HAL open science

From temporal network data to the dynamics of social relationships

Alain Barrat, Valeria Gelardi, Didier Le Bail, Nicolas Claidière

To cite this version:

Alain Barrat, Valeria Gelardi, Didier Le Bail, Nicolas Claidière. From temporal network data to the dynamics of social relationships. Proceedings of the Royal Society B: Biological Sciences, 2021, 288 (1959), pp.20211164. 10.1098/rspb.2021.1164 . hal-03357923

HAL Id: hal-03357923

https://hal.science/hal-03357923v1

Submitted on 29 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

From temporal network data to the dynamics of social relationships

Valeria Gelardi[a,b], Didier Le Bail[a], Alain Barrat[a,c], Nicolas Claidiere[b,d]

a Aix Marseille Univ, Université de Toulon, CNRS, CPT, Marseille, France
b Aix Marseille Univ, CNRS, LPC, FED3C, Marseille, France
c Tokyo Tech World Research Hub Initiative (WRHI), Tokyo Institute of Technology, Tokyo, Japan
d Station de Primatologie-Celphedia, CNRS UAR846, Rousset, France

Abstract

Networks are well-established representations of social systems, and temporal networks are widely used to study their dynamics. However, going from temporal network data, i.e., a stream of interactions between individuals, to a representation of the social group?s evolution, remains a challenge. Indeed, the temporal network at any specific time contains only the interactions taking place at that time and aggregating on successive time-windows also has important limitations. Here, we present a new framework to study the dynamic evolution of social networks based on the idea that social relationships are interdependent: as the time we can invest in social relationships is limited, reinforcing a relationship with someone is done at the expense of our relationships with others. We implement this interdependence in a parsimonious two-parameter model and apply it to several human and non-human primates? data sets to demonstrate that this model detects even small and short perturbations of the networks that cannot be detected using the standard technique of successive aggregated networks. Our model solves a long-standing problem by providing a simple and natural way to describe the dynamic evolution of social networks, with far-reaching consequences for the study of social networks and social evolution.

Introduction

Social relationships are created and maintained through interactions between individuals, which can last and be repeated over a variety of timescales. Social networks provide convenient representations for the resulting human and non-human animal social structures, where individuals are the nodes of the networks and links (ties) are summaries of their social interactions [Granovetter, 1973, Hinde, 1976, Wasserman and Faust, 1994, Brent et al., 2011]. Recently, the availability of temporally resolved data on interactions between individuals, from various types of communication [Eckmann et al., 2004, Kossinets and Watts, 2006, Onnela et al., 2007, Karsai et al., 2011, Miritello et al., 2013] to face-to-face interactions [Cattuto et al., 2010, Salathé et al., 2010, Stopczynski et al., 2014, Toth et al., 2015] has fueled the development of the field of temporal networks [Holme and Saramäki, 2012, Holme, 2015], which replaces static ties by information on the actual series of interactions on each tie, allowing researchers to further the study of social networks. For instance, aggregating temporal information over successive time windows has made it possible to follow the evolution of ties over larger timescales [Saramäki et al., 2014, Fournet and Barrat, 2014, Gelardi et al., 2019, Aledavood et al., 2015]. Taking into account the temporal features of each tie during a certain time window can also shed light on their strength and persistence [Navarro et al., 2017, Ureña-Carrion et al., 2020].

Finally, researchers have identified temporal structures with no static equivalent [Kovanen et al., 2011, Kobayashi et al., 2019, Galimberti et al., 2018] that reveal interesting patterns of relevance to the analysis of social phenomena or dynamic processes in a social group [Kovanen et al., 2013, Ciaperoni et al., 2020].

Despite this progress, moving from a stream of interactions within a group of individuals, represented by a temporal network, to a meaningful quantification of the strength and evolution of their social relationships, remains a challenge. Indeed, the temporal network seen at any specific time t contains by definition only the interactions taking place at t, while the state of a relationship between two individuals at t depends potentially on the whole history of their previous interactions, both mutual and with others. Temporal aggregation over successive time windows is a commonly used approach to address this issue, but a number of properties of the networks obtained by temporal aggregation on successive windows depend on the window length and placement [Sulo et al., 2010, Krings et al., 2012, Psorakis et al., 2012, Kivelä and Porter, 2015]. Aggregating over increasingly long time windows also averages out relevant temporal information by treating in the same way old and recent interactions, and by not taking into account possible temporal correlations between successive interactions, nor the impact of a single interaction on multiple ties. Moreover, no single natural time scale for aggregation can be defined, as relevant dynamics occur on multiple timescales [Holme, 2013, Saramäki and Moro, 2015, Darst et al., 2016, Masuda and Holme, 2019].

Here, we put forward a new systematic way to transform the stream of interactions between individuals into a continuously evolving representation of the social structure, i.e., a network with time-varying weights, taking into account the temporal ordering of interactions in a non-trivial way. The evolving weight $w_{ij}(t)$ of the tie between nodes i and j represents a quantification of the strength of their relationship at t. Moreover, our framework goes beyond the few such dynamic network models proposed to date [Ahmad et al., 2018, Zuo and Porter, 2021, Jin et al., 2001, Palla et al., 2007], that are based on the idea that the weight of a tie between two individuals strengthens when they interact, and that in the absence of interaction, the tie's weight decays exponentially with time (the timescale of the decay is the model's parameter): these rules of evolution assume that the links between distinct pairs of individuals are independent, while the interdependence of social relationships is instead often well justified. For instance, in the complex social groups formed by humans and other primates [Dunbar and Shultz, 2007, Mitani, 2009, Silk et al., 2010], investing in a social relationship is a costly strategic decision that requires specific cognitive skills [Cheney et al., 1986] and the quality of an individual's social relationships depends on the time invested in them [Dunbar, 2020, Dunbar et al., 2009, Borgeaud et al., 2021]. Thus, the occurrence of a social interaction between two individuals not only reinforces their mutual relationship, but it also weakens the relationships they have with others: the time and energy spent to maintain the tie with an individual is taken from a finite interaction capacity and thus is time that is not spent with others. The framework that we put forward here takes into account this interdependence of social relationships to transform a stream of interactions into an evolving weighted network: with each interaction between two individuals, the weight of their tie increases, while the weights of the ties they have with other individuals decrease. In contrast to other recent temporal network representations [Ahmad et al., 2018, Zuo and Porter, 2021], time itself is not explicit, and the weight of a tie remains unchanged if the corresponding individuals do not interact with anyone. Our framework is therefore linked to the Elo rating method [Elo, 1978] used to rank chess players and analyze animal hierarchies: the dynamics of the system are determined by the pace of interactions between individuals, not by the absolute time between events.

In the following, we define a parsimonious model for the evolution of social ties based on these concepts, with two parameters quantifying respectively the increase in the weight

of a tie $i-j$ when an interaction occurs between i and j, and its decrease when another interaction involving either i or j (but not both) takes place. We then show the relevance of the model by applying it to several data sets describing interactions in groups of human and non-human primates and by using it to automatically detect naturally occurring changes in the groups' dynamics and artificially generated perturbations in the data.

Results

Framework

The concepts highlighted above can be translated in various ways to transform a stream of interactions into evolving tie weights of an evolving directed network $G(t)$. The nodes of the network represent the individuals and the weight $w_{ij}(t)$ represents the strength of the social relationship from i to j at time t. Here we consider a model depending on two parameters, α and β, with the following rules:

- We start from an empty network with uniform weights initialized to zero, i.e., $w_{ij}(0) = 0 \quad \forall i, j$;
- At each time t, we denote by $E(t)$ the set of interacting ties at t. Then, for each tie $(i, j) \in E(t)$, the weights of the ties in which i and j are involved are updated according to

$$\begin{aligned} w_{ij}(t^+) &= w_{ij}(t^-) + \alpha(w_{max} - w_{ij}(t^-)) \\ w_{ji}(t^+) &= w_{ji}(t^-) + \alpha(w_{max} - w_{ji}(t^-)) \end{aligned} \tag{1}$$

  and

$$\begin{aligned} w_{ik}(t^+) &= (1-\beta)w_{ik}(t^-) \;\; \forall k \neq j, (i,k) \notin E(t) \\ w_{jk}(t^+) &= (1-\beta)w_{jk}(t^-) \;\; \forall k \neq i, (j,k) \notin E(t) \; . \end{aligned} \tag{2}$$

The weights of all ties interacting at t thus increase according to (1), while the weights of the neighbouring ties that do not interact at t decrease according to (2). These rules of evolution can be applied to temporal network data expressed either in continuous time (i.e., an interaction between two individuals can occur at any time) or in discrete time (when the data itself has a finite temporal resolution): in the former case, t^- and t^+ stand respectively for the times immediately before and after the interaction; in the latter case, t^- is simply replaced by t and t^+ by $t+1$ in Eqs (1)-(2). The parameter $0 < \alpha < 1$ quantifies how much a tie strength is reinforced by each interaction, while $0 < \beta < 1$ accounts for the weakening of the strength of the ties with other individuals. $w_{max} > 0$ represents the maximum possible value of the weights, which we set to $w_{max} = 1$ without loss of generality. These rules ensure that the weights remain bounded between 0 and w_{max} and, if a tie's weight is zero, it remains so unless there is an interaction involving that tie. Moreover, the weights of the ties between individuals who interact often increase towards w_{max}. The weights thus represent a quantification of the strength of the social ties at each time, taking into account the history of interactions as well as the impact of each interaction of an individual on all its ties, strengthening some and weakening others. Interestingly, in a simple case of random and uncorrelated interactions, the long-time limit of the weight between two individuals can be shown to correspond to their probability of interaction (see Supplementary material).

