Maximum entropy temporal networks

Paolo Barucca
Department of Computer Science, University College London, London, United Kingdom

(Received 15 October 2025; accepted 25 February 2026; published 17 April 2026)

Temporal networks consist of time-stamped directed interactions that may appear continuously in time, yet few studies have directly tackled the continuous-time modeling of networks. Here, we introduce a maximum-entropy approach to temporal networks and with basic assumptions on constraints, the corresponding network ensembles admit a modular and interpretable representation: a set of global time processes and a static maximum-entropy edge, e.g., node pair or probability. This time-edge labels factorization yields closed-form log-likelihoods, degree, clustering, and motif expectations, and yields an entire class of effective generative models. We provide the maximum-entropy derivation for the nonhomogeneous Poisson process (NHPP) intensities governing the probability of directed edges in temporal networks via the functional optimization over path entropy, connecting NHPP modeling to maximum-entropy network ensembles. NHPPs consistently improve log-likelihood over generic Poisson processes, while the maximum-entropy edge labels recover strength constraints and reproduce expected unique-degree curves. We discuss the limitations of this framework and how it can be integrated with multivariate Hawkes calibration procedures, renewal theory, and neural kernel estimation in graph neural networks.

DOI: 10.1103/78vv-hs72

I. INTRODUCTION

Networks are widely used to describe the architecture of social, biological, and economic systems. While static network models have been studied extensively, many real-world systems are better represented by temporal networks, where interactions are time-stamped and evolve continuously over time. In such systems, links appear as discrete events on a fixed or evolving node set, producing temporal structures that cannot be captured by static snapshots alone. A central challenge is to develop models for temporal networks that are both interpretable and predictive, while resting on principled statistical foundations.

In the static setting, maximum-entropy ensembles have become a cornerstone of network science. Inspired by Jaynes' work on information-theoretic inference [1], maximum-entropy network ensembles reproduce prescribed node- or network-level statistics while remaining otherwise maximally random. These ensembles have been successfully applied to reproduce degree and strength distributions, higher-order constraints, and mesoscale structures, thereby providing a baseline against which empirical patterns can be tested [2,3]. Similarly, the identification of communities and mesoscopic organization through block models has led to principled clustering frameworks that combine interpretability with prediction accuracy [4–6]. Together, these works have established maximum-entropy principles and related exponential-family ensembles as indispensable for static network modeling [2].

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

At the same time, temporal dynamics display signatures that cannot be explained by static constraints alone. A wide body of research has shown that interaction streams exhibit burstiness, memory, and correlation patterns across scales, leading to non-Poissonian interevent time distributions [7–9]. Such deviations from Poissonian timing can often be modeled through renewal processes [10], where interevent times are independent but follow heavy-tailed distributions such as power-law, lognormal, or Weibull forms [11–14]. Renewal processes naturally reproduce the alternation of rapid bursts and long inactive periods observed in empirical data, providing a minimal, interpretable baseline for bursty temporal networks that captures temporal heterogeneity without invoking explicit event-event dependencies. In contrast, self-exciting point processes, introduced in Hawkes' seminal work [15], describe causal or feedback-driven event generation, reproducing clustered activity and contagion-like excitation effects in contexts ranging from seismology to financial order flows [16–20].

Temporal networks naturally combine both aspects: structural constraints in the marks (who interacts with whom) and temporal patterns in the timing of events. Empirical studies have shown that such systems exhibit distinctive motifs, heterogeneous attention, and limited memory [7,21]. Recent work introduced maximum-entropy approaches for discrete-time temporal networks [21], while advances in fitting multivariate Hawkes processes have produced highly calibrated generative models [18,22,23]. However, despite these parallel developments, there is no principled maximum-entropy ensemble for temporal networks in continuous time. In particular, there is no framework that jointly encodes structural constraints and realistic bursty dynamics, while remaining analytically tractable and interpretable.

Here we introduce such a framework. By maximizing path entropy for nonhomogeneous Poisson processes under structural and temporal constraints, we derive a family of maximum-entropy ensembles that factorize into a *time* component (modeled by nonhomogeneous Poisson intensities) and a *mark* component (given by static maximum-entropy edge weights subject to strengths, masks, or block constraints). This separation yields closed-form log-likelihoods, analytical expectations for degrees and unique edges, and motif statistics that combine mark probabilities with temporal cross dependencies. Empirically, we show on a benchmark dataset (Enron email dataset) that adding a time layer consistently improves likelihoods, while not necessarily achieving optimality in terms of information criteria, and motif fits over Poisson baselines, while maximum-entropy marks reproduce degree and clustering. Our contribution unifies maximum-entropy ensembles with bursty temporal processes, filling a key gap in the modeling of temporal networks and offering a tractable foundation for future extensions involving calibrated kernels, motif analysis, and neural inference methods.

