



HAL
open science

Statistical models of complex brain networks: a maximum entropy approach

Vito Dichio, Fabrizio De Vico Fallani

► **To cite this version:**

Vito Dichio, Fabrizio De Vico Fallani. Statistical models of complex brain networks: a maximum entropy approach. Reports on Progress in Physics, 2023, 86 (10), pp.102601. 10.1088/1361-6633/ace6bc . hal-03781966v2

HAL Id: hal-03781966

<https://inria.hal.science/hal-03781966v2>

Submitted on 18 Jan 2024

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

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



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

Statistical models of complex brain networks: a maximum entropy approach

Vito Dichio¹, Fabrizio De Vico Fallani¹

¹ Sorbonne Universite, Paris Brain Institute - ICM, CNRS, Inria, Inserm, AP-HP, Hopital de la Pitie Salpêtriere, F-75013, Paris, France

E-mail: fabrizio.de-vico-fallani@inria.fr

14 August 2023

Abstract. The brain is a highly complex system. Most of such complexity stems from the intermingled connections between its parts, which give rise to rich dynamics and to the emergence of high-level cognitive functions. Disentangling the underlying network structure is crucial to understand the brain functioning under both healthy and pathological conditions.

Yet, analyzing brain networks is challenging, in part because their structure represents only one possible realization of a generative stochastic process which is in general unknown. Having a formal way to cope with such intrinsic variability is therefore central for the characterization of brain network properties.

Addressing this issue entails the development of appropriate tools mostly adapted from network science and statistics. Here, we focus on a particular class of maximum entropy models for networks, *i.e.* exponential random graph models (ERGMs), as a parsimonious approach to identify the local connection mechanisms behind observed global network structure. Efforts are reviewed on the quest for basic organizational properties of human brain networks, as well as on the identification of predictive biomarkers of neurological diseases such as stroke.

We conclude with a discussion on how emerging results and tools from statistical graph modeling, associated with forthcoming improvements in experimental data acquisition, could lead to a finer probabilistic description of complex systems in network neuroscience.

Keywords: Statistical modeling, Complex systems, Exponential random graph model, Brain networks, Inference, Maximum entropy principle.

Contents

1	Introduction	2
2	Brain networks in a nutshell	4
2.1	Building brain networks from experimental data	4
2.2	The rationale of a statistical approach	5
3	Exponential random graph model (ERGM)	6
3.1	Model parameter estimation	7
3.2	Interpretation of ERGMs	9
3.3	Graph statistics for ERGMs	10

4 ERGMs in neuroscience 14

4.1 Minimal model of brain networks 15

4.2 Identification of discriminant brain network features 19

4.3 Predicting states in temporal brain networks 21

5 On the interpretation of ERGMs 24

5.1 ERGMs and the maximum entropy principle (MEP) 24

5.2 Description and prediction: the choice of graph statistics 26

5.3 Explanation: ERGMs and Stat.Mech. 26

6 Conclusion and perspectives 27

Acknowledgements 28

References 28

1. Introduction

The human brain is a biological system of tremendous complexity. At different scales of neuronal organization, the paradigm of a system “made up of a large number of parts that interact in a nonsimple way” [1] turns out to be an apt abstraction. Notably, it suits neurons interacting through synapses at the microscale as well as brain regions’ activity coordinating at the macroscale and resulting in the rich spectrum of mind’s functional states. The study of the brain as a complex system has flourished in the last 20 years, drawing analogies from the physics of disordered systems, graph theory, dynamical systems and fueled by the modern deluge of data upon nearly all scientific fields [2].

By explicitly representing the interactions between the system’s components, networks, or graphs, constitute a natural and powerful way to inquire its organizing principles. The description of the brain under the lens of network science has lead to a number of fundamental results. Topological network properties, such as node centrality, modular organization and global efficiency, play a fundamental role in the emergence of basic physiological functions, as well as on the apperance of many brain diseases [3, 4, 5].

Beyond purely descriptive analyses, important questions on the network structure include what is the underlying generative process and how local wiring rules result in the observed large-scale properties [6]. To this end, network models based on random edge rewiring or nodal preferential attachment rules have been initially investigated and allowed to assess nontrivial properties of brain networks, such as integration and segregation of information or the presence of few nodes with a high number of connections, *i.e.* the so-called hubs [7]. In addition to purely topological models, brain networks can be more finely modeled by explicitly taking into account the spatial position of the nodes so as to penalize the cost of having long distance connections and minimize the associated metabolic consumption [8, 9].

While these models allow to identify putative mechanisms involved in generating the observed brain networks, they implicitly make the assumption that the connections have a physical meaning, *i.e.* they are tangible quantities. However, it is important to remind that almost all brain networks are currently inferred from experimental data through advanced tools from image and signal processing (**Fig.1a**) [10]. In particular,




Figure 1: **Statistical analysis of brain networks inferred from experimental data.** (a) Neuroimaging techniques provide images (*e.g.*, diffusion tensor imaging) and signals (*e.g.*, electroencephalography) informing about the anatomical structure and functional dynamics of the brain, respectively. (b) Methods from image and signal processing are then used to build the brain network G^* , where nodes correspond to different brain regions and links correspond to structural connections or functional interactions. (c) The aim of statistical modeling is to describe $P(G)$, the probability distribution in the space of all possible graphs G . It encodes the architectural principles of the observed brain network G^* . Typically, the latter has intermediate properties between a totally random G^{random} and a totally ordered $G^{ordered}$ network.

links are statistical estimates of possibly existing anatomical pathways or functional interactions between different brain areas and they are affected by some uncertainty (**Fig.1b**).

Statistical modeling has been introduced in an effort to identify the probability distribution of all the possible network realizations associated with an observed one (**Fig.1c**). Approaches of this kind include generating *a-posteriori* network surrogates preserving some properties (*e.g.*, node degree distribution [11]), or *a-priori* reproducing, for example, the underlying modular structure of networks via stochastic block modeling [12, 13]. These approaches have effectively improved the characterization of brain networks with respect to random null-models and provided a statistical framework to estimate confidence intervals. Nevertheless, they remain quite limited in terms of the number of local network properties that can be simultaneously tested and modeled.

This is a crucial aspect as in general we don't know how many and which are the local connection properties that have generated the global observed structure. Exponential random graph models (ERGMs) represent an intriguing solution to this limitation, as demonstrated by the recent development in the field. Broadly speaking, ERGMs belong to the family of maximum entropy (Max.Ent.) models [14, 15]. Given an observed network G^* , the ERGM probability distribution is the one with the largest

information entropy that statistically reproduces some key structural properties, or features, of G^* . The latter are associated with the model parameters, which can be inferred from the data.

A fortunate alignment among availability of experimental brain data, development of network theory and of powerful computational tools - mainly borrowed from social sciences [16] - has allowed researchers to study the brain network structure and dynamics under the lens of the ERGMs. The goal of the present report is to track down these efforts and present them in a self-contained review with an unified accessible language. The ultimate ambition is to show their usefulness to a broad audience and encourage their exploitation to tackle current open questions in basic and clinical neuroscience.