It is important to stress that, while instantaneous interactions may be undirected (no source nor target individuals such as in face-to-face interaction data), the evolution rules (1)-(2) result in a directed network. Upon an interaction between i and j, w_{ij} and w_{ji}

evolve in the same way; however, when i interacts with other individuals than j, w_{ij} decreases while w_{ji} is not affected. For instance, if j interacts only with i but i interacts with many other individuals, w_{ji} can only increase (upon each interaction with i), while w_{ij} will increase at each interaction of i and j and decrease at each interaction of i with $k \neq j$: w_{ji} thus becomes larger than w_{ij}, reflecting the fact that i is more important to j than j to i. Naturally, the evolution rules could easily be modified in the case of directed interactions, such as in an exchange of text messages or on online social media: for instance, if i sends a message to j, the weights w_{ij} and w_{ik} could be affected more strongly than the weights w_{ji} and w_{jk}. However, this would require the introduction of additional parameters.

Application to empirical data

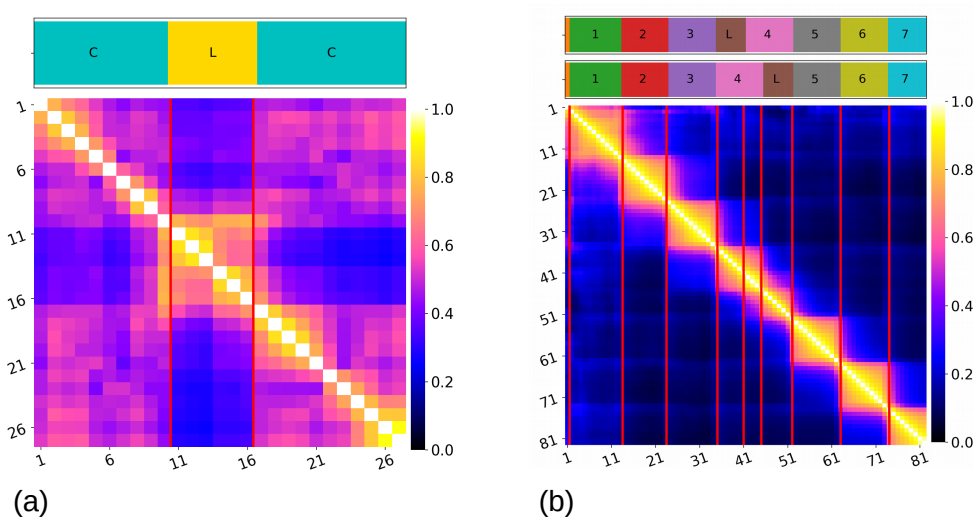

Figure 1. Similarity matrices and school schedules for a day in a French elementary school (a) and in a US middle school (b). Here we use $\alpha = \beta = 0.1$, and the evolving networks are observed every $\Delta = 20$ minutes for the French school and every $\Delta = 5$ minutes for the US school. The horizontal bars give information about the schedule of a school day. The different colors correspond to the class times (indicated by the letter C in (a) and with different numbers in (b)) and lunchtimes (indicated by the letter L). In (b) there are two bars because the students were split into two groups.

Let us first consider empirical data describing face-to-face proximity interactions collected by wearable devices in two schools, namely a French elementary school [Stehlé et al., 2011b] and a US middle school in Utah [Toth et al., 2015, Leecaster et al., 2016], with a temporal resolution of approximately 20 seconds in both cases (the devices collected data on the relative proximity of individuals, and not on their location, see Materials and Methods for more details). Although both cases involve school contexts, the classes were organized very differently, as described in [Stehlé et al., 2011b, Leecaster et al., 2016]: the elementary school students remained in the same classroom for their different classes, while the middle school students changed classrooms between classes.

In each case, we transformed the temporal network data into a network $G(t)$, with the weights evolving according to the rules (1)-(2). For simplicity, we used $\alpha = \beta$ and considered various values of α. We then stored the network $G(t)$ and the tie weights every Δ time steps (i.e., we store $G(n\Delta)$ for $n = 0, 1, 2, \cdots$) and computed the similarity

between each pair of the stored networks $G(n\Delta)$ and $G(n'\Delta)$ (see Materials and Methods). We thus obtained a matrix of similarity values [Masuda and Holme, 2019, Gelardi et al., 2019] for each value of α, shown in Figure 1 for $\alpha = 0.1$.

For instance, in the case of the US school, the large values of similarity found in the diagonal blocks (in yellow) indicate periods in which the network $G(t)$ remains stable, and lower values (off-diagonal) indicate that these periods of stability are different from each other; as seen from the comparison with the school schedule, each diagonal block (period of stability of the network) corresponds to a specific class period. In the French school, the organization in blocks correspond to the class and lunch periods. These matrices thus clearly highlight that the two contexts correspond to different schedules and organizations of interactions and reflect the temporal organization and the periods of importance in the school schedules. We show in Figure S1 in the Supplementary material that, at small α, the weights evolve too slowly and the distinction between the various periods is blurred: the distinction between the various periods becomes clearer as α increases.

Detection of a perturbation

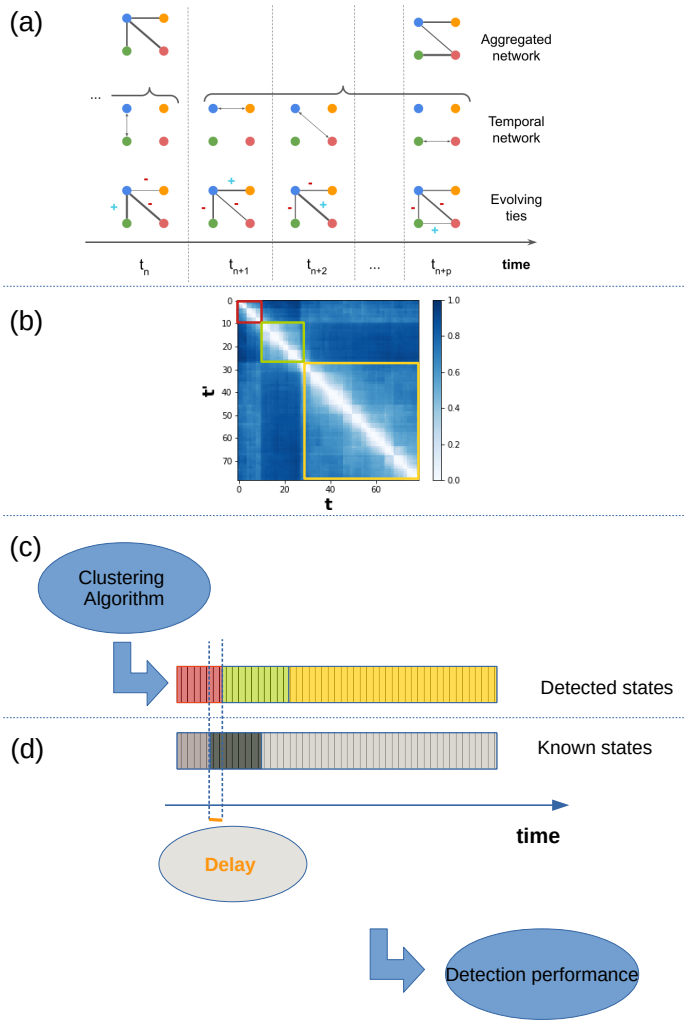

Figure 2. Workflow used to detect discrete states and change points in temporal networks and to estimate the performance of the detection. (a) Creation of a sequence of networks, either by temporal aggregation over successive time windows of p time steps: namely aggregating the p snapshots from t_{mp+1} to t_{mp+p} (m being an integer), and using as weight w^a_{ij} of a link ij in the aggregated network the number of interactions between i and j in this time range, or by transforming the data into an evolving network observed every p time steps. (b) Computation of the similarity between all pairs of networks using the global cosine similarity (see Methods). (c) Classification of the networks into discrete states using a hierarchical clustering algorithm on the distance matrix (the distance between two graphs being simply defined as 1 minus their similarity). (d) To estimate the performance of the classification, we use (i) the Jaccard index between the known and detected time frame of the perturbation and (ii) the delay of the detected perturbation, defined as the number of timestamps between the actual starting time of the perturbation and the smallest timestamp of the following cluster detected.

To go beyond a mere visual inspection of the similarity matrices, we considered a more systematic analysis of the capacity of a temporal network representation, obtained either

by temporal aggregation or through our framework, to detect perturbations in a social group's interaction patterns. To this aim, we first introduced a synthetic perturbation of controled intensity and duration in the temporal network data, for instance by switching the identity of some nodes for a certain duration. We then followed the steps outlined in Fig. 2. First, we used our framework to transform the perturbed temporal network into an evolving weighted graph according to the evolution rules (1)-(2). This weighted graph was observed every p time steps (if the real time duration of one time step is δ, this means that we observed the graph every $\Delta = p\delta$). As a baseline, we also aggregated the temporal network data on successive time windows of duration Δ (Fig. 2a), i.e., we considered at time $t = (m+1)p$ (m being an integer) the aggregation of the p snapshots $t_{mp+1}, t_{mp+2}, \ldots, t_{mp+p}$, using as weights of the aggregated links the number of interactions in that time range. We then followed Masuda et al.'s procedure for detecting states in a temporal network [Masuda and Holme, 2019]. Namely, we computed the cosine similarity matrix between graphs observed at different times (Fig. 2b), transformed it into a distance matrix, and applied a hierarchical clustering algorithm (see Methods) to detect discrete states of the network. As the ground truth perturbation is known, we added a validation step to compare the states obtained by the clustering algorithm to the known perturbation. In this step we quantified the detection performance through two indicators (Fig. 2d), namely the Jaccard index between the sets of timestamps of the actual perturbation and the timestamps of the perturbed state detected, and the delay between the start time of the actual perturbation and the corresponding value obtained through the clustering algorithm (see Methods).

To illustrate the procedure, we considered proximity data from a group of baboons (see Material and Methods). We introduced a small perturbation in the data, namely the exchange of two individual's identities in the data during a certain period. In Figure 3 we use a perturbation duration of 2 hours and show the resulting similarity matrices between the weighted evolving networks obtained for three values of $\alpha = \beta$ and observed every 30 minutes. We also measure and show the detection performance as a function of α. Strikingly, even such a small and short perturbation is well detected over a wide range of α values, excepting the smallest and largest. Notably, the perturbation is instead not detected when using temporal aggregation over successive windows of 30 minutes. The perturbation is not detected for small α values, as the resulting network dynamics is too slow: Fig. 3(a) shows that the network remains very similar to itself during the whole explored time range. However, we observe a sharp increase in detection performance as soon as the resulting dynamics are fast enough. At very large α values, the detection becomes impossible again because each single interaction induces large changes in the weights, leading to rapidly changing dynamics.

