



Dose–response functions and surrogate models for exploring social contagion in the Copenhagen Networks Study

Jonathan F. Donges^{1,2,a} , Jakob H. Lochner^{1,3} , Niklas H. Kitzmann^{1,4} , Jobst Heitzig¹ , Sune Lehmann^{5,6} , Marc Wiedermann^{1,7,8}, and Jürgen Vollmer³ 

¹ Earth System Analysis and Complexity Science, Potsdam Institute for Climate Impact Research, Member of the Leibniz Association, Potsdam, Germany

² Stockholm Resilience Centre, Stockholm University, Stockholm, Sweden

³ Institute for Theoretical Physics, University of Leipzig, Leipzig, Germany

⁴ Institute for Physics and Astronomy, University of Potsdam, Potsdam, Germany

⁵ Department of Applied Mathematics and Computer Science, Technical University of Denmark, Lyngby, Denmark

⁶ Center for Social Data Science, University of Copenhagen, Copenhagen, Denmark

⁷ Robert Koch-Institut, Berlin, Germany

⁸ Institute for Theoretical Biology, Humboldt University of Berlin, Berlin, Germany

Received 15 March 2021 / Accepted 3 September 2021 / Published online 1 October 2021
© The Author(s) 2021

Abstract Spreading dynamics and complex contagion processes on networks are important mechanisms underlying the emergence of critical transitions, tipping points and other non-linear phenomena in complex human and natural systems. Increasing amounts of temporal network data are now becoming available to study such spreading processes of behaviours, opinions, ideas, diseases and innovations to test hypotheses regarding their specific properties. To this end, we here present a methodology based on dose–response functions and hypothesis testing using surrogate data models that randomise most aspects of the empirical data while conserving certain structures relevant to contagion, group or homophily dynamics. We demonstrate this methodology for synthetic temporal network data of spreading processes generated by the adaptive voter model. Furthermore, we apply it to empirical temporal network data from the Copenhagen Networks Study. This data set provides a physically-close-contact network between several hundreds of university students participating in the study over the course of 3 months. We study the potential spreading dynamics of the health-related behaviour “regularly going to the fitness studio” on this network. Based on a hierarchy of surrogate data models, we find that our method neither provides significant evidence for an influence of a dose–response-type network spreading process in this data set, nor significant evidence for homophily. The empirical dynamics in exercise behaviour are likely better described by individual features such as the disposition towards the behaviour, and the persistence to maintain it, as well as external influences affecting the whole group, and the non-trivial network structure. The proposed methodology is generic and promising also for applications to other temporal network data sets and traits of interest.

1 Introduction

Spreading and complex contagion processes shape the dynamics of diverse complex ecological, societal and technological systems studied in many fields of research [1–3]. Examples include biological infections [4, 5] such as the spreading of the COVID-19 pandemic [6]; cascading failures in interdependent infrastructure systems [7]; diffusion of innovations and technologies [8–10]; evolutionary processes [11, 12]; social norms [13], behaviours [14], and other social, political and tech-

nological innovations relevant for sustainability transition and rapid decarbonisation [15–18]; political changes [19]; or religious missionary work [20, 21]. These spreading processes on complex networks often give rise to non-linear dynamics and the emergence of macroscopic phenomena, such as phase transitions and tipping points that separate qualitatively different dynamical regimes [22]; for example, a transition between regimes where a local infection or innovation is locally contained, and those where it spreads globally to a large part of the network [1, 2, 10, 23, 24]. Furthermore, spreading processes can interact with the underlying complex network structures, e.g. through the process of homophily, giving rise to complex coevolutionary feedbacks between dynamics on and structure of these

Jonathan F. Donges and Jakob H. Lochner share the lead authorship.

^a e-mail: donges@pik-potsdam.de (corresponding author)

networks [25–28]. Better understanding of such complex spreading processes, based on improved methods for data analysis and modelling, is highly relevant for finding robust approaches to identify, analyse, influence or govern their dynamics. This way, harmful impacts may be avoided, or desirable outcomes reached, e.g. for containing pandemic outbreaks [6, 29, 30], preventing cascading failures in power grids [7, 31], or fostering the spreading of social-cultural-technological innovations towards a rapid sustainability transformation [15–17, 22].

In recent years, temporal network data has become more abundantly available from social media platforms such as Facebook [32] and Twitter [33], or long-term health studies such as the Framingham Heart Study [34] that have been leveraged for studying spreading and contagion processes, e.g. in the dynamics of obesity [35], smoking [36], happiness [37], loneliness [38], alcohol consumption [39], depression [40], divorce [41], emotional contagion [42] and political mobilisation [43]. So far such studies of empirical temporal network data mainly relied on standard statistical methods such as generalised linear models, generalised estimating equations or spatial autoregressive models [3]. However, these methods are typically not well equipped to deal with network dependencies [44]. Furthermore, analogous to the problem of identifying causal associations in multivariate time series data [45, 46], there are challenges in extracting possible causal effects induced by contagion processes, and in separating their imprints from other mechanisms such as homophilic rewiring of network structure, common external forcing from the system's environment and other confounding effects. After all, most studies rely on observational data and not on controlled experiments [44].

Here, we contribute to this field by developing a methodology for the analysis of complex spreading processes in temporal network data sets based on dose–response functions (DRFs) that have been used in the theoretical description of simple and complex contagion processes [2, 23]. Among others, they have been applied to the study of behavioural contagion in animal systems such as startling cascades in fish schools [47] and the spread of information on social media networks [48]. Dose–response functions encode a network nodes' probability of being infected with a new trait, given the level of exposure to this trait in its network neighbourhood. We propose an algorithm including Gaussian filtering to robustly estimate DRFs from synthetic and empirical temporal network data, including the possibility of propagating various types of uncertainties. To test for the possibility of an actual causal spreading process being involved in generating the data, and to identify confounding effects, we also develop a hierarchy of temporal network surrogate models. These models comprise a family of methods that rely on partial data randomisation to analyse specific features of (networked) processes without assuming particular underlying mechanisms and have been proven highly useful in exploratory data analyses [49, 50]. In particular, they have been used extensively to investigate temporal net-

works [51, 52], including epidemic and social contagion processes [53, 54]. A conceptually related application for surrogate models is the study of time series data [55, 56]. Here, we combine methods from both temporal network and time series surrogate models. This enables us to investigate which features and structures in the data are possibly sufficient to explain the obtained dose–response functions.

We apply our methodology to synthetic data from the adaptive voter model as a proof of concept, and to empirical observational temporal network data from the Copenhagen Networks Study. Based on the latter we analyse the spreading dynamics of the illustrative behaviour of “regularly going to the fitness studio” on a physically-close-contact network between university students participating in the study over the course of 3 months with daily time resolution. We do not find robust evidence of a causal spreading process underlying the observed dynamics. This suggests that possible social contagion effects in this context are limited, and dominated by other factors or shadowed by excessive noise. This is in agreement with findings from health behaviour psychology [57]. Hence, this first application study suggests that the proposed methodology is generic and promising for investigations of other data sets and possibly spreading traits of interest.

This paper is structured as follows: we first introduce the synthetic and empirical temporal network data sets, obtained from the adaptive voter model and the Copenhagen Networks Study, respectively (Sect. 2). In a next step, we describe the methodology developed here for data analysis, including estimating dose–response functions and generating surrogate data sets for testing hypotheses on underlying data generating processes (Sect. 3). Finally, we report results obtained for the synthetic and empirical data sets (Sect. 4), discuss these findings and conclude (Sect. 5).

2 Data

Here, we describe the data sets used in this study to test our proposed dose–response function methodology. The data has the form of temporal networks (Sect. 2.1), it includes synthetic temporal network data generated by the adaptive voter model (Sect. 2.2) and empirical temporal network data from the Copenhagen Networks Study (Sect. 2.3).

2.1 Temporal social networks

The data sets investigated in this work are structured as temporal networks $\mathcal{G}(t)$ with a fixed number of nodes N and a time-dependent set of links described by the adjacency matrix $A_{ij}(t)$, where $i, j \in \{1, \dots, N\}$ [52], sampled at discrete time steps t . In addition, node traits $o_i(t)$ are time-dependent as well, for example encoding different opinions or behaviours.

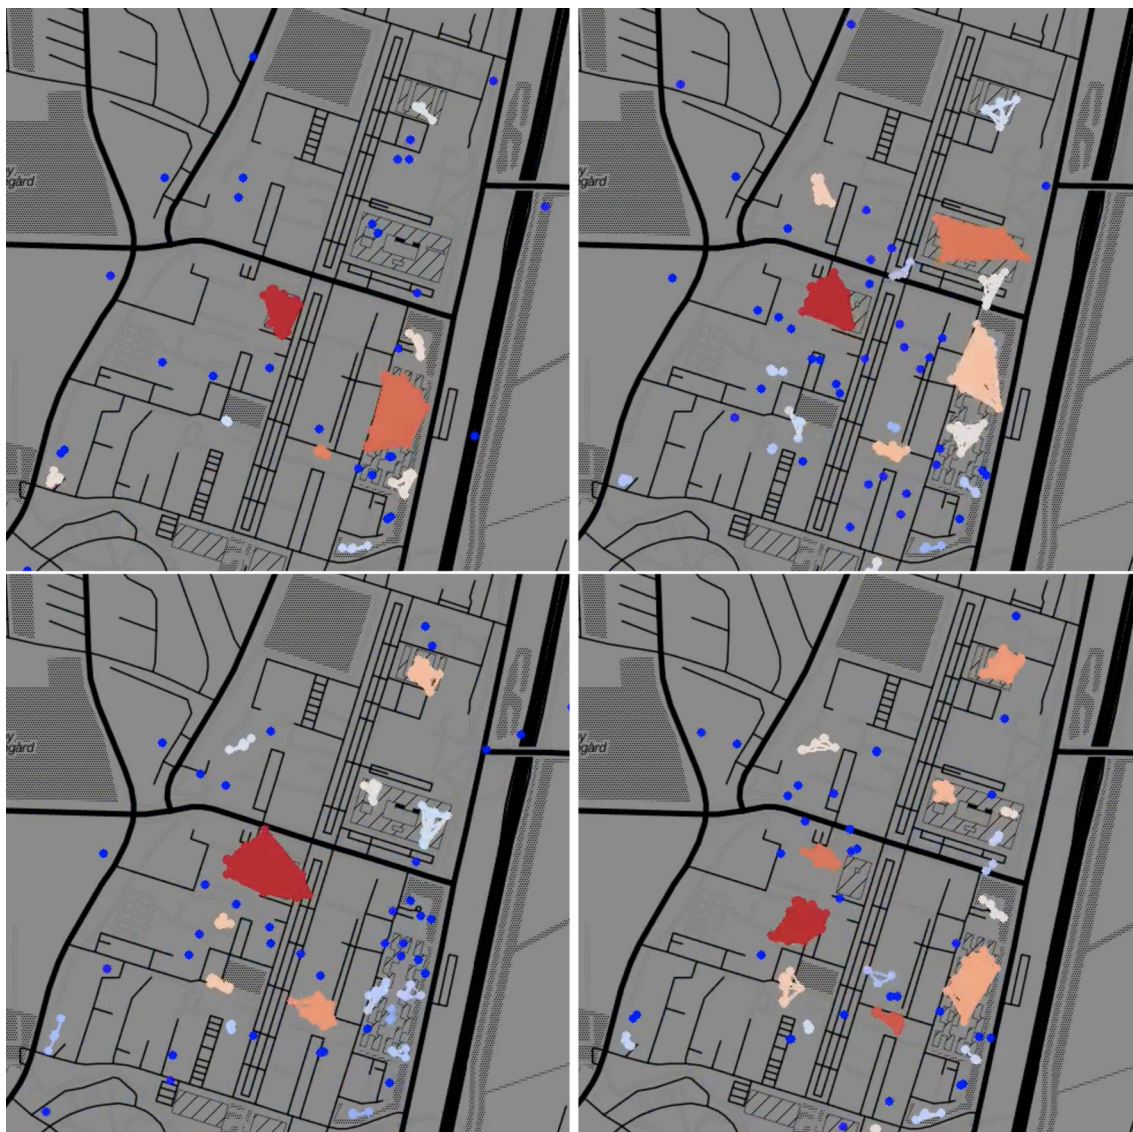


Fig. 1 Temporal network snapshots throughout a typical day during the first semester of the Copenhagen Networks Study. Each dot represents an individual, colour coded according to cluster size from single nodes (dark blue) to

large clusters (dark red). Node clusters evident in the snapshots correspond to students engaging in joint activities, such as lectures or eating lunch in a cafeteria

2.2 Synthetic temporal network data: adaptive voter model

One prototypical model of temporal network dynamics is the adaptive voter model (AVM) [25] that incorporates core processes in social systems, i.e. homophily [58] and social learning of traits [59]. As such, the AVM can be interpreted as a straightforward generalisation of the so-called voter model [60] to any prescribed initial social network topology and the ability of the represented individuals to deliberately change their neighbourhood structure. It thereby aims to explain the emergence of like-minded communities within a larger social network and the extent to which individuals (i) become like-minded because of shared social ties or (ii) form such social ties because they are like-minded.

We use an AVM to generate synthetic temporal network data that resembles the experimental data from the Copenhagen Networks Study. This choice has several motivations: first, it matches our initial hypothesis that a quasi-symmetric social learning process underlies the spread of “active” and “passive” behaviours of individuals. Under this hypothesis, individuals can equally imitate active or passive behaviour occurring in their network neighbourhood. This is in contrast to standard SI(S/R)-type models [61, 62], where only one trait spreads infectiously, and a spontaneous recovery process is assumed. Furthermore, the AVM also includes both the processes of social learning and homophilic social network rewiring that we hypothesise to be present in the empirical data. Finally, the AVM is one

of the simplest and best understood models that has these desired properties [27, 61].

Specifically, the AVM considers a temporal network $\mathcal{G}(t)$ with a fixed number of N nodes and M links. Each node v_i holds one of Γ opinions or traits o_i that are initially distributed at random among them. The M links are initially distributed uniformly at random as well, thus mimicking the configuration of an Erdős–Rényi graph. At each discrete time step t , a single node v_i with opinion or trait o_i is randomly chosen. If its degree k_i , i.e. the number of directly connected neighbours, is non-zero, either of two processes takes place:

1. *Homophilic rewiring.* With fixed probability φ , we select one of the edges that are attached to v_i and move its other end to a randomly selected node v_k that holds the same trait o_k as v_i , and is not connected to v_i yet. v_i thereby *adapts* its neighbourhood structure to align more with its own trait o_i .
2. *Social learning:* Otherwise, with fixed probability $1 - \varphi$, we pick a random neighbour v_j of v_i and set v_i 's trait equal to that of v_j , i.e. $v_i \leftarrow v_j$. Hence, v_i *imitates* the trait o_k of v_k to become more alike to its immediate neighbourhood.

The model reaches a steady state once only one trait per connected network component remains. In this case, no additional updates to the nodes' states or their neighbourhood structure are possible. The fixed probability φ is a model parameter that allows to scale the relative frequencies of imitation and adaptation events. For $\varphi = 0$, only imitation, and for $\varphi = 1$, only adaptation takes place. The model displays a phase transition at intermediate values of φ where the system's steady state qualitatively shifts from a large connected component of a single remaining trait to a fractionalized configuration of multiple disconnected components that each show distinct predominant traits [25].

In our specific study, we set the number of nodes to $N = 619$, the number of edges to $M = 5724$ and the number of traits to $\Gamma = 2$ to ensure consistency with the (filtered) empirical data from the Copenhagen Networks Study (CNS), see below.

2.3 Empirical temporal network data: Copenhagen Networks Study

In the following, we present the Copenhagen Networks Study as our main empirical data source (Sect. 2.3.1) and describe the methodology used for extracting a temporal social network with time-dependent node traits from this data set (Sect. 2.3.2).

2.3.1 Description of data sources

The data analysed here originates from the Copenhagen Networks Study (CNS) [63, 64]. CNS was carried out from 2012–2016 and focussed on collecting temporal network and demographic data on a densely interconnected cohort of nearly 1000 individuals. To collect

the temporal network information, the study handed out state-of-the-art smartphones to consenting freshman students at the Technical University of Denmark. Specifically the study collected information on networks of physical proximity (using Bluetooth signals), phone calls, text messages, and online social networks. In addition to the network data, the study also collected information on the participants' mobility, using the phones' GPS sensors—and demographic and personality data, using questionnaires. The study was approved by the Danish Data Protection agency, the appropriate legal entity in Denmark. In terms of research, data from CNS have been used in a number of contexts e.g. epidemiology [65–67], mobility research [68, 69], network science [70, 71], studies of gender-related behaviour [72], and education research [73, 74].

