Temporal Network Analysis

(tnet_analysis toolkit helper)

Simachew Mengiste

2025

1 Introduction

Many complex systems, from brains to societies, produce time-varying patterns of interaction. In neuroimaging, for instance, functional magnetic resonance imaging (fMRI) yields multivariate time series reflecting the activity of N brain regions. A common way to study their interactions is through *dynamic functional connectivity (DFC)*, where temporal windows of the time series are correlated to yield evolving connectivity matrices.

From correlation matrices to temporal networks. By thresholding DFC matrices at levels chosen to match a target average degree, we obtain a sequence of binary adjacency matrices. This sequence is naturally interpreted as a *temporal network (TNet)*:

$$\mathcal{T} = \{B^{(1)}, B^{(2)}, \dots, B^{(S)}\}, \qquad B^{(t)} \in \{0,1\}^{N\times N}. \tag{1}$$

Edges in $\mathcal{T}$ represent strong statistical associations at a given time window. Importantly, these edges do not imply causal influence or physical transmission; instead, they encode *patterns of co-fluctuation* that may persist, shift, or recur across time.

Why analyze temporal networks built from correlations? Even though the edges reflect statistical association, organizing them as a TNet allows us to leverage a powerful toolbox of graph-theoretic and information-theoretic measures:

- **Dynamism vs. statism:** quantify how rapidly association patterns change (transition rates, entropy, similarity, mutual information).
- **Latency and circulation:** assess how quickly time-respecting chains of associations appear between regions, and how often they recur back to their source, revealing temporal persistence or recurrence of co-fluctuation.
- **Integration vs. segregation:** analyze clustering, modularity, and partner stability/diversity to understand whether regions form cohesive groups or shift partners over time.

Conceptual interpretation. In this framework, TNets capture not physical flows but *the temporal organization of statistical dependence*. A short latency between i and j indicates that strong association pathways align across consecutive windows, making j predictable from i over a short lag. High circulation rates mean that local association neighborhoods tend to recur, even if shuffled across partners. Entropy and similarity indices quantify the balance between stability (redundancy) and variability (flexibility) of associations. Together, these measures characterize the temporal landscape of correlation structure in a principled way.

Roadmap. The remainder of this article is structured as follows:

1. From time series to temporal networks (Section 2).
2. Null models and temporal network generators (Section 3).
3. Measures of dynamism and statism (Section 4).
4. Latency and circulation metrics (Section 5).
5. Segregation and cohesion measures (Section 6).

2 From Time Series to Temporal Networks

2.1 From Time Series to Dynamic Functional Connectivity

In many applications, such as brain imaging or multivariate physiological recordings, we begin with a time series $X(t)$ of dimension $T \times N$, where T is the number of time points and N is the number of regions or nodes.

The notion of **dynamic functional connectivity (DFC)** captures the fact that correlations between nodes are not static, but may change over time. A standard way to estimate DFC is the *sliding window correlation method*:

$$A^{(s)} = \operatorname{corr}\big(X_{[s:s+W]}^{\top}\big), \quad s = 0, \ell, 2\ell, \dots \tag{2}$$

where W is the window length, ℓ is the step size (lag), and $A^{(s)}$ is the connectivity matrix for the snapshot starting at time s.

This produces a sequence of weighted, symmetric adjacency matrices:

$$\mathcal{A} = \{A^{(1)}, A^{(2)}, \dots, A^{(S)}\}, \tag{3}$$

where each entry $A_{ij}^{(t)} \in [-1, 1]$ reflects correlation between nodes i and j in snapshot t.

2.2 Thresholding and Binarization

While DFC matrices capture continuous correlation values, most network-based measures require binary adjacency matrices. This raises the question: *at what threshold θ should we binarize?*

A naive fixed threshold (e.g. keeping correlations $r > 0.5$) may produce networks with widely varying densities across subjects or conditions, making comparisons difficult. Instead, a principled approach is to control the **average degree** of the resulting networks.

The degree of a node i in snapshot t after thresholding is

$$k_i^{(t)}(\theta) = \sum_{j \neq i} \mathbf{1}\{A_{ij}^{(t)} > \theta\}, \tag{4}$$

and the average degree across all nodes and snapshots is

$$\bar{d}(\theta) = \frac{1}{SN} \sum_{t=1}^{S} \sum_{i=1}^{N} k_i^{(t)}(\theta). \tag{5}$$

We then search for a threshold θ^* such that $\bar{d}(\theta^*) \approx d_{\text{target}}$, a target average degree chosen by the researcher.