In the supplementary material we also considered daily and weekly time scales of observation with perturbations lasting days or weeks (Figures S5 and S6) At such timescales our framework resulted in a perfect or almost perfect detection of the perturbation for a wide range of values of the parameter α (values of the Jaccard index close or equal to 1), while the perturbation was rarely detected when using daily aggregated networks. Note that, at this stage, no single optimal value of the parameter emerges: rather, the existence of the perturbation can be assessed with some degree of confidence by the fact that the procedure detects the same perturbed state over a range of parameter values.

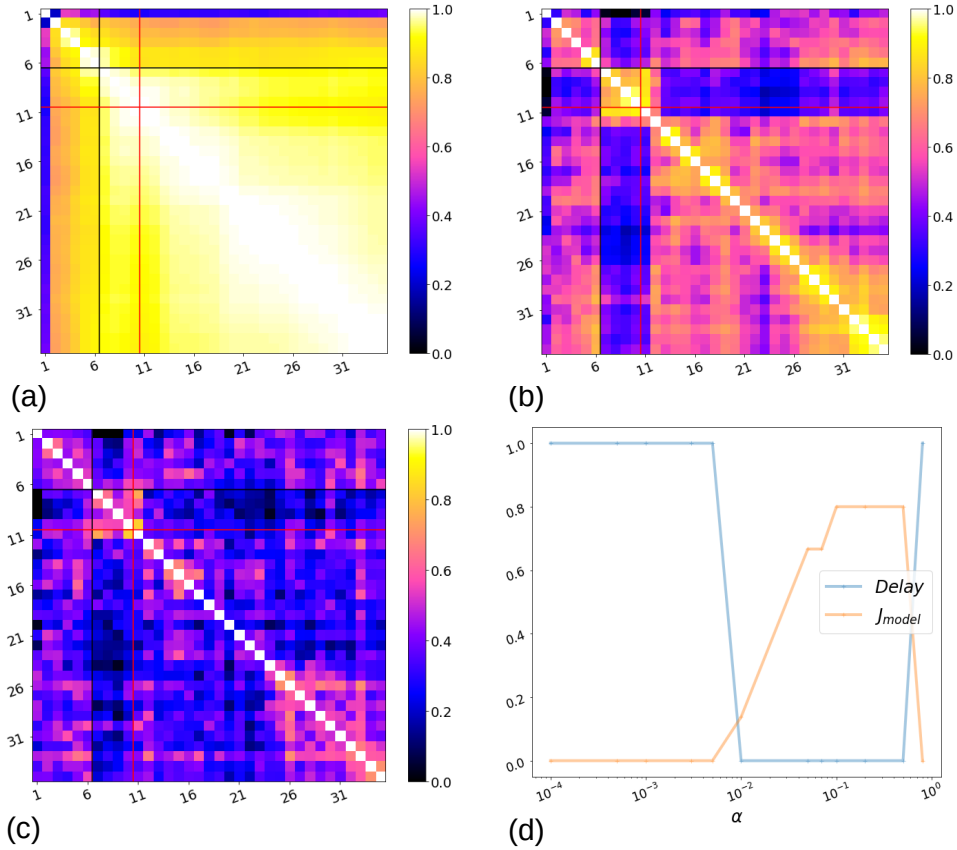

Figure 3. Detection of a simulated perturbation in a temporal network data set. Here we consider one day of proximity data collected from a group of 13 baboons (see Material and Methods). The data, with a temporal resolution of 20 seconds, are artificially perturbed by exchanging the identity of two nodes for 2 hours. The resulting temporal network is transformed into a weighted evolving network as described in the text, and this network is observed every 30 minutes. Panels (a), (b), (c) represent the resulting cosine similarity matrices for values of $\alpha = \beta = 0.001$, 0.1, 0.5, respectively. The black and red lines correspond to the (known) start and end times of the perturbation. Panel (d) shows the performance detection of network states (see Fig. 2), computed from the hierarchical clustering analysis applied to the distance matrices, with the number of clusters fixed to $C = 3$. The blue line represents the relative delay in the detection of the perturbation, i.e. the difference between the known beginning of the perturbation (black line) and the detection of a new network state, divided by the total length of the perturbation. The orange line indicates the Jaccard index between the known perturbation and the perturbation detected by the clustering algorithm. The detection performance relative to the aggregated network is not presented because no cluster detected by the algorithm corresponded to the simulated perturbation. The similarity matrices for the aggregated network with different time window lengths are shown in figure S4 of the Supplementary material.

We further investigated whether using different values for the parameters α and β could lead to an improvement in the detection performance. We show the results in Figure 4 for the same data and perturbation as for Figure 3 (see also Supplementary Figure S7). We found that the detection performance worsened for $\beta < \alpha$, while it increased for $\beta > \alpha$. This can be understood as follows: at small values of $\alpha = \beta$, the weights' increase

and decrease are too slow upon a brutal change in the interactions, and the perturbation is not well detected; this can be compensated by a larger β that induces a fast decrease of the weights of non-interacting ties. For instance, if a node i was repeatedly interacting with a node j before the perturbation, but interacts more with another one k during the perturbation, w_{ij} decreases quickly as soon as the perturbation starts, and this can be easily detected even if the small value of α makes w_{ik} increase only slowly.

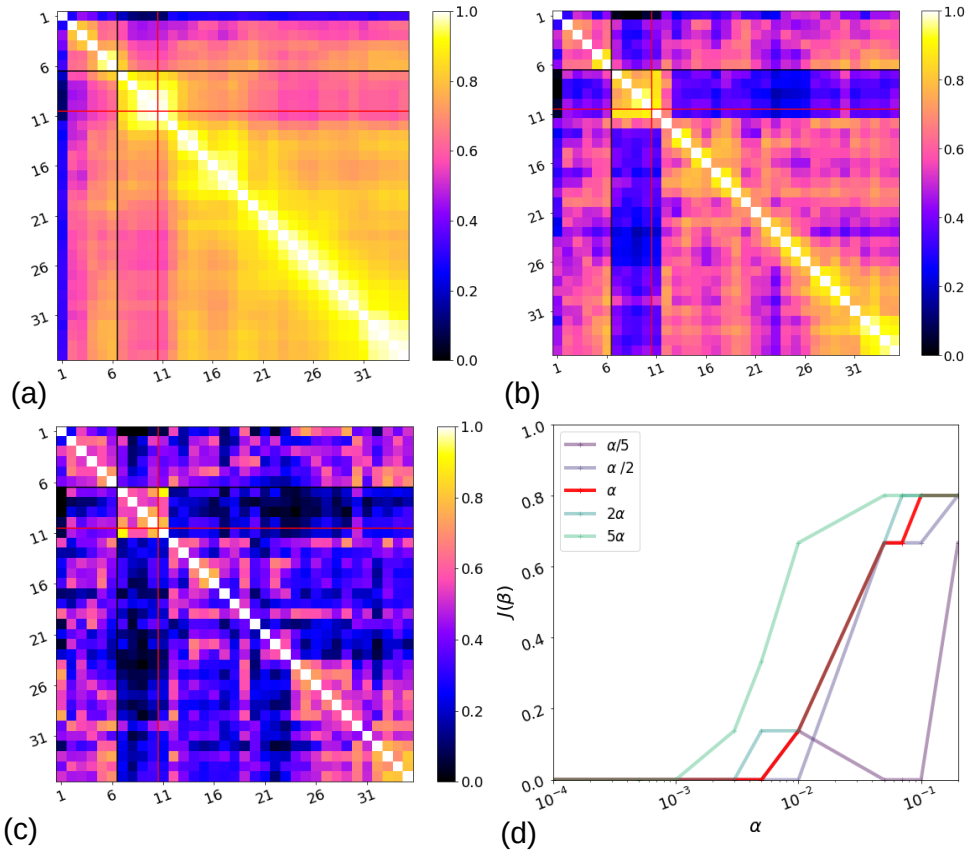

Figure 4. Performance of the detection of simulated changes when varying β. Panels (a), (b), (c) represent the cosine similarity matrices for $\alpha = 0.1$ and values of $\beta = \alpha/5, \alpha, 5\alpha$, respectively, using the same simulated perturbation as in Fig. 3. Panel (d) shows the performance detection, namely the Jaccard index between the real and detected perturbations, as a function of α and for different values of β.

Discussion

How can we represent the evolution of social relationships? Temporal aggregation procedures have provided in-depth knowledge on the dynamics of social networks at various timescales [Aledavood et al., 2015, Saramäki et al., 2014, Fournet and Barrat, 2014, Saramäki et al., 2014] and are also used for data-driven numerical simulations of dynamic processes of networks [Stehlé et al., 2011a], possibly with aggregation schemes adapted to the specific process under study [Holme, 2013]. They however lose information on the temporal ordering of interactions and do not take into account the impact of interactions on neighbouring ties.