In addition to the data from the Copenhagen Networks Study, and in view of our aim to investigate the illustrative behaviour “regularly going to the fitness studio”, a data set was generated with the locations of fitness studios in the vicinity of Copenhagen. The studios were selected from the locations provided by Open Street Map [75] and listed with the keys ‘leisure=fitness_center’ or ‘sport=fitness’. A comprehensive list of all considered studios can be found in Appendix C.

2.3.2 Generation of empirical temporal social network

The empirical temporal social network is generated as a physically-close-contact network between the study's participants. A network edge is created when two participants are in close proximity to each other once during day t . The network's adjacency matrix $A_{ij}(t)$ is then defined as

$$A_{ij}(t) = \begin{cases} 1, & |s_{ij}(t)| > 80 \text{ dBm} \\ 0, & \text{otherwise} \end{cases}, \quad (1)$$

where time t is in units of days and $s_{ij}(t)$ is the maximum Bluetooth signal strength between participants i and j measured during day t , while measurements were performed every five minutes. The threshold 80 dBm corresponds to a distance of about 2 m and maximises the ratio of social interactions to transient and unimportant connections [76].

To minimise noise from the beginning and end periods of data collection, i.e. noise due to participants joining late or dropping out early, in this study we focus on the period from the first of February 2014 to the end of April 2014, which corresponds to the spring semester and is in the middle of the “SensibleDTU 2013” data collection, the second deployment of CNS.

Much of human behaviour proceeds in weekly cycles [77]. To account for this periodicity in the data, we define a time window $T(t, t')$ using a Gaussian kernel:

$$T(t, t') = e^{-(t-t')^2/(2t^2)}, \quad (2)$$

$$X(t) = \sum_{t'=0}^t x(t') \cdot T(t, t'), \quad (3)$$

where $t_c = 7$ days is the characteristic time. Equation 3 illustrates how $T(t, t')$ is functioning as temporal weight in a sum over an arbitrary time-dependent variable $x(t')$. We suppose that $t_c = 7$ days introduces the least additional assumptions as it coincides with the typical seven-day rhythm of study, work, leisure and exercise activities and behaviours (e.g. a university student would attend a particular lecture at a particular day of the week, visit the fitness studio on another particular day etc.). The Gaussian kernel is a preferable choice to a rectangular kernel, as the latter can produce artefacts due to discontinuities. It is also a preferable choice to an exponential kernel because it decreases slowly for $t - t' < t_c$ and then tends to zero quickly. In contrast, an exponential kernel quickly falls towards zero and is, therefore, not suitable for a time window that represents typical horizons of human short-term activity.

The raw data contain students with no or fluctuating social interaction. Reasons might be that they have left campus or spend time with people not participating in the study. To minimise their influence onto this study's results, two filters were applied to the data. The first sorts out participants who had no or very few contacts over the whole study period by setting a lower limit for the average degree $\bar{k}_i \geq k_{\min} = 4$. Variations of k_{\min} in the interval $1 \leq k_{\min} \leq 5$ were tested, and showed no significant influence on this study's results. The second filter compensates for the fluctuating contact behaviour of the participants. Some participants have a regular number of contacts on average, but occasionally this number drops to only a few or no contacts (e.g. illness could be a plausible explanation). These absences could confound the results of the study. Therefore, we only consider students who had at least one contact in the last week. For this purpose, the participants were filtered according to their average node degree in the past week:

$$\tilde{k}_i(t) = \frac{\sum_{t'=0}^t k_i(t') \cdot T(t, t')}{\sum_{t'=0}^t T(t, t')}. \quad (4)$$

Here, k_i is the node degree and $T(t, t')$ is the time window defined in Eq. 2. We, therefore, interpret $\tilde{k}_i(t)$ as the average number of daily contact events in past week, and we consider only students in our analysis that had in the order of one contact in the last week, i.e. we set the lower bound to $\tilde{k}_i(t) \geq \tilde{k}_{\min} = 1/7$. Variations of \tilde{k}_{\min} in the interval $1/7 \leq \tilde{k}_{\min} \leq 1$ were tested, and showed no significant influence on this study's results.

To investigate possible spreading dynamics of the illustrative behaviour “regularly going to the fitness studio”, we match stop-locations with the locations of fitness studios (Appendix C). Here, stop-locations are coordinates generated from the GPS data, where the participants spent at least 15 min [78]. The accuracy chosen for matching is 10 m, which corresponds to the

precision of GPS [79]. Hence, we record for each node i at the time t the behaviour:

$$b_i(t) = \begin{cases} 1, & \text{if node } i \text{ visited a studio at day } t \\ 0, & \text{otherwise} \end{cases}. \quad (5)$$

To distinguish between students who go to the studio occasionally and students who go regularly, we introduce the past-week behaviour:

$$\bar{b}_i(t) = \sum_{t'=0}^t b_i(t') \cdot T(t, t'), \quad (6)$$

with $T(t, t')$ the 1-week time window defined in Eq. 2. We interpret $\bar{b}_i(t)$ as typical behaviour during the last week.

Finally, for each point in time t , we split the participants into two groups: (i) students going occasionally or not at all to the fitness studio, and (ii) students going more often to the studio. A typical behaviour of regularly going into the fitness studio would be to go once a week. This suggests to select $\bar{b}_i(t) = 1$ as a threshold criterion, and to explore the following time-dependent trait $o_i(t)$ for each node in the network:

$$o_i(t) = \begin{cases} 1, & \bar{b}_i(t) \geq 1 \\ 0, & \text{otherwise} \end{cases}. \quad (7)$$

Indeed, there is a clear boundary in the cumulative distribution of $\bar{b}(t)$ plotted in Fig. 2 for $\bar{b}(t) \approx 1$ and for all t . The boundary indicates that $\bar{b}(t) > 1$ is occurring less frequently than $\bar{b}(t) < 1$. This supports the choice to separate participants with the threshold $\bar{b}(t) = 1$. In the following, the students going to gyms at least once in the last week ($o_i(t) = 1$) are referred to as “active” nodes, while the others ($o_i(t) = 0$) are referred to as “passive” nodes.

The procedure presented here generates a social network consisting of 619 nodes with an average degree of $\bar{k}_i = 19$. The nodes change their trait $o_i(t)$ on average 5.94 times over the course of the considered 3-month period.

3 Methods

In this section, we describe the methodologies used to estimate empirical dose–response functions from temporal network data (Sect. 3.1) and for generating surrogate data sets to test hypothesis on the processes and structures underlying specific features of the empirical dose–response functions (Sect. 3.2).

3.1 Estimating dose–response functions from temporal network data

Dose–response functions (DRFs) represent the functional dependence between the probability of changing

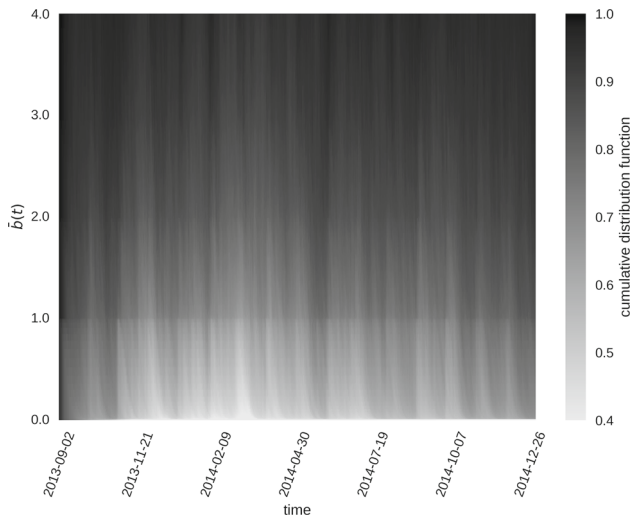


Fig. 2 Cumulative distribution of the past-week behavioural function $\bar{b}(t)$ plotted as a heat map over the period of the entire “SensibleDTU 2013” data collection. Our study analyses the 3-month subperiod from February to April 2014. A clear boundary is visible at $\bar{b}(t) \approx 1$ for all t , with values of $\bar{b}(t) > 1$ being much less frequent than $\bar{b}(t) < 1$. Therefore, $\bar{b}(t) = 1$ is a reasonable choice to separate the participants into two groups. Members of the group with $\bar{b}(t) \geq 1$, who visited the fitness studio at least once in the past week are referred to as active nodes, while individuals with $\bar{b}(t) < 1$ are referred to as passive nodes

a trait $p_{o \rightarrow o'}$ and the exposure K , which is defined as the joint influence of all contacts with a given trait, or more formally as the superposition of all received doses from neighbouring nodes. We assume that the influence of each node is equal and that the recent influence from the last week has a greater impact on the decision-making process than the influence from the distant past, i.e. it contributes more to the exposure K . To measure the exposure to which a single node i is subjected, we put

$$K_i(o, t) = \sum_{t'=0}^t \mathcal{N}_i(o, t') \cdot T(t, t'), \tag{8}$$

where $\mathcal{N}_i(o, t')$ is the number of neighbouring nodes with trait o at time t' and $T(t, t')$ is the weight of the encounter as defined in Eq. 2, which down-weights the influences from encounters from further back than 1 week.

From the time series of each node’s traits $o_i(t)$, the received exposures $K_i(o, t)$ can be computed, allowing us to estimate the DRFs as relative frequencies as

$$p_{o \rightarrow o'}(K) \approx \frac{C(K)}{N(K)}. \tag{9}$$

Here, $C(K)$ is the number of nodes that have changed their trait between $t - 1$ and t and having experienced a certain level of exposure K . Furthermore, $N(K)$ is the

total number of nodes that have experienced exposure level K . $C(K)$ and $N(K)$ are the result of an aggregation over all time steps and are thus time-independent.

$p(K)$ is an estimator of the actual probability of changing trait when experiencing an exposure level of K . If the reactions (changing trait or not) to subsequent exposures are assumed to be independent, this estimator is simply the empirical success rate of an $N(K)$ times repeated Bernoulli experiment, and its standard error can thus be estimated by

$$\sigma_p = \sqrt{\frac{p(K)(1 - p(K))}{N(K)}} = \sqrt{\frac{C(K)(N(K) - C(K))}{N(K)^3}}. \tag{10}$$

In the present study, we adopt

$$\sigma_p^c = \sqrt{C(K)(N(K) + C(K))/N(K)^3} \tag{11}$$

as a conservative upper bound to this error. Where multiple data sets are used for one result, as is the case when multiple simulation runs or surrogate model realisations are computed using the same parameters, the data are considered as one ensemble for further analysis. The error estimation in Eq. 11 is thus performed on these pooled data sets where applicable.

3.2 Generating surrogate data sets for hypothesis testing

To probe the empirical data from the Copenhagen Networks Study for contagion effects relating to the studied behaviour, we use the method of surrogate data sets. The surrogate data approach is a statistical method for identifying non-linearity, such as contagion effects, in time series. This is achieved by performing hypothesis tests on data sets that are generated from the empirical data by using Monte Carlo methods [51, 52, 55, 56]. Surrogate data sets have been used in the past to study a wide range of time series [80–82] and network data [83–85]. The method is described in the following paragraph, followed by the description of the surrogate data studies examined in the present contribution.

First, a class of processes that may potentially be sufficient in explaining the empirical data, is specified as a composite null hypothesis \mathcal{H}_0 . To test this hypothesis, a new, “surrogate” data set is derived from the empirical data in a way that is consistent with \mathcal{H}_0 . Any structures that the null hypothesis excludes are destroyed in this process, while other features of the original data are retained.

One algorithm which can be used to produce such surrogate data sets is the creation of random permutations of the original data, for example by permuting the nodes’ time series or network connections. The product resembles the empirical data, but lacks the features excluded by the null hypothesis, such as contagion processes. This method, known as Constrained Realisations [86], represents a parameter-free way of producing surrogate data sets without the use of a specific model.

A discriminating statistic is then computed on the original data and surrogate data sets alike. If there is a significant difference between the value or distribution computed for the original data, and the ensemble of values or distributions computed for the surrogate data sets, the null hypothesis is rejected. Put simply, the empirical data are permuted in a way that is consistent with a composite null hypothesis, and if this substantially changes a statistical measure of interest, the null hypothesis can be rejected. Through the careful choice of iteratively more complex null hypotheses, preserving different sets of data properties, the nature of the true underlying non-linear process can be investigated.

Six surrogate data sets are produced for this analysis. The first four investigate the influence of different assumptions about the node dynamics on the dose–response functions, by permuting the node traits $o_i(t)$ and keeping the network component $A_{ij}(t)$ unchanged. The last two surrogate models address the effect of the network component, by permuting the network edges $A_{ij}(t)$ and keeping the node dynamics $o_i(t)$ unchanged. An overview of the investigated null hypotheses is displayed in Fig. 8B. In this figure, arrows from a surrogate test at a higher to one at lower location indicate a higher degree of randomisation in the former than in the latter. This illustrates the hierarchical nature of surrogate randomisation models. To describe the surrogate data sets P associated with the null hypotheses \mathcal{H}_0 , the canonical naming convention from [51] is used. This convention is based on defining surrogate data sets by the quantities they conserve with respect to the original data. In the following, the estimated DRF of the empirical data is referred to as the empirical DRF $p_{o \rightarrow o'}$, while the one estimated for surrogate data may be referred to as the surrogate DRF $\tilde{p}_{o \rightarrow o'}$. To reduce statistical uncertainties, ten surrogate data realisations are performed for each null hypothesis. They are considered as one ensemble to compute the dose–response functions and their error bars. The following surrogate data test were conducted:

1. $\mathcal{H}_0^1: P(A_{ij}(t), O)$. *The empirical DRF can be reproduced with a class of models that is based only on the global mean activity level $O = \overline{\langle o_i(t) \rangle_i}$. Here, the overline and brackets represent the time and ensemble average, respectively. This null hypothesis represents the most basic assumption, corresponding to an underlying process that is completely random. For this surrogate data set, all traits $o_i(t)$ are permuted randomly. Only the average activity level across the entire ensemble and observation period is conserved.*
2. $\mathcal{H}_0^2: P(A_{ij}(t), O_i)$. *The empirical DRF can be reproduced with a class of models that is based only on each node's individual activity level $O_i = \overline{o_i(t)}$. This null hypothesis leaves room for an activity factor unique to each individual node, while still assuming otherwise random node dynamics. For the corresponding surrogate data set, the activity levels are permuted in time, separately for each node.*
3. $\mathcal{H}_0^3: P(A_{ij}(t), \{\tau_{i;0,1}\})$. *The empirical DRF can be reproduced with a class of models that is based only on the distribution of time intervals for which the node stays in either activity state $\tau_{i;0,1}$, which implicitly conserves O_i and the number of activity level switches as well. This null hypothesis builds on the previous one by also conserving each node's overall persistence, defined as the inverse of a node's number of switches between behaviours, and the corresponding distribution of time intervals. This is realised by permuting the length of intervals with a constant activity level, separately for periods of active and passive behaviour, for each node. E.g. the sequence (active for 2 steps, inactive for 5 steps, active for 3 steps, inactive for one step) may be turned into (active for 3 steps, inactive for one step, active for 2 steps, inactive for 5 steps). The number of activity level switches is a constraint on the randomisation space for this surrogate model. However, the average number of activity level switches allows for sufficient randomisation in our data (see Appendix B).*
4. $\mathcal{H}_0^4: P(A_{ij}(t), O(t))$. *The empirical DRF can be reproduced with a class of models that is based only on the mean time-dependent activity level $O(t) = \langle o_i(t) \rangle_i$ of the ensemble. This null hypothesis assumes a non-stationary temporal dynamics of the ensemble's behaviour, while excluding any non-random individual node characteristics. The surrogate data set is produced by permuting the activity states of all nodes, separately for each time step.*
5. $\mathcal{H}_0^5: P(A, O_i(t))$. *The empirical DRF can be reproduced with a class of models that is based only on individual activity dynamics and the average network edge density $A = \overline{\langle A_{ij}(t) \rangle_{i,j}}$. In this case, the null hypothesis contains the assumption that the observed DRF is independent of the specific topology of the connection network, and arise solely based on the individual nodes' behaviour. The corresponding surrogate data set is produced by randomly permuting all edges across nodes and time.*
6. $\mathcal{H}_0^6: P(k_i(t), O_i(t))$. *The empirical DRF can be reproduced with a class of models that is based only on the individual node dynamics, and each node's time-dependent network degree $k_i(t) = \sum_{j=0}^N A_{ij}(t)$. This null hypothesis builds on the previous one by randomising the neighbourhood of the nodes, but preserving each nodes connectivity in the network. This can serve as a check for homophilic effects in the network dynamics. To produce the surrogate data set, we use the random link switching algorithm [87, 88]. Pairs of connections (i, j) and (k, l) are drawn randomly, and are transformed into the connections (i, k) and (j, l) . This procedure ensures that each node's degree remains unchanged.*