The rest of the review is organized as follows. In (Sec.2) we introduce brain networks, briefly describe how to build them from experimental data and which are their main topological properties. In (Sec.3) we focus on a minimal version of the ERGM, we present both its theoretical and computational aspects, we linger on the interpretation of its parameters and introduce some widely used ERGM graph statistics. In (Sec.4) we review the existing literature of ERGM-based methods for neuroscience. We show how they have been used to characterize brain states, discriminate between different experimental conditions and model dynamic brain networks. In (Sec.5) the reader will find a discussion of the theoretical subtleties of the ERGMs with a focus on their applications to the brain. Finally, in (Sec.6) we conclude with a perspective on the future of the approach, and point out potentially useful alternatives.

2. Brain networks in a nutshell

The study of the brain from a network perspective is nowadays regarded as a vibrant interdisciplinary field of research, referred as network neuroscience [4]. Such efforts have initially focused on quantifying the organizational properties of the nervous system by means of different network metrics [17, 2]. For example, it has been shown that brain networks tend to exhibit a pronounced modular structure *i.e.* they are partitioned in densely inter-connected communities linked by sparse inter-community connections [18]. The high clustering together with the low average path length between pairs of nodes has suggested that neuronal interactions are organized in a small-world topology, optimizing the balance between segregation and integration of information [19, 20]. Moreover, there is mounting evidence on the presence of hubs [21, 22] and metabolic constraints of the brain's wiring, with spatially closer regions supporting stronger patterns of connections [8, 23]. These brain network properties tend to span different spatial and temporal scales, are in common to several species, correlate with individual behavior and can be significantly affected by neurological disorders [24, 25]. In the following, we will focus on brain networks derived from neuroimaging data resulting from the noninvasive recording of large-scale neural information in humans.

2.1. Building brain networks from experimental data

Neuroimaging broadly refers to a large set of techniques that allow to measure noninvasively the brain structure and function. The translation of these measurements into a network hinges on the definition of what nodes and edges are. In general, these

definitions depend on experimental constraints as well as on the specific scientific question. Here, we will provide a gentle overview of the main approaches adopted so far in an effort to offer a minimal understanding of the nature of the networks that we will later explore. Should the readers be interested in more detailed technical descriptions, we refer them to some recent books and reviews [10, 26, 27, 28].

A first type of brain networks aims at representing the structural wiring of the brain. Nodes are defined as different brain sites and edges represent their anatomical interconnections. To this purpose, diffusion tensor imaging (DTI) is a noninvasive technology that allows to infer the diffusion of water molecules through white matter tracts of the brain in the three-dimensional space [29]. The collected data are discretized into non-overlapping gray matter volumes (nodes) and the connection weights (edges) between two nodes are proportional to the anatomical properties of the fiber tracks (*e.g.* number, length).

The second type of brain networks aims instead at capturing the functional interactions between brain regions (nodes). These regions of interest (ROIs) typically correspond to cytoarchitectural landmarks known a-priori or more directly to sensors used for the measurements [30]. For functional networks, interactions (edges) are defined as statistical similarities between the dynamics taking place in different nodes. Many technologies exist to record nodal dynamics as signals with different temporal and spatial resolutions. Among others, electroencephalography (EEG) and functional magnetic resonance imaging (fMRI) allow for noninvasive recordings of electrical activity through sensors placed over the scalp [31] and blood oxygen levels from 3D images of the brain, respectively [32, 33]. Functional interactions are then inferred from the recorded signals using related measures such as Pearson or Spearman correlations, mutual information or Granger causality [10, 30]. Thus, the study of functional networks is primarily intended to reveal the brain organizational properties from an information processing point of view.

At this stage, both structural and functional brain networks are characterized by an adjacency matrix A whose generic entry a_{ij} is a scalar number, typically normalized between 0 and 1, representing the magnitude of the connection between nodes i and j . Thresholding procedures can be eventually adopted to filter out irrelevant links and obtain sparse unweighted networks. More details on why and how to filter information in complex brain networks can be found in [34, 35, 36].

2.2. The rationale of a statistical approach

The rationale for statistical modeling is based on the assumption that the observed system is a realization of a stochastic process, and that by characterizing the properties of this process we can gain insight into the underlying mechanisms governing the system's behavior. More specifically, the probability of observing each possible realization depends on endogenous system constraints, which can be either postulated or inferred from the data.

This type of statistical uncertainty is different from that associated with exogenous experimental errors - here, the network construction process. Statistical uncertainty is due to the inherent variability of a system and is quantified using statistical measures. Uncertainty due to experimental errors arises from factors such as measurement noise and reflects the limitations of our measurement apparatus. In this review we deal with the former, discussion on the latter can be found elsewhere [10].

In network neuroscience, a common first approach to assess the statistical significance of a network’s properties has been to compare it *a-posteriori* to a null model. The latter is a simplified version of the network that preserves some basic structural features, but removes other non-random ones [37]. By comparing a network to its null model, we can determine whether its observed properties are due to chance, or whether they reflect underlying mechanisms or constraints on the system. This comparison can be done using a variety of statistical tests, such as comparing the network’s properties to those of a large ensemble of randomized networks, or comparing it to a null model that is specifically designed to test a particular hypothesis.

Such an approach however do not allow to test more specific hypotheses about the mechanisms that govern the network organization, *e.g.* the likelihood that *a-priori* defined structures are at the origin of the observed network. This kind of questions can be targeted by using instead an inferential approach. This involves (i) constructing a statistical model that describes the relationships between different variables in the system, and (ii) using this model to estimate the parameters that govern the behavior of the system. By comparing the predictions of the model to the observed data, we can determine whether the model accurately captures the underlying mechanisms that drive the system.

Furthermore complex systems typically arise from multiple organizing principles across different scales. Statistical (inferential) frameworks that can simultaneously account for these effects are crucial to understanding the global properties of the system. In network neuroscience, where the focus is on understanding the brain’s structure and function, this approach is particularly useful [24, 38, 39]. By using inferential methods, we can identify the combinations of relevant network properties that give rise to the observed large-scale structures and predict how these properties will change under different conditions or perturbations.

Recent endeavors in exponential random graph models (ERGMs) have increasingly attracted the interest of the network science and neuroscience community due to their ability to match the aforementioned desiderata within a coherent and unifying framework. Hence, it is timely to discuss these emerging developments, and to seek to put them together into a common theoretical ground that can be used to tackle current open questions in modern neuroscience.

3. Exponential random graph model (ERGM)

The interest in the exponential class of probability distributions dates back to the dawn of modern statistics [40, 41] and was originally motivated by its properties with respect to sufficient statistics, as introduced by R.A. Fisher in the early ’20s [42, 43]. A number of theoretical and computational properties makes them very appealing for the problem of statistical inference [44, 45] and justifies their ubiquity in all branches of scientific research. In graph theory, the exponential family, here referred to as exponential random graph models (ERGMs), came on stage in the early ’80s [46, 47, 48] building on the seminal work of J. Besag [49] and further developing since then [50, 51]. Analogies with well-known methods of statistical mechanics have been recently explored in simple cases in conjunction with the outbreak of network science [52, 53, 15]. ERGMs have recently become popular in several scientific fields *e.g.* epidemiology [54], sociology [55, 56, 57] as well as neuroscience. As a consequence an increasing number of publicly available related softwares has appeared, such as the R language packages collected in the `statnet` suite [58, 59].

Our reference implementation throughout this section is the `ergm` package [60].

Here, we give the theoretical minimum to familiarize the reader with ERGMs and their estimation methods. Notably, we will consider unweighted and undirected graphs as they represent the most common scenario in the current literature. Later in Sec.4 we will extend the discussion to more sophisticated cases, when appropriate.