Here, we have presented a new framework to go from a stream of interactions to a

quantification of the strength of ties in a social network and to study their dynamic evolution, based on the idea that social relationships are interdependent: since time resources to invest in social relationships are limited [Borgeaud et al., 2021], reinforcing a relationship with someone is necessarily done at the expense of the relationships with others. While this idea can be translated in various ways into specific rules of evolution, here we have focused on a parsimonious two-parameter model rather than on more complex alternatives. We have applied this model to several data sets of interest, showing its ability to highlight changes in the dynamics of the networks and differences between data representing interactions in different contexts. Moreover, we have systematically tested its ability to detect a perturbation in the network at different timescales. Notably, our results show that this simple model yields a high detection performance even for small and short perturbations that cannot be detected by the dynamics of successive aggregated networks. Overall, our framework is able to detect perturbations in a broad range of conditions spanning different data sets and various timescales and perturbations. This point is particularly important as real-world variations in social relationships can occur on a broad range of timescales, from hours to days to months. For instance, despite decades of research, the timescale of the exchange of favors in primates (e.g., grooming in exchange for other commodities) is still very uncertain [Sánchez-Amaro and Amici, 2015]. Our framework does not require a specification of the timescale of changes to be detected a priori and in our current study, the temporal organization relevant in the school data sets and the artificial perturbations introduced were detected in a broad range of parameter values. However, whether the same range of parameter values could be used to detect any kind of temporal patterns in an unknown data set remains an open question. To explore a new data set, a sensible path would be either to consider known changes in the network (if such information is available) or to simulate a plausible perturbation of the data, and scan parameter values to check when such changes or perturbations are detected. A natural hypothesis would then be that the temporal patterns present in the data should be detected using the same range of parameter values. Finding the same temporal patterns on a range of parameter values would also give confidence on the significance of these patterns. To investigate this point in more detail, further research will use a collection of temporal network models with tunable parameters and different levels of complexity and realism [Perra et al., 2012, Laurent et al., 2015]. Introducing perturbations of various types (e.g., changes in the community structure over time, changes in activity, etc), and of tunable intensity and duration, will allow us to systematically explore the detection capacities and limitations of the evolving weighted graph framework introduced here.

Our focus here has been on the issue of detecting when a perturbation occur in a social network, as it is a known challenging task. Once the existence of a perturbation is established, investigating more detailed quantities such as the distributions of local similarities between the neighborhoods of individuals, or individual trajectories of the similarity of each local neighborhood between successive times (see [Gelardi et al., 2019]) could also make possible to detect which nodes or links are involved in the perturbation. Alternatively, this might also be achieved by considering individual matrices of similarities between neighborhoods of individuals at different times, and applying clustering algorithms to each such matrix.

An interesting property of our framework is that, starting from a stream of undirected interactions, it yields directed ties, because individuals do not invest in their mutual relationship in the same way: for instance, one individual may spend 80% of her time with another, while the other spends only 50% of her time with the first). The weights on each tie can therefore be more or less symmetric, and it would be interesting to investigate this point with respect to the social relationships under study. To this aim, one would need to compare the directed network obtained from our framework to other independent measures, such as friendship surveys in a human group or grooming behavior

in non-human primates.

While we have limited our current study to a simple version of the model, several extensions could be of interest. In particular, directed interactions between individuals (such as phone or online messages) could be taken into account, with different impacts on the ties originating from the source of the interaction and on the ties originating from the interaction target. Moreover, one could take into account individual characteristics that are often important in relationships by introducing α and β coefficients that depend on individual characteristics such as age, sex, kinship or rank. This would be appropriate for instance when the costs and benefits of interactions differ between low ranking and high ranking individuals [Silk et al., 1999]. Our framework could also provide an extension of models of social contagion or consensus formation [Guilbeault et al., 2018, Rosenthal et al., 2015]: in the spirit of [Holme, 2013], it could help take into account that interactions with different individuals and at different times are not equivalent, by providing a way to dynamically weigh these interactions (an interaction along a currently strong tie would weigh more than along a weak tie).

Finally, our focus here has been on social relationships of primates in particular, but our conceptual contribution lies in taking into account the interdependence of ties in evolving networks. Thus, our framework may well apply to other systems where such interdependence is relevant, possibly with changes in the rules of evolution. In particular, we have considered that an interaction between two nodes reinforces the tie between them at the expense of ties with other nodes, but in other contexts, the increase of a tie's weight may in fact increase the importance of related ties. For instance, if a new flight route is created between two airports, passengers may take other flights to connect to other destinations, increasing the traffic on the corresponding routes [Barrat et al., 2004]. Taking these interactions into account might open up new perspectives to study the evolution of these types of infrastructure networks [Sugishita and Masuda, 2021].

Materials and Methods

Data Description and Aggregation

We used three data sets of time-stamped dyadic interactions between individuals corresponding to physical proximity events:

- A data set of contacts between students in an urban public middle school in Utah (USA) measured by an infrastructure based on wireless ranging enabled nodes (WRENs) [Toth et al., 2015, Leecaster et al., 2016]. The data, available in reference [Leecaster et al., 2016], involve 679 students in grades 7 and 8 (typical age range from 12 to 14 years old). Participants were recorded over two consecutive days.
- A data set gathered by the SocioPatterns collaboration (http://www.sociopatterns.org/) using radio-frequency identification devices in an elementary school in France. These sensors record face-to-face contacts within a distance of about 1.5m. The data were aggregated with a temporal resolution of 20 seconds (for more details see [Cattuto et al., 2010]). Contacts between 242 participants (232 elementary school children and 10 teachers) were recorded over two consecutive days [Stehlé et al., 2011b]. The data are publicly available at http://www.sociopatterns.org/datasets.
- Data of proximity contacts within a group of Guinea baboons (*Papio papio*), collected from June to November 2019. A subgroup of 13 baboons consisting only of juveniles and adults were equipped with leather collars fitted with the wearable proximity sensors developed by the SocioPatterns collaboration (see [Gelardi et al., 2020]).

Similarity between networks

To compare the weighted evolving networks (or aggregated networks) observed at different times, we chose the global cosine similarity between the two vectors formed by the list of all the weights

in each network (using a weight 0 if a link was not present).

A cosine similarity measure is generally defined between two vectors and is bounded between -1 and $+1$. It takes the value $+1$ if the vectors are proportional with a positive proportionality constant, a value of -1 if the proportionality constant is negative, and 0 if they are perpendicular. For positive weights, as in our case, it is bounded between 0 and 1.

In the case of two networks, G_1 and G_2, the global cosine similarity is defined as:

$$GCS_{G_1,G_2} = \frac{\sum_{i>j} w_{ij}^{(1)} w_{ij}^{(2)}}{\sqrt{\sum_{i>j} \left(w_{ij}^{(1)}\right)^2} \sqrt{\sum_{i>j} \left(w_{ij}^{(2)}\right)^2}} \ , \tag{3}$$

where the subscripts $^{(1)}$ and $^{(2)}$ denote the weights of the links in the networks G_1 and G_2, respectively.

Clustering method

To obtain discrete system states by hierarchical clustering, we used the "fcluster" function of the `scipy.hierarchy` library from the SciPy module in Python. The function is applied directly on the $t_{max} \times t_{max}$ distance matrix d, obtained by transforming the cosine similarity matrix elements for each pair of timestamps (t, t'): $d(t, t') = 1 - CS(t, t')$. To define the distance between clusters, we used the "average" method in the "linkage" function of the library. We set the number of clusters to $C = 3$, corresponding to the periods before, during and after the perturbation.

Detection performance

To assess the performance of our model, our rationale was that the temporal network representation should allow us to detect changes in the social structure, and the quality of the detection entails two aspects: it has to be detected (i) without delays and (ii) clearly, i.e., social changes have to be distinguished from the noise represented by "ordinary" variations in social activity. In particular, a perturbation is said to be well detected if one of the states found by the clustering algorithm includes all the timestamps of the perturbation and only those.

We first verified that one of the detected clusters could be associated with the perturbation in the data. To this end we determined that each cluster would correspond to a set of contiguous timestamps (thus forming an interval), with the smallest time equal to or larger than the initial timestamp of the perturbation, and largest time equal to or larger than the final timestamp of the perturbation. A first measure to evaluate the quality of the detection was then given by the "delay" between the actual and the detected perturbation (the number of timestamps between the actual starting time of the perturbation and the smallest timestamp of the second cluster detected; see Figure 2d). The second measure was given by the Jaccard index J between the set of time steps during which the actual perturbation takes place, $\mathcal{T}_{groundtruth}$, and the set of time steps of the state detected as a perturbation by the clustering procedure, $\mathcal{T}_{detected}$:

$$J = \frac{|\mathcal{T}_{groundtruth} \cap \mathcal{T}_{detected}|}{|\mathcal{T}_{groundtruth} \cup \mathcal{T}_{detected}|} \tag{4}$$

Acknowledgments

Many thanks to Yousri Marzouki for planting the seed of the idea for this article and to Clément Sire for interesting discussions and the suggestion of studying case $\beta \neq \alpha$ in the model. A.B. was supported by the ANR project DATAREDUX (ANR-19-CE46-0008) and JSPS KAKENHI Grant Number JP 20H04288.

References

Ahmad et al., 2018. Ahmad, W., Porter, M. A., and Beguerisse-Díaz, M. (2018). Tie-decay temporal networks in continuous time and eigenvector-based centralities. *arXiv preprint arXiv:1805.00193*.

Aledavood et al., 2015. Aledavood, T., Lehmann, S., and SaramÃ€ki, J. (2015). Digital daily cycles of individuals. *Frontiers in Physics*, 3:73.

Barrat et al., 2004. Barrat, A., Barthélemy, M., and Vespignani, A. (2004). Weighted evolving networks: Coupling topology and weight dynamics. *Phys. Rev. Lett.*, 92:228701.

Borgeaud et al., 2021. Borgeaud, C., Jankowiak, B., Aellen, M., Dunbar, R. I., and Bshary, R. (2021). Vervet monkeys socialize more when time budget constraints are experimentally reduced. *Ethology*, n/a(n/a).

Brent et al., 2011. Brent, L. J., Lehmann, J., and Ramos-Fernández, G. (2011). Social network analysis in the study of nonhuman primates: A historical perspective. *American Journal of Primatology*, 73(8):720–730.

Cattuto et al., 2010. Cattuto, C., Van den Broeck, W., Barrat, A., Colizza, V., Pinton, J.-F., and Vespignani, A. (2010). Dynamics of person-to-person interactions from distributed RFID sensor networks. *PLOS ONE*, 5(7):1–9.

Cheney et al., 1986. Cheney, D., Seyfarth, R., and Smuts, B. (1986). Social relationships and social cognition in nonhuman primates. *Science*, 234(4782):1361–1366.