We choose the dose–response function, introduced in Sect. 3.1, as the discriminating statistic used to compare empirical and surrogate data sets. The comparisons of surrogate DRFs $\tilde{p}_{p \rightarrow a}$ and empirical $p_{p \rightarrow a}$

DRFs are presented in Sect. 4.2. To test our methodology, we also create the hierarchy of surrogate models for the synthetic AVM data with realistic parameter choices (see Appendix (A)). To quantify the difference between $\tilde{p}_{p \rightarrow a}$ and $p_{p \rightarrow a}$, we use a test statistic ζ that combines the k many individual z -scores (denoted as $z_i, i = 1, \dots, k$) of the DRFs into a single score similar to Stouffer's z -score method [89,90], but using the sum of squared z -scores instead of their simple sum so that negative and positive deviations cannot cancel out. Since under the null hypothesis, that sum has a χ^2 -distribution with k degrees of freedom, which depends in a non-trivial way on k , we additionally normalise the sum of squares by dividing it by the 95th percentile of that distribution, so that a value of $\zeta \geq 1$ indicates a significant deviation from the null hypothesis:

$$\zeta = \sum_{i=1}^k z_i^2 / Q_{0.95}(\chi_k^2). \quad (12)$$

4 Results

Here, we report on the results obtained by applying our proposed dose–response function methodology. As a first step, we analyse synthetic data generated by the adaptive voter model as a proof of concept (Sect. 4.1). Building on these insights, we then investigate the empirical temporal network data obtained from the Copenhagen Networks Study (Sect. 4.2). Our findings are summarised in Sect. 4.3.

4.1 Synthetic data

As a first application of our methodology, we analyse synthetic temporal network data generated by the adaptive voter model (Sect. 2.2). Figure 3 shows the estimated DRFs for the AVM with $\varphi = 0$ (green dots), which includes only imitation dynamics, and with $\varphi = 0.6$ (blue crosses), involving both imitation and homophily dynamics. Two cases are simulated: In Fig. 3A, model parameters are chosen to align the average frequency of behaviour switches across the system, and the number of time steps, with the data from the CNS study. To display the effects of more progressed network adaptation, Fig. 3B displays the DRF of a similar simulation, where the model updates per time step, and the total number of simulated time steps, are significantly increased. Each plot contains data from ten independent model runs. The probabilities for the change of trait $p_{o \rightarrow o'}$ are generated for equally sized bins with a width of $K = 2$. Only bins with at least 30 data points were considered. For increasing K , the DRF $p_{o \rightarrow o'}$ is subject to increasing uncertainties, since exposures $K > 30$ are very rare in the network.

As suggested by the imitation rule in the model, we observe that $p_{o \rightarrow o'}$ depends monotonically, but

non-linearly, on K . Moreover, the plots for $\varphi = 0.6$ clearly show the impact on $p_{o \rightarrow o'}(K)$ of the additional homophily compared to the plot of $\varphi = 0$. For $K \gtrsim 15$, the DRF of these data is significantly larger than for those with $\varphi = 0$. For $K \gtrsim 30$, the difference between the DRFs is obscured by the increasing errors in case A, but it is still clearly showing for the longer simulations in panel B.

From this first proof of concept application, we can conclude that contagion dynamics such as the imitation rule in the model [2,23] leads to positive correlation of $p_{o \rightarrow o'}$ and K . However, from the estimated DRF for $\varphi = 0.6$, we learn that homophily is reflected in the DRFs as well. To distinguish between the different dynamics, we use a surrogate analysis in the following investigation of the empirical temporal network data (Sect. 3.2).

To validate our data analysis methodology, we computed the complete hierarchy of surrogate models (described in Sect. 3.2) on the synthetic AVM data set with CNS-aligned parameter choices. The details of this study are given in Appendix A, while the results are summarised in Fig. 8A. In line with our expectations, we find evidence for contagion effects in both the $\varphi = 0.0$ and $\varphi = 0.6$ cases. Significant homophilic effects are only found where the network adaptation process of the AVM was active ($\varphi = 0.6$), also confirming our expectations. This demonstrates the sensitivity and appropriateness of our methodology for detecting contagion and homophily in the studied empirical data set. A detailed exposition of the approach is now given for the empirical data on the Copenhagen network study. Subsequently, the results for both the synthetic and the empirical data are discussed in Sect. 4.3.

4.2 Empirical data

In the following, we apply our methodology to empirical temporal network data from the Copenhagen Networks Study (Sect. 2.3) to investigate possible spreading dynamics of the illustrative behaviour “regularly going to the fitness studio”. The DRF $p_{o \rightarrow o'}(K)$ is estimated for equal-sized bins with a width of $K = 5$. Only bins with at least 30 data points were considered. The resulting DRFs are shown in Fig. 4.

We observe that the probabilities for becoming active $p_{p \rightarrow a}$ (Fig. 4A) and for becoming passive $p_{a \rightarrow p}$ (Fig. 4B) do not behave in a symmetric way. Since the initiation and the maintenance of an activity represent two rather distinct phases [57], this is not necessarily surprising. To test whether we observe significant monotonic relationships of $p_{p \rightarrow a}(K)$ and $p_{a \rightarrow p}(K)$ with K , we calculate Spearman's rank correlation coefficient ρ [91]. For a perfect monotonic increase (decrease), the coefficient is equal to $\rho = 1$ ($\rho = -1$), while $\rho = 0$ indicates the absence of a monotonic relationship. For $p_{a \rightarrow p}$ a slight but significant monotonic decrease can be identified with $\rho = -0.89$ and a p value of $p = 3.5 \cdot 10^{-7}$. Going to the gym more often than contacts (large K) could potentially be an incentive to maintain active behaviour

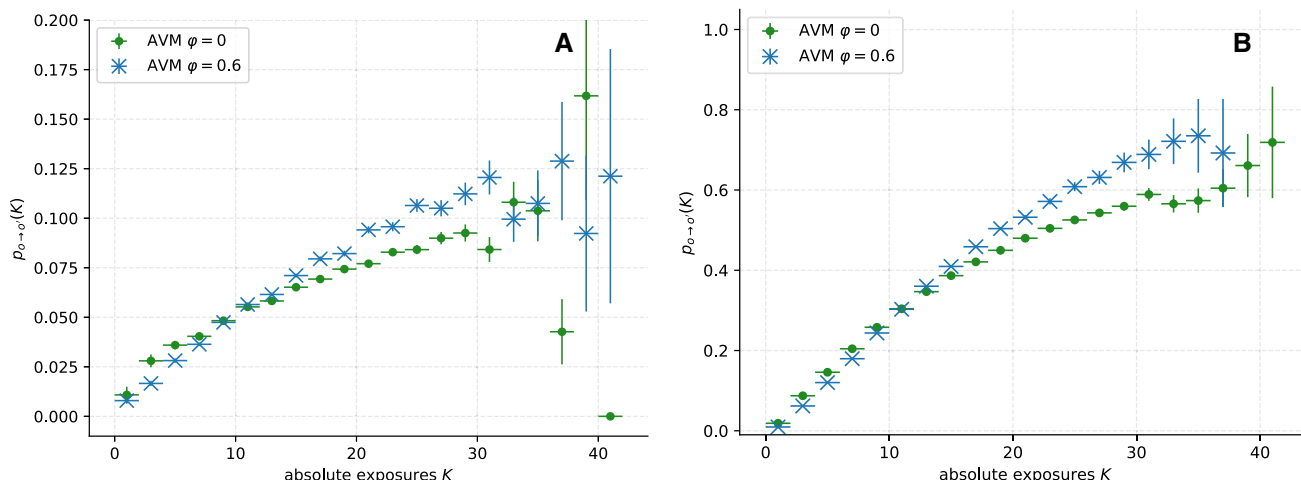


Fig. 3 Average estimated dose–response functions (DRFs) for synthetic temporal network data generated by ten runs of the adaptive voter model for rewiring probability $\varphi = 0$ and $\varphi = 0.6$. Error bars are computed as described in Sect. 3.1. In **A** the number of nodes $N = 619$, the average degree $\bar{k}_i = 19$ and the number of simulated time steps $\tau = 90$ were chosen analogously to the empirical temporal network from the Copenhagen Networks Study. The number of model updates per time step was adjusted to align with the average number of behaviour changes per time step with the CNS data. The error for the data point at $K \approx 41$ could not be estimated due to a lack of measurements; a large error is plausible. In **B**, the simulations are repeated for a larger network with $N = 851$ nodes, an

average degree of $\bar{k}_i = 13.5$ and significantly more model updates per time step. Here, the system was simulated until consensus (i.e. all nodes having the same trait for $\varphi = 0$, while for $\varphi = 0.6$ the model converges to two distinct groups with consensus each) was reached at $\tau = 190$ steps. Both **A** and **B** show a monotonic increasing relationship between $p_{p \rightarrow a}$ and K , while in **B** this trend is clearer due to the larger number of data points. The DRFs for $\varphi = 0$ differ significantly from those for $\varphi = 0.6$, owing to the more progressed network adaptation in the latter case. This difference shows that their form is not only influenced by contagion (imitation or social learning) effects, but also by homophily (network adaptation) dynamics

and lead to the observed monotonic decrease. However, we address in this study the switching between active and passive behaviour as a consequence of social contagion, and therefore, focus on the probability of becoming active $p_{p \rightarrow a}$ in the following analysis.

The probability $p_{p \rightarrow a}$ is subject to large errors for $K > 100$. The low occurrence of large K seems to be the main reason. However, we find a significant monotonic increase of $p_{p \rightarrow a}$, with Spearman’s rank correlation coefficient $\rho = 0.61$ and p value $p = 0.007$. This correlation could indicate contagion or homophilic dynamics. To pursue this indicator further, we examine the DRF using the surrogate data set method (Sect. 3.2). First, we investigate the possible influence of contagion dynamics (Sect. 4.2.1), then for group dynamics or external influences (Sect. 4.2.2) and finally for homophily dynamics (Sect. 4.2.3).

4.2.1 Investigation for contagion dynamics

For investigating the possible influence of contagion dynamics on the DRF, we employ the surrogate data tests \mathcal{H}_0^1 , \mathcal{H}_0^2 , and \mathcal{H}_0^3 introduced in Sect. 3.2, i.e. consider surrogate models in which explicitly no contagion takes place and we explore if they nevertheless reproduce the empirically observed DRF. To do so, we permute the traits of the nodes $o_i(t)$ and leave the network component $A_{ij}(t)$ unchanged. These permutations destroy possible temporal correlations of expo-

sure K with changes in traits and, thus, any trace of contagion dynamics. In three steps, we analyse the impact of different assumptions about the node dynamics on the dose–response functions and show step by step which assumptions are necessary to explain the observed DRF.

First data test. Hypothesis \mathcal{H}_0^1 : $P(A_{ij}(t), O)$. The empirical DRF can be reproduced with a class of models that is based only on the global mean activity level $O = \overline{\langle o_i(t) \rangle}_i$.

We test the most basic assumption of whether the empirical DRF can be explained by uncorrelated traits. To do so, all traits were uniformly permuted at random and only the global mean activity level $O = \overline{\langle o_i(t) \rangle}_i$, was conserved. Here, the overline and the brackets represent the time and ensemble mean, respectively. All possible contagion dynamics are destroyed in the model due to the random permutations.

Expectation. We expect to observe no correlation between the DRF $\tilde{p}_{p \rightarrow a}$ of the surrogate and K due to the permutations. Moreover, $\tilde{p}_{p \rightarrow a}(K)$ should be equal to the fraction of active states in the whole observed period.

Result. In Fig. 5A, the DRF $\tilde{p}_{p \rightarrow a}$ of the surrogate is contrasted with the empirical DRF $p_{p \rightarrow a}$. We find our expectations confirmed, $\tilde{p}_{p \rightarrow a}$ is quantitatively and qualitatively different from $p_{p \rightarrow a}$. Moreover, $\tilde{p}_{p \rightarrow a}$ is

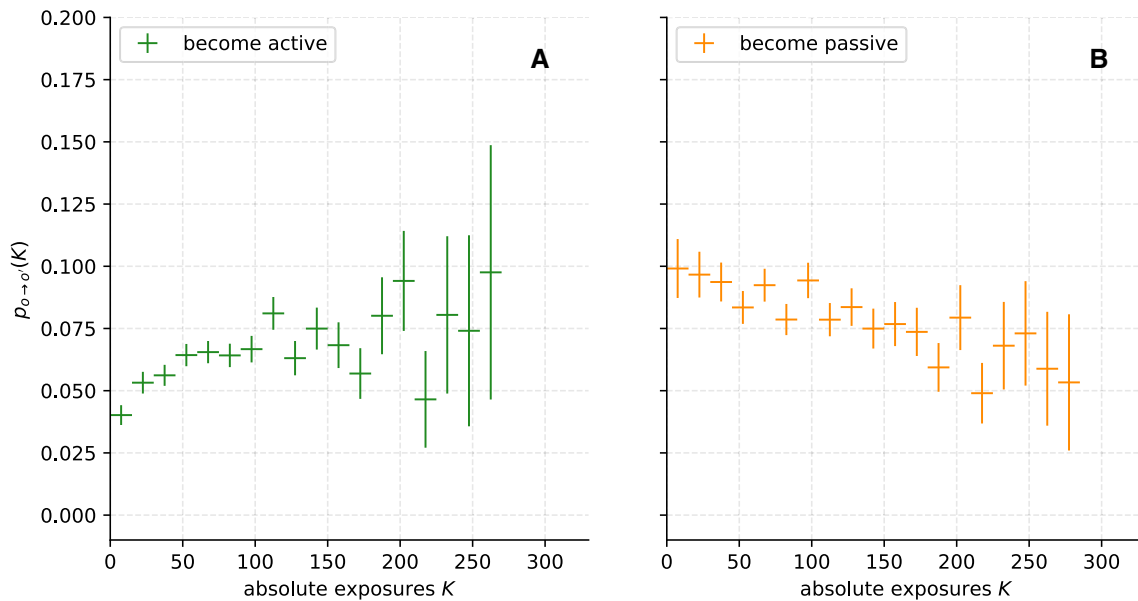


Fig. 4 Empirical dose–response functions computed from the Copenhagen Networks Study temporal network data, representing the probability to become active (**A**) or passive (**B**), as a function of the absolute exposure to these respective activity levels. Error bars are computed as described

approximately equal to the share of active states. We quantify the observed difference using the ζ test statistic introduced in Sect. 3.2. For the here discussed DRFs, the score is $\zeta = 328 \gg 1$. Therefore, the model is not sufficient to explain the empirical dynamics and we reject the first null hypothesis.

Second data test. Hypothesis \mathcal{H}_0^2 : $P(A_{ij}(t), O_i)$. *The empirical DRF can be reproduced with a class of models that is based only on each node’s individual activity level $O_i = \bar{o}_i(t)$.*

We test the effects of the individual activity level of each node $O_i = \bar{o}_i(t)$. Analogous to the previous model, the traits per node are randomly permuted in time, but this time only within each node’s time series. Therefore, O_i is conserved. As in the previous model, any possible contagion dynamics are destroyed due to the permutations.

Expectation. Due to the permutation in the surrogate, the individual probability of the node to change its trait is equal to O_i . In particular, this probability is independent of the exposure K . Therefore, we do not expect any correlation between $\tilde{p}_{p \rightarrow a}$ and K .