II. METHODS

Any temporal network observed over a time window $[0, T]$ can be represented as a marked point process. Each interaction is an event occurring at a time $t_k \in [0, T]$ and carrying a mark $m_k = (i \to j)$ that identifies the ordered pair of nodes involved. Thus the dataset consists of events $\{(t_k, m_k)\}_{k=1}^{K}$, where K is the total number of time-stamped edges. We model the event times as a nonhomogeneous Poisson process (NHPP). A NHPP on $[0, T]$ with intensity $\lambda(t) \geqslant 0$ is a counting process $N(t)$ such that the following occurs.

(i) For any measurable set $A \subset [0, T]$, the number of events $N(A)$ is Poisson distributed with mean $\Lambda(A) = \int_A \lambda(t)\, dt$.

(ii) For disjoint sets $A_1, \ldots, A_k$, the random variables $N(A_1), \ldots, N(A_k)$ are independent.

Equivalently, $\lambda(t)$ is the instantaneous event rate, so that $\lambda(t)dt$ gives the probability of observing an event in $[t, t + dt)$ to first order in dt. In the marked setting, we introduce edge-specific intensities $\lambda_{ij}(t) \geqslant 0$, describing the instantaneous rate at which the directed edge (i, j) occurs at time t. The global event rate is, therefore,

$$\Lambda(t) = \sum_{(i,j)\in E} \lambda_{ij}(t),$$

where E denotes the set of admissible directed edges on the node set V. Under this model, the log-likelihood of observing the events $\{(t_k, m_k)\}_{k=1}^{n}$, with $m_k = (i_k \to j_k)$, is

$$\ell(\{\lambda_{ij}\}) = \sum_{k=1}^{K} \log \lambda_{i_k j_k}(t_k) - \sum_{(i,j)\in E} \int_0^T \lambda_{ij}(t)\, dt. \quad (1)$$

The integrated intensity $\mu_{ij} = \int_0^T \lambda_{ij}(t)\, dt = N_{ij}$ corresponds to the expected number of interactions from node i to node j, i.e., the expected directed edge weight in the time-aggregated network. Node strengths are recovered as $s_i^{\text{out}} = \sum_j \mu_{ij}$ and $s_j^{\text{in}} = \sum_i \mu_{ij}$, linking the temporal model directly to standard weighted network quantities. Taking expectation with respect to the process yields the probability-mass adjusted Poisson entropy functional's

$$\mathcal{H}[\lambda] = -\int_0^T [\lambda(t) \log \lambda(t) - \lambda(t)]\, dt.$$

Maximizing entropy thus corresponds to selecting intensity functions that are as unstructured as possible subject to the chosen constraints on the expected values over a set of prescribed functions of the events, in direct analogy to the canonical ensemble in statistical mechanics.

The central methodological step is to impose both temporal and structural constraints on the intensities. Temporal constraints specify aggregate time profiles for sets of edges, e.g., the total rate of activity from a given sender, within a block pair, or globally. Structural constraints fix time-integrated quantities such as edge totals, node strengths, or block sums. Formally, if $\mathcal{P} = \{E_r\}$ is a partition of the edge set and $G_r(t)$ are prescribed profiles, then we require $\sum_{(i,j)\in E_r} \lambda_{ij}(t) = G_r(t)$ for all r and t. Similarly, integrated constraints take the form $\int_0^T \lambda_{ij}(t)\, dt = \mu_{ij} = N_{ij}$, for edge constraints, $\sum_j \mu_{ij} = s_i^{\text{out}}$, for node-strengths constraints, or $\sum_{i\in a, j\in b} \mu_{ij} = N_{ab}$, for block-pair flows. A block partition is defined as a mapping $g : V \to B$ assigning each node to one of $|B|$ groups. Block-pair quantities N_{ab} aggregate expected interactions between nodes in block a and block b. These linear restrictions capture both prior information and desired marginal statistics of the temporal network.

To enforce these conditions, we introduce a Lagrangian functional combining the entropy with time-dependent multipliers $\alpha_r(t)$ for the temporal partitions and static multipliers for the mark-side margins. Explicitly,

$$\begin{aligned}\mathcal{L} = &-\sum_{(i,j)\in E} \int_0^T (\lambda_{ij} \log \lambda_{ij} - \lambda_{ij})\, dt \\ &+ \sum_r \int_0^T \alpha_r(t) \left(\sum_{(i,j)\in E_r} \lambda_{ij}(t) - G_r(t) \right) dt \\ &+ \sum_{(i,j)} \Theta_{ij} \left(\int_0^T \lambda_{ij}(t)\, dt - N_{ij} \right), \end{aligned} \quad (2)$$

and similar terms for node strengths and block constraints. Stationarity with respect to each $\lambda_{ij}(t)$ yields

$$\log \lambda_{ij}(t) = \alpha_{r(i,j)}(t) + \Psi_{ij} - 1, \quad (3)$$

with Ψ_{ij} collecting all static multipliers (edge, sender, receiver, and block) and mask terms. Thus, the solution always factorizes as

$$\lambda_{ij}(t) = \phi_{r(i,j)}(t)\, w_{ij}, \quad \phi_r(t) = e^{\alpha_r(t)}, \quad w_{ij} = e^{\Psi_{ij}-1}. \quad (4)$$

