



OPEN

weg2vec: Event embedding for temporal networks

Maddalena Torricelli^{1,2}, Márton Karsai^{3,4,1} & Laetitia Gauvin¹✉

Network embedding techniques are powerful to capture structural regularities in networks and to identify similarities between their local fabrics. However, conventional network embedding models are developed for static structures, commonly consider nodes only and they are seriously challenged when the network is varying in time. Temporal networks may provide an advantage in the description of real systems, but they code more complex information, which could be effectively represented only by a handful of methods so far. Here, we propose a new method of event embedding of temporal networks, called *weg2vec*, which builds on temporal and structural similarities of events to learn a low dimensional representation of a temporal network. This projection successfully captures latent structures and similarities between events involving different nodes at different times and provides ways to predict the final outcome of spreading processes unfolding on the temporal structure.

An interacting group of people, the collectively active neurons in the brain, or the transportation system of a city are only a few examples of complex systems, which are all intrinsically dynamical¹. They can be commonly interpreted as a set of entities, which interact over time and form a network structure coding the architecture of the system in hand^{2,3}. This duality of the structure and temporal nature of interactions can be effectively captured by temporal networks, proposing a new and precise representation of complex systems as compared to earlier strategies^{4,5}. On the finest level, temporal networks consist of time-varying events between interacting nodes, and as a whole they appear as systems with high complexity and dimensionality. Events in real temporal structures, however, are not random but correlated with each other and arguably driven by several microscopic mechanisms leading to several generative characters of the network. Emerging properties like the heterogeneous number or strength of interactions^{6,7}, community structure^{8–10}, degree correlations, bursty temporal patterns^{11,12}, or temporal motifs of causally correlated events^{13,14} are all arguably induced by such mechanisms. In turn, temporal events and their correlations largely influence dynamical processes as well¹⁵, like they determine the speed and final outcome of information or epidemic spreading^{16–19}. The recognition of these impacts has put temporal networks in the focus point of several investigations recently, which yet struggle to propose efficient representations to capture the complex temporal/structural information coded in them, while reducing their dimensionality to ease their analysis.

Correlated patterns in the structure and dynamics of networks usually can be described by certain higher-order representations²⁰. For static networks, *line graphs* propose an efficient description^{21,22}, which in their simplest form identify static links as nodes and connect them if they are adjacent, i.e., share an ending node in common. Other technique is based on simplicial complexes²³ considering homology of the network topology to capture higher order structures. At the same time, recent *network embedding* methods propose inventive ways to obtain a reduced dimensional representation of static structures. Their common goal is to map complex networks to a low-dimensional space, while conserving certain similarities of nodes reflected by some distance metrics in the embedding. Most common methods use random walk sampling^{24,25} or graph convolution^{26,27} to capture the local structural context of network nodes. In case of temporal networks, the recently proposed *event graph* representation^{28,29} defines a higher-order description by identifying relations between events, which are adjacent, i.e. not simultaneous and share at least one ending node. Adjacent events then are connected by links with direction respecting the arrow of time, and with weight defined as the absolute time difference between the connected events. This type of description is very useful as it codes all time-respecting paths in a temporal network at once, while proposing an information lossless representation of the temporal structure as a static weighted directed acyclic graph. Recently a few *dynamical network embedding* methods^{30–33} have been developed to consider dynamical changes in the structure in the learned network representations. At the base of many methods there is the

¹ISI Foundation, Turin, Italy. ²Alma Mater Studiorum University of Bologna, Bologna, Italy. ³Department of Network and Data Science, Central European University, H-1051, Budapest, Hungary. ⁴Univ de Lyon, ENS de Lyon, INRIA, CNRS, UMR 5668, IXXI, 69364, Lyon, France. ✉e-mail: gauvinlaetitia@gmail.com




Figure 1. Schematic presentation of the methodological pipeline of the presented temporal network embedding method, which takes a temporal network to (a) project it into a weighted event graph; (b) to sample a set of environments for each event; (c) and uses it as input for a Skip-Gram model; (d) to obtain an event embedding of the original network.

modification of the standard representation of the temporal network, whether it is in the form of a list of events, a tensor⁹ or a supra-adjacency matrix³⁴. All of these methods like DyANE³⁴, Online Node2Vec³², STWalk³⁰, or the one proposed by Singer *et al.*³¹ commonly aim to solve a node embedding problem by locally sampling the temporal-structural neighbourhood of nodes to create contexts, which they feed to a Skip-Gram learning architecture borrowed from the text representation literature³⁵. As a solution, they build a sequence of correlated/updated embeddings of network snapshots, which consider short term history of the network backward in time. However, these methods have some common limitations; first of all, it can be hard to manage a high number of hyper parameters for the control of the sampling random walk process and the embedding itself. At the same time, the embedding of nodes may miss to reflect the dynamical changes of temporal interactions. Finally, taking into account only past and present interactions in the embedding can crucially limit the performance of the prediction, while the consideration of future events can significantly improve this task.

Here we propose a new temporal network embedding method, that we call *weg2vec* (weighted event graph to vector), which aims to tackle all these shortcomings. This is an event embedding method, which represents an entire temporal network in the same reduced dimensional abstract space. It is based on combined event contexts built by sampling locally a higher-order static representation of temporal networks, which in turn code the complex patterns characterising the structure and dynamics of real world networks. This is an unsupervised representation learning technique, which can consider the past and future context of an event simultaneously. It is sampling without using dynamical processes, thus it can be controlled by a handful of hyper parameters. It identifies similarity between different events/nodes, which may be active at different times, but influence a similar set of nodes in the future. To demonstrate the power of this representation, we showcase its utility through the prediction of the final outcome of modelled spreading processes on several real world temporal networks. This prediction task performs significantly better when it builds on our representation as compared to other dynamical network embeddings.

As follows, first we will present the pipeline to build our embedding method. We will show the characteristics of our representation, measuring its stability and its ability to capture temporal and structural information from the network. We will show then the results obtained in estimating an epidemic spreading outcome, as a showcase of the potentials of our embedding method in the analysis of dynamic processes. Finally, we will compare our results to similar computations performed with two other embedding methods. In the final Sections, we present the discussion of the results and the analysis of the methods.

Results

An embedding method of temporal networks may take a list of temporal interactions as input, and provide a lower dimensional representation, in which vectors corresponding to similar nodes or events in the original structure ideally point close to each other in the embedding. In our pipeline we solve this problem in three consecutive methodological steps. First, we turn the original temporal network into a higher-order representation, which captures pairs of adjacent and consecutive events, which may be in causal relationship. Second, we use this representation to generate environments for each event, sampled from their adjacent neighbours. Finally we obtain an embedding of events by training a Skip-Gram model on the obtained environments. These steps are schematically drawn in Fig. 1) and introduced next in the coming sections.