Result. Contrary to our expectations, in Fig. 5B, we find the probability $\tilde{p}_{p \rightarrow a}$ and K positively correlated, qualitatively similar to the correlation of $p_{p \rightarrow a}$ and K . However, for $K > 100$, the probability $\tilde{p}_{p \rightarrow a}(K)$ continues to increase, while $p_{p \rightarrow a}(K)$ appears to saturate. Furthermore, $\tilde{p}_{p \rightarrow a}$ and $p_{p \rightarrow a}$ differ quantitatively by a factor of about six. Thus, the conservation of O_i is not sufficient to explain the empirical DRF $\zeta = 309 \gg 1$, and we also reject the second null hypothesis.

in Sect. 3.1. For the probability to become active $p_{p \rightarrow a}$, a clear upward trend is noticeable ($\rho = 0.61$; $p = 0.007$), which might be caused by contagion. For the probability to become passive $p_{a \rightarrow p}$, a monotonic decrease can be identified ($\rho = -0.89$; $p = 3.5 \cdot 10^{-7}$)

In the second considered model, we found that the DRFs of the surrogate and the empirical data behave in a qualitatively similar way. This could be the result of pre-existing clustering in the data set: contacts j of nodes i would have similar activity values $O_j \approx O_i$ over the entire observation period. A node i with e.g. low O_i thus has contacts j with low O_j , and therefore, receives low exposure K . A positive correlation would be the result. Even without fully understanding the cause of the correlation found, it can be concluded that the individual activity level O_i is an essential feature in the empirical network. In addition to the correlation, we found a shift of the DRF $\tilde{p}_{p \rightarrow a}(K)$ by a factor of six compared to $p_{p \rightarrow a}$. We suspect the reason for this shift to be the non-preserved persistence of the nodes (inverse number of individual activity state changes). Due to the random permutations, the nodes change their trait more frequently than in the empirical network. In the following surrogate, this hypothesis is analysed in more detail.

Third data test. Hypothesis \mathcal{H}_0^3 : $P(A_{ij}(t), \{\tau_{i,0,1}\})$. *The empirical DRF can be reproduced with a class of models that is based only on each node’s individual activity level O_i , and its individual persistence (inverse number of individual activity state switches).*

In addition to O_i , the effect of individual persistence is tested. To achieve this, both the intervals with active trait $o_i(t) = 1$ and the intervals with passive trait $o_i(t) = 0$ were permuted at random. Hence, O_i and the persistence are conserved. Similar to the previous

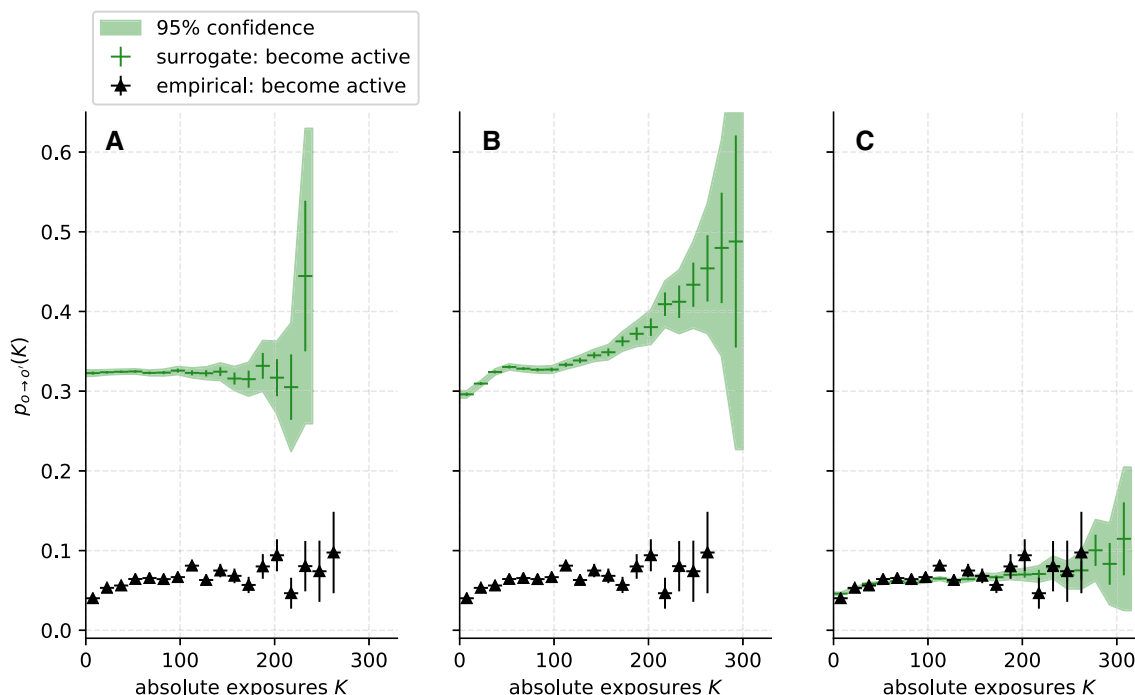


Fig. 5 Comparison of DRFs computed on empirical data (black triangles) and surrogates of the node traits (green crosses), corresponding to the null hypotheses \mathcal{H}_0^1 through \mathcal{H}_0^3 . It can be observed that neither **A** the preservation of the average trait O (\mathcal{H}_0^1), nor **B** the additional preservation of each individual node’s average trait O_i (\mathcal{H}_0^2) is sufficient to reproduce the data. **C** However, when the individual node

persistence, defined as the inverse of the number of trait switches, is also conserved (\mathcal{H}_0^3), the surrogate and empirical data show good agreement. Thus, we do not find sufficient evidence that contagion plays a significant role. Error bars are computed as described in Sect. 3.1. Confidence bounds for surrogate DRFs are the 95% confidence interval of the distribution of $p_{o \rightarrow o'}(K)$ over all surrogate realisations

models, the random permutations remove any possible contagion dynamics.

Expectation. Due to the additional conservation of individual persistence, we expect $\tilde{p}_{p \rightarrow a}$ to be qualitatively similar to $\tilde{p}_{p \rightarrow a}$ from the second model, but shifted closer to the empirical DRF on the y-axis.

Result. In Fig. 5C, we find, consistently with our expectations, that the DRF of the surrogate is shifted. Moreover, the probability $\tilde{p}_{p \rightarrow a}$ saturates for $K > 100$, analogous to the empirical DRF. Using the ζ test statistic, no significant deviation $\zeta = 0.79 < 1$ between $\tilde{p}_{p \rightarrow a}$ and $p_{p \rightarrow a}$ can be found. Therefore, we do not reject the third null hypothesis.

The third model showed that individual persistence is a main feature in the empirical network. Moreover, the model reproduces the empirical DRF in the model even without contagion. Thus, the third model shows that the data are not sufficient evidence that contagion plays a significant role in the empirical network, contrary to the hypothesis we formed when we first observed the correlation of $p_{p \rightarrow a}$ and K .

4.2.2 Investigation for group dynamics

In the previous section, we tested the effects of individual properties such as the individual activity level O_i or the individual persistence with our models. To

investigate the importance of group dynamics, in this section, we discard all individual properties and test the following null hypothesis:

Fourth data test. Hypothesis \mathcal{H}_0^4 : $P(A_{ij}(t), O(t))$. The empirical DRF can be reproduced with a class of models that is based only on the mean time-dependent activity level $O(t) = \langle o_i(t) \rangle_i$ of the ensemble.

We test the relevance of the mean time-dependent activity level $O(t) = \langle o_i(t) \rangle_i$ for the empirical dynamics. To do this, the traits between nodes were permuted at random for each time point separately, and only $O(t)$ is preserved.

Expectation. Given the permutations, both the probability of becoming active $\tilde{p}_{p \rightarrow a}$ and the exposure K depend on $O(t)$. Thus, a correlation between $\tilde{p}_{p \rightarrow a}$ and K is to be expected. Furthermore, we expect $\tilde{p}_{p \rightarrow a}(K) \gg p_{p \rightarrow a}(K)$ resulting from the destruction of the persistence of the nodes.

Result. Figure 6A compares the DRF $\tilde{p}_{p \rightarrow a}$ obtained from the surrogate data to the empirical DRF $p_{p \rightarrow a}$. Figure 6b shows the same DRFs, but the DRF of the surrogate (green, left y-axis) is offset by 0.25 to better compare the shape of the functions. In line with our expectations, $\tilde{p}_{p \rightarrow a}$ is correlated with K . For $K < 100$, the probability $\tilde{p}_{p \rightarrow a}(K)$ increases linearly. The empirical $p_{p \rightarrow a}(K)$ also increases for $K < 100$, but slightly

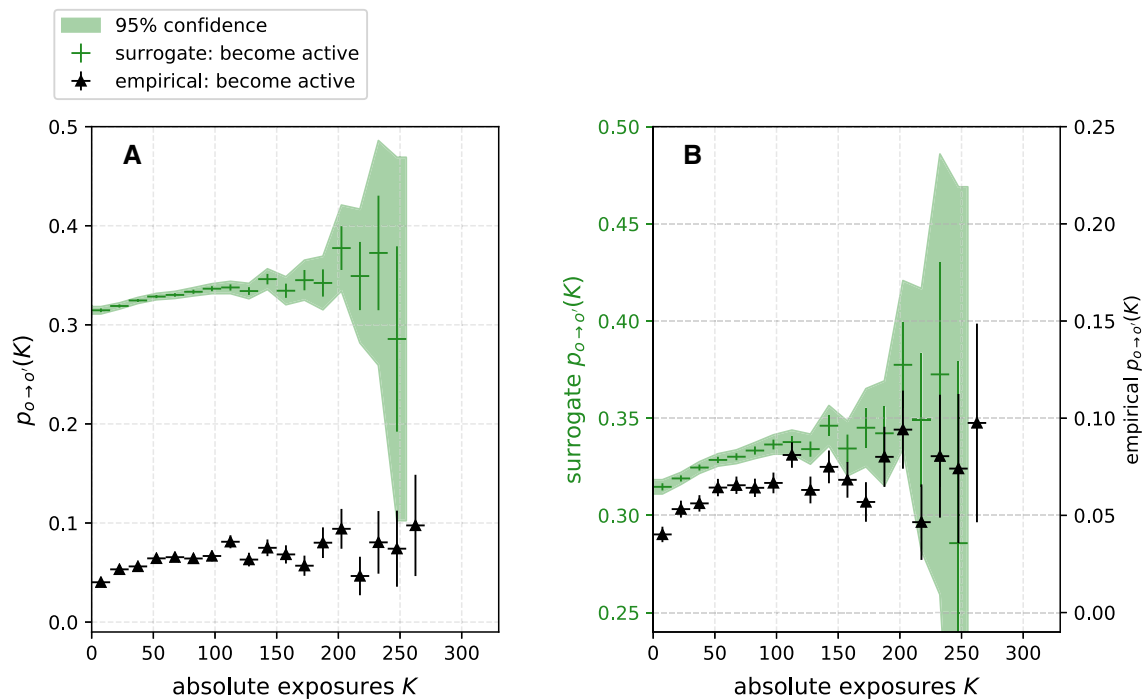


Fig. 6 Comparison of the DRF for empirical (black triangles) and surrogate (green crosses) data for null hypothesis \mathcal{H}_0^4 . To investigate external influences that affect all nodes simultaneously, the node traits were randomised in a way that conserves the time-varying mean activity level $O(t)$ of the group. The two figures contain the same data: **A** compares the absolute values of the data points, while in **B** the surrogate data y -axis (green, left side) is offset by 0.25

non-linearly. Quantitatively, we observe $\tilde{p}_{p \rightarrow a}(K) \gg p_{p \rightarrow a}(K)$. Thus, without individual traits, the model is not able to reproduce the empirical DRF $\zeta = 326 \gg 1$. Therefore, we reject the fourth null hypothesis.

Although the surrogate model DRF is quantitatively significantly different from the empirical DRF, the model predicts a qualitatively similar functional form. Temporal group dynamics thus seems to be another important feature in the empirical temporal network data. Apparently, participants change their behaviour collectively, as is also evident from the fluctuations observed in the mean activity level (Fig. 2). Such non-stationarities could emerge from internal collective dynamics or be due to external influences such as, for example, exam periods, weekends or holidays. A more detailed analysis is needed to distinguish these possible effects.

4.2.3 Investigation for homophily dynamics

Continuing our investigation, we look for homophily dynamics in the network. Analogously to the analysis testing for contagion effects, we create surrogate models in which explicitly no homophily takes place. With these, we attempt to reproduce the empirical dynamics. To this end, we permute the network edges $A_{ij}(t)$ and keep the properties of the nodes $o_i(t)$ unchanged.

to facilitate comparison of the functional forms. While the absolute values differ strongly, similarities in the functional forms are apparent, pointing to the importance of external influences on the collective group dynamics. Error bars are computed as described in Sect. 3.1. Confidence bounds for surrogate DRFs are the 95% confidence interval of the distribution of $p_{o \rightarrow o'}(K)$ over all surrogate realisations

This approach removes any homophily dynamics from the network, since the drawing and breaking of edges is randomised. The investigation is carried out in two steps, testing the following null hypotheses:

Fifth data test. Hypothesis \mathcal{H}_0^5 : $P(A, O_i(t))$. The empirical DRF can be reproduced with a class of models that is based only on individual activity dynamics and the average network edge density $A = \langle A_{ij}(t) \rangle_{i,j}$.

We test the most basic assumption that the empirical dynamics can be explained by a random network. For this purpose, all edges were permuted uniformly at random. Only the average temporal network edge density $A = \langle A_{ij}(t) \rangle_{i,j}$ was conserved. In this model, any homophily dynamics is removed, as the formation and breaking of edges is randomised.

Expectation. Since the traits have been kept unchanged, we expect the DRF of the model and the empirical DRF to be of the same order of magnitude. Due to the randomisation of the network, the neighbourhoods of the nodes are randomised as well. Thus, no correlation between the exposure K received from the neighbours and the probability $\tilde{p}_{p \rightarrow a}$ of changing the trait is to be expected.

Result. The DRF of the model and the empirical DRF are compared in Fig. 7A. Contrary to our expectation,

we can observe a correlation between $\tilde{p}_{p \rightarrow a}$ and K . Moreover, for the model, the case $\tilde{p}_{p \rightarrow a}(K)$ for $K > 100$ does not exist. Both DRFs have the same order of magnitude, which is in line with our expectations. However, only a few bins of the empirical DRF lie within the 95% confidence interval of the DRF from the surrogate and calculating the ζ test statistic gives $\zeta = 61 > 1$. Consequently, we reject the fifth null hypothesis.

When analysing our model based on a random network, we observed a positive correlation between $\tilde{p}_{p \rightarrow a}$ and K . This correlation was significantly different from the correlation found for the empirical DRF. Therefore, the non-trivial network structure and dynamics appear to be essential for reproducing the empirical dynamics. One explanation for the correlation found could be the external influences already described in Sect. 4.2.2. Nodes may change their traits in synchrony, independently of the network and caused by an external influence. This would affect K as well and could explain the correlation found. A further analysis is necessary here. Another feature of the surrogate model's DRF is that no large exposure $K > 100$ occurred. This is likely caused by a much smaller variance of the degree distribution in the random network than in the empirical one. In the following surrogate, this hypothesis is analysed in more detail.

Sixth data test. Hypothesis \mathcal{H}_0^6 : $P(k_i(t), O_i(t))$. *The empirical DRF can be reproduced with a class of models that is based only on the individual node dynamics, and each node's time-dependent network degree $k_i(t) = \sum_{j=0}^N A_{ij}(t)$.*

Building on the previous model, we test whether the time-dependent network degree of the nodes $k_i(t) = \sum_{j=0}^N A_{ij}(t)$ has a significant impact on the network dynamics. For this purpose, the edges of the network are permuted at random, but $k_i(t)$ is preserved. Analogous to the previous model, the homophily dynamics are removed by the permutations.

Expectation. For the correlation of $\tilde{p}_{p \rightarrow a}$ and K , we expect it to be similar to the one of the previous model. However, for this model we conserved the node's degree. Thus, the progression of the DRF should also extend over $K > 100$.