Let \mathcal{G} be a set of all finite graphs of N nodes with no self-loops, containing at most one single edge between two nodes. Each graph $G \in \mathcal{G}$ can be equivalently represented by a $N \times N$ adjacency matrix A containing Boolean values, *i.e.* $a_{ij} = \{1, 0\}$.

At the hearth of the ERGM approach there is the definition of an appropriate probability mass function $P(G)$ over the ensemble \mathcal{G} . The crucial idea of ERGMs is to encode the *relevant* information about the graph G in a vector $\mathbf{x}(G) \in \mathbb{R}^r$ of r statistics or metrics, where each element $x_\alpha(G)$ measures a different network property of interest.

Specifically, the probability of observing a generic graph G reads as

$$P(G|\boldsymbol{\theta}) = \frac{e^{\boldsymbol{\theta} \cdot \mathbf{x}(G)}}{\sum_{\tilde{G} \in \mathcal{G}} e^{\boldsymbol{\theta} \cdot \mathbf{x}(\tilde{G})}}, \quad (1)$$

where $\boldsymbol{\theta} \in \mathbb{R}^r$ are the model parameters weighting the graph statistics and \cdot is the dot product. A formal analogy with the canonical Boltzmann distribution is readily recognized by defining the Hamiltonian

$$\mathcal{H}(G) = -\boldsymbol{\theta} \cdot \mathbf{x}(G) = -\sum_{\alpha} \theta_{\alpha} x_{\alpha}(G) \quad (2)$$

Hence, the denominator of Eq.1 is a normalizing constant, which turns out to play the role of partition function \mathcal{Z} and we can rewrite $P(G|\boldsymbol{\theta}) = \frac{1}{\mathcal{Z}} e^{-\mathcal{H}(G)}$. Given a specific observed network G^* , one can infer the model parameters $\boldsymbol{\theta}^*$ so that the expected value of each graph statistics over the ensemble \mathcal{G}

$$\langle \mathbf{x} \rangle = \sum_{G \in \mathcal{G}} \mathbf{x}(G) P(G|\boldsymbol{\theta}^*) \quad (3)$$

statistically matches the observed value $\simeq \mathbf{x}(G^*)$.

Note that (Eq.3) sets the number of model parameters from $2^{\binom{N}{2}}$ - the number of possible edges, in the worst case - to a significantly smaller amount determined by the chosen graph statistics. By consequence, in general a tie-level matching between the networks generated with $P(G|\boldsymbol{\theta}^*)$ and the observed graph G^* is neither expected nor desired.

3.1. Model parameter estimation

Apart from very simple cases, ERGM parameters are hard to obtain in a closed form. In general, the nature and the number of graph statistics make the denominator of Eq.1 impossible to compute analytically. The evaluation of the partition function \mathcal{Z} is a very well-known problem in many situations, like for example when trying to infer parameters from Boltzmann-like distributions [61, 62].

In the last 30 years there has been a tremendous theoretical and computational effort in developing methods to address this issue and achieve efficient estimation algorithms. In the following, we present a brief introduction to the key-idea behind many of these methods [63, 64].

Given a graph G^* and a set of statistics $\mathbf{x}(G)$, we can take a Bayesian perspective and argue that our knowledge on the parameters $\boldsymbol{\theta}$ is better described by the posterior distribution

$$P(\boldsymbol{\theta}|G^*) = \frac{P(G^*, \boldsymbol{\theta})}{P(G^*)} = \frac{P(G^*|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(G^*)}. \quad (4)$$

In the case where no prior information is available for $\boldsymbol{\theta}$, we can take $P(\boldsymbol{\theta})$ to be uniformly distributed. From Eq.4 we see that the posterior distribution $P(\boldsymbol{\theta}|G^*)$ is directly proportional to the distribution of the data G^* given the parameters $\boldsymbol{\theta}$. Accordingly, our best guess on the parameters value is given by the Maximum Likelihood Estimator (MLE)

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} \mathcal{L}(G^*|\boldsymbol{\theta}), \quad (5)$$

where $\mathcal{L}(G^*|\boldsymbol{\theta}) = \log P(G^*|\boldsymbol{\theta})$ is the log-likelihood function.

The evaluation of (Eq.5) is hampered by the computation of the partition function \mathcal{Z} inside $P(G^*|\boldsymbol{\theta})$. As discussed above, this entails the need for non-exact numerical methods.

Rather than maximizing $\mathcal{L}(G^*|\boldsymbol{\theta})$ directly, the idea is to maximize a shifted log-likelihood $\bar{\mathcal{L}}(G^*|\boldsymbol{\theta}) = \mathcal{L}(G^*|\boldsymbol{\theta}) - \mathcal{L}(G^*|\boldsymbol{\theta}_0)$, where $\boldsymbol{\theta}_0$ is an arbitrarily parameter vector. ‡

A bit of algebra reveals that the new log-likelihood function can be written as

$$\bar{\mathcal{L}}(G^*|\boldsymbol{\theta}) = (\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}(G^*) - \log \langle e^{(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \cdot \mathbf{x}(G)} \rangle_{\boldsymbol{\theta}_0} \quad (6)$$

where the subscript $\langle \cdot \rangle_{\boldsymbol{\theta}_0}$ indicates the expectation value over the distribution $P(G|\boldsymbol{\theta}_0)$. Note that the new log-likelihood $\bar{\mathcal{L}}(G^*|\boldsymbol{\theta})$ is maximized by the same $\boldsymbol{\theta}^*$ that maximizes $\mathcal{L}(G^*|\boldsymbol{\theta})$.

In practice, the evaluation of the partition functions \mathcal{Z} is bypassed and we just need to calculate the expectation $\langle \cdot \rangle_{\boldsymbol{\theta}_0}$, which can be efficiently performed with standard Markov-chain Monte-Carlo (MCMC) approximations. Eventually, the whole procedure is repeated until the convergence to some stable solution $\boldsymbol{\theta}^*$ is reached (**Tab. 1**). In the following, we will refer to this method as MCMC-MLE.

3.1.1. ERGMs in practice The increasing diffusion of ERGMs has been driven by a concurrent technical effort to develop efficient algorithms for carrying out the inference task. By far the most popular implementation available to date is the **statnet** suite of R-based software packages [58, 67]. To our knowledge, no other implementation has reached the same level of maturity and extending these techniques outside the R language is an open challenge for the next future. The reference implementation for the vast majority of existing packages is the **ergm** package [60]. A number of other packages have been built on it, both for static [68, 69, 70, 71] and temporal networks [72, 73].