To demonstrate our method we used four different datasets all obtained from the SocioPatterns project³⁶. These open datasets record the time evolving physical proxy interactions of a large number of people in different settings like in *conference*, *high school*, *hospital*, or *primary school*. The data comes as a long sequence of network snapshots recording simultaneous interactions between any participants in every 20 seconds. While for demonstration most of the results in the paper are shown only for the conference and primary school settings, a detailed data description together with the analysis for the other networks are presented in the Supplementary Information.

Temporal networks as weighted event graphs. Let us consider a temporal network

$$G_T = (N, E_T, T), \quad (1)$$

where E_T denotes a set of events (temporal edges) among nodes in n at times $t \in T$. Specifically, we define an event $e = (i, j, t)$ as an interaction between two nodes $(i, j) \in N \times N$ at a given timestamp $t \in T$. The time aggregation of interactions in G_T over t maps the underlying structure into a static graph $G = (N, E)$ defined over the same set of nodes N , which are connected if they interacted at least once. For simplicity here we assume that events are undirected and no self events/links exist, i.e. for any event (i, j, t) or link $(i, j), i \neq j$.

We define two events $e_1 = (i, j, t_1)$ and $e_2 = (k, l, t_2)$ to be *adjacent* ($e_1 \rightarrow e_2$) if they share at least one node $(\{i, j\} \cap \{k, l\} \neq \emptyset)$ and they are time consecutive ($t_1 < t_2$). Note that adjacency in a static network G can be similarly defined between links $l_1 = (i, j)$ and $l_2 = (k, l)$, which share at least one node $(\{i, j\} \cap \{k, l\} \neq \emptyset)$. More restrictively, we called them δt -adjacent ($e_1 \xrightarrow{\delta t} e_2$) if they are adjacent and $|t_2 - t_1| \leq \delta t$, thus follow each other within a given period of time ($0 \leq \delta t \leq \tau$) where τ corresponds to the total period of time of the interactions. Adjacency introduces a directed relation between events, with orientation respecting their order in time. Using this notion we can formally define a static directed network representation $D = (E_T, E_D)$ of any temporal network, where original events in E_T are defined as nodes and they are connected by directed links $e_D \in E_D$ if they are adjacent $e_D = e_1 \rightarrow e_2$. The obtained network is a directed acyclic graph called the *event graph*, defined earlier in^{28,29}. It can be interpreted as a temporal line graph, which provides a higher-order representation to map out simultaneously all time respecting paths of the original temporal network without any loss of information. Note, that to simplify our representation, for a given event if it has multiple future adjacent events with the same pair of nodes, we only consider the earliest one.

Event graphs can be easily enriched with various types of link weights reflecting some temporal or structural information coded in the original structure. Here, to better capture the strength of potential causal relationships, first we consider a weight defined as $w_{path} = \frac{1}{(1 + |t_2 - t_1|)}$, which is a measure inversely proportional to the absolute time difference between adjacent events at t_1 and t_2 . This definition of the weight allows us to include the temporality of interactions such as long decay in social activities. At the same time we define a second weight for adjacent events (links of the event graph), based on the total number of co-occurring events on the underlying adjacent links in the static network. More precisely, the $w_{co}(e_1, e_2)$ co-occurrence weight counts the number of δt -adjacent events in G_T appearing on a given pair of adjacent links $l_1(i, j)$ and $l_2(k, l)$ in the static graph G . Note that adjacent events connected in d , which corresponds to the same links in the underlying network G , will have the same w_{co} values. Datasets analysed in this paper are defined as sequences of snapshots aggregating temporal interactions in consecutive time windows of size Δt . In these systems we compute w_{co} for adjacent links as the number of co-occurrence of corresponding events within a single snapshot. This definition may slightly underestimate the real co-occurrence (if $\delta t \leq \Delta t$), but provides the best plausible solution due to the un-ambiguity of timings of events within a single snapshot.

Neighbourhood sampling strategy. In the same spirit of recent node embedding techniques^{25,35} based on the Skip-Gram model, we propose an event embedding method, which samples neighbourhoods for events from the weighted event graph representation to map them to a lower dimensional space. To assign an environment to an event e_k , we sample its local neighbourhood set N_k , which consists of the set of its first in- (past) and out- (future) neighbours (also called its *predecessors* and *successors* from now on). The sampling is done according to probabilities determined by the two types of weights of the links that connect the actual event to its neighbours.

In order to consider not only the past but the future of an event in its environment, during sampling we use a combined set of its predecessor and successor events. Further, to balance the contribution of causality and temporal co-occurrence, we introduce a neighbourhood sampling strategy such that the probability $p(e_l)$ of picking an event e_l from the combined neighbourhood set N_k of the central event e_k is given by:

$$p(e_l) = \alpha F(w_{path}(e_k, e_l)) + (1 - \alpha) F(w_{co}(e_k, e_l)) \quad (2)$$

where α is a coefficient between 0 and 1 scaling the contribution of the two types of weights and F is a normalised weighted function defined as:

$$F(w) = \frac{w(e_k, e_l)}{\sum_{n \in N_k} w(e_k, e_n)}. \quad (3)$$

Using such probabilities computed for each neighbour, we sample nb number of environments randomly for each event for an effective training of the model explained next. Each environment contains s events, we call the number s , the length of the environment.

Embedding of empirical temporal networks. Once the environments have been created, they are passed as inputs to the Skip-Gram model, with parameters fixed to different values according to the analysis to conduct. The result is a d -dimensional vector space in which events are represented by Cartesian coordinates. As an illustration, Fig. 2 shows a 3-dimensional embedded representation of two of the empirical networks we used for the analysis, recorded in a conference and a primary school settings. Our aim with this setting was to investigate the performance of low-dimensional embedding on the one hand, and on the other taking into account equally the effects of causal temporal paths and co-occurrences by setting $\alpha = 0.5$. The environment parameters