Result. In Fig. 7B, we compare the DRF of the model with the empirical one. In agreement with our expectation, we find $\tilde{p}_{p \rightarrow a}(K)$ for $K > 100$. However, the correlation of $\tilde{p}_{p \rightarrow a}$ and K is different from the previous model (Fig. 7A). No significant difference $\zeta = 0.31 < 1$ to the empirical DRF can be found anymore, using the ζ test statistic. Therefore, we cannot reject the sixth null hypothesis.

With this final surrogate model, we were able to reproduce the empirical DRF by conserving the node degree sequence in the temporal network data. Accordingly, node degree $k_i(t)$, the number of social contacts a student has at a given time t within the student population covered by the study, seems to be an important feature in the empirical data set. Furthermore, the

reproduction succeeded without including the dynamics of homophily. Thus, we do not detect a significant influence of contagion (see the results for \mathcal{H}_0^3 reported above), but neither a significant influence of homophily.

4.3 Summary

In Sects. 4.1 and 4.2, we presented the results of our methodology, which we applied first to synthetic data from the Adaptive Voter Model (AVM) and second to empirical data from the Copenhagen Networks Study (CNS). For both the synthetic and the empirical DRF, we found a monotonic functional dependency. In the synthetic case, it arises from the dynamics of the model: homophilic rewiring and social learning. To investigate whether contagion and homophily are the main driver for the empirical DRF, six null hypotheses \mathcal{H}_0^1 to \mathcal{H}_0^6 were tested. The tests were conducted by analysing two classes of surrogate models. In one, the traits $o_i(t)$ and in another, the edges $A_{ij}(t)$ were randomly permuted. Each class consists of a hierarchy of surrogate models. Starting with the most basic model, in which all traits resp. edges are randomly permuted, we gradually conserve parts of the system until the surrogate DRF $\tilde{p}_{p \rightarrow a}(K)$ and the empirical DRF $p_{p \rightarrow a}(K)$ are considered equal within an error margin. As proof of concept, this methodology was applied to the synthetic DRF of an adaptive voter model (see Appendix A for detailed results). In Fig. 8, we present a result compilation of the test hierarchy for the synthetic data (A) of the AVM ($\varphi = 0.6$) as well as for the empirical data (B). The red and the blue branches give the class of surrogate tests with permuted traits $o_i(t)$, while for the yellow branches the edges A_{ij} were permuted at random. An arrow from a surrogate test at a higher location to a lower one indicates that the former shuffles more than the latter. The differences between $\tilde{p}_{p \rightarrow a}(K)$ and $p_{p \rightarrow a}(K)$ are displayed on the horizontal axis and was quantified using a test statistic ζ introduced in Sect. 3.2. For the synthetic data (A), the yellow and the red branches end with \mathcal{H}_0^3 and \mathcal{H}_0^6 outside the grey area ($\zeta \geq 1$), indicating a significant difference between $\tilde{p}_{p \rightarrow a}(K)$ and $p_{p \rightarrow a}(K)$. Since we test with H_0^3 (H_0^6) whether the DRF can be explained without contagion (homophily), but both are core dynamics in the underlying model, this result was expected. In contrast, for the empirical data (B) \mathcal{H}_0^3 and \mathcal{H}_0^6 lie within the grey area, indicating no significant difference between $\tilde{p}_{p \rightarrow a}(K)$ and $p_{p \rightarrow a}(K)$. Consequently, this leads to the conclusion that we find neither significant evidence for an influence of contagion nor significant evidence for homophily in the CNS data. Considering all the tests performed on the empirical data, individual activity level, individual behavioural persistence, the effects of a possibly externally forced collective group dynamic and the individual number of social contacts (the node degree sequence) are sufficient to explain the estimated empirical DRF.

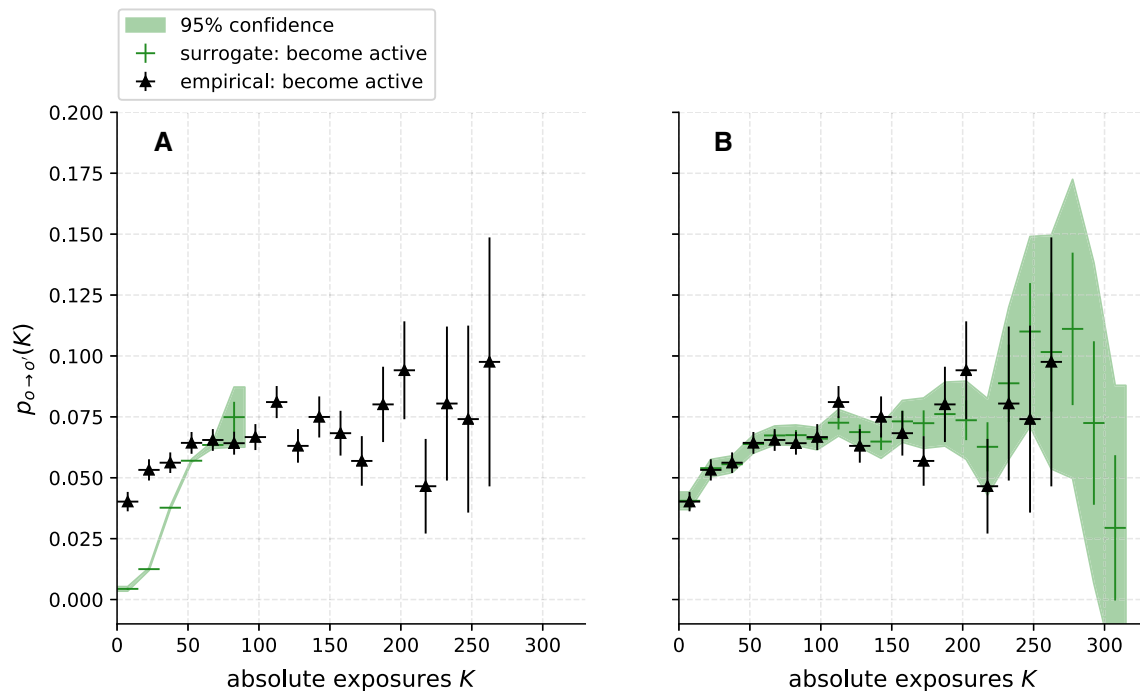


Fig. 7 Comparison of DRFs computed on empirical data (black triangles) and surrogates of the network topology (green crosses) for null hypotheses \mathcal{H}_0^5 and \mathcal{H}_0^6 . In **A**, only the mean node degree k is conserved (\mathcal{H}_0^5), leading to a significant difference between empirical and surrogate data. In **B**, each node's time-varying degree $k_i(t)$ is conserved as well (\mathcal{H}_0^6), corresponding to a test for homophily in the

network, with good agreement between the DRFs. It can be concluded that, while the non-trivial network structure appears to be of importance, no significant evidence for homophilic dynamics can be found. Error bars are computed as described in Sect. 3.1. Confidence bounds for surrogate DRFs are the 95 % confidence interval of the distribution of $p_{o \rightarrow o'}(K)$ over all surrogate realisations

5 Discussion and conclusion

In this paper, we proposed a methodology for estimating dose–response functions (DRFs) from temporal network data. We developed a hierarchy of surrogate data models to evaluate to what degree the observed DRFs can be explained by underlying processes such as social contagion, collective group dynamics and homophily. These surrogate models test the effects of distinct data features, such as overall and individual node activity levels, individual node trait persistence, overall network link density and individual node degrees. We applied this methodology to empirical temporal network data from the Copenhagen Networks Study, focussing on the illustrative health-related behaviour “regularly going to the fitness studio” in a physically-close-contact network of 619 university students, observed over the course of 3 months. We find neither significant evidence for an influence of contagion, nor significant evidence for homophily. The individual activity level, individual behavioural persistence, effects of possibly externally forced collective group dynamics, and individual number of social contacts (the node degree sequence) are sufficient to explain the estimated empirical dose–response function. These findings are underlined by a validation study performed using synthetic data, in which the sensitivity of our methodology to contagion and homophilic effects is demonstrated.

In the context of the application case considered in the present study, our findings contradict the perspective that social interactions influence adopted behaviour, for example via subjective norms [92], as supported by psychological research [93]. In particular, the ability of social norms to influence individual decision-making has been identified previously as a potential tool for large-scale group behaviour transformations [13, 94]. However, in the present context of exercise behaviour a person may only be susceptible to social influence during particular stages of their decision process, while being almost “immune” at other times [57, 95]. At any time, too few people may be in this socially susceptible state to rise above the noise threshold in the data.

Overall, our results demonstrate that care needs to be taken in interpreting dose–response functions obtained from empirical temporal network data; in particular when considering observational data that did not emerge from experiments in more controlled environments [42, 43]. Even pronounced positive correlations between exposure to a trait and the probability to adopt this trait can arise from structures in the temporal network data that do not need to be related to contagion and spreading processes, or homophily. Applying and further developing methodologies based on hierarchies of surrogate models, such as the one proposed in this article, provides a way forward to discern the specific imprints of complex spreading processes in tem-

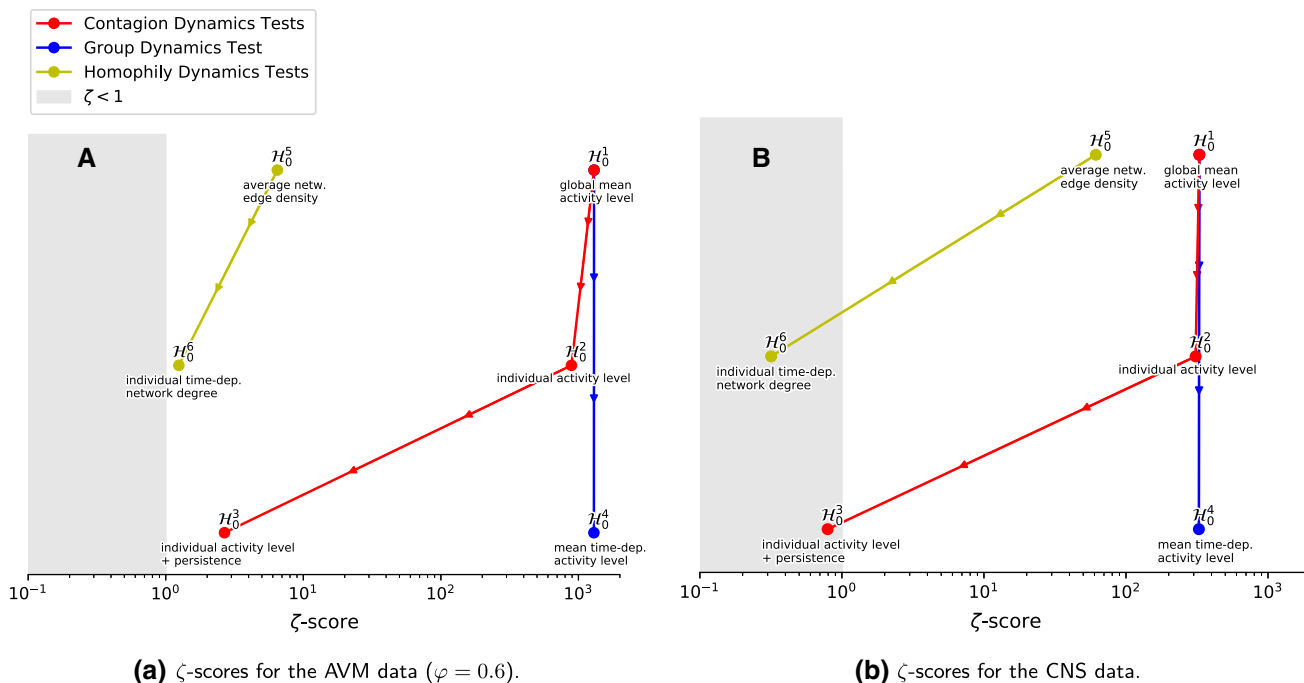


Fig. 8 Comparison of ζ -scores for a hierarchy of surrogate tests, for **A** the synthetic AVM data ($\varphi = 0.6$) and **B** the empirical CNS data. Each circle in the figure represents a single surrogate data test. The horizontal location of the circle reports the ζ -score of the tested hypothesis. An arrow from a surrogate test at a higher location to a lower one indicates that the former randomises more structure in the data than the latter. The null hypothesis name of each test

is given above each circle, and the conserved features of the surrogate model below it (see Sect. 3.2). The link and circle colour indicate which dynamics were investigated with the tests. The red and the blue branches give the class of surrogate tests with permuted traits $o_i(t)$, while for the yellow branches the edges A_{ij} were permuted at random. The grey rectangle marks the area where the empirical DRF does not differ significantly from the surrogate DRF

poral network data. Cases where the presence of such processes is not supported by the data can thus be excluded.

Our analysis has limitations in several dimensions that should be considered. First, in terms of data limitations, the empirical temporal network data set extracted from the Copenhagen Networks Study depends on multiple assumptions on thresholds and other parameter values. The definition of social contacts as links in a physically-close-contact network could be too unspecific for discerning social contagion effects. Social contagion might be expected to require a more permanent and intense social relationship such as friendship to be effective. Likewise, the chosen 1-day timescale of the contact network may need to be reconsidered, as clustering in the CNS data has been shown to disappear at time scales greater than 1 h [70]. Furthermore, the definition of node traits as active or passive may suffer from noise and missing data issues, since most likely some fitness studios and other relevant exercise institutions (e.g. university gyms, swimming pools etc.) are missing from our list. Also, using GPS coordinates to determine whether a student is visiting a fitness studio introduces uncertainties: in a densely populated urban area like the city of Copenhagen, a café or a library might be located right next to, or even above or below a fitness studio, introducing additional noise into our data set.

Second, considering methodological limitations, DRFs are a highly aggregate statistical indicator describing a complex temporal network data set. They might not be specific enough to detect subtle spreading processes or to discriminate different types of complex contagions. Arguably this calls for higher order statistics with larger statistical power. Moreover, the proposed methodology based on a hierarchy of surrogate data sets is limited in that it allows only for indirect inference on the possible presence of spreading or contagion processes. In this respect, it is desirable to augment the present analysis with more direct investigations including generative models of complex network spreading processes.

In summary, we suggest that our methodology is promising for applications to other systems and temporal network data sets. This can, among other applications, possibly aid our understanding of the social dynamics, spreading potentials and possible social tipping points in behaviours and social norms relevant for the adoption of healthy and sustainable diets [96] that can help to feed the world within planetary boundaries [97]. Efforts should be directed towards providing high-quality empirical temporal network data sets that can be leveraged for understanding complex spreading processes in these relevant domains. Promising directions of methodological developments include higher order statistics such as multi-node correlations for discerning the effects of longer contagion chains, spreading conta-

gion waves, or the imprints of network motifs on complex spreading processes. Astute surrogate data models can provide detailed insights into such spreading processes. Connecting empirical network data to generative statistical and dynamical adaptive network models more directly, e.g. via maximum likelihood methods, appears similarly promising. Hence, one can open new perspectives to predict future spreading dynamics. Ultimately, this research thus aids in designing targeted interventions for fostering desirable or suppressing unwanted contagions in diverse complex systems including pandemics, the brain, traffic and sustainability transformations.

Acknowledgements The authors would like to thank Franziska Gutmann and Michaela Schinkoeth of the Sport and Exercise Psychology research group at University of Potsdam for a helpful discussion. JFD, JH, JHL and MW are thankful for financial support by the Leibniz Association (project DominoES). JFD acknowledges support from the European Research Council project Earth Resilience in the Anthropocene (743080 ERA). NHK is grateful to the Geo.X Young Academy for financial support. SL acknowledges support by the Danish Research Council and the Vilum Foundation.

Funding Open Access funding enabled and organized by Projekt DEAL.

Author contribution statement

SL conducted the Copenhagen Networks Study, and provided the experimental data. JFD, JH, and JV conceived this study. JHL curated the data, implemented most of the methods and simulations, and visualised the results, with supervision by JFD and JV. NHK implemented the AVM-based method validation, with support from JHL. All the authors interpreted and discussed the results. JFD, JHL, and NHK wrote the manuscript, with support from JH, SL, JV and MW. All the authors give their final approval of the article version to be published.

