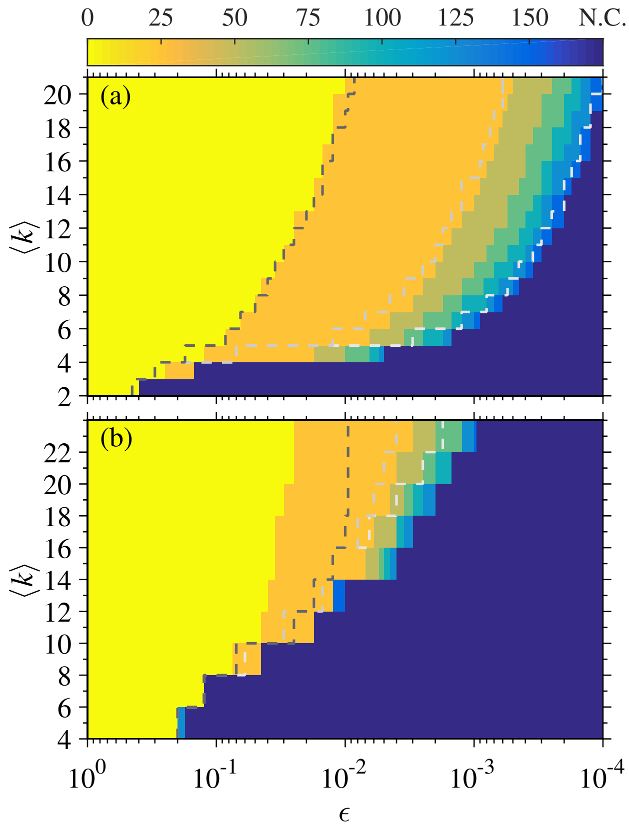

Fig. 4 $(\epsilon, \langle k \rangle)$ cascade diagram for different cycles (coded by color), with fixed $\omega = 3$. Results are obtained for synthetic (a) Erdös-Rényi, and (b) and scale-free (with $\gamma = 3$) uncorrelated networks of size $N = 10^4$. Vertical axis and each dashed line define a confined region in which global cascades might occur according to Eq. 6 and for a specific cycle (here we show only the expected zones for $c = 0$ –dashed white– and $c = 150$ –dashed gray). A cascade is considered "macroscopical" if the synchronized cluster $S_c \geq 0.25N$. Color codes indicate the existence of at least one cascade $S > S_c$ in numerical simulations; analytical predictions are averaged over 200 networks with random initial conditions. Note that the cascade condition in (a) often underestimates the actual cascade regions because it does not take into account second order interactions; the same applies in the lower panel (b), except for $c = 0$ where the analytical prediction overestimates the results because the inclusion of the hub into the cascade is improbable starting from a uniform distribution. Adapted from [30].

contagion. The aim now is to show that the model gives room as well to a certain degree of formalization. To gain analytical insight, we depart from Watts' model, using Eq. 1 to derive the cascade condition in the MRA framework. Note that now the distribution of activity is governed by $\rho(t) = 1 - \int_0^{1-\epsilon} g(m,t)\, dm$ where $g(m,t)$ corresponds to the states' probability distribution at a certain time t. For an initial uniform distribution of motivation m and a fixed ϵ, the condition for the emergence of cascades at $t = 0$ reads

$$\epsilon \sum_k k(k-1)p_k = \langle k \rangle \quad (3)$$

(see inset (a) in Fig. 1). In general, for any time t

$$\rho(t) \sum_k k(k-1)p_k = \langle k \rangle, \quad (4)$$

which implies that the cascade condition depends on time in our proposed framework. Clearly, in this scenario $\rho(t)$ is not a function of the node degree k, as opposed to ρ_k in Watts' proposal.