This ensures **comparability**: all temporal networks produced have approximately the same density, independent of absolute correlation values in the data.

2.3 From DFC to Temporal Network (TNet)

Once the threshold is found, each weighted snapshot is binarized:

$$B_{ij}^{(t)} = \mathbf{1}\{A_{ij}^{(t)} > \theta^*\}. \tag{6}$$

Diagonal entries are set to 1 (to allow self-waiting in temporal paths), yielding the final temporal network:

$$\mathcal{T} = \{B^{(1)}, B^{(2)}, \dots, B^{(S)}\}. \tag{7}$$

This sequence of binary matrices forms the **temporal network (TNet)** on which all subsequent analyses (integration, segregation, dynamism) can be applied.

2.4 Purpose

- **Normalization:** Controlling density via target average degree eliminates trivial effects of sparsity/density on higher-level metrics.
- **Comparability:** Networks across subjects, sessions, or conditions become directly comparable.
- **Interpretability:** Many temporal network measures (e.g., shortest paths, clustering, circulation) are defined in binary settings and assume consistent density.
- **Flexibility:** The choice of target average degree d_{target} allows tailoring networks to the resolution appropriate for the study.

3 Temporal Network Generators

3.1 Why Generate Null Models?

To understand the significance of empirical temporal network structure, it is essential to compare against *null models* that preserve some aspects of the data while randomizing others. Null models allow us to separate the effect of trivial properties (e.g. density, activation counts) from meaningful structural organization.

We distinguish between two broad classes:

1. **Static temporal networks**, where a single frame is repeated across all time steps.
2. **Dynamic temporal nulls**, where temporal evolution is reshuffled or randomized in controlled ways.

3.2 Static Temporal Networks

A static temporal network is obtained by generating a single adjacency matrix of a chosen topology and repeating it across all time steps T. The average degree is matched to that of the empirical temporal network for comparability.

- **Static ER (Erdős–Rényi):** a random graph with M edges chosen uniformly among $\binom{N}{2}$ possible pairs.
- **Static SW (Small-World):** a Watts–Strogatz small-world network with rewiring probability p_{rand}, combining clustering with short paths.
- **Static SF (Scale-Free):** a scale-free network generated with preference weights, in four variants:
  1. Linear preference
  2. Exponential preference
  3. Power-law preference
  4. Hybrid (convex combination of linear and power-law)

Since these networks are repeated identically in time, they have *no temporal variability*. They serve as baselines to assess the role of temporal change.

3.3 Dynamic Temporal Nulls

Dynamic nulls explicitly manipulate the time-varying structure. They preserve some statistics of the empirical temporal network but shuffle or randomize others.

- **Empirical:** the observed temporal network, serving as ground truth.
- **Time-shuffler:** permutes the order of snapshots, destroying temporal correlations while preserving per-frame structure.
- **Edge randomizer:** for each snapshot, replaces the frame with a graph of one of the three network topologies (random, small-world, and scale-free) with the same number of links. Preserves density dynamics but destroys topological structure.
- **Link activation model:** preserves the total number of activations of each edge across time but redistributes them randomly across snapshots.
- **Density variability model:** Generates a temporal network where the number of links per snapshot follows a prescribed mean and variance (free parameters). The total number of links across all snapshots is fixed to that of the empirical temporal network, while the distribution of links across time is adjusted to match the chosen density statistics.

Together, these models provide a principled set of baselines to test hypotheses about the structure and dynamics of empirical temporal networks.

Table 1: Summary of temporal network null models and the properties they preserve. "Temporal order" indicates whether the original sequence of snapshots is preserved. "Density distribution" refers to whether the number of edges per snapshot is preserved.

Model	**Temporal order**	**Density distribution**	**Edge activation**
Static	No	No	No
Time-shuffler	No	Yes	Yes
Edge randomizer	Yes	Yes	No
Link activation	No	No	Yes
Density variability	Free	Free	No

4 Temporal Dynamism vs. Statism

We summarize how much a temporal network changes ("dynamism") versus remains similar ("statism") by combining complementary quantities computed from a binary temporal network $\mathcal{T} = \{B^{(1)}, \dots, B^{(S)}\}$ of size N across S snapshots. All edgewise computations use the lower triangle ($i < j$).