Code availability The model and data analysis scripts used in this study are available as open source code at www.github.com/pik-copan/pydrf, pydrf v1.0, <https://doi.org/10.5281/zenodo.5526641>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need

to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Surrogate method validation with synthetic data

To evaluate how well the surrogate model method performs, we apply it to the two synthetic data sets created with CNS-aligned parameters for Fig. 3A. This data set is generated using the Adaptive Voter Model (AVM), once with ($\varphi = 0.6$) and once without ($\varphi = 0.0$) the network adaptation process. Other model parameters are chosen to align with the filtered data extracted from the Copenhagen Networks Study (CNS): the number of nodes $N = 619$, the average degree $k_i = 19$ and the number of simulated time steps $\tau = 90$. The number of model updates per time step is determined empirically, to align the average number of behaviour switches per time step across the entire system with the value found in the CNS data (40.24 ± 0.96 behaviour switches per time step). To maintain the comparability to the CNS data, a single simulation run of the AVM model is used, based on which ten surrogate model realisations are computed.

Using AVM-generated data to test the surrogate methods is a natural choice; when compared with e.g. SI(R) models, the AVM can best describe the processes and conditions of the system. For example, the behaviour is already rather common in the population; there is no “patient zero.” Furthermore, we assume that contact with “infected” (high activity level) individuals may increase infection probability—but also vice versa, that contact with “uninfected” (low activity level) individuals makes “recovery” more likely. However, even when aligning the model parameters to the CNS data, it should be noted that the AVM model does not necessarily represent a “best guess” for the real-world dynamics, but only an over-simplified stand-in.

In the following, we create the hierarchy of surrogate models, which is described in detail in Sect. 3.2. In this chapter, “AVM data” refers to the synthetic data set generated by the Adaptive Voter Model with the parameters described above, and “surrogate data” refers to surrogate data sets created using the AVM data. To quantify the difference between surrogate DRF and AVM DRF, we calculate the ζ -score (Eq. 12). A graphical presentation of the test hierarchy can be found in Fig. 12 for $\varphi = 0.0$, while the corresponding figure for $\varphi = 0.6$ is presented in Sect. 4.3.

First AVM test. Hypothesis $\mathcal{H}_0^1 : P(A_{ij}(t), O)$: Displayed in Fig. 9A and D for $\varphi = 0.0$ and $\varphi = 0.6$, respectively. As could be expected in this complete randomisation of activity states, the DRF becomes flat in both cases, at a level corresponding to the fraction of active nodes in the network. The ζ -score for the run with $\varphi = 0.0$ is $\zeta = 1281$ and the score for $\varphi = 0.6$ is $\zeta = 1299$.

Second AVM test. Hypothesis $\mathcal{H}_0^2 : P(A_{ij}(t), O_i)$: Displayed in Fig. 9B and E for $\varphi = 0.0$ and $\varphi = 0.6$, respectively. In this randomisation that conserves the individual node's activity levels, the surrogate DRF is still much higher than the AVM DRF. A likely explanation for the rising trends in the surrogate DRFs is the formation of net-

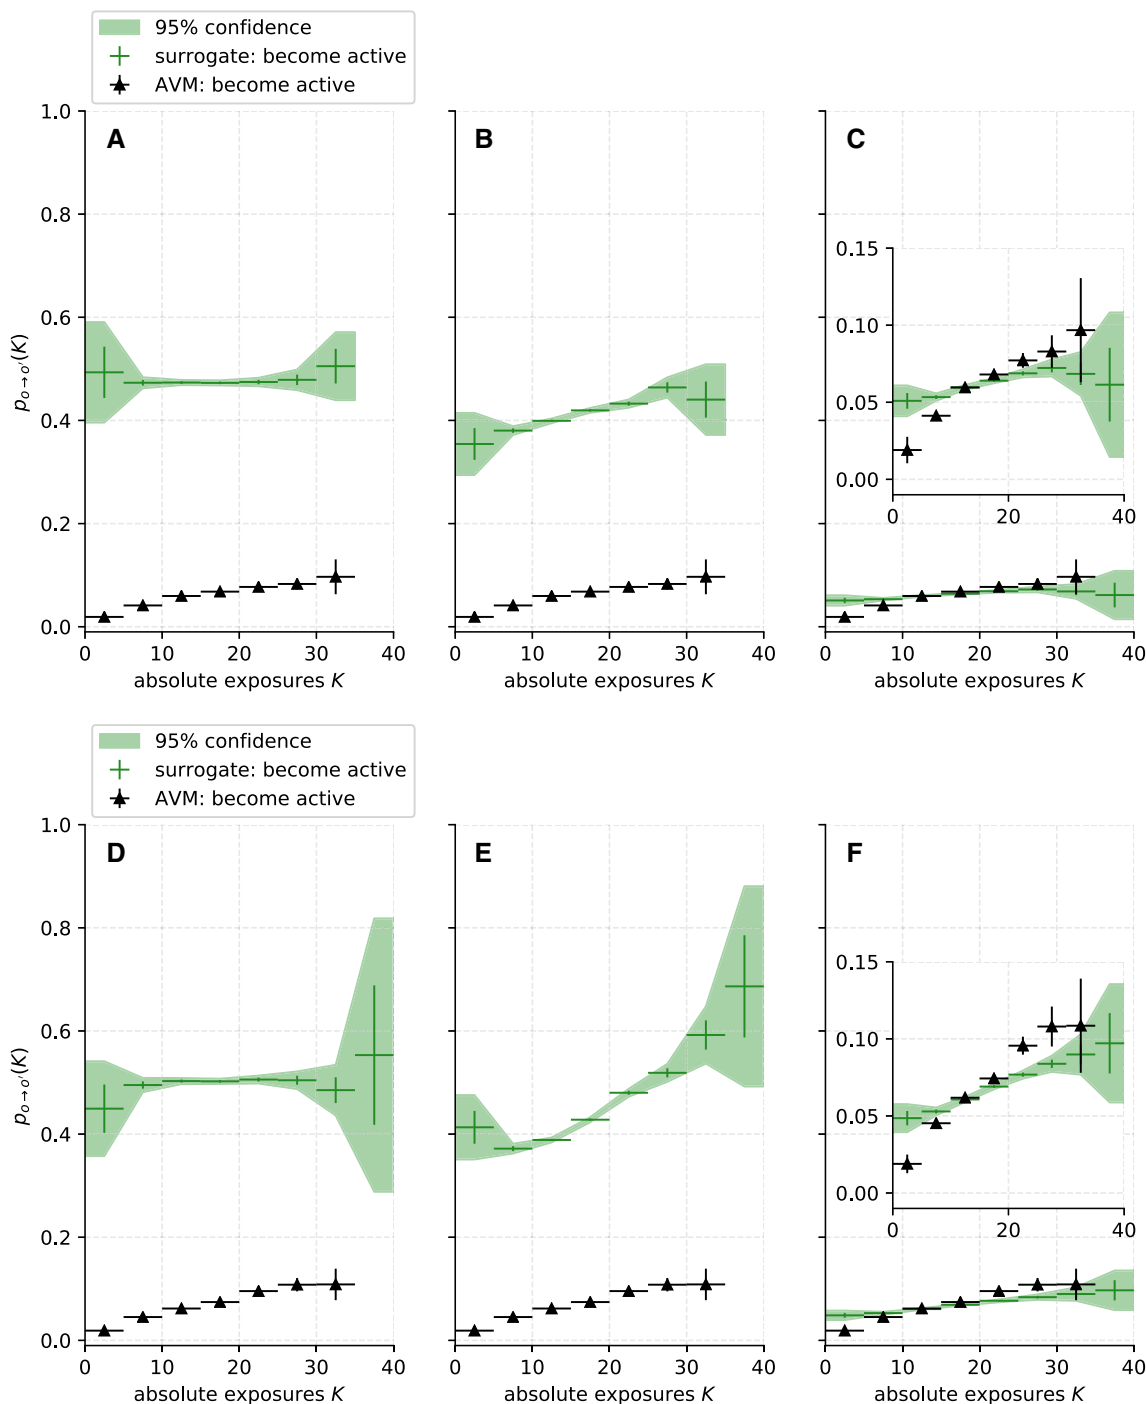


Fig. 9 Comparison of DRFs computed on AVM data (black triangles) and surrogates of the node traits (green crosses), corresponding to the null hypotheses \mathcal{H}_0^1 through \mathcal{H}_0^3 . **A–C** are for the $\varphi = 0.0$ case, while **D–F** correspond

to the $\varphi = 0.6$ case. Error bars are computed as described in Sect. 3.1. Confidence bounds for surrogate DRFs are the 95% confidence interval of the distribution of $p_{o \rightarrow o'}(K)$ over all surrogate realisations

work regions that have relatively homogeneous activity levels through the AVM process. Such regions, which consist of nodes that lean towards one activity level and whose neighbourhood comprise a majority of nodes with the same activity level, are not destroyed by the \mathcal{H}_0^2 shuffling. This effect can be expected to be stronger for the $\varphi = 0.6$ case, where homophilic rewiring is an additional driver in the formation

of such regions. The greater slope in Fig. 9E supports this. The ζ -score for the run with $\varphi = 0.0$ is $\zeta = 955$ and the score for $\varphi = 0.6$ is $\zeta = 890$.

Third AVM test. Hypothesis $\mathcal{H}_0^3 : P(A_{ij}(t), \{\tau_{i;0,1}\})$: Displayed in Fig. 9C and F for $\varphi = 0.0$ and $\varphi = 0.6$, respectively. As expected, when conserving the number of behaviour switches, the average switching probability dis-

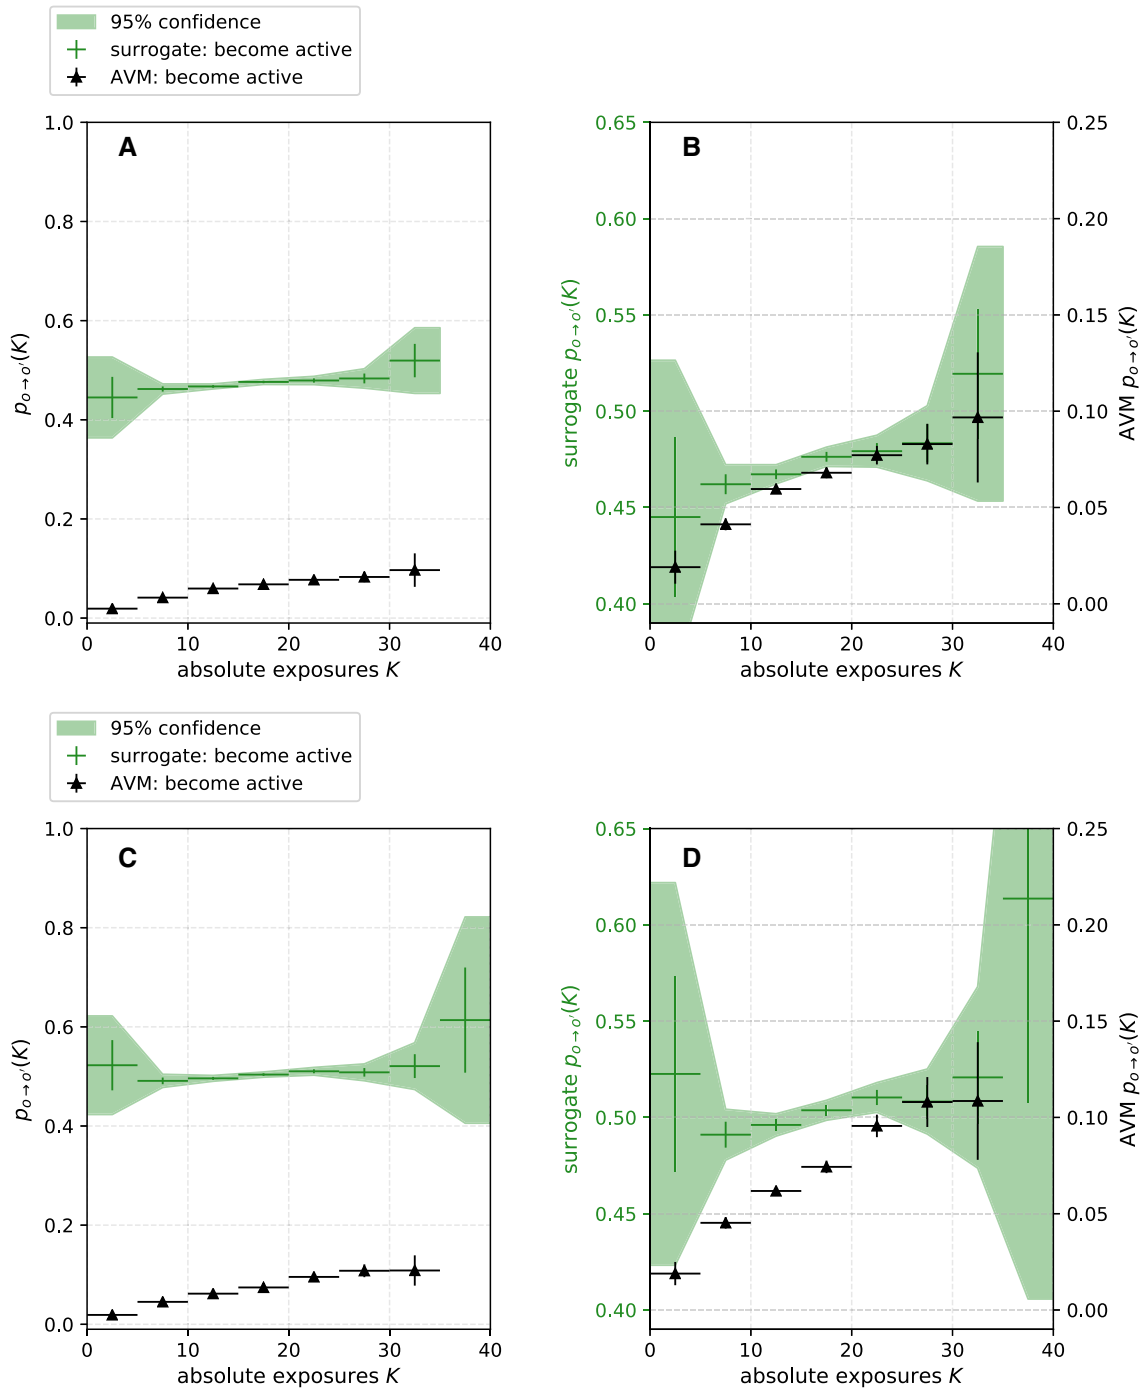


Fig. 10 Comparison of DRFs computed on AVM data (black triangles) and surrogates of the node traits (green crosses), corresponding to the null hypothesis \mathcal{H}_0^4 . **A** and **B** are for the $\varphi = 0.0$ case, while **C** and **D** correspond to the $\varphi = 0.6$ case. **B** and **D** display the same data as **A** and **C**, respectively, but with the range of the vertical axes indepen-

dently shifted for AVM and surrogate data, to facilitate the direct comparison of the DRFs. Error bars are computed as described in Sect. 3.1. Confidence bounds for surrogate DRFs are the 95% confidence interval of the distribution of $p_{o \rightarrow o'}(K)$ over all surrogate realisations