Assuming a supercritical regime, the states of the nodes become increasingly correlated as the dynamics evolve in time and, consequently, the distribution of states changes dramatically. The evolution of the states distribution is outlined in Figure 1. The initially uniform distribution $g(m, 0)$ (inset (a)) evolves towards a Dirac δ function (inset (d)) as the network approaches global synchronization, i.e. global cascade. So far, there is not a closed analytical expression for the consecutive composition of the function $g(m, t)$ after an arbitrary number of time steps to reveal the evolution of $\rho(t)$. Nonetheless, it can be solved numerically, and Eq. 4 reduces to

$$\rho(t)(\langle k^2 \rangle - \langle k \rangle) = \langle k \rangle. \quad (5)$$

The cascade condition is then

$$\frac{\rho(t)}{1 + \rho(t)} = \frac{\langle k \rangle}{\langle k^2 \rangle}, \quad (6)$$

which exactly corresponds to the bond percolation critical point on uncorrelated networks [36–38]. For the case of random Poisson networks $\langle k^2 \rangle \sim \langle k \rangle^2$, then

$$\frac{\rho(t)}{1 + \rho(t)} = \frac{1}{\langle k \rangle} \quad (7)$$

This result, as exposed in [30], represents an advance in our understanding of non-linear, pulse-coupled dynamics, regardless of our (social) interpretation, and could be further developed in different settings beyond a minimal MRA (see next section).

5 Empirical evidence

So far, we have reviewed a minimal MRA as a generalization of networked threshold models, inasmuch MRA include socially relevant mechanisms of the latter –complex contagion, mainly–, but incorporate as well additional desirable ingredients: temporal evolution, endogenous and exogenous motivation, etc. The validity and convenience of MRA is however subject to its capacity to

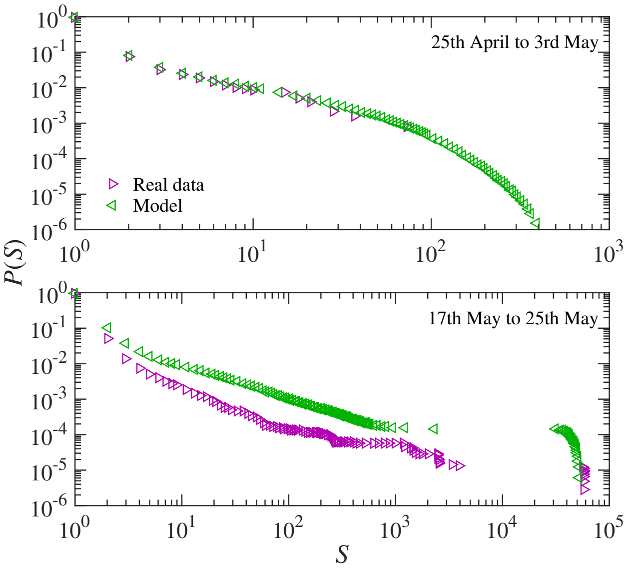

Fig. 5 Cascade size cumulative distributions $P(S)$ of real data (green triangles) and the model counterpart (purple triangles). We have considered two time windows which significantly differ: first eight days (top) for which we have set $\omega = 0.1$; last eight days (bottom) for which we have $\omega = 30.0$. Note that $\epsilon_c \approx 10^{-3}$. The model performs well in both periods, the relative error of the slope in the linear region is $< 1\%$ (see slope values on the figure). Real data distributions are measured taking into account the definition of behavior-based activity cascades, see also [17, 24]. Adapted from [30].

reproduce and explain observed phenomena. As a first step, the model should be able to reproduce observed cascades as they happened at different times during a single, unfolding event. In Figure 5 we see an example of this. These results correspond to a dataset of the civil protests occurred in Spain ("15M movement") that resonated on Twitter, in the period April-May 2011 [27]. The collection comprises 521,707 tweets produced by $N = 87,569$ users, and the observation period ranges from the 25th of April at 00:03:26 to the 26th of May at 23:59:55, 2011.

Figure 5 compares empirical *vs.* synthetic cascade size distributions for different periods of the protests: a 'slow-growth" phase (the "brewing period" of the mobilization, 25th April to 3rd May; green triangles in the upper panel), when the protest is limited to some online activists; and an "explosive" phase (19th to 25th May; green triangles in the lower panel), which comprehends the most active interval. MRA dynamics run here on the empirical following-follower topology for different ω values, with remarkable success (purple triangles), though the bottom panel does not show so good of an agreement as the top one. Remarkably, the values that ω takes in simulations are not chosen arbitrarily as a best fit to empirical data, but are inversely proportional to the "level of excitation" of the system, quantified by the inter-event time distribution, see [30] for details.