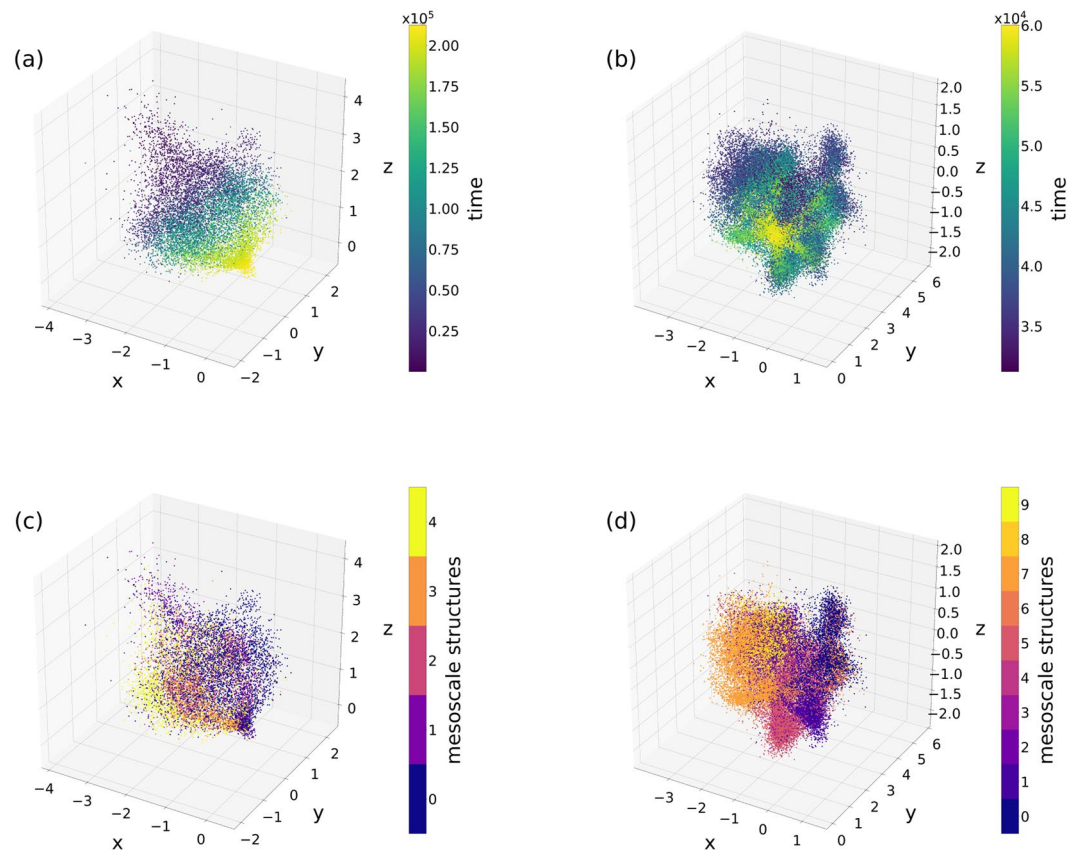


Figure 2. 3-dimensional embeddings of the conference and primary school networks. x , y and z axes indicate event coordinates, while colour in panels (a and b) shows the time at which the event occurs and in panels (c and d) mesoscale structure membership found using a tensor factorisation method (see Section Tensor Factorisation for mesoscale structure extraction) respectively set to find 5 and 10 of these structures. Hyper parameters were set to $\alpha = 0.5$, $nb = 10$ and $s = 10$.

nb and s have been set both to 10 as an example. In Fig. 2(a,b) each event is represented as a point in the embedded space with colour indicating the time at which they occurred in the original temporal network. Interestingly, the gradient change of colours indicates that these embeddings capture in large part the time ordering of the events. At the same time, Fig. 2(c,d) shows the same 3-dimensional embedded representations, but with colours representing the membership to mesoscale structures detected by tensor factorisation methods applied on the original temporal network⁹ (see Section on Tensor Factorisation for mesoscale structure extraction). Evidently, as colours are not distributed randomly but similar colours are somewhat grouped together in space, it suggests that our embedding is able to capture some of these mesoscale structures as well.

We propose an additional microscopic scale analyses of our embedding method in the Supplementary Information. There we studied pairs of events and we were interested in the relation between their time difference measured in the temporal network and their euclidean distance observed in the embedding. Indeed, we found clear correlation between these quantities and demonstrated that the euclidean distances among linked events in the temporal network are significantly smaller than the distances measured between randomly selected events pairs. These observations demonstrate that our method simultaneously captures structural and temporal vicinity of events.

Effects of the dimension and of the neighbourhood sampling. One of the most important hyper parameter of our method is the number of dimensions of the embedding vector space. Lower than optimal number of dimensions may lead to neglected but otherwise relevant latent correlations in the temporal structure, while overestimation of this number may give us a highly redundant embedded space. We test here the consistency and robustness of our embedding technique in terms of this parameter. We argue that as we increase the number of dimensions, once it reaches and overpasses an optimal number, it starts coding increasingly redundant information in the embedding. As a consequence, further dimensions would not alter the relative positions of embedded events and the pairwise euclidean distances among them would stabilise. To check this assumption, we use an entropy measure capturing the fluctuations of pairwise euclidean distances over several realisations with the same number of embedding dimensions. More precisely, for selected event pairs, we measured the probability distribution of their pairwise euclidean distances over 10 realisations (see Section Entropy Computation) and used it to compute an entropy score. Averaging these scores over all the event pairs provided us an indicator of the stability of the embedding.




Figure 3. Entropy values with respect to d number of embedding dimensions for the conference (a) and the primary school (b) networks at $\alpha = 0.5$. The dash line represents the value ($d = 20$ and 24 respectively for (a and b)) of the optimal embedding dimension in which stability is reached. The blue line and the shaded area represent respectively the average and the variance among the samples we used for the analysis.

r^2					
Data		Original	Snapshot	Timeline	Link
Conference	($d = 20$)	0.79 ± 0.01	0.53 ± 0.04	0.66 ± 0.03	0.57 ± 0.01
Hospital	($d = 14$)	0.53 ± 0.03	0.11 ± 0.02	0.35 ± 0.06	0.50 ± 0.04
High School	($d = 26$)	0.56 ± 0.02	0.23 ± 0.01	0.53 ± 0.02	0.76 ± 0.04
Primary School	($d = 24$)	0.68 ± 0.02	0.12 ± 0.01	0.31 ± 0.02	0.55 ± 0.02

Table 1. R-squared values, r^2 obtained between the simulated and predicted epidemics outcomes using embedding of the real empirical temporal networks and of the randomised model. We set the environment parameters s and nb both to 10. The optimal embedding dimension were chosen as found in Effects of the dimension and of the neighbourhood sampling.