An illustrative workflow based on the **ergm** package is implemented in [74], consisting of six fundamental steps. They include: i) creating a network object using the **network** package from the adjacency matrix and nodal/edge covariates (if any), ii) computing ERGM statistics for the experimental network (optional), iii) specifying

‡ In practice $\boldsymbol{\theta}_0$ should be close enough to $\boldsymbol{\theta}$ to ensure convergence in a reasonable computational time. One popular choice is to take $\boldsymbol{\theta}_0$ as the maximum pseudo-likelihood estimation (MPLE) of $\boldsymbol{\theta}$ under the additional hypothesis that the edges are mutually independent [16, 65].

```

 $\theta^{(0)} = \theta^{PL}$ 
 $\tau = 0$ 
while (conv = FALSE) do
   $\tau += 1$ 
  generate  $n$  graphs  $G_k \sim P(G|\theta^{(\tau-1)})$  by MCMC
   $\theta^{(\tau)} = \arg \max_{\theta} \left\{ (\theta - \theta^{(\tau-1)}) \cdot \mathbf{x}(G^*) - \log\left(\frac{1}{n} \sum_{k=1}^n e^{(\theta - \theta^{(\tau-1)}) \cdot \mathbf{x}(G_k)}\right) \right\}$ 
  if (convergence criterion) then
    conv = TRUE
  end if
end while

```

Table 1: MCMC-MLE pseudo-code. The initial choice of parameters $\theta^{(0)}$ is arbitrary, here the pseudo-likelihood value θ^{PL} . Several convergence criteria are possible, from the simplest $\|\theta^{(\tau)} - \theta^{(\tau-1)}\| < \epsilon$ for some $\epsilon > 0$, up to more refined choices [16, 66].

and estimating the model, iv) assessing MCMC convergence (optional), v) evaluating the model fit using goodness of fit methods (GoF), and vi) simulating the model using estimated parameters (optional).

The computational time required for the model estimation might vary from a few seconds to several hours, depending on a number of factors including the type and number of graph statistics, the basic properties of the networks (*e.g.* size, sparsity) and several estimation settings that can be possibly adjusted. Because of the intrinsic characteristics of the current neuroimaging technology (Sec.2), brain networks of interest in this review have a relatively small number of nodes (at most hundreds). For such sizes, the ERGM estimation is in general doable. For further information regarding the computational aspects and a comprehensive overview of the latest technical enhancements of the `ergm` 4.0 package, we refer to [75, 66] and references therein.

3.2. Interpretation of ERGMs

A crucial, often overlooked, aspect in ERGM is the interpretation of the parameters θ in Eq.1. Let us address the question starting with a simple example [47, 76, 77]. Consider an ERGM based on a single graph metric $x(G)$ corresponding to the number of edges in the network, *i.e.* $x(G) = \sum_{i<j} a_{ij}$. We first consider the *forward* problem in which there is no observed graph but instead the parameter θ is given and the goal is to evaluate the expected value for $x(G)$.

In such a simple case, the mass probability function can be computed exactly:

$$P(G|\theta) = \frac{e^{\theta x(G)}}{\prod_{i<j} \sum_{a_{ij}=0}^1 e^{\theta a_{ij}}} = \frac{e^{\theta x(G)}}{(1 + e^{\theta})^{\binom{N}{2}}}. \quad (7)$$

Notably, by defining

$$p(\theta) = \frac{e^{\theta}}{1 + e^{\theta}} \quad (8)$$

we can rewrite $P(G|\theta) = p^{x(G)}(1 - p)^{\binom{N}{2} - x(G)}$, which is the probability of a Bernoulli graph with N nodes and $x(G)$ links [78]. The expected value for the number of edges in the graph is then simply $\langle x \rangle = \binom{N}{2} p$.

In practice, the most common scenario is quite the opposite. One does not know the value of θ but rather has an observed graph G^* with, say, M edges. In these situations, the goal is to solve the *inverse* problem by inferring θ^* from $x(G^*) = M$. Taking $p^* = M/\binom{N}{2}$ and inverting Eq.8:

$$\theta^*(p^*) = \log \frac{p^*}{1-p^*} \quad (9)$$

so that, consistently, $\langle x \rangle = \binom{N}{2}p = M$. Notably, Eq.9 tells us that if we start from a Bernoulli graph with $p^* = 0.5$, then $\theta = 0$. This is the maximally random case since each dyad corresponds to the toss of a fair coin. But if we start from a denser ($p^* > 0.5$) or sparser ($p^* < 0.5$) graph, then $\theta^* > 0$ or $\theta^* < 0$, respectively.

In such a simple case we have therefore a complete understanding of both the *forward* and *inverse* problems. In real situations, however, the interpretation of the parameters becomes trickier since ERGMs may include several statistics, often exhibiting some degree of dependence. As a consequence, there might be interactions among the model parameters that are in general difficult to compute analytically.

Yet, inference is still possible via approximate numerical methods such as the MCMC-MLE. The interpretation of the inferred parameters can be given by quantifying their effect on the likelihood that an edge between two nodes exists or not in the network. For example, consider a single edge toggle between nodes i and j . Let us call $P(G_{a_{ij}=1, \mathbf{a}_{\setminus ij}} | \boldsymbol{\theta}^*)$ the probability of having an edge between the two nodes given the rest of the graph $\mathbf{a}_{\setminus ij}$, and $P(G_{a_{ij}=0, \mathbf{a}_{\setminus ij}} | \boldsymbol{\theta}^*)$ its complementary.

From Eq.1 one easily finds that

$$\frac{P(G_{a_{ij}=1, \mathbf{a}_{\setminus ij}} | \boldsymbol{\theta}^*)}{P(G_{a_{ij}=0, \mathbf{a}_{\setminus ij}} | \boldsymbol{\theta}^*)} = \exp \left(\boldsymbol{\theta}^* \cdot \Delta_{ij}^x(G) \right), \quad (10)$$

where the $\Delta_{ij}^x(G) = \mathbf{x}(G_{a_{ij}=1, \mathbf{a}_{\setminus ij}}) - \mathbf{x}(G_{a_{ij}=0, \mathbf{a}_{\setminus ij}})$ are the so-called *change statistics*. From Eq. 10 is clear that the probability of a link to exist depends on the magnitude and signs of both $\boldsymbol{\theta}^*$ and $\Delta_{ij}^x(G)$.

Let us consider a very simple case where the presence of a new link between i and j increases the value of a given graph statistic *i.e.* $\Delta_{ij}^x(G) > 0$. If $\theta^* > 0$, then this change statistic favors the probability of this edge to exist in the graph ($a_{ij} = 1$). Instead, if $\theta^* < 0$ the same change statistic penalizes the existence of the connection, favouring instead ($a_{ij} = 0$). Overall, the magnitude of this effect is given by the term $\exp[\theta^* \Delta_{ij}^x(G)]$. In other words, one can disentangle and quantify the relative influence of each graph statistic on the probability to observe, or not, an edge between two nodes i and j .

Note that the above micro-level interpretation of the parameters is precisely at the core of the MCMC routines used to sample networks from ERGM probability distributions. Markov chains of graphs whose stationary distribution is $P(G|\boldsymbol{\theta}^*)$ are obtained by sequentially proposing changes on a graph, evaluating the probability of the new configuration by Eq.10 and accepting it according to the Metropolis-Hastings recipe. We refer to recent reviews in the field for a more detailed description of the MCMC methods used in ERGMs [16].

3.3. Graph statistics for ERGMs

A broad spectrum of graph statistics, or metrics, $x(G)$ have been formulated and included in Eq.1. Here, we present those that have been more frequently adopted in

neuroscience.

A first group of graph statistics includes the so-called dyadic independence terms, *i.e.* combinations of only single-dyad terms. The most intuitive metric is the **edges** term which measures the number of edges in the graph (**Fig. 2a**):

$$x_e(G) = \sum_{i < j} a_{ij} . \quad (11)$$

A straightforward extension is the so-called **edge covariate**

$$x_{dc}(G) = \sum_{i < j} a_{ij} \zeta_{ij} . \quad (12)$$

where ζ_{ij} is an attribute associated defined for each dyad, such as the Euclidean distance in a spatial network.