This is the key structural result: time and marks separate cleanly. Each temporal partition has its own profile $\phi_r(t)$, while all static constraints determine the weights w_{ij}. The separation is reminiscent of exponential-family factorization and ensures convexity of the inference problem.

The factorization allows direct enforcement of constraints. For each partition, $\phi_r(t)$ is pinned by $G_r(t)$ up to a normalization, while the integrated quantities $\mu_{ij} = \Phi_{r(i,j)} w_{ij}$ with $\Phi_r = \int_0^T \phi_r(t)\, dt$ must match the prescribed margins. The

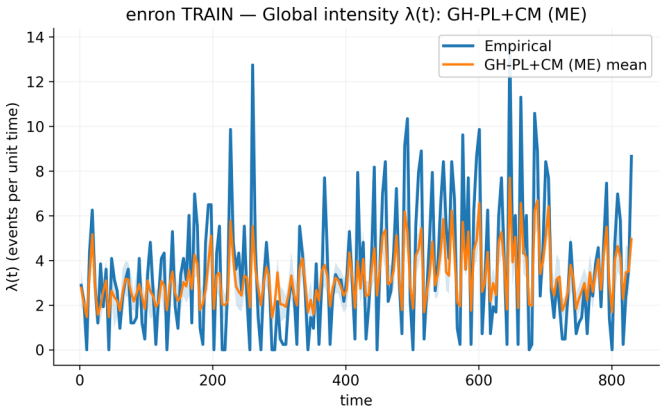


FIG. 1. Enron. Intensity over time obtained by fitting one global Hawkes process to the entire sequence of time stamps.

latter reduces to masked biproportional scaling (iterative proportional fitting) on w_{ij}, a well-studied convex optimization that generalizes the construction of static maximum-entropy ensembles. Importantly, this procedure remains feasible under masks that restrict admissible edges or enforce block quotas, provided the margin sums are consistent.

Several notable cases emerge. If only the edge totals are fixed and no temporal profile is prescribed, one obtains the edge-Poisson model with constant rates $\lambda_{ij}(t) = \mu^{\star}_{ij}/T$. If instead a global Hawkes intensity $\lambda_H(t)$ is imposed as an nonhomogeneous Poisson process, the model factorizes as $\lambda_{ij}(t) = (\mu^{\star}_{ij}/\Lambda_H)\,\lambda_H(t)$, effectively rescaling the Hawkes process across edges. In Fig. 1 we show the empirical and fitted global Hawkes intensity for the Enron dataset. Adding node strengths with a mask yields a temporal analog of the directed weighted configuration model, while imposing block-degree quotas enforces a fixed number of unique edges per block pair. Finally, partitioning by sender and assigning sender-specific profiles recovers a flexible class of nonhomogeneous but still convex ensembles. In each case, sampling reduces to simulating the partition-level time processes and assigning marks according to w_{ij}, thereby generalizing classical shuffling methods to bursty, constraint-aware temporal null models.

A. Network properties

The separation $\lambda_{ij}(t) = \phi_{r(i,j)}(t)\,w_{ij}$ permits a wide range of statistics to be computed in a closed form, providing a tractable description of structural, temporal, and higher-order properties of temporal networks. At the structural level, the integrated intensity $\mu_{ij} = \Phi_{r(i,j)} w_{ij}$ defines the expected number of events along edge (i, j) in $[0, T]$, with $\Phi_r = \int_0^T \phi_r(t)\,dt$.

Summing over targets or sources yields the expected out strength $s_i^{\text{out}} = \sum_j \mu_{ij}$ and in strength $s_j^{\text{in}} = \sum_i \mu_{ij}$, while block sums $\mu_{ab} = \sum_{i\in a, j\in b} \mu_{ij}$ provide mesoscopic statistics. The expected number of distinct edges activated is

$$\mathbb{E}[U_{\to}] = \sum_{i\neq j} (1 - e^{-\mu_{ij}}), \tag{5}$$

i.e., 1 minus the probability of the edge never being active in the time-window considered, generalizing the coverage formulas of static ensembles and allowing direct comparison with observed unique-edge counts. These statistics show how structural constraints can be consistently enforced on average in the ensemble, similarly to the corresponding static ensembles.