We show results in Fig. 3 for two empirical networks using a balanced embedding with $\alpha = 0.5$ in both cases. As we expected, the entropy decreases as we increase the number of dimensions. This is due to stabilisation of the distribution of pairwise euclidean distances, which in turn gives us a hint on the optimal dimension at which the local neighbourhoods (as defined by the sampling) are well captured by the embedding. To select this optimal dimension, we identified the lower bound at which entropy starts to fluctuate around a constant value (see Methods 4 and Supplementary Information). Note that to identify the optimal number of dimensions we have taken into account other algorithms³⁷ as explained in the Supplementary Information. However, our final choice fell on the entropy based method we introduced above as it maintained a good trade-off between the compactness (i.e. low dimensionality) and stability of the embedding. Other tested methods suggested an unrealistically high number of dimensions to be optimal, probably due to their incompatibility with the actual setting, as they were developed for word embedding problems.

Spreading process prediction with embedding events. Beyond the demonstrated capacities of our model to capture the temporal ordering and the underlying mesoscopic structures, it may provide further useful information about the embedded events. As a temporal network embedding, it positions events in proximity with similar neighbourhoods. In other words, it can help to identify similar events maybe involving different nodes at different times, but influencing a similar set of other nodes via overlapping temporal paths. As a consequence, this information can be used to predict the outcome of information diffusion processes on temporal networks.

To explore this problem, we model a Susceptible-Infected (SI) process, which is the simplest schematic model of epidemic or information spreading (see Section Spreading Process in Methods). Defined on networks, this model assumes that each node can be in one of two mutually exclusive states (susceptible (S) or infected (I)) at a given time. While initially each node is susceptible, infection can spread from a selected infected seed node/event via temporal interactions and can reach all other nodes via connected valid temporal paths. To obtain the expected outcome of SI process on a temporal network we took each event as the seed and simulated the spreading on the empirical temporal network to measure the final epidemic size in each case. Note that our aim to investigate the final outcome of an epidemic differs from the one pursued with DyANE³⁴. In our case, we are not interested in the status of the node time by time, but in the final outcome of the epidemic originated by a specific event. To test the versatility of our embedding method we trained a model using the embedded coordinates of events for epidemic size predictions and compared results directly to the corresponding simulated outcomes. We used linear regression to approximate the correspondence between the embedding coordinates of each event and the size of epidemic initiated from them. As the goodness of the prediction we simply computed the r^2 scores between the predicted and simulated epidemic sizes. Note, that we tested more complicated non-linear models but obtained lower performance in prediction (not shown here). We report our results in Table 1, where we fixed the environment parameters s and nb both to 10 and chose the optimal embedding dimension for each real network detected as we explained in Section Effects of the dimension and of the neighbourhood sampling.

Correlation		Local temporal	Weight-structural	Higher-order temporal	Structural
RRM					
Original		✓	✓	✓	✓
Snapshot	$P[f(t), p_T(\Gamma)]$	×	✓	×	✓
Timeline	$P[L, f(p_L(\Theta))]$	✓	×	×	✓
Link	$P[f(L), p_L(\Theta)]$	✓	×	×	×

Table 2. Summary of preserved and eliminated structural and temporal correlations (Local temporal, Weight-structural, Higher-order temporal, and Structural) in the Original and different random reference models (Snapshot, Timeline and Link shuffling) of temporal networks. For further explanation see text.

Real temporal networks are interwoven by several temporal and structural correlations. First there are *local temporal correlations* induced by intrinsic or environment effects emerging between events on the same link leading to bursty behaviour, event trains or circadian activity patterns. Another type is *higher-order temporal correlations* leading to the emergence of causal temporal motifs. *Structural correlations* are responsible for the emerging assortative patterns, communities, or any non-random connection pattern in a social network, while *weight-structural correlations* induce the non-random distribution of strong and weak ties inside and between communities. In Table 2 we summarise which RRM preserves and eliminates which type of correlations.

To get a better sense about the effects of these correlations, we used three types of randomised reference models (RRM)³⁸ to eliminate combinations of temporal and structural correlations, and to identify which of them are determinant for the prediction task. These RRMs were

- *Snapshot shuffling*, which randomises the timestamps of events in order to eliminate any temporal correlation between them inducing burstiness, causal motifs, or group activation etc. This model is assigned as $P[f(t), p_T(\Gamma)]$ using the notation convention introduced in³⁸.
- *Timeline shuffling* takes the complete timeline of events between a connected pair of nodes in the temporal network and swap it with the timeline of another randomly selected connected pair of nodes. This shuffling method, noted as $P[L, f(p_L(\Theta))]$ in³⁸, eliminates all correlations between the underlying structure and timelines while also vanish any casual correlations between events on adjacent links.
- *Link shuffling* method (noted as $P[f(L), p_L(\Theta)]$ in³⁸) randomises links of the underlying aggregated static network first to obtain a Bernoulli random structure, and then reassign randomly the original timelines of events to the new links without replacement. In this way, it destroys any structural and structural-temporal correlations in the network, while keeping local temporal correlations like burstiness unaltered.

For a summary of present and eliminated correlations in the different RRMs see Table 2.

The prediction results for the original and the RRM networks are summarised in Table 1, where we depict the observed average r^2 values with their standard deviation computed over the embedding realisations. These results suggest that in general the randomised model embeddings perform worse in predicting the final epidemic size with respect to the original network embeddings. In one way this is straightforward as some correlations have been eliminated from RRMs, which might be determinant for the prediction task. On the other hand, RRMs also appear with a less complex structure and limited irrelevant dependencies and noise, which in turn may help the prediction. It is the snapshot shuffling method, which leads consistently to a significant drop in performance, suggesting that temporal correlations (local or higher-order) are very important for the spreading process and that our embedding can capture these dependencies successfully. Timeline shuffling, which destroys weight-structural and higher-order temporal correlations but conserve the dynamics on links and the underlying network seems to perform better as compared to the snapshot shuffling method but yet worse than the original network. This suggests that while structural correlations can be captured by the embedding, local temporal correlations might be better predictors than weight-structural correlations. Interestingly, the link shuffling method, which conserves only local timeline dynamics on links, performs the best among the RRMs, sometimes even better than the original dataset. Consequently, indeed local temporal event dynamics is the most important feature of the temporal network, while removing structural correlations the system becomes homogeneous and easier to predict (for a supporting analysis see Supplementary Information). Although these general conclusions seem to be consistent over the several analysed datasets, results computed in different settings may have some fluctuations. As explained in the Supplementary Information, this can be partially explained by the variance of the epidemic size distribution reflecting the fluctuation of epidemic size started from different times and events. In most of the cases smaller variance of epidemic size correlates with higher predictive performance except for the link shuffling method, as explained in the Supplementary Information.