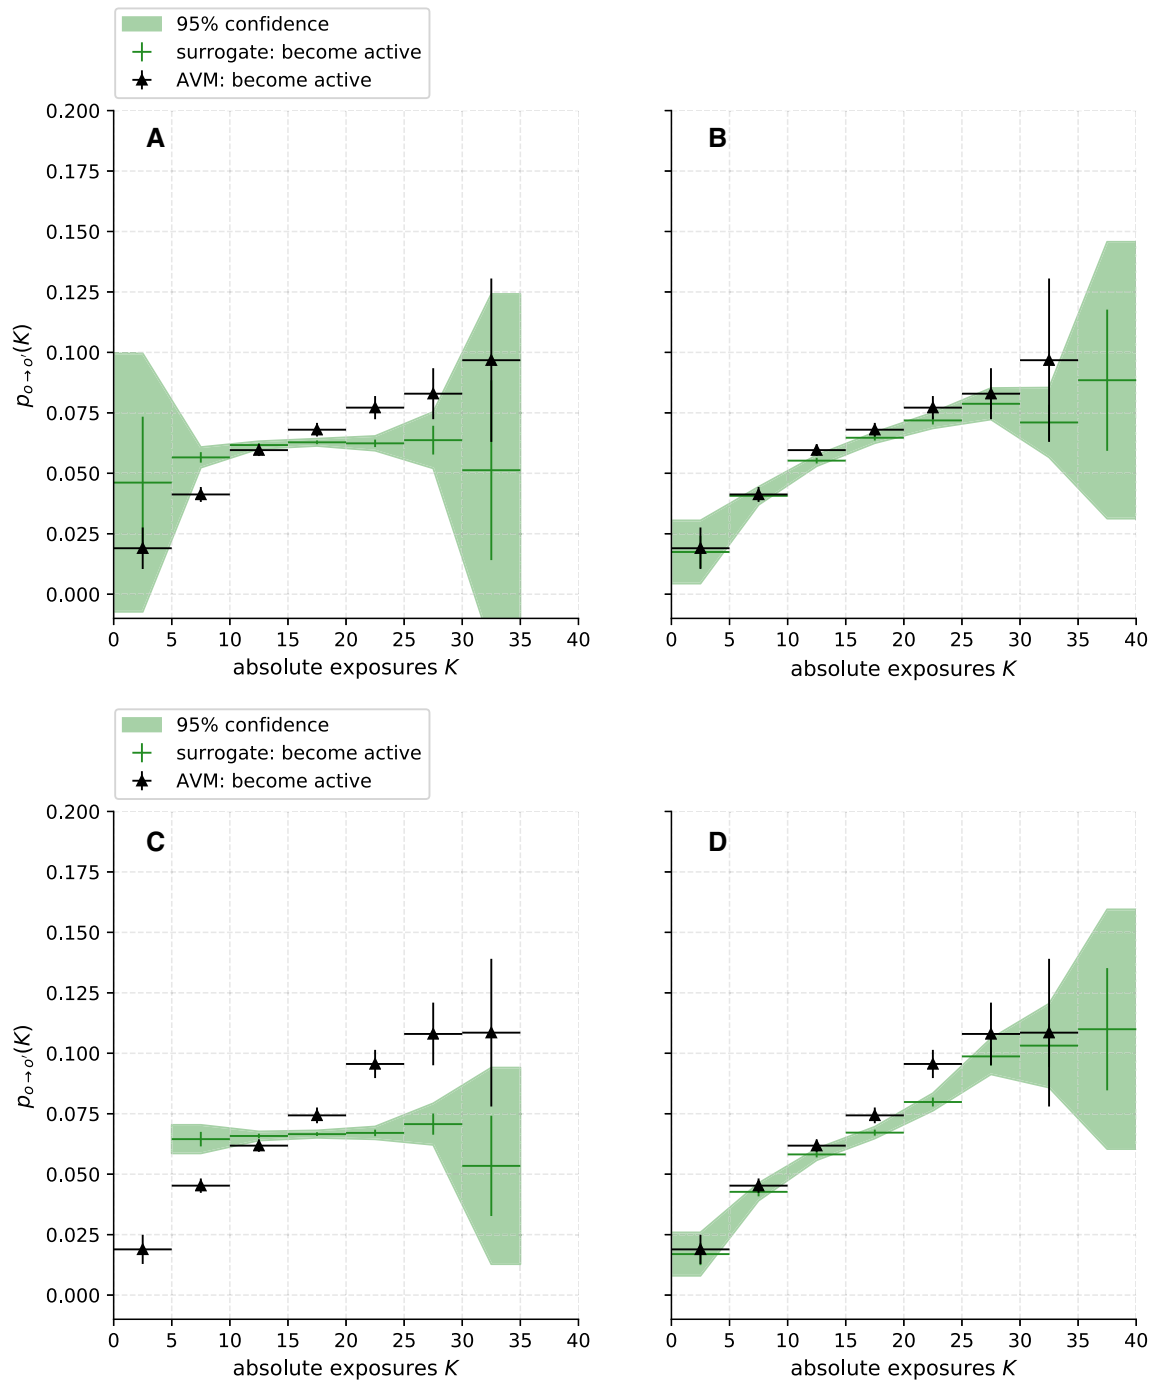


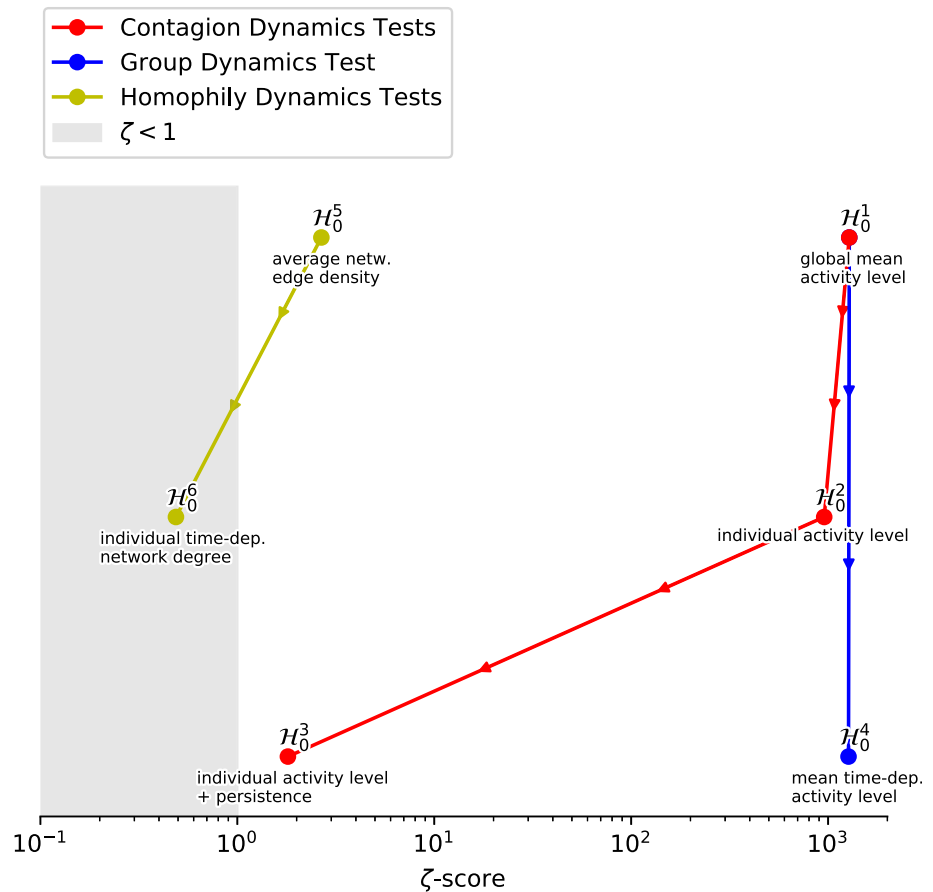
Fig. 11 Comparison of DRFs computed on AVM data (black triangles) and surrogates of the node traits (green crosses), corresponding to the null hypotheses \mathcal{H}_0^5 (A, C) and \mathcal{H}_0^3 (B, D). A and B are for the $\varphi = 0.0$ case, while

C and D correspond to the $\varphi = 0.6$ case. Error bars are computed as described in Sect. 3.1. Confidence bounds for surrogate DRFs are the 95% confidence interval of the distribution of $p_{o \rightarrow o'}(K)$ over all surrogate realisations

played in the DRF is very similar for the AVM and surrogate data. However, clear differences between the AVM and surrogate DRFs can be discerned. The ζ -score for the run with $\varphi = 0.0$ is $\zeta = 1.8$ and the score for $\varphi = 0.6$ is $\zeta = 2.7$, indicating a significant difference $\zeta > 1$. The upward trend of the AVM data DRFs is significantly greater than in the surrogate in both the $\varphi = 0.0$ and $\varphi = 0.6$ cases.

This is consistent with the true contagion process underlying the AVM simulation data. This shows the method to be sensitive to contagion effects, implying that the inability to reject \mathcal{H}_0^3 in the empirical data (see Fig. 5C) is likely due to a lack of dominant contagion dynamics in the studied behaviour. It should be noted that the surrogate DRFs do not become completely flat, but retain a more moder-

Fig. 12 Comparison of ζ -score for the hierarchy of surrogate tests on AVM data, for the $\varphi = 0.0$ case. Each circle in the figure represents a single surrogate data test. The horizontal location of the circle reports the ζ -score of the tested hypothesis. An arrow from a surrogate test at a higher location to a lower one means that the former shuffles more than the latter. The null hypothesis name of each test is given above each circle, and the conserved features of the surrogate model below it (see Sect. 3.2). The link and circle colour indicate which dynamics were investigated with the tests. The grey rectangle marks the area where the empirical DRF does not differ significantly from the surrogate DRF



ate upward trend. This can be explained analogously to the upward trend in the surrogate DRFs of the second AVM test, described above.

Fourth AVM test. Hypothesis \mathcal{H}_0^4 : $P(A_{ij}(t), O(t))$: Displayed in Fig. 10 (A,B) and (C,D) for $\varphi = 0.0$ and $\varphi = 0.6$, respectively. The surrogate and AVM DRFs have greatly differing y-scales. The ζ -score for the run with $\varphi = 0.0$ is $\zeta = 1269$ and the score for $\varphi = 0.6$ is $\zeta = 1295$. However, in the $\varphi = 0.0$ case, the surrogate DRF retains an upward trend, albeit smaller than the AVM DRF. Since \mathcal{H}_0^4 is essentially the mean-field approximation of the system, this demonstrates how the network is densely, and relatively homogeneously connected in this case. In the $\varphi = 0.6$ case, the randomisation destroys any significant slope. Here, the original AVM data apparently differs more strongly from the mean-field approximation, which can be explained by the greater degree of homophilic clustering in this case. The network structure, with its additional rewiring mechanism, thus appears more important in this case. The behaviour seen in the evaluation of \mathcal{H}_0^4 in the empirical CNS data (Fig. 6) resembles the $\varphi = 0.0$ case in AVM data, which can be interpreted as an absence of clustering in the CNS data.

Fifth AVM test. Hypothesis \mathcal{H}_0^5 : $P(A, O_i(t))$: Displayed in Fig. 11A and C for $\varphi = 0.0$ and $\varphi = 0.6$, respectively. As expected, after completely randomising the network, the surrogate model gives a nearly constant DRF. The ζ -score for the run with $\varphi = 0.0$ is $\zeta = 2.7$ and the score for $\varphi = 0.6$ is $\zeta = 6.5$. The difference between surrogate and AVM DRFs is less significant for the $\varphi = 0.0$ case than for the $\varphi =$

0.6 case, which can be explained by the additional network processes at work in the latter case: the randomisation has a larger effect here.

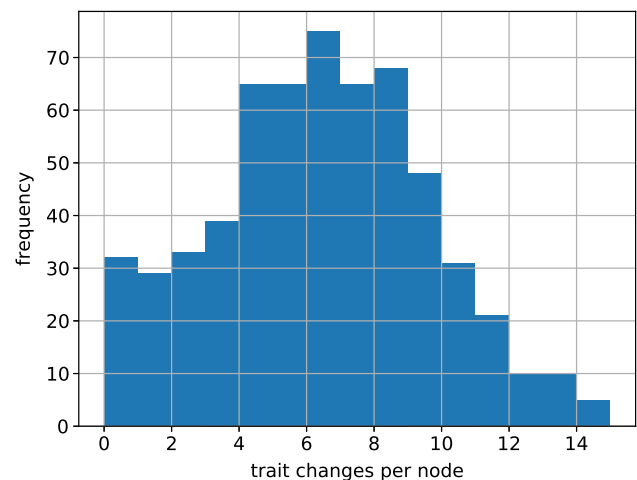


Fig. 13 Distribution of the total number of activity level changes per node during the studied time interval

Sixth AVM test. Hypothesis \mathcal{H}_0^6 : $P(k_i(t), O_i(t))$: Displayed in Fig. 11B and D for $\varphi = 0.0$ and $\varphi = 0.6$, respectively. The ζ -score for the run with $\varphi = 0.0$ is $\zeta = 0.49$ and the score for $\varphi = 0.6$ is $\zeta = 1.25$. The difference between the

surrogate and original AVM DRFs is not nearly as big as in many of the other surrogate tests, pointing to an effect of homophily that is moderate at most. For the $\varphi = 0.6$ case, hints for homophily effects can be observed, since the surrogate and original AVM curves are significantly separated here. For the $\varphi = 0.0$ case, the curves are not significantly separated (see also Fig. 12). This is consistent with our expectations, since homophilic clustering through preferential attachment is present, but not dominant in the $\varphi = 0.6$ model (see Fig. 3)

Figure 12 shows, analogously to Fig. 8, the significance of the deviations between surrogate and AVM DRFs for $\varphi = 0.0$. The case $\varphi = 0.6$ was already presented in Fig. 8A. For the $\varphi = 0.0$ case (Fig. 12), only \mathcal{H}_0^5 cannot be rejected based on the ζ test statistic. For the $\varphi = 0.6$ case (Fig. 8A), none of the hypothesis tests can be rejected. The difference in the rejection of \mathcal{H}_0^3 to the empirical case (Fig. 8B) appears to show that our method can detect contagion created by the social learning within the AVM. Moreover, the difference in the rejection of \mathcal{H}_0^5 between the $\varphi = 0$ and the $\varphi = 0.6$ cases suggests that our method can detect the small amount of homophily created by the adaptive rewiring.

B Permutation space for \mathcal{H}_0^3

For the surrogate method to work, the shuffling algorithms must provide sufficient randomisation, creating data sets with significant differences to the original data. This is eas-

ily achieved for most of the proposed surrogate models. However, the randomisation space for \mathcal{H}_0^3 is the most constrained. Here, the number of possible permutations of the activity intervals is limited by the total number of activity level switches of each node. In this section, we demonstrate that this randomisation space is sufficient for the method to function.

Figure 13 displays the distribution of total activity level (“trait”) changes per node in the studied time interval. Nodes switch behaviour on average 5.94 times. Thus, on average, there are 3–4 active and 3–4 inactive intervals for each node. If a node has 3 active and 4 inactive intervals, the shuffling can produce $3!4! = 144$ different surrogates. More than 43 percent of agents switch behaviour at least 7 times, thus having at least 4 active and 4 inactive intervals and hence at least $4!4! = 576$ different surrogates for each of these nodes. From this, we conclude that there is sufficient randomisation in \mathcal{H}_0^3 . This is supported by the validation of the methodology using synthetic AVM data, which shows a deviation between AVM and surrogate DRFs for \mathcal{H}_0^3 (see Fig. 9C and F).

C List of considered fitness centres in Copenhagen

See Table 1.