ERGM terms based on nodal attributes also belong to this group. If η_i is a categorical property of the nodes, then it is possible to define the so-called **nodematch** term as

$$x_{nm}(G) = \sum_{i < j} \delta_{\eta_i \eta_j} a_{ij} , \quad (13)$$

where δ is the Kronecker delta and x_{nm} counts the number of edges whose nodes are labeled by the same categorical attribute. For instance, $\eta_i = \text{R, L}$ could indicate whether a node belongs to the right (R) or left (L) hemisphere.

More in general, real networks are characterized by the presence of complex connectivity structures leading to dependencies between dyads. Therefore, we now turn into presenting graph statistics that involve products of two or more dyadic variables a_{ij} . Following the seminal work of O. Frank and D. Strauss on Markov Graphs [47], we assume that only dyads that share a node can be dependent or, equivalently, that nonincident dyads are conditionally independent.

Under this general definition, several statistics may be introduced based on the count of local connection patterns. For instance, the presence of hubs, documented in many real-world networks [79, 77], can be measured in an ERGM by the **k-stars** term:

$$x_{st}^{(k)}(G) = \frac{1}{k!} \sum_{i_0} \cdots \sum_{i_k} a_{i_0 i_1} \cdots a_{i_0 i_k} \quad (14)$$

with $k = 1, \dots, N - 1$.

The global tendency of networks to form clusters, can be instead measured by the presence of triplets, or triads, of connected nodes. A **k-triangles** term can be then defined as:

$$x_{tr}^{(k)}(G) = \frac{1}{\xi(k)} \sum_{i_0} \cdots \sum_{i_{1+k}} a_{i_0 i_1} a_{i_0 i_2} a_{i_1 i_2} \cdots a_{i_0 i_{1+k}} a_{i_1 i_{1+k}} \quad (15)$$

where the symmetry factor is $\xi(k) = 3!$ for $k = 1$ and $\xi(k) = 2k!$ for $k = 2, \dots, N - 2$. The simplest case $x_{tr}^{(1)}(G) = \frac{1}{6} \sum_{i_0 i_1 i_2} a_{i_0 i_1} a_{i_0 i_2} a_{i_1 i_2}$ counts the number of triangles in the graph.




Figure 2: **Illustration of graph statistics for ERGMs.** a) x_e counts the total number of edges in the graph; b) $x_d^{(k)}$ counts the number of nodes with degree k ; c) $x_{nsp}^{(k)}$ is the number of non-connected dyads sharing k common partners; finally, d) $x_{esp}^{(k)}$ counts the connected dyads sharing exactly k common partners.

3.3.1. Degeneracy of simple ERGMs. Starting from the earliest numerical investigations, it has been observed that for some basic ERGM specifications, MCMC simulations hardly "mix", meaning they do not reach a stationary distribution within a reasonable amount of time [48, 80]. The lack of mixing hampers the estimation task. This behavior is due to the geometry of ERGM stationary distributions. As it turns out, there exist choices of statistics that result in "degenerate" models, meaning that assign a significant probability only to a few graph configurations, often with a very different structure, such as empty and fully connected graphs. [81, 82]. A general, thorough discussion of the causes and consequences of degeneracy in the context of ERGMs can be found in [83].

For a few simple ERGMs, such as the 2-star model ($\mathcal{H} = x_e + x_{st}^{(2)}$) or the Strauss clustering model ($\mathcal{H} = x_e + x_{tr}^{(1)}$), it has been demonstrated through analytic solutions that this phenomenon corresponds to a standard phase separation observed in statistical mechanics between a low- and high-density phase, representing almost empty and full networks, respectively [84, 85]. For such model specifications, there is no parameters combination that is able to reproduce intermediate densities of edges, two-stars and triangles as observed in real-world networks.

These examples demonstrate that degeneracy issues arise when the model includes unstable sufficient statistics \S . Such models show excessive sensitivity in the sense that

\S According to the definition in [83], $x(G)$ is stable if there exist $C > 0$ and $M_c > 0$ such that $U_x \equiv \max_{G \in \mathcal{G}} x(G) \leq CM \forall M > M_c$, where $M = \binom{N}{2}$ is the number of degrees of freedom in the network. Note that for instance $U_{x_e} = M$ implies that the edges term is stable while $U_{x_{st}^{(2)}} \sim NM$

small state changes can result in extremely large odds ratios eq.(10) [83].

The main conclusion of these theoretical investigations is that inference for simple ERGMs is simply not possible, ill-posed. To solve degeneracy issues, a possible solution is to add structure to the model specification by incorporating additional information such as i) edge-specific or node-specific covariates, ii) block structure, iii) temporal structure, iv) multilevel structure, or v) spatial structure. A recent comprehensive discussion of these cases can be found in [86].

However, the most common approach to ERGM degeneracy issues in neuroscience applications has been so far to include the so-called "curved statistics" in the model specification. In the next section, we will discuss this approach in detail.

3.3.2. Curved Statistics. The inclusion of curved statistics into ERGM model specifications has been found to be effective in alleviating degenerate behaviors [83, 86]. These statistics are linear combinations of the complete distribution of degree or shared partner statistics, where the coefficients are geometrically weighted. Therefore, these statistics not only enumerate patterns but also include additional non-linear constraints on graph structures through the use of geometric coefficients. [16, 87].

For instance, the geometrically weighted degree (**gwd**) term reads:

$$x_{gwd}(G|\tau) = e^\tau \sum_{k=1}^{N-1} \left\{ 1 - (1 - e^{-\tau})^k \right\} x_d^{(k)}(G) , \quad (16)$$

where $\tau > 0$ is a parameter and $x_d^{(k)}(G)$ is the number of nodes in the graph G whose degree is exactly equal to k (**Fig.2b**). Similarly, the geometrically weighted non-edgewise shared partner (**gwnsp**) statistic:

$$x_{gwnsp}(G|\tau) = e^\tau \sum_{k=1}^{N-2} \left\{ 1 - (1 - e^{-\tau})^k \right\} x_{nsp}^{(k)}(G) , \quad (17)$$

where $\tau > 0$ and $x_{nsp}^{(k)}(G)$ is the number of non connected dyads having k neighbors in common, (**Fig.2c**). Finally, the geometrically weighted edgewise shared partner (**gwesp**) statistics:

$$x_{gwesp}(G|\tau) = e^\tau \sum_{k=1}^{N-2} \left\{ 1 - (1 - e^{-\tau})^k \right\} x_{esp}^{(k)}(G) , \quad (18)$$

where $\tau > 0$ and $x_{esp}^{(k)}(G)$ is now the number of connected dyads that share exactly k neighbors (**Fig.2d**).