Notation. Let $E = \binom{N}{2}$ and let $\mathbf{b}^{(t)} \in \{0,1\}^E$ be the flattened lower-triangular edge vector at snapshot t:

$$\mathbf{b}^{(t)} = \big(B_{ij}^{(t)}\big)_{1 \le i < j \le N}.$$

Define the per-edge activation probability across time $p_e = \frac{1}{S}\sum_{t=1}^{S} b_e^{(t)} \in [0,1]$ for $e = 1, \dots, E$.

4.1 Transition probability between successive frames

We quantify how often edges flip state between consecutive snapshots:

$$p_{\text{trans}} = \frac{1}{(S-1)E} \sum_{t=1}^{S-1} \sum_{e=1}^{E} \mathbf{1}\big\{b_e^{(t)} \neq b_e^{(t+1)}\big\}.$$

This is the empirical probability that a randomly chosen edge changes between t and $t+1$. Higher values indicate faster switching dynamics (more volatile edges).

4.2 Global dynamism via entropy of edge activity

For edges that are not identically 0 or 1 (to avoid $\log(0)$), let

$$\mathcal{E}_\star = \{e :\ 0 < p_e < 1\}, \qquad H_e = -\big[p_e \log p_e + (1-p_e)\log(1-p_e)\big].$$

We aggregate and normalize by the maximal possible entropy assuming *all* E edges had $p_e = \frac{1}{2}$:

$$gEnt = \frac{\sum_{e \in \mathcal{E}_\star} H_e}{E \cdot \big(-\log \frac{1}{2}\big)}.$$

Thus $gEnt \in [0,1]$; it is large when many edges spend substantial time both on and off (i.e., temporally heterogeneous activity).

4.3 Successive cosine similarity (statism proxy)

We compare successive snapshots by cosine similarity of their edge vectors:

$$\cos(t, t+1) = \frac{\langle \mathbf{b}^{(t)}, \mathbf{b}^{(t+1)} \rangle}{\|\mathbf{b}^{(t)}\|_2 \, \|\mathbf{b}^{(t+1)}\|_2}.$$

If both $\mathbf{b}^{(t)}$ and $\mathbf{b}^{(t+1)}$ are all-zeros, we set $\cos(t, t+1) = 1$ (identical emptiness); if only one is all-zeros, we set the cross-similarity to 0 elsewhere. The *mean* successive similarity is

$$\text{mnSim} = \frac{1}{S-1} \sum_{t=1}^{S-1} \cos(t, t+1) \in [0, 1].$$

Large mnSim indicates strong statism (snapshots remain similar).

4.4 Combined dynamism index

We combine (2) and (3) multiplicatively,

$$\text{DynH} = gEnt \cdot (1 - \text{mnSim}),$$

which is high only when the network has high global activation entropy *and* successive frames are dissimilar (low cosine similarity). It down-weights cases with high entropy but little frame-to-frame change, or with large change but low global entropy.

4.5 Mutual information between successive frames

As a complementary measure of redundancy (statism), we compute the *normalized mutual information (NMI)* between consecutive snapshots.

For two binary edge vectors $\mathbf{b}^{(t)}$ and $\mathbf{b}^{(t+1)}$ with E entries, we compute their joint distribution $p(x, y)$ over the states $\{(0,0), (0,1), (1,0), (1,1)\}$, and the marginals $p(x)$ and $p(y)$. The mutual information is

$$I(\mathbf{b}^{(t)}; \mathbf{b}^{(t+1)}) = \sum_{x,y \in \{0,1\}} p(x, y) \, \log \frac{p(x, y)}{p(x)\, p(y)}.$$

The normalized form is

$$\text{NMI}(\mathbf{b}^{(t)}, \mathbf{b}^{(t+1)}) = \frac{I(\mathbf{b}^{(t)}; \mathbf{b}^{(t+1)})}{\max\{H(\mathbf{b}^{(t)}), H(\mathbf{b}^{(t+1)})\}},$$

with Shannon entropy $H(\mathbf{b})$. $\text{NMI} \in [0, 1]$. We report the average across all successive pairs:

$$\text{MI} = \frac{1}{S-1} \sum_{t=1}^{S-1} \text{NMI}\big(\mathbf{b}^{(t)}, \mathbf{b}^{(t+1)}\big) \,.$$

Interpretation.