Turning to the temporal dimension, the partition profiles $\phi_r(t)$ govern interevent times and burstiness. For a homogeneous Poisson process with rate $\bar{\lambda}$, interevent times are exponentially distributed with variance $1/\bar{\lambda}^2$. In contrast, when $\phi_r(t)$ follows a Hawkes process with branching ratio $\eta < 1$, the stationary rate is $\bar{\lambda}_r = \mu_r/(1-\eta)$ but the variance is inflated by the kernel-dependent autocovariance $\gamma_r(u) = \text{Cov}[\lambda_r(t), \lambda_r(t+u)]$. This produces overdispersion in interevent intervals, a hallmark of bursty temporal dynamics. Within the ensemble, each partition inherits this property, so burstiness is naturally captured by the time-layer while marks remain maximally random subject to constraints. Thus, our framework reproduces both static marginals and temporal dispersion statistics in a unified manner.

Most importantly, the simple combination of mark probabilities and temporal covariances enables explicit computation of motif statistics. Consider ordered event pairs within a lag window $(0, \Delta]$ and an observation interval $I = [a, b]$. For each partition r, define the mass of the rows and columns of the edge weights $r_i^{(r)} = \sum_j \Pi_{ij}^{(r)}$ and $c_j^{(r)} = \sum_i \Pi_{ij}^{(r)}$, where $\Pi_{ij}^{(r)} = w_{ij}/W_r$ is the normalized edge distribution, where $W_r := \sum_{(i,j)\in E_r} w_{ij}$.

Then we can evaluate the pure-edge probabilities of two-edge motifs, in particular, we show the repeat case and reciprocity case, and the broadcast and convergence case

$$p_{\text{REP}}^{(r,s)} = \sum_{i,j} \Pi_{ij}^{(r)}\,\Pi_{ij}^{(s)}, \tag{6}$$

$$p_{\text{REC}}^{(r,s)} = \sum_{i,j} \Pi_{ij}^{(r)}\,\Pi_{ji}^{(s)}, \tag{7}$$

$$p_{\text{BCAST}}^{(r,s)} = \sum_{i} r_i^{(r)}\,r_i^{(s)} \;-\; p_{\text{REP}}^{(r,s)}, \tag{8}$$

$$p_{\text{CONV}}^{(r,s)} = \sum_{j} c_j^{(r)}\,c_j^{(s)} \;-\; p_{\text{REP}}^{(r,s)}, \tag{9}$$

where "repeat" considers the repetition of the same edge, "reciprocity" the appearance of both the connection from i to j and from j to i, the broadcast case shows the appearance of two (distinct) edges with the same source i, and finally in the convergence case, node j receives two (distinct) edges. This is the full list of two-edge motifs involving two and three nodes in which time ordering is irrelevant: if we account for time ordering we could consider the chain and reverse-chain (see Fig. 2).

These quantities are purely structural, i.e., edge dependent yet time independent, reflecting the allocation of the probability mass across edges. Temporal amplification is captured by the cross moment

$$C_{r,s}(\Delta; I) = \int_a^b \int_0^{\Delta} \mathbb{E}[\lambda_r(t)\lambda_s(t+u)]\,du\,dt, \tag{10}$$

where $\lambda_r(t) = \sum_{(i,j)\in E_r} \lambda_{ij}(t)$. For $r \neq s$ independent, $C_{r,s} = |I|\Delta\bar{\lambda}_r\bar{\lambda}_s$, while for $r = s$ one obtains $C_{r,r} = |I|[\bar{\lambda}_r^2\Delta + \int_0^{\Delta}\gamma_r(u)\,du]$. Combining both pieces yields the general motif

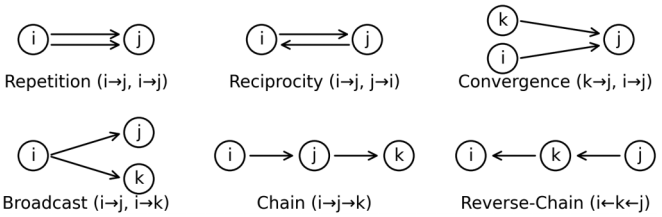


FIG. 2. Directed two-event temporal motifs involving, at most, three nodes. From left to right: *repeat* ($i \to j$, $i \to j$), *reciprocity* ($i \to j$, $j \to i$), *convergence* ($i \to j$, $k \to j$), *broadcast* ($i \to j$, $i \to k$), *chain* ($i \to j$, $j \to k$), and *reverse-chain* ($i \to j$, $k \to i$). These six configurations form a mutually exclusive and exhaustive set of two-event motifs including at most three distinct nodes.

master formula

$$\mathbb{E}\big[M^{(\Delta)}_{\mathrm{TYPE}}(I)\big] = \sum_{r,s} p^{(r,s)}_{\mathrm{TYPE}}\, C_{r,s}(\Delta; I). \tag{11}$$

This decomposition explicitly shows that temporal autocorrelation γ_r amplifies short-lag motifs, while mark probabilities control the balance among motif types. In particular, self-exciting Hawkes kernels enhance broadcast and convergence motifs in line with empirical bursty communication, while static constraints can enforce degree and clustering profiles.