Notably, MRA will be a useful and predictive tool as long as it focuses on this type of unfolding events, see [39] for collections along these lines. Political mobilizations, elapsed online discussions, pre-announced events are examples of MRA-like social activity, while unexpected or extremely short cascades fit better other modelling approaches.The accumulation of empirical evidence for MRA models is one of the biggest challenges for scientists interested in unveiling the dynamical evolution of information cascades, starting from an event or from multiple events. Some recent studies point out the usefulness of focussing on resharing [40], or in the dynamic changes in the topology from information cascades [41], and several others, see [42] for a review.

6 Conclusions and outlook

Our understanding of social spreading in networked systems has been shifting paradigms over the last two decades, as new challenges were raised at an accelerating rate. First, network scientists acknowledged that complex contagion –which had been brought forward much earlier in sociology– was a better mechanism to explain many of the spreading phenomena occurring on social networks. This translated into an update of the threshold model, followed by a decade of analytical developments and variations of the model itself, in a quest for better insights and stronger accommodation to empirical facts. Similarly, the advent of new communication platforms has provided evidence that a large fraction of activity calls for another twist –one that retains the desirable features of the threshold model, but incorporates naturally the unfolding nature of recurrent activity.

We claim here that a good starting point to address this new challenge comes under the form of a model for recurrent activation (MRA) on social systems. In this line, we have reviewed a set of works which bring about a minimal framework, which we understand as a starting point rather than a destination itself. This minimal model already contains many convenient features –complex contagion, endogenous and exogenous mechanisms that trigger activity, recurrence, temporal evolution, etc.–, and some evidence that it captures some empirical observations. We wish, however, to outline here some open challenges that lay ahead.

On one hand, there are many aspects of MRA that call for a better understanding. The observation of the existence of critical values in the system (either ϵ_c or ω_c) suggests that the model has room for analytical development –which, as a side-effect, points at a functional inter-dependency between ϵ and ω. This would have an impact on the study of which conditions need to be met so as to observe global cascades, be them on real topologies or synthetic ones (small world, scale-free, etc.). Similarly, it remains a relevant question to understand how synchronization builds up in this MRA –the path to synchronization [43]– under different scenarios (and specially in modular networks, which abound in empirical settings).

On the other hand, the minimal configuration outlined here is insufficient to attain sociological relevance. If only numerically, MRA has already been studied under different heterogeneous contexts (e.g. ω_i vs. homogeneous ω), see [44]. This line of increasing sophistication needs to continue, introducing further enrichments that bring MRA from its promising minimal version, towards a more realistic one. This includes perhaps the introduction of memory, or the inclusion of refractory periods –mimicking aspects such as fatigue.

7 Acknowledgements

AA acknowledges financial support from MINECO (projects FIS2015-71582-C2), from the ICREA Academia and the James S. McDonnell Foundation.