- High p_{trans}, high $gEnt$, low mnSim, low MI $\Rightarrow$ *volatile dynamics*: edges switch often and snapshots differ.

- Low p_{trans}, low *gEnt*, high mnSim, high MI $\Rightarrow$ *static structure*: edges rarely switch and snapshots are redundant.
- DynH serves as a single scalar that is large only under jointly high entropy and low similarity.

5 Latency and Circulation Measures

Temporal networks allow us to study how fast information, influence, or random walkers can spread through time-respecting paths. In this group we capture two complementary notions:

- **Latency**: how long it takes for a node to reach another (or return to itself) via time-respecting connections.
- **Circulation**: the rate and latency of paths that eventually come back to their source.

Notation. A temporal path from node i to j is a sequence of edges $(i = v_0, v_1), (v_1, v_2), \ldots, (v_\ell = j)$ occurring at times $t_1 < t_2 < \cdots < t_\ell$. The *latency* of the path is $t_\ell - t_0$, i.e. the elapsed time from departure to arrival.

5.1 Latency via temporal shortest paths (Smart Walker)

We first consider deterministic earliest-arrival paths. The **Smart Walker** algorithm propagates reachability forward in time: at each snapshot t, we update the reachability matrix by multiplying with $B^{(t)}$ (the adjacency at time t). The first time a node j becomes reachable from i is recorded as the latency d_{ij}.

This yields the *latency matrix* $D = (d_{ij})$, where

$$d_{ij} = \min\{\, \tau : \exists \text{ time-respecting path } i \to j \text{ within } \tau \text{ steps} \,\}.$$

The matrix encodes minimum latencies between all pairs of nodes.

5.2 Latency via random walks (Random Walker)

While the Smart Walker yields deterministic earliest-arrival latencies, we can also study stochastic exploration. A **Random Walker** moves at each time step t by choosing uniformly among the neighbors of its current node in $B^{(t)}$. The first-passage time from i to j is the time when j is visited for the first time. Repeating across many trials gives an empirical latency distribution, summarized by mean latency and reachability rates.

5.3 Circulation latency

To assess how quickly activity returns to its source, we define the **circulation latency** for node i as the minimum $\tau > 0$ such that a time-respecting path starting and ending at i exists within τ steps. Averaging across start times and nodes yields distributions of return latencies.

5.4 Circulation rate

In addition to latency, we quantify how frequently circulation occurs. For node i, the circulation rate is the proportion of trials or start times in which a return path is found within the observation horizon T. This gives an overall measure of how "self-circulating" the network is.

Interpretation.

- **Latency measures** capture the efficiency of temporal reachability (how fast signals can propagate).
- **Circulation latency** measures how quickly activity tends to come back to its source (a notion of temporal recurrence).
- **Circulation rate** quantifies how likely circulation is at all, separating frequent recurrences from rare events.

6 Segregation and Cohesion Measures

Beyond global dynamism and latency, temporal networks also reveal whether nodes form *stable partnerships*, *diversify* their associations, or *cluster* into cohesive neighborhoods. These measures emphasize **segregation**, in the sense of how local neighborhoods evolve in time, complementing global integration measures.

6.1 Node persistence

Persistence captures whether a node tends to maintain links with the *same set of neighbors* across time. For node i, we compute the persistence probability

$$P_i = \frac{1}{|\mathcal{N}_i^\star|} \sum_{j \in \mathcal{N}_i^\star} \frac{1}{S} \sum_{t=1}^{S} B_{ij}^{(t)},$$

where $\mathcal{N}_i^\star$ is the set of neighbors that ever connect to i. This is compared to a theoretical chance threshold (based on mean edge density) to identify nodes that are *significantly persistent*. Persistence is useful for detecting "core" regions in correlation networks whose associations are consistently present, potentially indicating stable communication backbones.