Together, these results demonstrate the analytical tractability of the ensemble. Edge-level expectations follow from μ_{ij}, burstiness is quantified by interevent variance and autocovariance, and motif statistics emerge from a combination of structural and temporal factors.

III. RESULTS

A central contribution of our framework is to provide a principled benchmark against which empirical temporal network properties can be evaluated for statistical significance. By specifying both structural constraints (degrees, strengths, and block totals) and temporal constraints (global or partitioned profiles, Hawkes kernels, and renewal processes), the maximum-entropy ensembles define a family of null models that preserve low-order features while randomizing higher-order patterns.

This allows us to compute the expected values of edge statistics, burstiness, and motifs under the null model and to compare them directly with observed values. Any significant deviation between empirical and expected quantities can thus be interpreted as evidence for structure beyond the imposed constraints, analogously to the role of canonical ensembles in equilibrium statistical mechanics.

We tested this principle across four widely studied temporal network datasets and report results on the Enron mail corpus. The Enron corpus records time-stamped internal communications among employees of the Enron Corporation over a period of years. This dataset has become a benchmark for studies of temporal communication patterns because it exhibits both strong structural heterogeneity (hubs of activity, communities of practice) and highly bursty temporal dynamics. Figure 3 provides an analysis of the temporal and degree heterogeneity of the Enron dataset, where edges are sorted by total strength $s_i^{\mathrm{out}} + s_i^{\mathrm{in}}$. The left panel shows the logarithm of event counts per directed edge, revealing strong structural heterogeneity: a small set of high-activity nodes concentrate a large fraction of interactions, while most dyads remain sparsely populated. The right panel reports the burstiness index for each edge, highlighting substantial temporal variability even among edges with comparable total counts. Highly active edges are not uniformly bursty; rather, burstiness varies across dyads, indicating that structural intensity and temporal irregularity are partially decoupled.

The challenge is to separate the features of the dataset (such as degree heterogeneity or daily cycles) from significant temporal signatures, such as anomalous motifs, excess

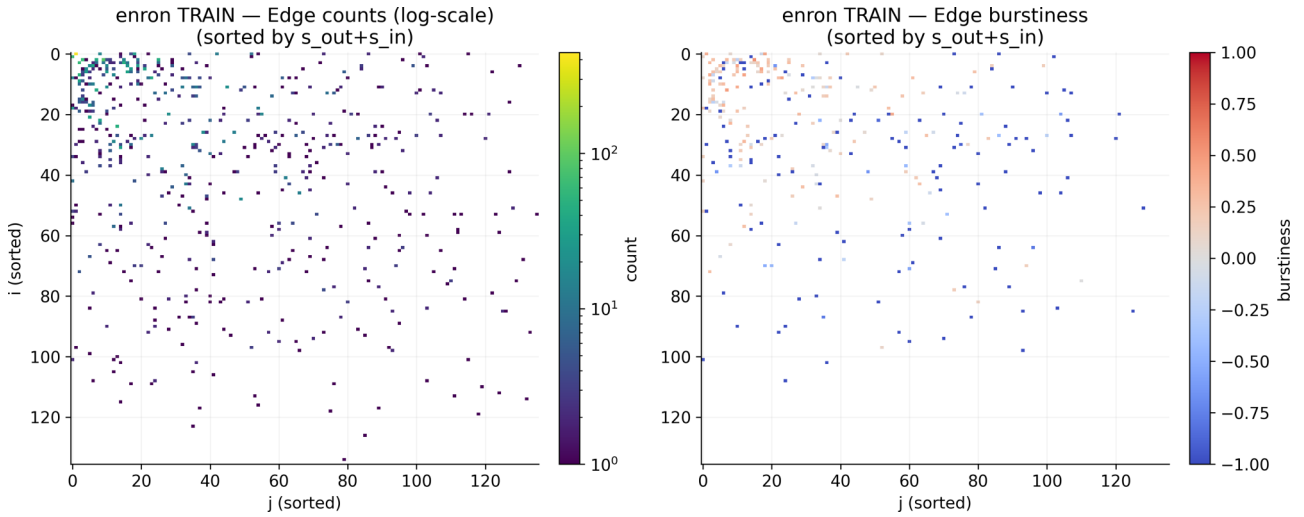


FIG. 3. Log of event counts per directed edge (left) and burstiness index per edge (right) for the Enron dataset. Nodes are sorted by total strength $s_i^{\mathrm{out}} + s_i^{\mathrm{in}}$, and edges are shown in the resulting order. The train split denotes the temporal segment of the data used for parameter calibration.