To provide an interpretation of curved statistics and better understand the underlying modeling assumptions, let's examine the case of **gwesp** and proceed as in (Sec.3.2). Namely, let G be a graph and consider two connected nodes i, j with k shared partners. Let G' be the same graph as G but with an additional common neighbor between the nodes i, j , this implies

$$\begin{aligned} x_{esp}^{(k)}(G') &= x_{esp}^{(k)}(G) - 1 \\ x_{esp}^{(k+1)}(G') &= x_{esp}^{(k+1)}(G) + 1 , \end{aligned} \quad (19)$$

implies that the 2-star term is unstable.

where $x_{esp}^{(k)}(G)$ is the number of connected dyads that share exactly k neighbors. For an ERGM including a **gwerp** term, we find

$$\begin{aligned} \frac{P(G'|\boldsymbol{\theta})}{P(G|\boldsymbol{\theta})} &\propto \exp \left[\theta_{gwerp} \left(x_{gwerp}(G'|\tau) - x_{gwerp}(G|\tau) \right) \right] \\ &\propto \exp \left[\theta_{gwerp} (1 - e^{-\tau})^k \right] \equiv \rho . \end{aligned} \quad (20)$$

We start by noting that since $\tau > 0$, we have $1 - e^{-\tau} \in (0, 1)$. Moreover:

- (i) the scaling factor $\theta_{gwerp} > 0$ implies $\rho > 1$, hence preference for adding shared partners; $\theta_{gwerp} < 0$ on the contrary implies a preference for deleting shared partners. This interpretation is akin to that of all standard ERGM parameters as discussed in (Sec.3.2).
- (ii) $\lim_{k \rightarrow \infty} \rho = 1$ *i.e.* adding/removing shared partners to pairs of connected nodes with already many of them ($k \gg 1$) has little effect on the ratio $P(G')/P(G)$. This implies an effective cutoff for the order k of $x_{esp}^{(k)}$ statistic that is relevant for the system.
- (iii) The speed of the geometric decay is controlled by an external parameter τ , small values imply a rapid decay while high values imply a slow decay.

Hence, the **gwerp** statistic is a refined version of the simple sum of triangles. When the parameter θ_{gwerp} is positive, the general tendency for clustering will not push the system into a fully connected state. In fact, its implementation through the **gwerp** statistic implies a strong advantage for the addition of shared partners for poorly connected nodes, while almost no effect for nodes that already have many connections. This is indeed a more realistic behavior of a triadic closure, which in the case of brain networks it is expected to be constrained by the inner physical and functional resources [9].

The same reasoning applies to other curved statistics, such as **gwd** and **gwnsp**. These two statistics are alternatively used to test the hypothesis that shortest paths results from the presence of hubs or two-paths. The former would affect the node degree distribution, the latter the distribution of non-edgewise shared partners. Thus, via the **gwd** and **gwnsp** statistics, the presence of short paths in the network will result in positive θ_{gwd} and θ_{gwnsp} values, without the system ending up in a fully connected/empty state.

4. ERGMs in neuroscience

ERGMs have started to be exploited as statistical framework to study brain networks since 2010 [34]. The typical workflow of an ERGM analysis in neuroscience is summarized in **Fig.3**. In this section, we present an overview of the most recent results obtained in the last decade. A schematic summary of the existing literature can be found in **Tab.2**.

While ERGMs are mathematical abstractions and can be applied to any network, most research has so far focused on functional brain connectivity derived from **fMRI** and **EEG** data. In the following, we structure the presentation of the current literature around three main scientific questions: *i*) what are the basic network mechanisms of healthy brain functioning (Sec.4.1), *ii*) how to statistically compare brain networks between different states (Sec.4.2), and *iii*) how to explicitly characterize the temporal network evolution following a pathological event (Sec.4.3).

Reference	Data (nodes)	Method/Package
[88] <i>Simpson et al. (2011)</i>	fMRI(90)	ERGM / <code>ergm</code> [60]
[89] <i>Simpson et al. (2012)</i>	fMRI(90)	ERGM / <code>ergm</code> [60]
[90] <i>Sinke et al. (2016)</i>	DTI(96)	BERGM / <code>Bergm</code> [68]
[91] <i>Obando et al. (2017)</i>	EEG(56)	ERGM / <code>ergm</code> [60]
[92] <i>Stillman et al. (2017)</i>	fMRI(20)	cGERGM/ <code>gergm</code> [69]
[93] <i>Dell'Italia et al. (2018)</i>	fMRI(148)	sTERGM / <code>stergm</code> [72]
[94] <i>Stillman et al. (2019)</i>	fMRI(9 – 25)	cGERGM/ <code>gergm</code> [69]
[95] <i>Lehmann et al. (2021)</i>	fMRI(90)	BERGM / <code>Bergm</code> [68]
[96] <i>Obando et al. (2022)</i>	fMRI(81.13*)	TERGM / <code>btergm</code> [73]

Table 2: **Recent ERGM studies of human brain networks.** We summarize each paper by its reference, the kind of experimental method for the data and the average number of nodes of the resulting networks, ERGM-method employed and its R implementation.

4.1. Minimal model of brain networks

At the core of many ERGM implementations for neuroscience there is the model first proposed by *Simpson et al.* in two consecutive works [88, 89], where the \mathcal{H} -function reads as

$$-\mathcal{H}(G|\boldsymbol{\theta}) = \theta_1 x_e(G) + \theta_2 x_{gwesp}(G|\tau) + \theta_3 x_{gwnsp}(G|\nu) \quad (21)$$

The **edges** term $x_e(G)$ characterizes the density of the network. This statistic is typically used to ensure that ERGMs also reproduce the actual number of connections in the observed network. The **gwersp** term $x_{gwersp}(G|\tau)$ is meant to capture the overall network clustering, a property that is crucial for the segregation of information in the brain [19, 20]. The **gwnsp** term $x_{gwnsp}(G|\tau)$ is related to the presence of two paths in the network and can be regarded as a measure of integration of information in the brain [98]. Hence, Eq.21 defines a parsimonious model of the fundamental properties of brain networks, *i.e.* connectedness, clustering, and global-efficiency.

In [88, 89] the authors applied this model to fMRI brain networks from 10 healthy subjects, fitted separately per each individual. The same model was also used in [95] to characterize fMRI brain networks from 200 healthy individuals. All these works agreed on assigning negative values to the parameters associated with **edges** and **gwnsp** statistics and positive values to **gwersp** statistics, while the decay parameter for the last two cases was fixed to $\tau = \nu = 0.75$ [88]. The **edges** parameter $\theta_1 < 0$ confirms the tendency of the network to be sparse. The **gwnsp** parameter $\theta_3 < 0$ indicated that any tendency for global integration of information in the brain is unlikely to be supported by shortest paths with length equal to 2. The **gwersp** parameter $\theta_2 > 0$ showed that brain networks are statistically inclined to form clusters, which is a basic hallmark of brain segregation of information.

In subsequent studies, [92, 94] investigated whether and how the aforementioned




Figure 3: **ERGMs workflow for brain networks.** a) A graph G^* is given, for instance a structural or functional brain network. b) The information stored in the graph G^* is condensed in a set of statistics $\mathbf{x}(\cdot) \in \mathbb{R}^r$, whose choice is up to the modeller. The statistics evaluated on G^* represent the input of the ERGM inference methods. c) The goal of an ERGM is the estimation of the set of parameters $\boldsymbol{\theta}^* \in \mathbb{R}^r$ associated with $\mathbf{x}(\cdot)$. d) The inferred parameters can be used to simulate n new brain networks G_1, \dots, G_n that statistically reproduce the observed network i.e. $\frac{1}{n} \sum_{i=1}^n \mathbf{x}(G_i) \sim \mathbf{x}(G^*)$. Goodness of fit (GoF) methods can be used to assess the performance of the estimated model in reproducing other features of the input network [97].

whole-brain statistical network properties were also present in specific subsystems associated with basic behavioral functions. Resting-state fMRI networks were then constructed and fitted separately in 21 healthy subjects and for 8 different subsystems, namely auditory, subcortical, dorsal attention, ventral attention, salience, cingulo-opercular task control, fronto-parietal task control, default mode [99] (**Fig.4**).