As a general conclusion we showed that the embedding successfully captures a combination of temporal and structural features of the network. On the other hand, the fact that temporal and structural features can be entangled has an impact on embedding performance but not on what the embedding is able to learn about them. For instance, the model can under-perform in a community-rich network where nodes of the same community have totally uncorrelated activities, while can provide precise predictions in shuffled datasets if structural and temporal correlations code redundant information.

In Methods in Section Parameter dependencies we also analyse the impact of changing the environment parameters nb and s , the dimension of the embedding d and the hyper parameter α on the predictability of the epidemic spreading outcome (i.e. with respect to the r^2 score). The tuning of the parameters for the best prediction gives a hint about the sensitivity of the prediction task on local properties of the networks and the sampling parameters.

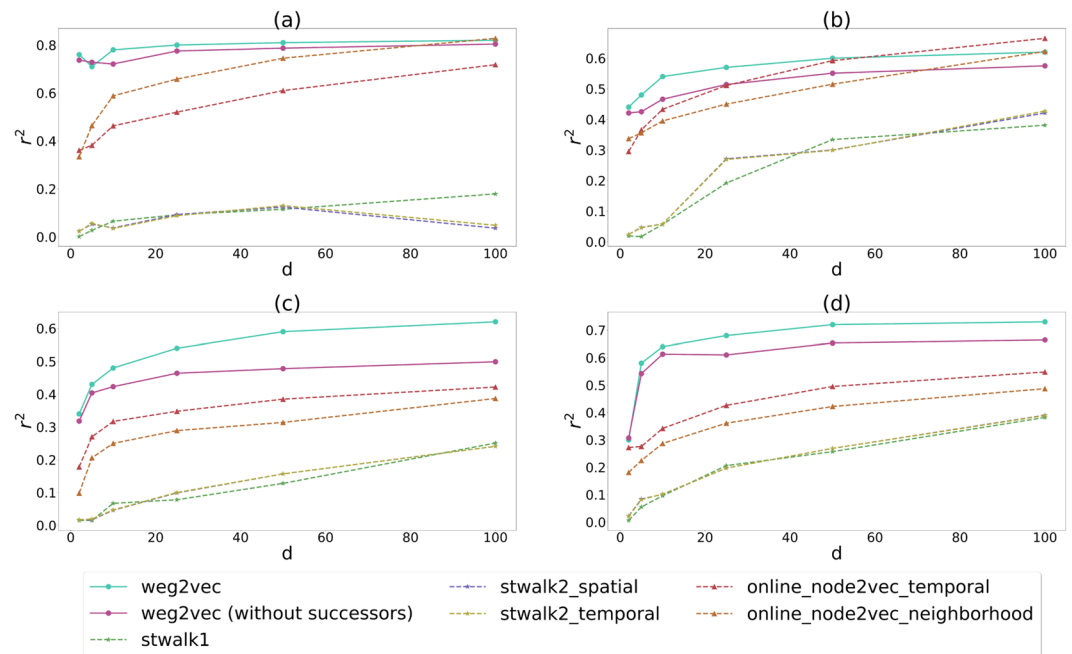


Figure 4. Comparison of STWalk, Online-Node2vec and our embedding methods in predicting spreading outcomes on empirical networks in different settings as (a) conference, (b) hospital, (c) high school, and (d) primary school. Results shown are r^2 scored obtained from linear regression on coordinates in embedding spaces with various dimensions computed for each method and empirical temporal networks.

Comparison with other methods. There are a few other recently proposed dynamical network embedding methods, which can be used for the prediction of spreading outcome. Here we consider two of the most promising ones, the STWalk^{30,39}, and the Online-Node2vec embedding methods^{40,41} to compare their predictive performances to weg2vec. Both methods are thought to build node embeddings for dynamic graphs using the Skip-Gram model, which introduces a significant difference to our event embedding method.

STWalk is designed to learn trajectory representations of nodes in temporal graphs by operating with two graph representations, a graph at a given time step and a graph from past time steps. It performs random walks respectively called space-walk and time-walk, to sample environments to input for the Skip-Gram embedding. The authors propose two variants of STWalk, different in the way the environment is built. In STWalk1, space-walk and time-walk are performed as part of a single step on a combined graph, while in STWalk2, space- and time-walks are done separately.

Online-Node2vec is a node embedding method updating coordinates each time a new event appears in a temporal network. It also applies random walks to generate environments possibly using two strategies, the Temporal Walk algorithm and the Temporal Neighbourhood algorithm. In the Temporal Walk algorithm⁴² a temporal path based centrality metric is used to capture similarity between nodes on the same temporal path close to each other in the embedding. In the Temporal Neighbourhood algorithm⁴³, node similarity is inferred via a fingerprinting method, which projects nodes with similar neighbourhoods close to each other.

To compare the performance of the different methods, we test all of them on the four empirical networks introduced earlier. The environment parameter nb and s have been set to 10 and 10 for all cases to give them the same amount of information to learn and for a fair comparison of outcome. Further, we fix the balance parameter α to 0.5. We then compute the average r^2 scores of simulated spreading outcomes as we vary the number of embedding dimensions. Since STWalk and Online-Node2vec use only the past and the present as basis for the nodes environment, we run the simulation for our methods using only the predecessors for each event as well (see Section Neighbourhood sampling strategy). Finally, as previously, we estimate the epidemic size by using the coordinates of the actual embedding in a linear regression model (see Section Spreading Process).

According to the results in Fig. 4, our method outperforms all the other methods on any of the networks for a broad range of dimensions. The performance improves if we also consider the successors and not only the predecessors in building the environment, as expected. The exception is the hospital network, where our method gets lower scores with respect to Online-Node2vec for dimensions 50 or larger. In general, the difference in the scores can be explained due to the advantage of event embedding instead of node embedding. Indeed if we are looking at epidemic spreading mediated by temporal interactions, it becomes more natural to work with events. In the case of STWalk, the lower scores can be partly explained by the selection of the environments that are allowed to include higher-order correlations among nodes. This more complex information coded in the environments can appear less relevant or noisy for the learning task here. In case of Online-Node2vec, the relative under-performance can be due to the fact that information of the temporal and neighbourhood information are considered separately instead. Missing to join these two aspects can lead to limited information and prediction capacities.