Ciaperoni et al., 2020. Ciaperoni, M., Galimberti, E., Bonchi, F., Cattuto, C., Gullo, F., and Barrat, A. (2020). Relevance of temporal cores for epidemic spread in temporal networks. *Scientific Reports*, 10(1):12529.

Darst et al., 2016. Darst, R. K., Granell, C., Arenas, A., Gómez, S., Saramäki, J., and Fortunato, S. (2016). Detection of timescales in evolving complex systems. *Scientific Reports*, 6:39713.

Dunbar, 2020. Dunbar, R. (2020). Structure and function in human and primate social networks: Implications for diffusion, network stability and health. *Proceedings of the Royal Society A*, 476(2240):20200446.

Dunbar et al., 2009. Dunbar, R. I., Korstjens, A. H., Lehmann, J., and Project, B. A. C. R. (2009). Time as an ecological constraint. *Biological Reviews*, 84(3):413–429.

Dunbar and Shultz, 2007. Dunbar, R. I. and Shultz, S. (2007). Evolution in the social brain. *science*, 317(5843):1344–1347.

Eckmann et al., 2004. Eckmann, J. P., Moses, E., and Sergi, D. (2004). Entropy of dialogues creates coherent structures in e-mail traffic. *PNAS*, 101:14333–14337.

Elo, 1978. Elo, A. E. (1978). *The rating of chessplayers, past and present*. Arco Pub.

Fournet and Barrat, 2014. Fournet, J. and Barrat, A. (2014). Contact patterns among high school students. *PLoS ONE*, 9(9):e107878.

Galimberti et al., 2018. Galimberti, E., Barrat, A., Bonchi, F., Cattuto, C., and Gullo, F. (2018). Mining (maximal) span-cores from temporal networks. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 107–116. ACM.

Gelardi et al., 2019. Gelardi, V., Fagot, J., Barrat, A., and Claidière, N. (2019). Detecting social (in)stability in primates from their temporal co-presence network. *Animal Behaviour*, 157:239–254.

Gelardi et al., 2020. Gelardi, V., Godard, J., Paleressompoulle, D., Claidière, N., and Barrat, A. (2020). Measuring social networks in primates: wearable sensors versus direct observations. *Proceedings of the Royal Society A*, 476:20190737.

Granovetter, 1973. Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380.

Guilbeault et al., 2018. Guilbeault, D., Becker, J., and Centola, D. (2018). Complex contagions: A decade in review. In *Complex Spreading Phenomena in Social Systems*, pages 3–25. Springer.

Hinde, 1976. Hinde, R. A. (1976). Interactions, relationships and social structure. *Man*, 11(1):1–17.

Holme, 2013. Holme, P. (2013). Epidemiologically optimal static networks from temporal network data. *PLoS Comput Biol*, 9(7):e1003142.

Holme, 2015. Holme, P. (2015). Modern temporal network theory: a colloquium. *The European Physical Journal B*, 88(9):234.

Holme and Saramäki, 2012. Holme, P. and Saramäki, J. (2012). Temporal networks. *Physics reports*, 519(3):97–125.

Jin et al., 2001. Jin, E. M., Girvan, M., and Newman, M. (2001). Structure of growing social networks. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 64 4 Pt 2:046132.

Karsai et al., 2011. Karsai, M., Kivelä, M., Pan, R. K., Kaski, K., Kertész, J., Barabási, A.-L., and Saramäki, J. (2011). Small but slow world: How network topology and burstiness slow down spreading. *Physical Review E*, 83(2).

Kivelä and Porter, 2015. Kivelä, M. and Porter, M. A. (2015). Estimating interevent time distributions from finite observation periods in communication networks. *Physical Review E*, 92(5):052813.

Kobayashi et al., 2019. Kobayashi, T., Takaguchi, T., and Barrat, A. (2019). The structured backbone of temporal social ties. *Nature Communications*, 10(1):220.

Kossinets and Watts, 2006. Kossinets, G. and Watts, D. J. (2006). Empirical analysis of an evolving social network. *Science*, 311:88–90.

Kovanen et al., 2011. Kovanen, L., Karsai, M., Kaski, K., Kertész, J., and Saramäki, J. (2011). Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(11):P11005.

Kovanen et al., 2013. Kovanen, L., Kaski, K., Kertész, J., and Saramäki, J. (2013). Temporal motifs reveal homophily, gender-specific patterns, and group talk in call sequences. *Proceedings of the National Academy of Sciences*, 110(45):18070–18075.

Krings et al., 2012. Krings, G., Karsai, M., Bernhardsson, S., Blondel, V. D., and Saramäki, J. (2012). Effects of time window size and placement on the structure of an aggregated communication network. *EPJ Data Science*, 1(1):4.

Laurent et al., 2015. Laurent, G., Saramäki, J., and Karsai, M. (2015). From calls to communities: a model for time-varying social networks. *The European Physical Journal B*, 88(11):301.

Leecaster et al., 2016. Leecaster, M., Toth, D. J. A., Pettey, W. B. P., Rainey, J. J., Gao, H., Uzicanin, A., and Samore, M. (2016). Estimates of social contact in a middle school based on self-report and wireless sensor data. *PLOS ONE*, 11(4):1–21.

Masuda and Holme, 2019. Masuda, N. and Holme, P. (2019). Detecting sequences of system states in temporal networks. *Scientific reports*, 9(1):1–11.

Miritello et al., 2013. Miritello, G., Lara, R., and Moro, E. (2013). Time allocation in social networks: correlation between social structure and human communication dynamics. In *Temporal networks*, pages 175–190. Springer.

Mitani, 2009. Mitani, J. C. (2009). Male chimpanzees form enduring and equitable social bonds. *Animal Behaviour*, 77(3):633–640.

Navarro et al., 2017. Navarro, H., Miritello, G., Canales, A., and Moro, E. (2017). Temporal patterns behind the strength of persistent ties. *EPJ Data Science*, 6(1):31.

Onnela et al., 2007. Onnela, J.-P., Saramäki, J., Hyvönen, J., Szabó, G., Lazer, D., Kaski, K., Kertész, J., and Barabási, A.-L. (2007). Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences*, 104(18):7332–7336.

Palla et al., 2007. Palla, G., Barabási, A.-L., and Vicsek, T. (2007). Community dynamics in social networks. *Fluctuation and Noise Letters*, 07(03):L273–L287.

Perra et al., 2012. Perra, N., Gonçalves, B., Pastor-Satorras, R., and Vespignani, A. (2012). Activity driven modeling of time varying networks. *Scientific Reports*, 2(1):469.

Psorakis et al., 2012. Psorakis, I., Roberts, S. J., Rezek, I., and Sheldon, B. C. (2012). Inferring social network structure in ecological systems from spatio-temporal data streams. *Journal of the Royal Society Interface*, 9(76):3055–3066.

Rosenthal et al., 2015. Rosenthal, S. B., Twomey, C. R., Hartnett, A. T., Wu, H. S., and Couzin, I. D. (2015). Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion. *Proceedings of the National Academy of Sciences*, 112(15):4690–4695.

Salathé et al., 2010. Salathé, M., Kazandjieva, M., Lee, J. W., Levis, P., Feldman, M. W., and Jones, J. H. (2010). A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences*, 107(51):22020–22025.

Saramäki et al., 2014. Saramäki, J., Leicht, E. A., López, E., Roberts, S. G. B., Reed-Tsochas, F., and Dunbar, R. I. M. (2014). Persistence of social signatures in human communication. *Proceedings of the National Academy of Sciences*, 111(3):942–947.

Saramäki and Moro, 2015. Saramäki, J. and Moro, E. (2015). From seconds to months: an overview of multi-scale dynamics of mobile telephone calls. *The European Physical Journal B*, 88(6):164.

Silk et al., 1999. Silk, J., Cheney, D., and Seyfarth, R. (1999). The structure of social relationships among female savanna baboons in moremi reserve, botswana. *Behaviour*, 136:679–703.

Silk et al., 2010. Silk, J. B., Beehner, J. C., Bergman, T. J., Crockford, C., Engh, A. L., Moscovice, L. R., Wittig, R. M., Seyfarth, R. M., and Cheney, D. L. (2010). Female chacma baboons form strong, equitable, and enduring social bonds. *Behavioral Ecology and Sociobiology*, 64(11):1733–1747.

Stehlé et al., 2011a. Stehlé, J., Voirin, N., Barrat, A., Cattuto, C., Colizza, V., Isella, L., Régis, C., Pinton, J.-F., Khanafer, N., Van den Broeck, W., et al. (2011a). Simulation of an seir infectious disease model on the dynamic contact network of conference attendees. *BMC medicine*, 9(1):87.

Stehlé et al., 2011b. Stehlé, J., Voirin, N., Barrat, A., Cattuto, C., Isella, L., Pinton, J.-F., Quaggiotto, M., Van den Broeck, W., Régis, C., Lina, B., et al. (2011b). High-resolution measurements of face-to-face contact patterns in a primary school. *PloS one*, 6(8):e23176.

Stopczynski et al., 2014. Stopczynski, A., Sekara, V., Sapiezynski, P., Cuttone, A., Madsen, M. M., Larsen, J. E., and Lehmann, S. (2014). Measuring large-scale social networks with high resolution. *PloS one*, 9(4):e95978.

Sugishita and Masuda, 2021. Sugishita, K. and Masuda, N. (2021). Recurrence in the evolution of air transport networks. *Sci Rep*, 11:5514.

Sulo et al., 2010. Sulo, R., Berger-Wolf, T., and Grossman, R. (2010). Meaningful selection of temporal resolution for dynamic networks. In *Proceedings of the Eighth Workshop on Mining and Learning with Graphs*, pages 127–136.

Sánchez-Amaro and Amici, 2015. Sánchez-Amaro, A. and Amici, F. (2015). Are primates out of the market? *Animal Behaviour*, 110:51–60.

Toth et al., 2015. Toth, D. J., Leecaster, M., Pettey, W. B., Gundlapalli, A. V., Gao, H., Rainey, J. J., Uzicanin, A., and Samore, M. H. (2015). The role of heterogeneity in contact timing and duration in network models of influenza spread in schools. *Journal of The Royal Society Interface*, 12(108):20150279.