Table 1 List of the fitness centres in Copenhagen considered in this study, with their respective coordinates, as extracted from Open Street Maps [75]

Name	Longitude [° E]	Latitude [° N]
Fresh Fitness Hvidovre	12.4691961	55.6415696
Fitness.dk	12.5618214	55.6614733
FitnessDK	12.5114098	55.6647699
Fresh Fitness	12.5404751	55.6975516
Fresh	12.4199488	55.6493081
Fitness World	12.4418141	55.7231967
Fitness World Ballerup	12.3579672	55.7296181
Fitness World Brøndby	12.4383494	55.6673030
Fitness World Farum Park	12.3513120	55.8172970
Fitness World Frederiksberg Bernhard Bangs Alle	12.5104671	55.6844058
Fitness World Frederiksberg Forum	12.5524718	55.6830906
Fitness World Frederiksberg Peter Bangs Vej	12.5131680	55.6795400
Fitness World Gentofte	12.5378949	55.7386120
Fitness World Glostrup	12.4008395	55.6640800
Fitness World Greve Hundige Storcenter	12.3274148	55.5987709
Fitness World Greve	12.2984612	55.5905648
Fitness World Herlev	12.4160534	55.7253403
Fitness World Husum	12.4810239	55.7095419
Fitness World København Baron Boltens Gård	12.5848511	55.6820125
Fitness World København Ellebjergvej	12.5108247	55.6507568
Fitness World København Emdrup Station	12.5409464	55.7218740
Fitness World København Englandsvej	12.6043943	55.6569690
Fitness World København Gasværksvej	12.5570237	55.6708078
Fitness World København Jagtvej	12.5509410	55.6964980
Fitness World København Lyngbyvej	12.5604444	55.7116463
Fitness World København Lyongade	12.6099453	55.6613686

Table 1 Continued

Name	Longitude [° E]	Latitude [° N]
Fitness World København Nordre Fasanvej	12.5364747	55.6985181
Fitness World København Strandvejen	12.5777058	55.7219712
Fitness World København Vester Farimagsgade	12.5623173	55.6782088
Fitness World København Århusgade	12.5872772	55.7067752
Fitness World Lyngby	12.5039072	55.7688801
Fitness World Måløv	12.3187172	55.7485909
Fitness World Søborg	12.4932893	55.7395909
Fitness World Taastrup	12.3017208	55.6529634
Fitness World Valby Mosedalvej	12.5134815	55.6674858
Fitness World Værløse	12.3615021	55.7821745
fitnessdk	12.4392816	55.7249089

References

- D.J. Watts, A simple model of global cascades on random networks. *Proc. Natl. Acad. Sci.* **99**, 5766–5771 (2002)
- P.S. Dodds, D.J. Watts, Universal behavior in a generalized model of contagion. *Phys. Rev. Lett.* **92**, 218701 (2004)
- S. Lehmann, Y.-Y. Ahn, *Complex Spreading Phenomena in Social Systems* (Springer, New York, 2018)
- J.D. Murray, *Mathematical Biology?: I. An Introduction* (Springer-Verlag, New York, 2002)
- D.J. Daley, J. Gani, *Epidemic Modelling* (Cambridge University Press, Cambridge, 1999)
- B.F. Maier, D. Brockmann, Effective containment explains subexponential growth in recent confirmed COVID-19 cases in China. *Science* **368**, 742–746 (2020)
- S.V. Buldyrev, R. Parshani, G. Paul, H.E. Stanley, S. Havlin, Catastrophic cascade of failures in interdependent networks. *Nature* **464**, 1025–1028 (2010)
- J.S. Coleman, E. Katz, H. Menzel, *Medical Innovation: A Diffusion Study* (Bobbs-Merrill Co, Indiana, 1966)
- T. Valente, Network models of the diffusion of innovations. *Comput. Math. Organ. Theory* **2**, 163–164 (1996)
- F.W. Geels, B.K. Sovacool, T. Schwanen, S. Sorrell, Sociotechnical transitions for deep decarbonization. *Science* **357**, 1242–1244 (2017)
- V. Capraro, M. Perc, Mathematical foundations of moral preferences. *J. R. Soc. Interface* **18**, 20200880 (2021)
- P. Turchin, T.E. Currie, E.A.L. Turner, S. Gavrillets, War, space, and the evolution of Old World complex societies. *Proc. Natl. Acad. Sci.* **110**, 16384–16389 (2013). ISBN: 9781308825113
- K. Nyborg et al., Social norms as solutions. *Science* **354**, 42–43 (2016)
- M. Tsvetkova, M.W. Macy, The social contagion of generosity. *PLOS One* **9**, e87275 (2014)
- J.D. Tàbara et al., Positive tipping points in a rapidly warming world. *Curr. Opin. Environ. Sustain.* **31**, 120–129 (2018)
- J.D. Farmer et al., Sensitive intervention points in the post-carbon transition. *Science* **364**, 132–134 (2019)
- I.M. Otto et al., Social tipping dynamics for stabilizing earth’s climate by 2050. *Proc. Natl. Acad. Sci.* **117**, 2354–2365 (2020)
- S. Sharpe, T.M. Lenton, Upward-scaling tipping cascades to meet climate goals: plausible grounds for hope. *Clim. Policy* **21**, 421–433 (2021)
- S. Lohmann, The dynamics of informational cascades: the Monday demonstrations in Leipzig, East Germany, 1989–91. *World Politics* **47**, 42–101 (1994)
- R. Stark, Why religious movements succeed or fail: a revised general model. *J. Contemp. Religion* **11**, 133–146 (1996)
- R.L. Montgomery, *The Diffusion of Religions: A Sociological Perspective* (University Press of America, Maryland, 1996)
- R. Winkelmann, et al. Social tipping processes towards climate action: a conceptual framework. *Ecol Econ* (in press). [arXiv:2010.04488](https://arxiv.org/abs/2010.04488) (2020)
- P.S. Dodds, D.J. Watts, A generalized model of social and biological contagion. *J. Theor. Biol.* **232**, 587–604 (2005). [arXiv:1705.10783](https://arxiv.org/abs/1705.10783)
- M. Wiedermann, E.K. Smith, J. Heitzig, J.F. Donges, A network-based microfoundation of Granovetter’s threshold model for social tipping. *Sci. Rep.* **10**, 11202 (2020)
- P. Holme, M.E. Newman, Nonequilibrium phase transition in the coevolution of networks and opinions. *Phys. Rev. E* **74**, 056108 (2006)
- T. Gross, C.J.D. D’Lima, B. Blasius, Epidemic dynamics on an adaptive network. *Phys. Rev. Lett.* **96**, 208701 (2006)
- T. Gross, H. Sayama, *Adaptive Networks* (Springer, New York, 2009)
- M. Wiedermann, J.F. Donges, J. Heitzig, W. Lucht, J. Kurths, Macroscopic description of complex adaptive networks coevolving with dynamic node states. *Phys. Rev. E* **91**, 052801 (2015)
- S. Hsiang et al., The effect of large-scale anti-contagion policies on the COVID-19 pandemic. *Nature* **584**, 262–267 (2020)
- F. Schlosser et al., Covid-19 lockdown induces disease-mitigating structural changes in mobility networks. *Proc. Natl. Acad. Sci.* **117**, 32883–32890 (2020)
- P.J. Menck, J. Heitzig, J. Kurths, H.J. Schellnhuber, How dead ends undermine power grid stability. *Nat. Commun.* **5**, 3969 (2014)
- K. Lewis, J. Kaufman, M. Gonzalez, A. Wimmer, N. Christakis, Tastes, ties, and time: a new social network dataset using Facebook.com. *Social Netw.* **30**, 330–342 (2008)

33. B. Suh, L. Hong, P. Pirolli, E.H. Chi, in *2010 IEEE Second International Conference on Social Computing*, pp. 177–184 (2010)
34. M. Feinleib, W.B. Kannel, R.J. Garrison, P.M. McNamara, W.P. Castelli, The Framingham offspring study. Design and preliminary data. *Prev. Med.* **4**, 518–525 (1975)
35. N.A. Christakis, J.H. Fowler, The spread of obesity in a large social network over 32 years. *New Engl. J. Med.* **357**, 370–379 (2007)
36. N.A. Christakis, J.H. Fowler, The collective dynamics of smoking in a large social network. *New Engl. J. Med.* **358**, 2249–2258 (2008)
37. J.H. Fowler, N.A. Christakis, Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ (Online)* **337**, a2338 (2008)
38. J.T. Cacioppo, J.H. Fowler, N.A. Christakis, Alone in the crowd: the structure and spread of loneliness in a large social network. *J. Person. Social Psychol.* **97**, 977–991 (2009)
39. J.N. Rosenquist, J. Murabito, J.H. Fowler, N.A. Christakis, The spread of alcohol consumption behavior in a large social network. *Ann. Intern. Med.* **152**, 426–433 (2010)
40. J.N. Rosenquist, J.H. Fowler, N.A. Christakis, Social network determinants of depression. *Mol. Psychiatry* **16**, 273–281 (2011)
41. R. McDermott, J.H. Fowler, N.A. Christakis, Breaking up is hard to do, unless everyone else is doing it too: social network effects on divorce in a longitudinal sample. *Social Forces* **92**, 491–519 (2013)
42. A.D. Kramer, J.E. Guillory, J.T. Hancock, Experimental evidence of massive-scale emotional contagion through social networks. *Proc. Natl. Acad. Sci.* **111**, 8788–8790 (2014)
43. R.M. Bond et al., A 61-million-person experiment in social influence and political mobilization. *Nature* **489**, 295–298 (2012)
44. E.L. Ogburn, in *Complex Spreading Phenomena in Social Systems*, 47–64 (Springer, 2018)
45. J. Runge et al., Identifying causal gateways and mediators in complex spatio-temporal systems. *Nat. Commun.* **6**, 8502 (2015)
46. J. Runge, Causal network reconstruction from time series: from theoretical assumptions to practical estimation. *Chaos* **28**, 075310 (2018)
47. M.M. Sosna et al., Individual and collective encoding of risk in animal groups. *Proc. Natl. Acad. Sci.* **116**, 20556–20561 (2019)
48. N.O. Hodas, K. Lerman, The simple rules of social contagion. *Sci. Rep.* **4**, 4343 (2014)
49. R. Vicente, M. Wibral, M. Lindner, G. Pipa, Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *J. Comput. Neurosci.* **30**, 45–67 (2011)
50. M. Casdagli, Chaos and deterministic versus stochastic non-linear modelling. *J. R. Stat. Soc. B Method.* **54**, 303–328 (1992)
51. L. Gauvin, et al. Randomized reference models for temporal networks. [arXiv:1806.04032](https://arxiv.org/abs/1806.04032) [physics, q-bio] (2020)
52. P. Holme, J. Saramäki, Temporal networks. *Phys. Rep.* **519**, 97–125 (2012)
53. M. Génois, C.L. Vestergaard, C. Cattuto, A. Barrat, Compensating for population sampling in simulations of epidemic spread on temporal contact networks. *Nat. Commun.* **6**, 8860 (2015)
54. F. Karimi, P. Holme, Threshold model of cascades in empirical temporal networks. *Physica A* **392**, 3476–3483 (2013)
55. J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, J. Doyne Farmer, Testing for nonlinearity in time series the method of surrogate data. *Physica D* **58**, 77–94 (1992)
56. T. Schreiber, A. Schmitz, Surrogate time series. *Physica D* **142**, 346–382 (2000)
57. B.H. Marcus, L.R. Simkin, The transtheoretical model: applications to exercise behavior. *Med. Sci. Sports Exercise* **26**, 1400–1404 (1994)
58. M. Boguná, R. Pastor-Satorras, A. Díaz-Guilera, A. Arenas, Models of social networks based on social distance attachment. *Phys. Rev. E* **70**, 056122 (2004)
59. C. Castellano, D. Vilone, A. Vespignani, Incomplete ordering of the voter model on small-world networks. *EPL* **63**, 153 (2003)
60. R.A. Holley, T.M. Liggett, Ergodic theorems for weakly interacting infinite systems and the Voter model. *Ann. Prob.* **3**, 643–663 (1975)
61. T. Gross, B. Blasius, Adaptive coevolutionary networks: a review. *J. R. Soc. Interface* **5**, 259–271 (2008)
62. N. Perra, B. Gonçalves, R. Pastor-Satorras, A. Vespignani, Activity driven modeling of time varying networks. *Sci. Rep.* **2**, 469 (2012)
63. A. Stopczynski et al., Measuring large-scale social networks with high resolution. *PloS One* **9**, e95978 (2014)
64. P. Sapiezynski, A. Stopczynski, D.D. Lassen, S. Lehmann, Interaction data from the Copenhagen Networks Study. *Sci. Data* **6**, 315 (2019)
65. E. Mones, A. Stopczynski, A.S. Pentland, N. Hupert, S. Lehmann, Optimizing targeted vaccination across cyber-physical networks: an empirically based mathematical simulation study. *J. R. Soc. Interface* **15**, 20170783 (2018)
66. A. Stopczynski, S. Lehmann et al., How physical proximity shapes complex social networks. *Sci. Rep.* **8**, 17722 (2018)
67. S. Kojaku, L. Hébert-Dufresne, E. Mones, S. Lehmann, Y.-Y. Ahn, The effectiveness of backward contact tracing in networks. *Nat. Phys.* **17**, 652–658 (2021)
68. L. Alessandretti, P. Sapiezynski, V. Sekara, S. Lehmann, A. Baronchelli, Evidence for a conserved quantity in human mobility. *Nat. Hum. Behav.* **2**, 485–491 (2018)
69. L. Alessandretti, U. Aslak, S. Lehmann, The scales of human mobility. *Nature* **587**, 402–407 (2020)
70. V. Sekara, A. Stopczynski, S. Lehmann, Fundamental structures of dynamic social networks. *Proc. Natl. Acad. Sci.* **113**, 9977–9982 (2016)
71. A. Mollgaard et al., Measure of node similarity in multilayer networks. *PloS One* **11**, e0157436 (2016)
72. I. Psylla, P. Sapiezynski, E. Mones, S. Lehmann, The role of gender in social network organization. *PloS One* **12**, e0189873 (2017)
73. V. Kassarnig, A. Bjerre-Nielsen, E. Mones, S. Lehmann, D.D. Lassen, Class attendance, peer similarity, and academic performance in a large field study. *PloS One* **12**, e0187078 (2017)
74. V. Kassarnig et al., Academic performance and behavioral patterns. *EPJ Data Sci.* **7**, 10 (2018)

75. OpenStreetMap contributors. Planet dump retrieved from <https://planet.osm.org>. <https://www.openstreetmap.org> (2019)
76. V. Sekara, S. Lehmann, The strength of friendship ties in proximity sensor data. *PLoS One* **9**, e100915 (2014)
77. J. Zuzanek, R. Mannell, Leisure behaviour and experiences as part of everyday life: the weekly rhythm. *Loisir Soc.* **16**, 31–57 (1993)
78. A. Cuttone, J.E. Larsen, S. Lehmann, in *UbiComp 2014-Adjunct Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 995–1004 (Association for Computing Machinery, Inc, New York, NY, 2014)
79. United States Department Of Defense. *Global positioning system standard positioning service performance standard*. Tech. Rep. 4th Edition (2008)
80. V. Venema, S. Bachner, H.W. Rust, C. Simmer, Statistical characteristics of surrogate data based on geophysical measurements. *Nonlinear Process. Geophys.* **13**, 449–466 (2006)
81. J.A. Scheinkman, B. LeBaron, Nonlinear dynamics and stock returns. *J. Bus.* **62**, 311–337 (1989)
82. W.S. Pritchard, D.W. Duke, K.K. Kriehle, Dimensional analysis of resting human EEG II: surrogate-data testing indicates nonlinearity but not low-dimensional chaos. *Psychophysiology* **32**, 486–491 (1995)
83. M. Wiedermann, J.F. Donges, J. Kurths, R.V. Donner, Spatial network surrogates for disentangling complex system structure from spatial embedding of nodes. *Phys. Rev. E* **93**, 042308 (2016)
84. S. Maslov, K. Sneppen, A. Zaliznyak, Detection of topological patterns in complex networks: correlation profile of the internet. *Physica A* **333**, 529–540 (2004)
85. S. Maslov, K. Sneppen, Specificity and stability in topology of protein networks. *Science* **296**, 910–913 (2002)
86. J. Theiler, D. Prichard, Constrained-realization Monte-Carlo method for hypothesis testing. *Physica D* **94**, 221–235 (1996)
87. G. Zamora-López, V. Zlatić, C. Zhou, H. Štefančić, J. Kurths, Reciprocity of networks with degree correlations and arbitrary degree sequences. *Phys. Rev. E* **77**, 016106 (2008)
88. Y. Artzy-Randrup, L. Stone, Generating uniformly distributed random networks. *Phys. Rev. E* **72**, 056708 (2005)
89. S.A. Stouffer, E.A. Suchman, L.C. Devinney, S.A. Star, R.M. Williams Jr., *The American Soldier: Adjustment During Army Life (Studies in Social Psychology in World War II)*, vol. 1 (1949)
90. M.C. Whitlock, Combining probability from independent tests: the weighted Z-method is superior to Fisher's approach. *J. Evolution. Biol.* **18**, 1368–1373 (2005)
91. C. Spearman, The proof and measurement of association between two things. *Am. J. Psychol.* **15**, 72 (1904)
92. I. Ajzen, The theory of planned behavior. *Organ. Behav. Hum. Decision Process.* **50**, 179–211 (1991)
93. A. Bandura, *Handbook of Personality*, 2nd edn. (Guilford Publications, New York, 1999), pp. 154–196
94. H.P. Young, The evolution of social norms. *Ann. Rev. Econ.* **7**, 359–387 (2015)
95. J.O. Prochaska, B.H. Marcus, *Advances in Exercise Adherence, 161–180* (Human Kinetics Publishers, Champaign, IL, England, 1994)
96. W. Willett et al., Food in the Anthropocene: the EAT–Lancet Commission on healthy diets from sustainable food systems. *Lancet* **393**, 447–492 (2019)
97. D. Gerten et al., Feeding ten billion people is possible within four terrestrial planetary boundaries. *Nat. Sustain.* **3**, 200–208 (2020)