Discussion

Embedding of networks has recently drawn a lot of attention as it both provides lower dimensional representations of networks and proves to be efficient to resolve task such as link prediction, node classification or anomaly detection. Here we proposed an embedding technique of temporal networks, a domain still little explored as most low dimensional representations of networks have been introduced for static networks. Moreover, instead of generating node embedding, we are creating link embedding, which we show to be much more efficient when we try to solve a task linked to spreading process compared to node embedding techniques.

The embedding method we introduced has the advantage to be very simple, it relies on the sampling of neighbourhood on a higher order static representation of the temporal networks and the use of the Skip-Gram model, largely developed for word embedding. The neighbourhood sampling was built so that it takes into account both the notions of causal temporal paths and co-occurrences, meaning that events that are on the same temporal paths and that tend to co-occur are projected close in space. We have shown that an embedding based on this neighbourhood sampling is particularly efficient to provide compact representations of temporal networks that retain essential features of the networks such as the time ordering and the organisation of networks in meso-scale structures. Here, we also provide a way to choose a relevant dimension for the embedding. Along with this, we show that the learned representations retain enough information of the original network to get a relevant estimate of the outcome of spreading process. Interestingly, this observation remains true even if the sampling strategy uses only information from the past for each event. This means that the technique introduced can be also used as an online method taking into account temporal events on the run. Moreover, tuning the neighbourhood sampling, i.e. playing on the trade-off between including causal temporal paths and co-occurrences, to get the best performance for the prediction of the outcome of the spreading process can be used to detect the relevant properties of the original network for the spreading process.

For future works, it would worth exploring other sampling strategies that decouples the purely structural properties, i.e. the presence of the communities in the aggregated network, from the temporal properties. Another important follow-up of this work would be the application of this embedding technique to solve questions such as the detection of key events in misinformation spreading.

Methods

Entropy Computation. To measure the entropy over the distributions of the euclidean distances between pairs of event coordinates in each of the temporal network embedding, we build 10 embeddings for 50 different dimensions (from 2 to 100 at step 2), setting $\alpha = 0.5$. The hyper parameters nb and s were set to 10 and 10 respectively. For each embedding, we then divide our dataset in 10 consecutive samples of 1000 consecutive events each, both to avoid to compute all the pairwise distances (which would be very costly computationally) and at the same time to have a representative set of events.

We compute the euclidean distance of each pair of events in the samples for each of the 10 different embeddings. We then bin the distances into $k = 10$ bins ranging from the global maximum and the global minimum values over all the possible dimension and realisations of the embedding for the same network, and measure the entropy over these sets of distances as

$$H = -\sum_k p_k \log p_k, \quad (4)$$

where p_k represents the probability associated with the k_{th} bin. In Fig. 3 (see also Supplementary Information) the blue curve represents the entropy values with respect to the number of embedding dimensions, averaged on the 10 samples as described above and the shaded surrounding area shows the variance among the 10 samples. The vertical dash line corresponds to the dimension at which the embedding stabilises. To determine this point we looked for the best fit of a horizontal line on the average entropy curve and took the value of the first interception of the curve with its fit.

Spreading Process. The simplest epidemic models are based on the assumption that the population can be divided into compartments, each representing a phase of the disease^{44–48}. The one we used for our analysis, the Susceptible-Infected model, which foresees that once a healthy node (in state S) is exposed to the infection, it will become infected (state I) with a given rate β and will never return to the original healthy state. In our specific case, the SI model has been implemented in such a way that $\beta = 1$, which defines a deterministic process from the spreading point of view.

Taking each event as the starting of the epidemic, we simulate the epidemic spreading on the temporal network and we assign the final epidemic size to each seed event. We build a dataset containing embedding coordinates of each event, the associated epidemic size that will be the target to be predicted, the square of each coordinate and the euclidean distance from each event in the network and the first event in time, that will be used as regressors. In other terms, we assume a linear relationship between the epidemic size (y) and the embedding coordinates (\vec{x}) defined by the so-called *regression equation*

$$y = \beta_0 + \sum_{i=1}^r \beta_i x_i \quad (5)$$

where x_1, \dots, x_r are the predictors (the embedding coordinates, their square and the euclidean distance) and β_0, \dots, β_r are the regression coefficients. For training we operate with a 10-fold cross-validation, i.e. we first randomly partitioned the original sample into 10 equal sized sub-samples and retain a single sub-sample as the validation data for testing the model, while using the remaining 9 sub-samples for the actual training. We repeat this

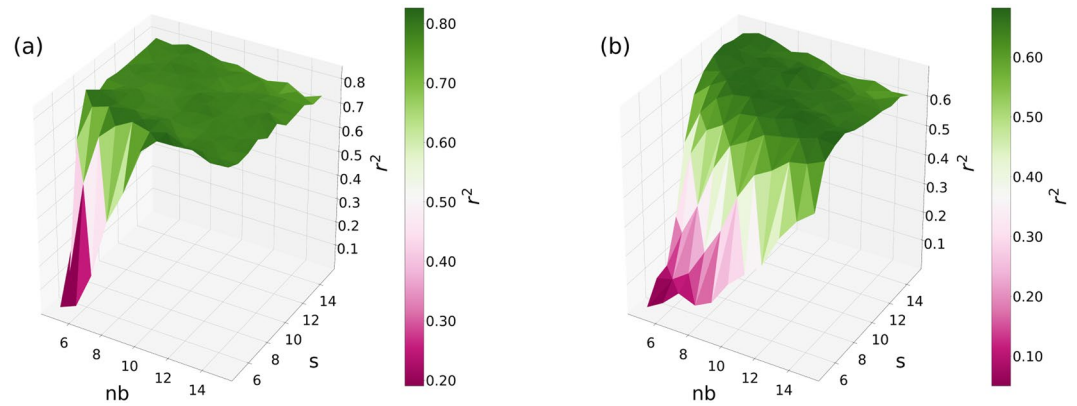


Figure 5. R-squared values, r^2 , dependency on the nb number (x-axis) and s size (y-axis) of sampled environments. Figure (a) shows results for the conference network and Figure (b) for the primary school network. Colours and z-axis code the obtained average r^2 score values for given nb and s parameter pairs computed over 10 realisations. α was fixed to 0.5; we set $d = 20$ for Figure (a) and $d = 24$ for Figure (b) - see Fig. 3.

process 10 times to train the embedding to best learn the coordinates of each event in the network. We use the coefficient of determination, denoted as r^2 , to understand which amount of variation in y can be explained by the dependence on \vec{x} using the particular regression model. Larger r^2 indicates a better fit and means that the model can better explain the variation of the output with different inputs.