Ureña-Carrion et al., 2020. Ureña-Carrion, J., Saramäki, J., and Kivelä, M. (2020). Estimating tie strength in social networks using temporal communication data. *EPJ Data Science*, 9(1):37.

Wasserman and Faust, 1994. Wasserman, S. and Faust, K. (1994). *Social Network Analysis: Methods and applications*. Cambridge University Press, Cambridge.

Zuo and Porter, 2021. Zuo, X. and Porter, M. A. (2021). Models of continuous-time networks with tie decay, diffusion, and convection. *Phys. Rev. E*, 103:022304.

Supporting Information

Supplementary Note 1. Interpretation of the weights in a simple case

We consider the simple example of N nodes among which interactions occur independently and at random, with uniform probability p at each timestep. We denote by $w_e(t)$ the weight of a directed edge $e = (i, j)$. At each time step, this weight can either remain the same (if i does not interact with any other node), increase (if i and j interact) or decrease (if i and j do not interact together but i interacts with $k \neq j$). Thus: $w_e(t+1) =$

(i) $(1-\alpha)w_e(t) + \alpha$ with probability p;

(ii) $w_e(t)$ with probability $P = (1-p)^{N-1}$;

(iii) $(1-\alpha)w_e(t)$ with probability $1 - p - P$.

It is possible to rewrite this recursion relation under the form :

$$w_e(t+1) - w_e(t) = \alpha\left(-w_e(t) + \epsilon_e(t)\right)\eta_e(t)$$

where $\eta_e(t)$ and $\epsilon_e(t)$ are independent Bernoulli variables of respective parameters $1 - P$ and $\frac{p}{1-P}$. Indeed, the case $\eta_e(t) = 1, \epsilon_e(t) = 1$ occurs with probability p and corresponds to (i); the cases $\eta_e(t) = 0$ correspond to (ii) (and happen with probability P); the case $\eta_e(t) = 1, \epsilon_e(t) = 0$ occurs with probability $(1-P)(1-p/(1-P)) = 1 - p - P$ and corresponds to (iii).

As $w_e(t)$ changes only when $\eta_e(t) = 1$, we denote by $(t_k)_{k\in\mathbb{N}}$ the sequence of times $t_k < t_{k+1}$ such that $\eta_e(t) \neq 0$, and we define $\tilde{w}_e(k) = w_e(t_k)$ and $\tilde{\epsilon}_e(k) = \epsilon_e(t_k)$. Then we have :

$$\tilde{w}_e(k+1) - \tilde{w}_e(k) = \alpha\left(-\tilde{w}_e(k) + \tilde{\epsilon}_e(k)\right) \tag{1}$$

As $t \to \infty$, the relation between w and $\tilde{w}$ reads $w_e(t) \simeq \tilde{w}_e(t(1-P))$ (as the probability for η to be 1 is $1 - P$), so in particular they have the same limit at large times. We can thus restrict ourselves to the study of $\tilde{w}_e$. We can obtain he expression for $\tilde{w}_e(k)$ by writing (1) for $\ell = k-1, k-2, \ldots, 0$, multiplying each equation by $(1-\alpha)^{k-\ell}$ and summing them. We then obtain

$$\tilde{w}_e(k) = (1-\alpha)^k w_e(0) + \frac{\alpha}{1-\alpha}\sum_{l=0}^{k-1}\tilde{\epsilon}_e(l)(1-\alpha)^{k-l},$$

where $w_e(0)$ is the initial condition for the weight w_e.

As the $\tilde{\epsilon}_e(l)$ are identically distributed variables of average $\frac{p}{1-P}$, and using $\sum_{l=0}^{\infty}(1-\alpha)^l = \frac{1}{\alpha}$, we obtain for the long time limit of the average of $\tilde{w}_e$, $\mu_e = \lim_{t\to\infty}\langle w_e(t)\rangle = \lim_{k\to\infty}\langle\tilde{w}_e(k)\rangle$:

$$\mu_e = \mu_e(\epsilon) = \frac{p}{1-P}$$

.

For $p = d/N$ with d finite for $N \to \infty$, i.e., with each node interacting with a finite number of other nodes in the thermodynamic limit at each time, $P \sim e^{-d}$ and $\mu_e = \frac{p}{1-e^{-d}}$.

As $\mu_e \propto p$, we can interpret the long-time limit of the edge weight $w_{i,j}(t)$ as the probability that i interacts with j at time t, modulated by a factor depending on the properties of the neighbourhood of i.

Supplementary Note 2. Supplementary figures

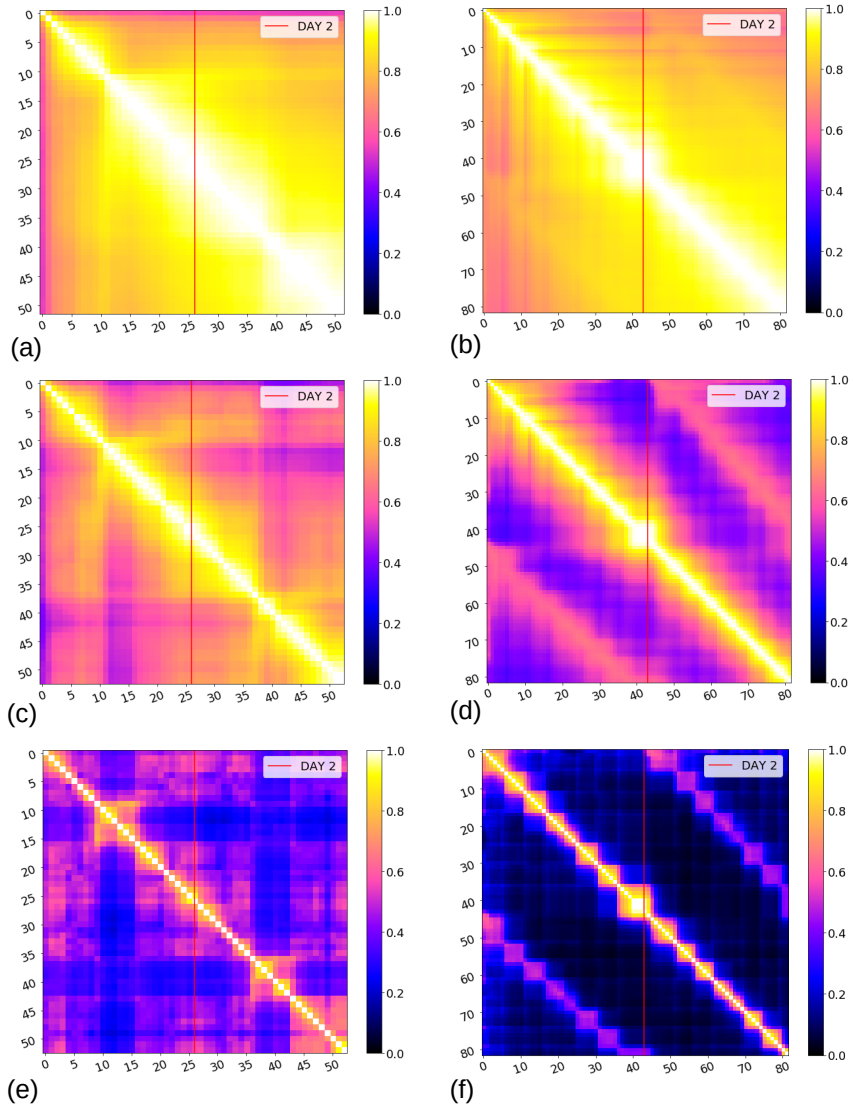

Figure S1. Cosine Similarity matrices for the French elementary school and Utah middle school, using different values of α. The first column (a),(c),(e) refers to the French elementary school data, sampling the network every 20 minutes; the second column refers to the Utah middle school data, sampling the network every 10 minutes. The global cosine similarity is calculated between every pair of observed networks. The vertical red lines indicate the time between the first and the second day of data collection (the data of the two days are concatenated). The different rows refer to three different values of the parameter α, namely $\alpha = 0.001$ (a) (b), $\alpha = 0.01$ (c) (d) and $\alpha = 0.1$ (e)(f). For $\alpha = 0.1$ the different structures of the schools schedules emerge more clearly. We observe for the French school that the structure of the networks during the two lunches are different, and that the structure during the classes remain similar in the two days. For the Utah middle school, we observe some similarity between the corresponding class periods in the two different days, indicating that the seating arrangements in each class are probably similar in different days.