References

1. H.W. Hethcote, SIAM review **42**(4), 599 (2000)
2. Y. Moreno, M. Nekovee, A. Pacheco, Physical Review E **69**(6), 066130 (2004)
3. D. Centola, M. Macy, American Journal of Sociology **113**(3), 702 (2007)
4. M. Granovetter, American Journal of Sociology pp. 1420–1443 (1978)
5. T.C. Schelling, *Micromotives and Macrobehavior* (Norton, London, 1978)
6. M. Granovetter, R. Soong, Journal of Mathematical sociology **9**(3), 165 (1983)
7. T.W. Valente, Social networks **18**(1), 69 (1996)
8. D. Watts, Proceedings of the National Academy of Sciences **99**(9), 5766 (2002)
9. P.S. Dodds, D.J. Watts, Physical Review Letters **92**(21), 218701 (2004)
10. R. Pastor-Satorras, A. Vespignani, Physical Review Letters **86**(14), 3200 (2001)
11. J. Gleeson, D. Cahalane, Physical Review E **75**(5), 056103 (2007)
12. A. Galstyan, P. Cohen, Physical Review E **75**(3), 036109 (2007)
13. J. Gleeson, Physical Review E **77**(4), 046117 (2008)
14. J. Gleeson, D. Cellai, J.P. Onnela, M.A. Porter, F. Reed-Tsochas, Proceedings of the National Academy of Sciences **111**(29), 10411 (2014)
15. D. Guilbeault, J. Becker, D. Centola, arXiv preprint arXiv:1710.07606 (2017)
16. J. Borge-Holthoefer, Y. Moreno, T. Yasseri (eds.), *At the Crossroads: Lessons and Challenges in Computational Social Science At the Crossroads: Lessons and Challenges in Computational Social Science* (Frontiers Media, 2016)
17. S. González-Bailón, J. Borge-Holthoefer, A. Rivero, Y. Moreno, Scientific Reports **1**, 197 (2011)
18. D. Liben-Nowell, J. Kleinberg, Proceedings of the National Academy of Sciences **105**(12), 4633 (2008)
19. E. Bakshy, J. Hofman, W. Mason, D. Watts, in *Proceedings of the 4th ACM International Conference on Web search and Data Mining* (ACM, 2011), pp. 65–74
20. K. Lerman, R. Ghosh, T. Surachawala, arXiv:1202.3162 (2012)
21. H. Kwak, C. Lee, H. Park, S. Moon, in *Proceedings of the 19th International Conference on World Wide Web* (ACM, 2010), pp. 591–600
22. E. Sun, I. Rosenn, C. Marlow, T. Lento, Proc. ICWSM **9** (2009)
23. G. Kossinets, J. Kleinberg, D. Watts, in *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining* (ACM, 2008), pp. 435–443
24. J. Borge-Holthoefer, A. Rivero, Y. Moreno, Physical Review E **85**, 066123 (2012)
25. R. Baños, J. Borge-Holthoefer, Y. Moreno, arXiv:1303.4629 (2013)
26. A. Hackett, S. Melnik, J. Gleeson, Physical Review E **83**(5), 056107 (2011)

27. J. Borge-Holthoefer, A. Rivero, I. García, E. Cauhé, A. Ferrer, D. Ferrer, D. Francos, D. Iñiguez, M. Pérez, G. Ruiz, et al., PloS One **6**(8), e23883 (2011)
28. P. Gai, S. Kapadia, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **466**(2120), 2401 (2010).
29. T.R. Hurd, J.P. Gleeson, Journal of Complex Networks **1**(1), 25 (2013).
30. P. Piedrahita, J. Borge-Holthoefer, Y. Moreno, A. Arenas, EPL (Europhysics Letters) **104**(4), 48004 (2013)
31. R. Mirollo, S. Strogatz, SIAM Journal on Applied Mathematics **50**, 1645 (1990)
32. A. Corral, Pérez-Vicente, Díaz-Guilera, A. Arenas, Physical Review Letters **74**, 118 (1995)
33. M. Timme, F. Wolf, T. Geisel, Physical Review Letters **89**(25), 258701 (2002)
34. A. Roxin, H. Riecke, S.A. Solla, Physical Review Letters **92**, 198101+ (2004)
35. S.A. Golder, M.W. Macy, Science **333**(6051), 1878 (2011)
36. M. Molloy, B. Reed, Random structures & algorithms **6**(2-3), 161 (1995)
37. M. Molloy, B. Reed, Combinatorics probability and computing **7**(3), 295 (1998)
38. S. Dorogovtsev, A. Goltsev, J. Mendes, Rev. Mod. Phys. **80**(4), 1275 (2008)
39. J. Borge-Holthoefer, N. Perra, B. Gonçalves, S. González-Bailón, A. Arenas, Y. Moreno, A. Vespignani, Science advances **2**(4), e1501158 (2016)
40. K. Subbian, B.A. Prakash, L. Adamic, in *Proceedings of the 26th International Conference on World Wide Web* (International World Wide Web Conferences Steering Committee, 2017), WWW '17, pp. 597–605
41. B. Baingana, G.B. Giannakis, IEEE Trans. Signal Processing **65**(4), 985 (2017)
42. M. Jalili, M. Perc, Journal of Complex Networks **5**(5), 665 (2017)
43. J. Gómez-Gardenes, Y. Moreno, A. Arenas, Physical Review Letters **98**(3), 034101 (2007)
44. P. Piedrahita, J. Borge-Holthoefer, Y. Moreno, S. González-Bailón, Social Networks (in press) (2017)