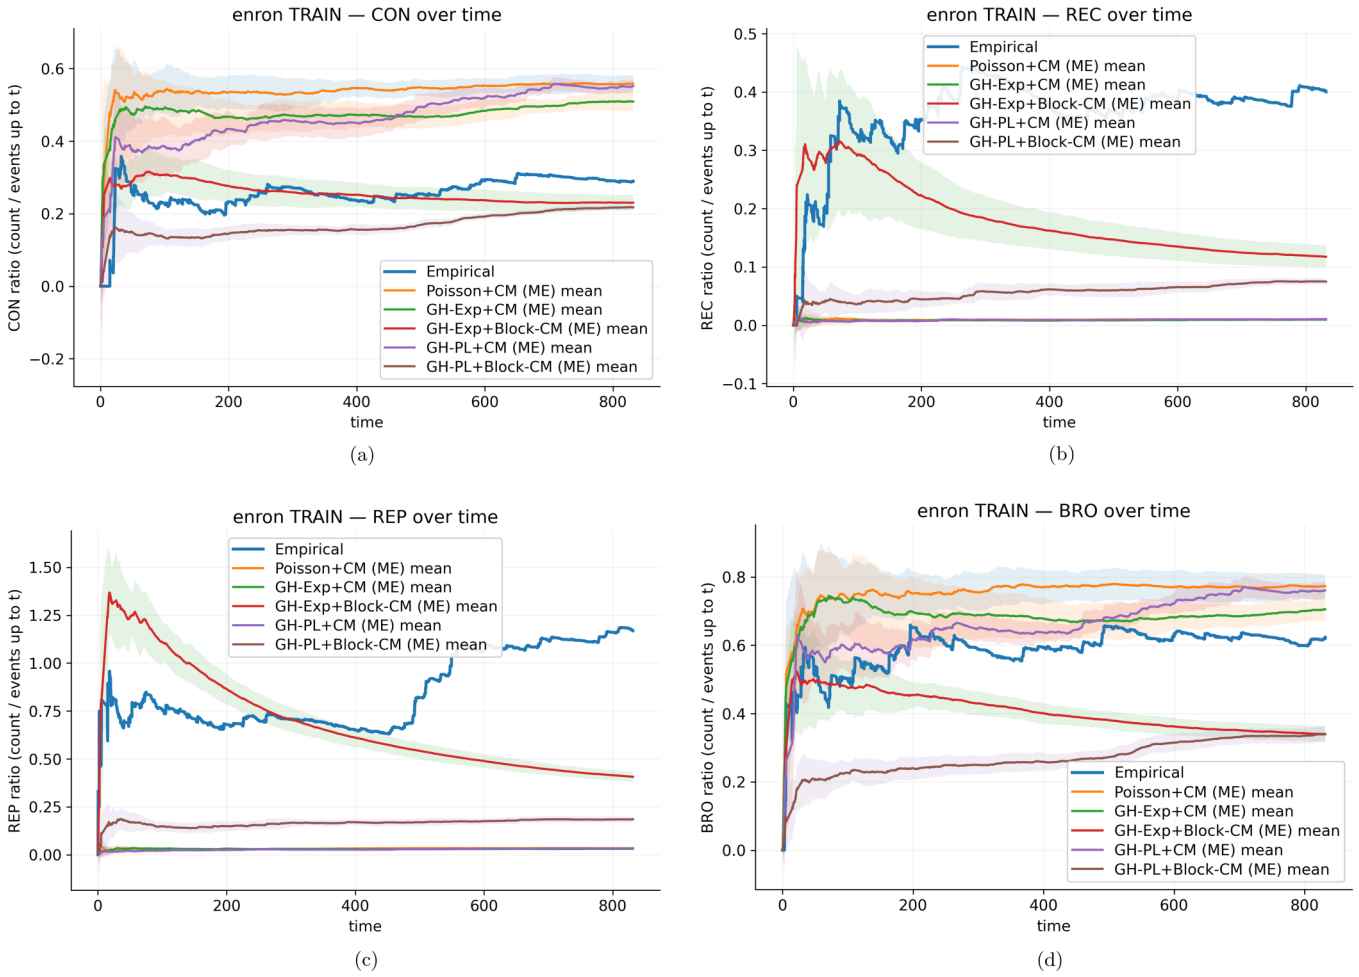


FIG. 4. Time-dependent motif ratios for the Enron dataset versus the maximum-entropy temporal network ensembles obtained combining Poisson and Exp-Hawkes (GH-Exp) and PL-Hawkes (GH-PL) temporal profiles and configuration model (CM) and block-configuration model (Block-CM) edge structures. (a) CON, Δ_{large}, (b) REC, Δ_{large}, (c) REP, Δ_{large}, (d) BRO, Δ_{large}.