Tensor Factorisation for mesoscale structure extraction. A temporal network can be fully described by a time-dependent adjacency matrix, where its entries are either one or zero depending on the presence or absence of interactions of pair of nodes at a given time. This matrix can be seen as a three-way tensor, whose size is $n \times n \times t$ (n indicates the number of nodes and t the number of snapshots in the temporal network). This tensor can then be factorised into a sum of r rank-one tensors, i.e. number of mesoscale structures we search into the temporal network. Using this technique we can group events into mesoscale structures. Indeed with the decomposition, we now have a value for each link for each time for each rank-1 tensor, we assign a mesoscale structure to each event based on these values. Basically the mesoscale structure assigned to each event is the one for which the corresponding link at the corresponding time has the highest value. To find the optimal number of mesoscale structures, we used the core consistency metric⁴⁹. It is based on scrutinising the *appropriateness* of the structural model based on the data and the estimated parameters of gradually augmented models. A model is called *appropriate* if adding other combinations of the same components does not improve the fit considerably. In practice, we operate different tensor decompositions for different value of the rank (ranging from 2 to 20, for all the networks) in order to estimate the best value for it.

Parameter dependencies. Here, we explain the way we investigate how hyper parameters of the environment sampling may impact the prediction score on different real networks. Figure 5 shows the r^2 scores computed for the conference and primary school networks with respect to the length s and number nb of environments sampled for each event. For these computations we fixed $\alpha = 0.5$ and the embedding dimensions to their optimal values. According to Fig. 5 on the conference network, except for very small number or length of contexts we see an emerging plateau of r^2 values. This means that the environment size compensates for the number of environments (or vice versa) when we measure the embedding performance. In other terms, increasing the length of the environment has the same effect than increasing the number of environments on the r^2 score. For the primary school network we observe a similar but somewhat weaker compensation effect (Fig. 5), while we observed the same behaviour even for the hospital and the high school networks (see Supplementary Information). These results suggest that these two parameters are highly redundant, thus we can effectively reduce our parameter set by fixing both to a large enough value.

In order to investigate the influence of the embedding dimension and the α sampling balance parameter, next we fix the environment parameters to $s = 10$ and $nb = 10$ based on our evaluation above. As shown in Fig. 6 (and in Supplementary Information for other networks) both increasing the number of embedding dimensions and α lead to better performances in predicting the spreading outcome. As the function of the number of dimensions, each case reaches a plateau in accordance with our earlier results presented in Section Effects of the dimension and of the neighbourhood sampling. On the other hand we observe somewhat stronger dependencies on α . While for the conference and the hospital networks the more one increases α the better the prediction gets, for the primary school and the high school networks the score reaches a plateau and become less sensitive to the change of α (see Fig. 6 and Supplementary Information). If we consider lower values of α , the similarity we capture between the event between adjacent events is mainly based on the co-occurrences, which are more relevant in school networks where participants might be active at the same time (e.g. in breaks between classes). This argument only

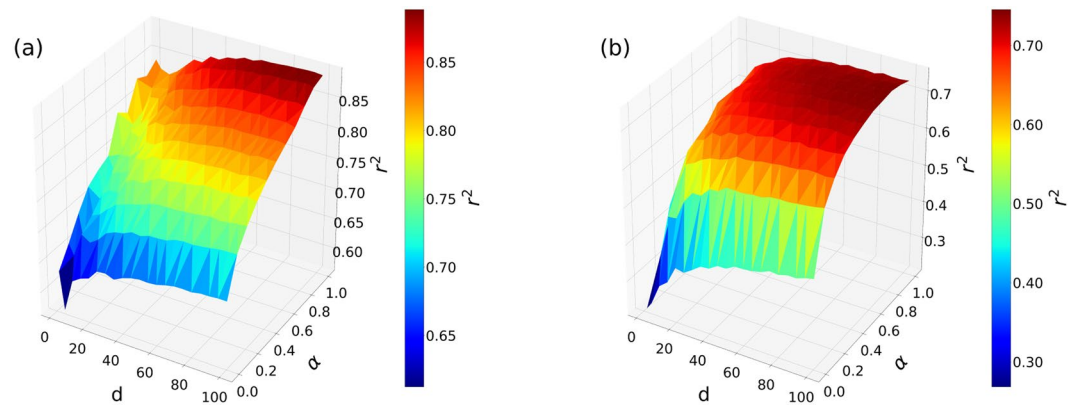


Figure 6. R-squared values, r^2 , as the function of the d number of embedding dimensions (x-axis) and α sampling balance (y-axis) parameters. Figure (a) shows results for the conference network and Figure (b) for the primary school network. Colours and z-axis code the obtained average r^2 score values for a given d and α parameter pairs computed over 10 realisations. Other hyper parameters were fixed to $nb = 10$ and $s = 10$.

moderately applies to a conference or hospital where simultaneous interactions typically happen in smaller groups or not at all. Higher values of α imbalance the sampling to contain more information about temporal paths, which actually indirectly codes co-occurrence frequencies as well. This gives the advantage to the model to learn both type of similarities and to predict the epidemic outcomes with higher precision.

Received: 18 November 2019; Accepted: 11 March 2020;