To carry out this analysis, the authors introduced the so-called correlation generalized ERGM (cGERGM). This formulation of the model is still based on the ERGM in Eq.21, with a slightly different geometrical weighting strategy of the graph statistics [69, 94]. However, cGERGM differs from the standard approach in two substantial ways. First, cGERGM was designed to handle full weighted correlation matrices, instead of dealing with unweighted sparse networks. Secondly, additional geometric information on the nodes were modeled externally to the ERGM by using a beta regression of the group-averaged link weights on the graph statistics [92]. In




Figure 4: **Differences in ERGM parameters across brain subsystems.** Data from [94]. For each subsystem and parameter, the mean value and st.dev. are reported. Asterisks indicate that the estimation was significantly different from 0 for all of the 21 subjects involved in the study. In general, larger networks (cingulo-opercular, salience, default mode, fronto-parietal) yield more robust estimations *i.e.* significant for a larger number of subjects. The corresponding brain subsystems are shown along the x-axis for illustrative purposes, as defined in [100].

particular, two kinds of geometric information were included: a `nodematch` term (Eq.13) to measure the tendency of connected nodes to lie in the same hemisphere, and a `edge covariate` term (Eq.12) to measure the total Euclidean distance between all connected nodes. To emphasize the difference from the ERGM parameters, we shall call β_4, β_5 the ones controlling for these effects.

Across all sub-systems, and almost all subjects, authors found consistent results with parameter values statistically different from 0 especially for the larger subsystems, *i.e.* cingulo-opercular, salience, default mode, fronto-parietal. The sign of the $\theta_2 > 0$ and $\theta_3 < 0$ parameters associated respectively with `gwesp` and `gwnsp` statistics, were in line with those reported for the entire brain network. Instead, the `edge` statistics was associated with a positive parameter $\theta_1 > 0$. This result is consistent with the fact that brain subsystems are more densely connected internally than between each other [101, 18].

The estimation of the regression parameters associated with the spatial geometry of the brain structural networks turn out to be much more inconsistent across subsystems and subjects. The significantly positive hemispheric parameter ($\beta_4 > 0$) indicates that the nodes in the fronto-parietal and default-mode subnetwork are more likely to be connected within hemispheres. Instead, there is a general weak tendency of all subnetworks, but the auditory one, to penalize long distant connections ($\beta_5 < 0$, albeit significant for only between 20% and 62% of the individuals).

The overall emerging picture is that of a modular brain network reflecting a significant clustering behavior together with a strong intrahemispheric connectivity. Note that the negative `gwnsp` parameters do not necessarily imply the absence of

2-paths in the networks. Instead, they indicate that short paths of length 2 are relatively underrepresented with respect to what could be expected when considering other statistics alone *e.g.*, triangles.

A related question is to what extent integration of information is statistically supported by short paths in the brain network. From a biological perspective, it would be more plausible that information flows through hubs and not via short paths, although both mechanisms are not mutually exclusive [102]. To address this question, [91] considered EEG resting-state networks in a group of 108 healthy subjects. In particular, they compared two different ERGMs, one including `gwesp` and `gwnsp` statistics, and the other including `gwesp` and `gwd` statistics, that is

$$-\mathcal{H}(G|\boldsymbol{\theta}) = \theta_2 x_{gwesp}(G|\tau) + \theta_3 x_{gwd}(G|\nu); \quad (22)$$

while holding fixed the number of edges in the network $x_e(G) \equiv x_e(G^*)$ and $\tau = \nu = 0.75$.

Results showed that the second ERGM was able to better reproduce brain network properties not included in the model such as modularity, global- and local-efficiency. Notably, both θ_2 and θ_3 values were in average positive and significant confirming their statistical relevance for brain networks. Taken together, these findings support the hypothesis that while information segregation emerges from triangles, information integration could be actually mediated by hubs rather than shortest paths.

4.1.1. Group-representative networks. Thus far, ERGMs have been used to characterize brain networks obtained from different individuals, separately. The resulting parameters distribution from many subjects can be then used to assess statistical properties at the group-level via standard tools such as hypothesis testing and regression analysis. Alternatively, assuming that brain networks from different individuals are stochastic realizations of the *same* system, they can be simultaneously modeled so as to obtain a unified group description. This is the rationale behind the construction of a group representative network (GRN) *i.e.* a network that summarizes the statistical properties of a given ensemble (**Fig.5**).

In network neuroscience, this question has been mainly addressed by aggregating the connectivity matrices of different individuals to obtain a group-average (*mean-GRN*) or group-median (*median-GRN*) representative network [22, 103]. Alternatively, methods to assess the statistical significance of the pooled connection values for each pair of nodes can be adopted [104]. Despite being computationally straightforward, the major limitation of these methods is that they treat edges independently and cannot directly inform on complex connection properties involving more than 2 nodes. An interesting possibility is therefore to incorporate the inter-subject variability directly into the ERGM formulation.

In [89], the authors proposed to generate a GRN starting from the ERGMs in (Eq.21) fitted separately on different subjects. Then, a group representative network is simulated via a new ERGM with the \mathcal{H} -function

$$-\mathcal{H}(G|\boldsymbol{\theta}) = \mu_1 x_e(G) + \mu_2 x_{gwesp}(G|\tau) + \mu_3 x_{gwnsp}(G|\nu) \quad (23)$$

where $\tau = \nu = 0.75$ are fixed and the $\boldsymbol{\mu}$ are the group-averaged values of the individual model parameters. This procedure better reproduced several topological properties of the actual brain networks, as compared to standard ERGMs fitted on the group-averaged network. Note that this is a frequentist approach, the estimation of the

group level parameters are point-like averages. We refer to this GRN construction scheme as a *mean-ERGM*.

The idea of summarizing individual estimates in a vector of hyper-parameters was pushed further in [95]. Accordingly, the model parameter vector for each individual is a realization of a normal probability distribution $\theta \sim \mathcal{N}(\mu, \Sigma_\theta)$. The hyper-parameters (μ, Σ_θ) are estimated by extending to a multilevel hierarchical setting a Bayesian formulation of the ERGM (BERGM), developed in [68]. A GRN can then be obtained by sampling a network from the posterior distribution of parameters. This method is named *multi-BERGM*.




Figure 5: **Comparison between different GRN construction schemes.** Functional (fMRI) brain networks from 100 subjects aged 18 – 33. Gray histograms show the distributions of parameter estimates from a Bayesian ERGM (BERGM) across individuals. Vertical lines correspond to group-level parameters. In red, the *multi-BERGM* (red line). Blue dotted line for *mean-BERGM* (group-averaged BERGM individual estimates). Green dashed line for a single BERGM estimated on the *mean-GRN*; purple dashed line for a BERGM on the *median-GRN*. Both *multi-BERGM* and *mean-BERGM* show greater accuracy than simpler *mean/median-GRN* in capturing the correct distribution of parameters in the population. Illustration from [95].

4.2. Identification of discriminant brain network features

ERGMs can also be used to assess statistical differences between brain networks obtained under different experimental conditions or belonging to different groups of subjects.

In [91] the authors analyzed EEG networks with 56 nodes in 108 healthy subjects under two different conditions: eyes closed (EO) and eyes closed (EC) resting-states. To this end, they used the same ERGM defined in Eq.22. In particular, separate ERGMs were estimated from brain networks extracted from each individual, condition and frequency band *theta* (4 – 7 Hz), *alpha* (8 – 13 Hz), *beta* (14 – 29 Hz), *gamma* (30 – 40 Hz). The estimated parameter values were then compared between the two conditions. Results showed that θ_2 values, associated with the *gwesp* statistics, were significantly higher in EO than EC, for both *alpha* and *beta* bands. This result is consistent with the local increase of EEG activity widely reported in literature for those bands [105, 106]. Notably, the difference in the *beta* band could not be detected by a standard network analysis.