reciprocity, or collective bursts, by defining ensembles that enforce the known marginals (e.g., strengths, block structure) and temporal baselines (Poisson or Hawkes profiles) while leaving higher-order correlations unconstrained. By generating expectations and variances for the relevant statistics under these ensembles, we can define the associated empirical deviations.

We test that the frequency of reciprocated motifs largely exceeds the maximum-entropy expectation given node strengths and bursty Hawkes timing. This can be simply interpreted as people working at Enron replying to each other more than would be expected simply from being active and emailing in bursts, thus suggesting genuine conversational feedback between workers. Figure 4 reports the time evolution of the four directed motifs under Δ_{large} for the Enron TRAIN split, comparing empirical ratios with the maximum-entropy temporal ensembles. Reciprocity (REC) and repetition (REP) are substantially underestimated by configuration-only models, even when coupled with Hawkes timing, indicating that strength constraints and global burstiness alone are insufficient to reproduce dyadic memory effects. Introducing block structure partially corrects this gap at short times but still decays away from the empirical curve, suggesting that mesoscale structure explains part, but not all, of the excess reciprocity. For convergence (CON) and broadcast (BRO), Hawkes-based models improve the early time amplification relative to Poisson baselines, yet persistent discrepancies remain, especially at longer horizons where empirical ratios stabilize above the ensemble means. Overall, the figure highlights that while the time-mark factorization captures bursty temporal activity and degree constraints, motif persistence in the data reflects additional temporal correlations beyond those encoded by global Hawkes excitation and static block masks. The ensemble serves as the baseline against which significance is judged: if an observed property falls within the distribution generated by the maximum-entropy temporal model, it is explained by the enforced constraints; if not, it signals mechanisms beyond those constraints.

The methodology requires only event times and node pairs, and it is applicable to any temporal network dataset. The case study presented here serves as an example, but the same procedure can be extended to neuronal spike trains, epidemiological contact networks, and financial transactions. In this sense, our framework provides a temporal-network analog

of static network maximum-entropy ensembles: a baseline against which higher-order structure can be detected, quantified, and interpreted across systems.

IV. CONCLUSION

We introduced a maximum-entropy framework for temporal networks in continuous time, combining nonhomogeneous Poisson processes with static maximum-entropy edge distributions. The key consequence of the derivation is a clean factorization of time and marks, which yields closed-form likelihoods, tractable fitting procedures, and explicit formulas for edge coverage, burstiness, and motif statistics. The framework provides a benchmark for assessing which features are explained by basic structural and temporal constraints and which require additional mechanisms. In particular, the use of NHPP improves agreement with empirical interevent distributions and motif counts, while the static maximum-entropy edge modeling accounts for degree and block structure. The resulting ensembles thus serve as principled null models for temporal networks, offering a consistent basis for significance testing and model comparison. The present results allow immediate generalizations in several directions. The use of nonhomogeneous Poisson processes enables us to fit an observed bursty behavior in the data or to impose it altogether with a theoretical model. In the maximum-entropy construction here presented there is no direct self-exciting behavior in the dataset generation, i.e., the self-excitation is driven by the original time-stamped process. The introduction of this ensemble connecting temporal networks to nonhomogeneous Poisson processes provides a framework to naturally explore richer time profiles in networks: accounting for burstiness-degree correlations by introducing network-dependent factors in the intensity rates, including explicit heterogeneous calibration of the interevent time distribution based on renewal processes theory, fitting block- and node-specific multivariate Hawkes processes, and generalizing the kernel estimation to neural network functions [24], thus contributing to bridging the gap between maximum entropy network ensembles and temporal graph-neural networks [25]. The framework can be used to evaluate higher-order structures beyond two-event motifs, providing a systematic baseline for studying temporal clustering, cascades, or community dynamics. In all these cases, the maximum-entropy perspective clarifies what can be expected from constraints alone and where genuinely new patterns emerge in networked systems changing continuously over time.

This paper has Supplemental Material [26] that derives the maximum-entropy factorization of temporal networks into time and mark components, providing closed-form expressions for likelihoods, edge counts, strengths, and motif statistics under Poisson and Hawkes kernels. It also presents extended empirical results across four benchmark datasets (Enron, Militarized Interstate Disputes, Reality Mining, and Facebook), including motif ratio plots, log-likelihood tables, and model comparison statistics.