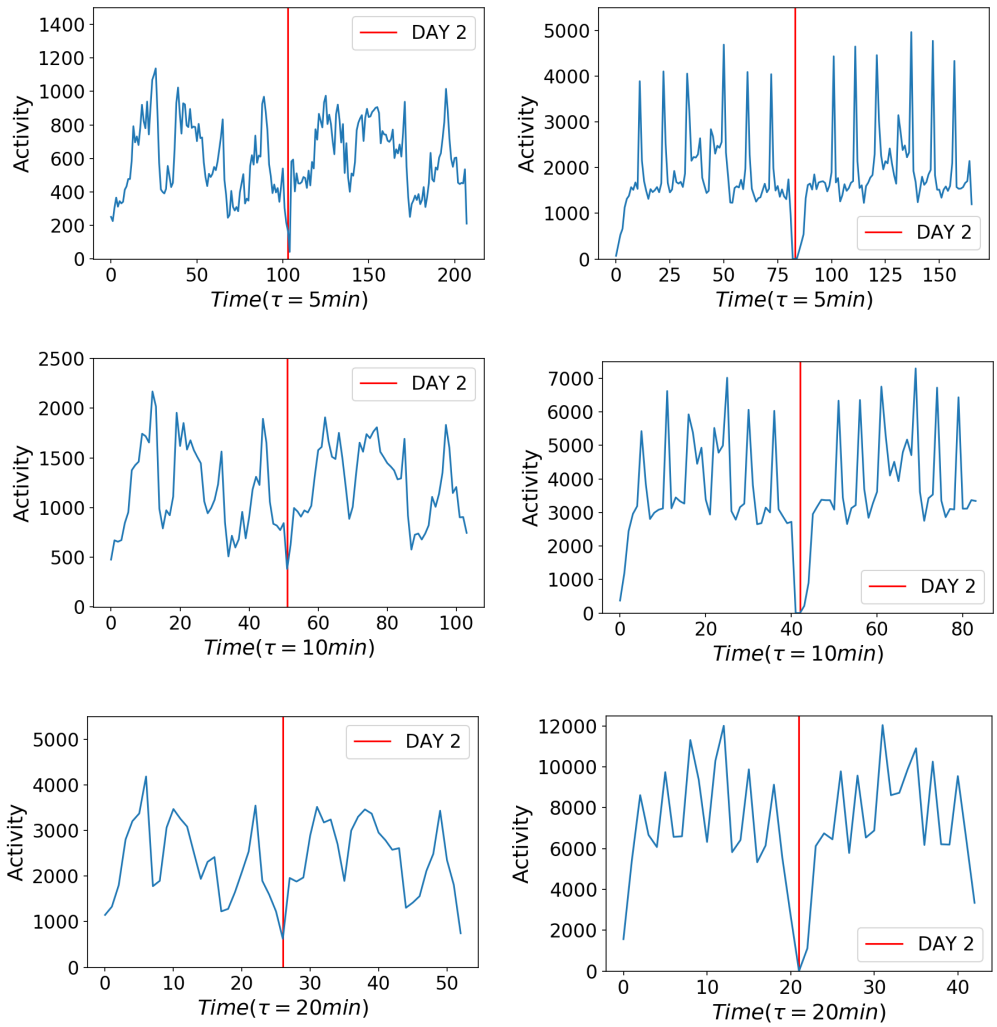

Figure S2. Event rate (number of temporal edges per timestamp) vs. time for the French elementary school data (left column) and the Utah middle school data (right column). Each row corresponds to a different time resolution (respectively 5, 10 and 20 minutes). While the activity timelines can highlight some timestamps of relevance, they do not shed information on the evolution of the structure of the system (i.e., whether the networks at different times have similar or different structures).

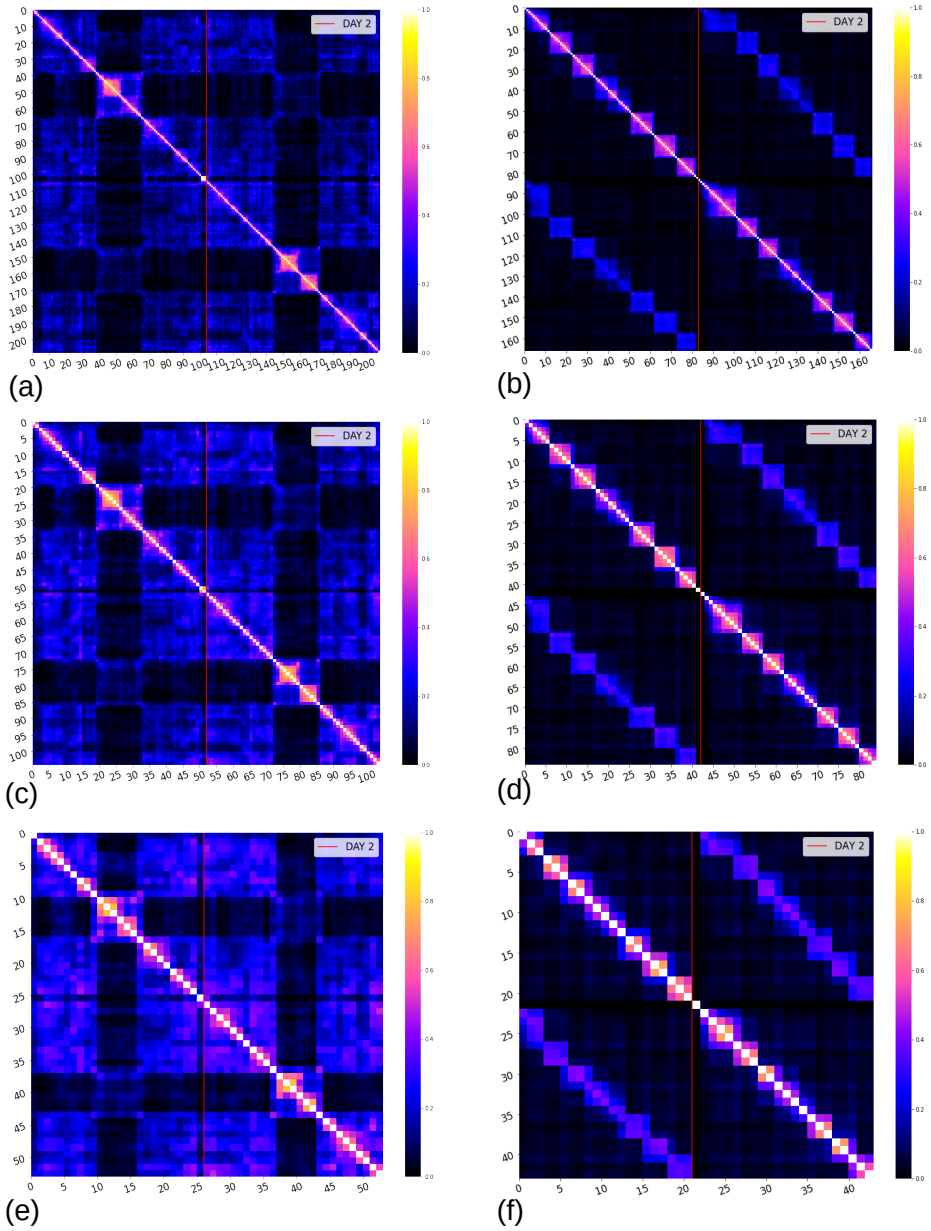

Figure S3. Cosine Similarity matrices for the French elementary school and Utah middle school, aggregating the network over different time window lengths. The first column (a),(c),(e) refers to the French elementary school data ; the second column refers to the Utah middle school data. The different rows refer to three different values of the aggregation time window length: respectively 5, 10 and 20 minutes.

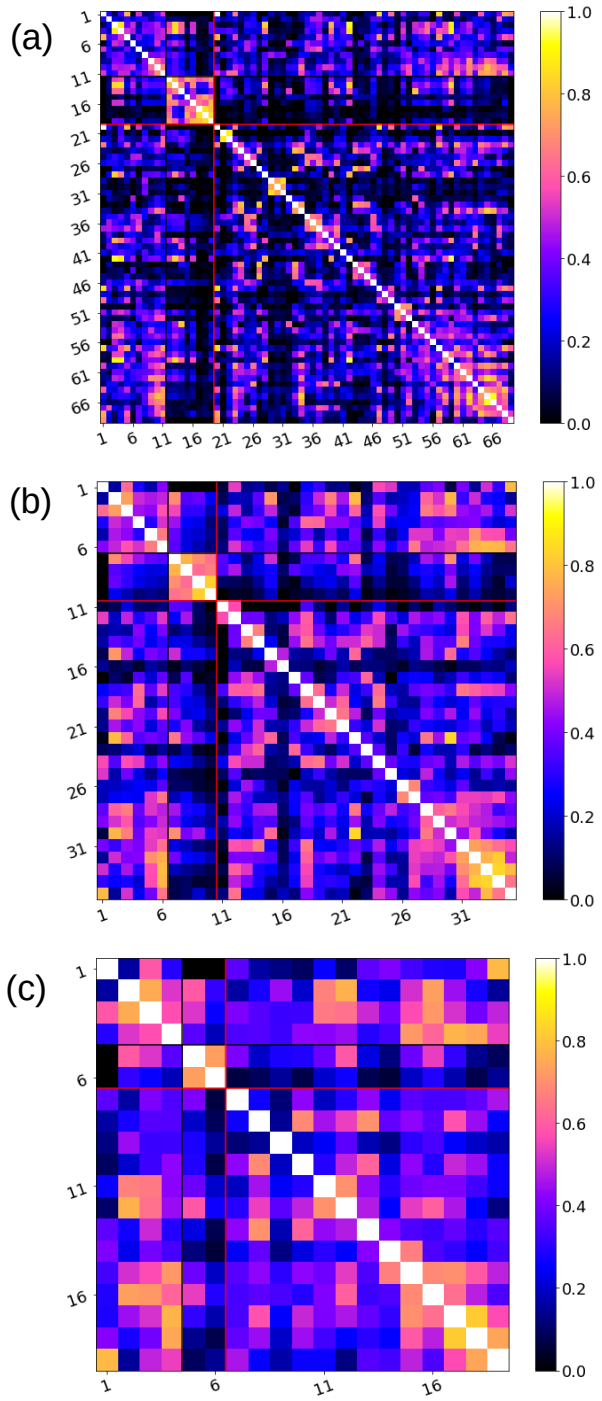

Figure S4. Cosine Similarity matrices for the baboon data, aggregating the network over different time scales. The three panels refer to the aggregated networks computed using the same data as in Fig. 3 for three different values of the aggregation time window, i.e. 15 minutes (a), 30 minutes (b) and 1 hour (c).