Published online: 28 April 2020

References

1. Bar-Yam, Y. *Dynamics of complex systems*, vol. 213 (Addison-Wesley Reading, MA, 1997).
2. Newman, M. E. The structure and function of complex networks. *SIAM review* **45**, 167–256 (2003).
3. Albert, R. & Barabási, A.-L. Statistical mechanics of complex networks. *Reviews of modern physics* **74**, 47 (2002).
4. Holme, P. & Saramäki, J. Temporal networks. *Physics reports* **519**, 97–125 (2012).
5. Masuda, N. & Lambiotte, R. *A Guide to Temporal Networks* (World Scientific, 2016).
6. Perra, N., Gonçalves, B., Pastor-Satorras, R. & Vespignani, A. Activity driven modeling of time varying networks. *Scientific reports* **2**, 469 (2012).
7. Karsai, M., Perra, N. & Vespignani, A. Time varying networks and the weakness of strong ties. *Scientific reports* **4**, 4001 (2014).
8. Laurent, G., Saramäki, J. & Karsai, M. From calls to communities: a model for time-varying social networks. *The European Physical Journal B* **88**, 301 (2015).
9. Gauvin, L., Panisson, A. & Cattuto, C. Detecting the community structure and activity patterns of temporal networks: a non-negative tensor factorization approach. *PLoS one* **9**, e86028 (2014).
10. Peixoto, T. P. & Rosvall, M. Modelling sequences and temporal networks with dynamic community structures. *Nature communications* **8**, 582 (2017).
11. Barabási, A.-L. The origin of bursts and heavy tails in human dynamics. *Nature* **435**, 207–211 (2005).
12. Karsai, M., Jo, H.-H. & Kaski, K. *Bursty human dynamics* (Springer, 2018).
13. Kovanen, L., Karsai, M., Kaski, K., Kertész, J. & Saramäki, J. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment* P11005 (2011).
14. Paranjape, A., Benson, A. R. & Leskovec, J. Motifs in temporal networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 601–610 (ACM, 2017).
15. Masuda, N. & Holme, P. *Temporal Network Epidemiology* (Springer, 2017).
16. Karsai, M. et al. Small but slow world: How network topology and burstiness slow down spreading. *Physical Review E* **83**, 025102 (2011).
17. Gauvin, L., Panisson, A., Cattuto, C. & Barrat, A. Activity clocks: spreading dynamics on temporal networks of human contact. *Scientific reports* **3**, 3099 (2013).
18. Delvenne, J.-C., Lambiotte, R. & Rocha, L. E. Diffusion on networked systems is a question of time or structure. *Nature communications* **6**, 7366 (2015).
19. Scholtes, I. et al. Causality-driven slow-down and speed-up of diffusion in non-markovian temporal networks. *Nature communications* **5**, 5024 (2014).
20. Benson, A. R., Gleich, D. F. & Leskovec, J. Higher-order organization of complex networks. *Science* **353**, 163–166 (2016).
21. Evans, T. & Lambiotte, R. Line graphs, link partitions, and overlapping communities. *Physical Review E* **80**, 016105 (2009).
22. Ahn, Y.-Y., Bagrow, J. P. & Lehmann, S. Link communities reveal multiscale complexity in networks. *nature* **466**, 761 (2010).
23. Petri, G., Scolamiero, M., Donato, I. & Vaccarino, F. Topological strata of weighted complex networks. *PLoS one* **8**, e66506 (2013).
24. Perozzi, B., Al-Rfou, R. & Skiena, S. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 701–710 (ACM, 2014).
25. Grover, A. & Leskovec, J. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 855–864 (ACM, 2016).
26. Hamilton, W., Ying, Z. & Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, 1024–1034 (2017).
27. Kipf, T. N. & Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
28. Kivelä, M., Cambe, J., Saramäki, J. & Karsai, M. Mapping temporal-network percolation to weighted, static event graphs. *Scientific reports* **8**, 12357 (2018).
29. Mellor, A. The temporal event graph. *Journal of Complex Networks* (2017).

30. Pandhre, S., Mittal, H., Gupta, M. & Balasubramanian, V. N. Stwalk: learning trajectory representations in temporal graphs. In *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*, 210–219 (ACM, 2018).
31. Singer, U., Guy, I. & Radinsky, K. Node embedding over temporal graphs. *arXiv preprint arXiv:1903.08889* (2019).
32. Béres, F., Pálovics, R., Kelen, D., Szabó, D. & Benczúr, A. Node embeddings in dynamic graphs. *7th International Conference on Complex Networks and Their Applications, Cambridge* 178–180 (2018).
33. Kumar, S., Zhang, X. & Leskovec, J. Learning dynamic embeddings from temporal interactions. *arXiv preprint arXiv:1812.02289* (2018).
34. Sato, K., Oka, M., Barrat, A. & Cattuto, C. DyANE: Dynamics-aware node embedding for temporal networks. *arXiv e-prints arXiv:1909.05976* (2019).
35. Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
36. Cattuto, C. *et al.* Dynamics of person-to-person interactions from distributed rfid sensor networks. *PloS one* 5, e11596 (2010).
37. Yin, Z. Understand functionality and dimensionality of vector embeddings: the distributional hypothesis, the pairwise inner product loss and its bias-variance trade-off. *arXiv preprint arXiv:1803.00502* (2018).
38. Gauvin, L. *et al.* Randomized reference models for temporal networks. *arXiv preprint arXiv:1806.04032* (2018).
39. Supriya Pandhre, M. G., Mittal, H. & Balasubramanian, V. N. Stwalk, <https://github.com/supriya-pandhre/STWalk> (Last Access: April 2019).
40. Ferenc Béres, D. M. K. D. S., Pálovics, R. & Benczúr, A. A. Node embeddings in dynamic graphs. In *Book of Abstracts of the 7th International Conference on Complex Networks and Their Applications*, 165–167 (2018).
41. Ferenc Béres, D. M. K. D. S., Pálovics, R. & Benczúr, A. A. Online node2vec, <https://github.com/mislam5285/online-node2vec> (Last Access: April 2019).
42. Béres, F., Pálovics, R., Oláh, A. & Benczúr, A. A. Temporal walk based centrality metric for graph streams. *Applied network science* 3, 32 (2018).
43. Fogaras, D. & Rácz, B. Scaling link-based similarity search. In *Proceedings of the 14th international conference on World Wide Web*, 641–650 (ACM, 2005).
44. Anderson, R. M. & May, R. M. *Infectious diseases of humans: dynamics and control* (Oxford university press, 1992).
45. Bailey, N. T. *et al.* *The mathematical theory of infectious diseases and its applications* (Charles Griffin & Company Ltd, 5a Crendon Street, High Wycombe, Bucks HP13 6LE., 1975).
46. Daley, D., Gani, J. & Yakowitz, S. An epidemic with individual infectivities and susceptibilities. *Mathematical and computer modelling* 32, 155–167 (2000).
47. Diekmann, O. & Heesterbeek, J. A. P. *Mathematical epidemiology of infectious diseases: model building, analysis and interpretation*, vol. 5 (John Wiley & Sons, 2000).
48. Barrat, A., Barthelemy, M. & Vespignani, A. *Dynamical processes on complex networks* (Cambridge university press, 2008).
49. Bro, R. & Kiers, H. A. A new efficient method for determining the number of components in parafac models. *Journal of Chemometrics: A Journal of the Chemometrics Society* 17, 274–286 (2003).

Acknowledgements

All authors are thankful for the insightful comments of M.Kivelä. MK and LG acknowledge support from the IXXI Complex System Institute. MK has been supported by the ACADEMICS (IDEX Lyon) and by the DataRedux ANR (ANR-19-CE46-0008) projects.

Author contributions

L.G. and M.K. designed research; L.G., M.T. and M.K. performed research; M.T. analyzed data; and L.G., M.T. and M.K. wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-63221-2>.

Correspondence and requests for materials should be addressed to L.G.

Reprints and permissions information is available at www.nature.com/reprints.

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



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 license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020