DATA AVAILABILITY

The data that support the findings of this article are openly available. All code used to generate the figures and reproduce the experiments is publicly available at [27]. All figures are fully reproducible from the repository.

[1] E. T. Jaynes, Information theory and statistical mechanics, Phys. Rev. **106**, 620 (1957).

[2] G. Cimini, T. Squartini, F. Saracco, D. Garlaschelli, A. Gabrielli, and G. Caldarelli, The statistical physics of real-world networks, Nat. Rev. Phys. **1**, 58 (2019).

[3] M. Newman, *Networks* (Oxford University Press, New York, 2018).

[4] B. Karrer and M. E. J. Newman, Stochastic blockmodels and community structure in networks, Phys. Rev. E **83**, 016107 (2011).

[5] K. Rohe, T. Qin, and B. Yu, Co-clustering directed graphs to discover asymmetries and directional communities, Proc. Natl. Acad. Sci. **113**, 12679 (2016).

[6] M. E. J. Newman, The structure of scientific collaboration networks, Proc. Natl. Acad. Sci. **98**, 404 (2001).

[7] N. Masuda and R. Lambiotte, *A Guide to Temporal Networks*, Complexity Science, Vol. 4 (World Scientific, Singapore, 2016).

[8] P. Holme and J. Saramäki, Temporal networks, Phys. Rep. **519**, 97 (2012).

[9] A. Vázquez, B. Rácz, A. Lukács, and A.-L. Barabási, Impact of non-Poissonian activity patterns on spreading processes, Phys. Rev. Lett. **98**, 158702 (2007).

[10] T. Hiraoka, N. Masuda, A. Li, and H.-H. Jo, Modeling temporal networks with bursty activity patterns of nodes and links, Phys. Rev. Res. **2**, 023073 (2020).

[11] K.-I. Goh and A.-L. Barabási, Burstiness and memory in complex systems, Europhys. Lett. **81**, 48002 (2008).

[12] M. Karsai, K. Kaski, A.-L. Barabási, and J. Kertész, Universal features of correlated bursty behaviour, Sci. Rep. **2**, 397 (2012).

[13] D. R. Cox, *Renewal Theory* (Methuen & Co. Ltd., London, 1962).

[14] S. Unicomb, G. Iñiguez, J. P. Gleeson, and M. Karsai, Dynamics of cascades on burstiness-controlled temporal networks, Nat. Commun. **12**, 133 (2021).

[15] A. G. Hawkes, Spectra of some self-exciting and mutually exciting point processes, Biometrika **58**, 83 (1971).

[16] Y. Ogata, Statistical models for earthquake occurrences and residual analysis for point processes, J. Am. Stat. Assoc. **83**, 9 (1988).

[17] E. Bacry, I. Mastromatteo, and J.-F. Muzy, Hawkes processes in finance, Market Microstructure and Liquidity **01**, 1550005 (2015).

[18] E. Bacry and J.-F. Muzy, First- and second-order statistics characterization of hawkes processes and non-parametric estimation, IEEE Trans. Inf. Theory **62**, 2184 (2016).

[19] V. Filimonov and D. Sornette, Quantifying reflexivity in financial markets: Toward a prediction of flash crashes, Phys. Rev. E **85**, 056108 (2012).

[20] I. Seabrook, P. Barucca, and F. Caccioli, Modelling equity transactions as bursty processes, *Data-Driven Modelling* (World Scientific, 2026).

[21] G. V. Clemente, C. J. Tessone, and D. Garlaschelli, Temporal networks with node-specific memory: Unbiased inference of transition probabilities, relaxation times, and structural breaks, Phys. Rev. Res. **6**, 043257 (2024).

[22] M. Achab, E. Bacry, S. Gaïffas, I. Mastromatteo, and J.-F. Muzy, Uncovering causality from multivariate hawkes integrated cumulants, J. Mach. Learn. Res. **18**, 1 (2017).

[23] H. Soliman, L. Zhao, Z. Huang, S. Paul, and K. S. Xu, in *Proceedings of the 39th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 162 (PMLR, Baltimore, 2022), pp. 20329–20346.

[24] H. Mei and J. M. Eisner, The neural hawkes process: A neurally self-modulating multivariate point process, *31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA*, Vol. 30 (ACM, 2017).

[25] X. Tian, X. Zhang, X. Du, and T. Lu, in *Proceedings of the 34th ACM International Conference on Information and Knowledge Management* (ACM, New York, 2025), pp. 2895–2904.

[26] See Supplemental Material at http://link.aps.org/supplemental/10.1103/78vv-hs72 for additional details.

[27] P. Barucca, Maximum-entropy-temporal-networks Public, Zenodo (2026), https://doi.org/10.5281/zenodo.18613086.