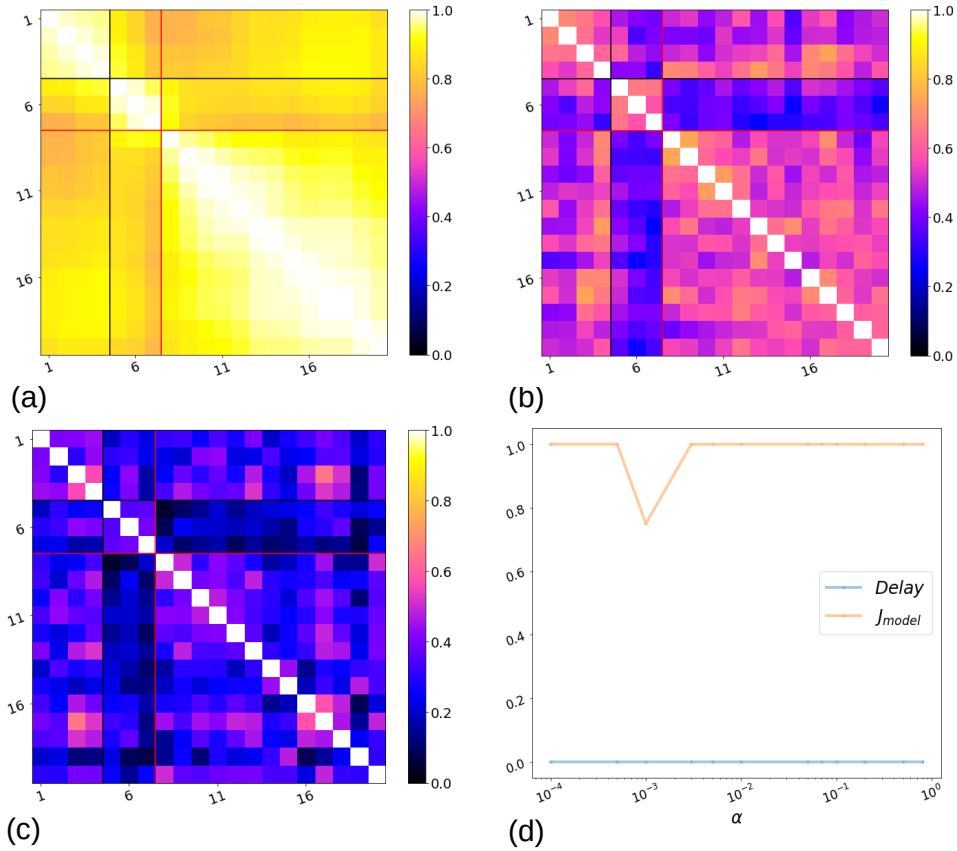

Figure S5. Detection of a simulated perturbation in a temporal network data set. Here we consider 20 days of proximity data collected in a group of 13 baboons (see Material and Methods). The data, whose temporal resolution is 20 seconds, is artificially perturbed by exchanging the identity of two nodes for 3 days. The resulting perturbed temporal network is transformed into a weighted evolving network as described in the text, and this network is observed on a daily basis. Panels (a), (b), (c) represent the resulting cosine similarity matrices for values of $\alpha = \beta = 0.001$, 0.1, 0.5, respectively. The black and red lines correspond to the (known) start and end times of the perturbation. Panel (d) shows the detection performance (see Fig. 2), computed from the hierarchical clustering analysis applied to the distance matrices, with the number of clusters fixed to $C = 3$. The blue line represents the relative delay in the detection of the perturbation, i.e. the difference between the known beginning of the perturbation (black line) and the detection of a new network state, divided by the total length of the perturbation. The orange line indicates the Jaccard index between the known perturbation timestamps and the perturbation detected by the clustering algorithm using different values of the parameter α. The detection performance relative to the aggregated network is not shown because no cluster detected by the algorithm could correspond to the simulated perturbation.

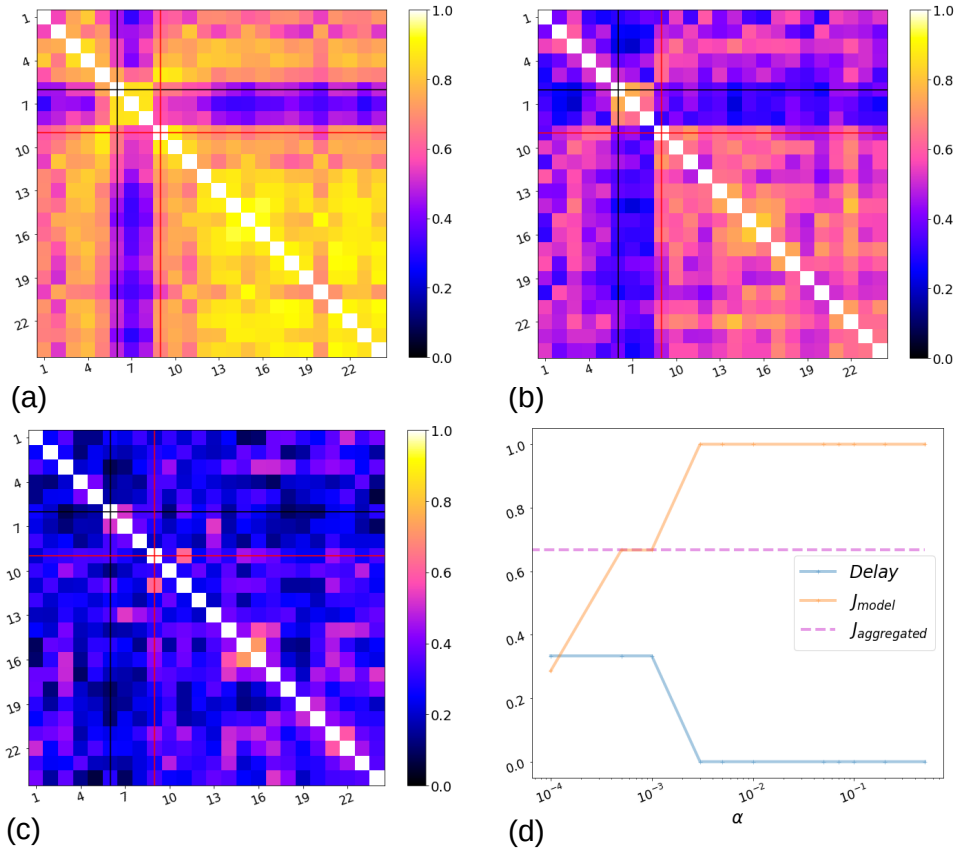

Figure S6. Detection of a simulated perturbation in a temporal network data set. Here we consider the proximity data collected in a group of 13 baboons over several months (see Material and Methods). The data, whose temporal resolution is 20 seconds, is artificially perturbed by exchanging the identity of two nodes during 15 days, affecting weeks 6 to 8. The resulting perturbed temporal network is transformed into a weighted evolving network as described in the text, and this network is here observed every 7 days. Panels (a), (b), (c) represent the resulting cosine similarity matrices for values of $\alpha = \beta = 0.001$, 0.1, 0.5, respectively. The black and red lines correspond to the (known) start and end times of the perturbation. Panel (d) shows the detection performance (see Fig. 2), computed from the hierarchical clustering analysis applied to the distance matrices, with the number of clusters fixed to $C = 3$. The blue line represents the relative delay in the detection of the perturbation, i.e. the difference between the known beginning of the perturbation (black line) and the detection of a new network state, divided by the total length of the perturbation. The orange line indicates the Jaccard index between the known perturbation timestamps and the perturbation detected by the clustering algorithm using different values of the parameter α; the dotted magenta line indicates the Jaccard index using networks aggregated over successive 7 day time windows.

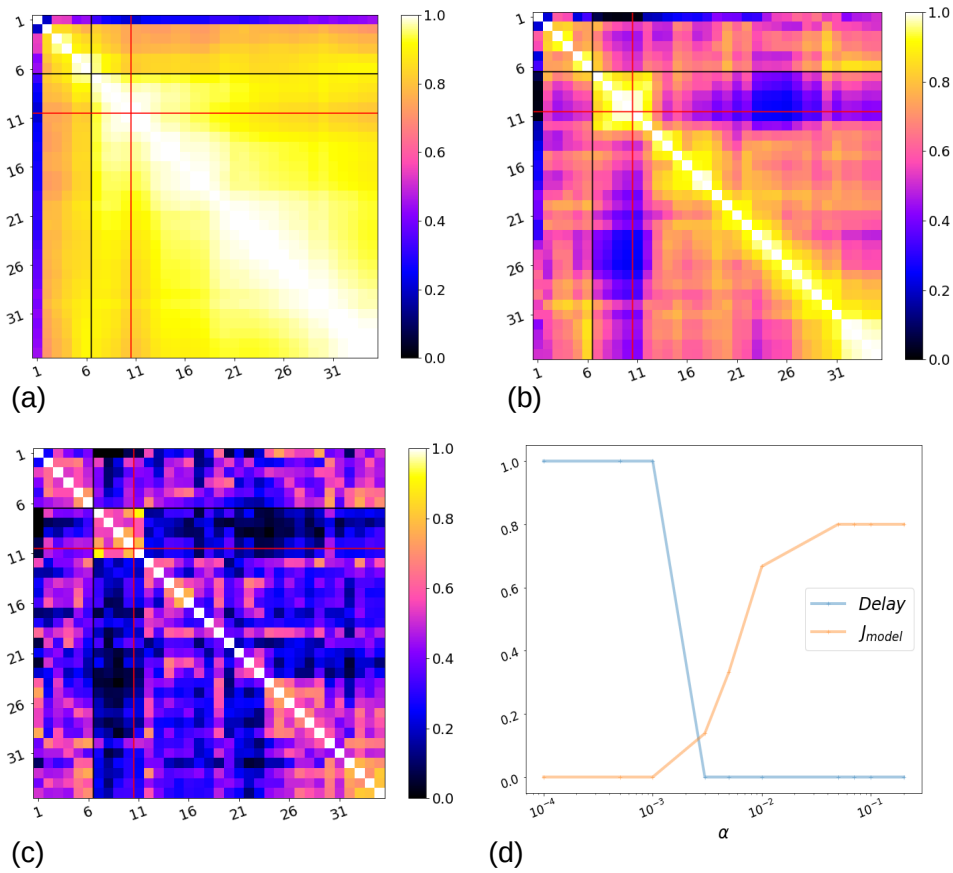

Figure S7. Detection of simulated perturbations in the baboon' data set, using the weighted evolving network with different values of the parameters α and β. Panels (a), (b), (c) represent the cosine similarity matrices for $\beta = 5\alpha$ and values of $\alpha = 0.001, 0.01$ and 0.1, respectively. The networks and the detection performance were computed using the same procedure as in Fig. 3.