6.2 Partner stability

While persistence tracks whether neighbors exist across time, **partner stability** emphasizes the *recurrence* of specific partners. For node i, we consider the time indices $\{t_k\}$ where i connects to neighbor j, and measure the distribution of inter-contact intervals $\Delta t = t_{k+1} - t_k$. The stability score is higher when these gaps are small on average:

$$S_i = 1 - \frac{\langle \Delta t \rangle}{S}.$$

Stability thus quantifies whether neighbors tend to *reappear frequently*. In correlation-derived TNets, high stability suggests that certain dyads retain synchronized fluctuations over repeated episodes.

6.3 Partner diversity

Nodes may either specialize with a few repeated partners or diversify across many. We measure diversity via normalized Shannon entropy. For node i, define the fraction of time its edge is active with partner j as p_{ij}, restricted to $\sum_j p_{ij} = 1$. Then

$$D_i = \frac{-\sum_j p_{ij} \log p_{ij}}{\log k_i},$$

where $k_i = |\{j : p_{ij} > 0\}|$ is the number of distinct partners. $D_i \in [0, 1]$; high values indicate broad and balanced partner usage, low values indicate specialization or repeated preference. This captures flexibility in temporal correlation structure.

6.4 Neighborhood memory

Neighborhood memory evaluates how similar a node's neighbor set is across time. For lag ℓ, define

$$M_i(\ell) = \frac{1}{S - \ell} \sum_{t=1}^{S-\ell} \frac{|\mathcal{N}_i^{(t)} \cap \mathcal{N}_i^{(t+\ell)}|}{|\mathcal{N}_i^{(t)} \cup \mathcal{N}_i^{(t+\ell)}|},$$

the average Jaccard similarity between neighbor sets $\mathcal{N}_i^{(t)}$. High memory implies structural inertia of neighborhoods. This measure is informative in correlation TNets because it tells us whether groups of co-fluctuating nodes remain intact across windows or reconfigure.

6.5 Clustering coefficients

Clustering describes whether a node's neighbors are themselves interconnected. We extend this to temporal networks in two ways:

- **Static clustering:** average over time of the usual clustering coefficient

$$C_i = \frac{\#\{\text{links among } \mathcal{N}_i^{(t)}\}}{\binom{k_i^{(t)}}{2}}.$$

- **Temporal clustering:** fraction of "time-respecting triangles", where i connects to j at t_1, to k at $t_2 > t_1$, and j connects to k at $t_3 \in [t_1, t_2]$.

Clustering measures are essential for identifying temporal communities: high values reflect segregation into tightly-knit modules of co-fluctuation.

6.6 Returnability

Returnability quantifies whether a node's neighborhood is likely to *return* over time. For node i, let $\mathcal{N}_i^{(t)}$ be its neighbors at t. Returnability is the fraction of neighbors at t that reappear at later times:

$$R_i = \frac{1}{S} \sum_{t=1}^{S} \frac{|\mathcal{N}_i^{(t)} \cap \bigcup_{u>t} \mathcal{N}_i^{(u)}|}{|\mathcal{N}_i^{(t)}|}.$$

High returnability indicates enduring associations, while low values suggest ephemeral or transient partners.

6.7 Link burstiness

Edges themselves can be irregular. For edge (i, j) with activation times $\{t_k\}$, define inter-event intervals $\tau_k = t_{k+1} - t_k$. The burstiness index is

$$B_{ij} = \frac{\sigma_\tau - \mu_\tau}{\sigma_\tau + \mu_\tau},$$

where μ_τ and σ_τ are mean and standard deviation of inter-event times. $B_{ij} = 0$ corresponds to Poisson-like regularity, $B_{ij} > 0$ to bursty clustering, $B_{ij} < 0$ to anti-bursty regular spacing. Burstiness captures irregular association dynamics, which are common in real-world correlation-derived TNets due to fluctuations in shared variance.

Interpretation. Together, these measures illuminate the *segregation profile* of temporal networks:

- High persistence, stability, memory, clustering, and returnability indicate long-lived, cohesive communities (statism).
- High diversity and burstiness reflect flexible, volatile, or exploratory association patterns (dynamism).

Thus, segregation metrics complement dynamism and latency measures by focusing on the mesoscopic organization of co-fluctuation patterns, clarifying whether the system is dominated by stable modules or dynamic partner switching.