ERGM-based analyses have been also adopted to characterize aging. To this end,

in [95] the authors compared **fMRI** brain networks in a group of young (Y: 18 – 33) and elderly (O: 74 – 89) human subjects [95]. Starting from the hierarchical Bayesian ERGM presented in (Sec.4.1.1), a further hierarchical layer is added to account for the population partitioned in two subgroups, yielding two different hyper-parameters $\phi^{(Y)} = (\mu^{(Y)}, \Sigma_{\theta})$ and $\phi^{(O)} = (\mu^{(O)}, \Sigma_{\theta})$ and same covariance structure is assumed for the two groups.

The results showed that the estimates for the parameter associated with the **gwnsp** statistics were significantly higher in the young group compared to the old one (**Fig.6**). These findings indicated that functional brain networks are prone to significant loss of integration of information across the lifespan, consistently with previous findings reporting a decreasing global-efficiency with age [22].




Figure 6: **Comparison of *multi-BERGM* parameters between young and old individuals.** Two groups of subjects are considered: young 18 – 33 (Y) and old 74 – 89 (O). 100 functional (**fMRI**) brain networks for each group. The young group is characterized by larger values for the **GWNSP** parameter while the edges and **GWESP** posteriors were almost identical between the two. Illustration from [95].

One might then wonder whether similar changes occur in structural brain networks, whose connectivity exhibits much slower changes [107]. *Sinke et al.* [90] studied **DTI** networks from 382 healthy subjects and considered 4 age-groups: 20 – 34, 35 – 50, 51 – 70, > 70 years (**Fig. 7**). Assuming a low variability between the individuals in the same group, authors constructed a *mean-GRN* for each group and used a Bayesian ERGM (**BERGM**) [68].

The \mathcal{H} -function in this case was slightly different from Eq.21 *i.e.*

$$-\mathcal{H}(G|\theta) = \theta_1 x_e(G) + \theta_2 x_{gwersp}(G|\tau) + \theta_3 x_{gwnsp}(G|\nu) + \theta_4 x_{nm}(G), \quad (24)$$

where $x_{nm}(G)$ is a **nodematch** term (Eq.13) measuring in this case the number of connections within the same hemisphere.

The signs of the fitted model parameters ($\theta_1 < 0$, $\theta_2 > 0$, $\theta_3 < 0$) confirmed the general tendencies observed in functional brain networks (Sec.4.1). The parameter θ_4 was instead negative for all groups, suggesting that the hemispheric nodematch is relatively low taking into account all the other included statistics. We stress that this does not necessary imply an absolute low intrahemispheric connectivity.




Figure 7: **Structural brain network changes across lifespan.** Top left: *mean-GRN* for the age range of 20 – 34. Other networks correspond to brain networks obtained in distinct older groups of individuals. Red lines represent "lost connections" and green lines represent "new connections". Brain network density changes across the lifespan with +1.6% (age 35–50), –4.4% (age 51–70) and –8.0% (age > 70). Illustration from [90].

Moreover, differently from previous studies [108, 109], only weak changes were found between the parameter values across age-groups. Future studies will elucidate whether this result is actually reflecting a more stable topology of structural brain networks or a result of the group-averaging of the original brain networks.

4.3. Predicting states in temporal brain networks

So far, ERGMs have not explicitly taken into account the fact that brain networks are intrinsically dynamic and their topology can change, fluctuate and evolve over multiple

time scales. Over short-time scales, brain networks rapidly reconfigure during motor or cognitive tasks to adapt to behavioral efforts or loads [110, 111, 112, 113]. Over longer time scales, brain networks can exhibit significant topological losses, such as in aging or in neurodegeneration [114, 115], as well as tentative restorative patterns in recovery after stroke or traumatic injuries [116].

In these cases, it is crucial to formulate statistical network models that explicitly include time, so as to capture time-varying connection properties that are predictive of the system behaviour [117, 118]. One possibility is to endow the ERGM with a time structure and estimate parameters from a time series, or sequence, of observed graphs G^1, \dots, G^T . One can start by assuming a first-order Markov time dependence, write the probability density function for the entire network sequence as

$$P(G^t, \dots, G^{t-1} | \boldsymbol{\theta}) = \prod_{t=2}^T P(G^t | G^{t-1}, \boldsymbol{\theta}) \quad (25)$$

and generalize Eq.1 in a straightforward fashion:

$$P(G^t | G^{t-1}, \boldsymbol{\theta}) = \frac{e^{\boldsymbol{\theta} \cdot \mathbf{x}(G^t, G^{t-1})}}{\sum_{G^* \in \mathcal{G}} e^{\boldsymbol{\theta} \cdot \mathbf{x}(G^*, G^{t-1})}} \quad (26)$$

where the \mathbf{x} statistics can be now defined over two consecutive graphs in the time-series. We shall refer to this model as temporal ERGM (TERGM) [119].

TERGMs have been recently adopted to model brain network reorganization after stroke *Obando et al.* [96]. The study was carried out on longitudinal **fMRI** resting-state networks measured in 49 patients at 2 weeks, 3 months and 1 year after their first ever unilateral stroke event [120, 116] (**Fig.8a**).

The model (Eq.26) was specified with the following temporal \mathcal{H} function:

$$\begin{aligned} -\mathcal{H}(G^t | G^{t-1}, \boldsymbol{\theta}) &= \theta_1 x_e(G^t) + \theta_2 x_{stab}(G^t, G^{t-1}) + \\ &+ \theta_3 x_{te}(G^t, G^{t-1}) + \theta_4 x_{ttr}(G^t, G^{t-1}) \end{aligned} \quad (27)$$

Besides the standard **edges** term (Eq.11), the three other temporal statistics were defined as follows.

The **stability** term x_{stab} simply counts the number of dyads that do not change in the time step $t-1 \rightarrow t$ and it has a fundamental role for the estimation since it sets the pace of allowed changes in the network sequence. The **temporal-edge** x_{te} term counts the number of new inter-hemispheric edges appearing at time t , while the **temporal-triangle** x_{ttr} term counts the number of triangle closures over time within the lesioned hemispheres (**Fig.8b**).

These two last statistics are entrusted with the goal of mimicking the brain plasticity process *i.e.* the recovery of between-hemisphere integration and within-hemisphere segregation [116, 121]. The TERGM analysis resulted in positive estimates ($\theta_3 > 0$, $\theta_4 > 0$) and significantly higher in stroke patients compared to healthy controls (**Fig.8c**). This suggested that both x_{te} and x_{ttr} are fundamental temporal processes of post-stroke brain reorganization. Notably, the same conclusions could not be reached when considering a static equivalent formulation of Eq.27 [96].

Moreover, the same parameters estimated from the temporal networks in the subacute phase (from 2 weeks to 3 months) significantly predicted the future language




Figure 8: **TERGM analysis of longitudinal brain networks after a stroke event.** a) Dynamic brain network of a stroke patient 2 weeks, 3 months and 1 year after the stroke event. Different colours emphasize the belonging to different brain systems. b) Graphical representation of temporal edges x_{te} and temporal triangles x_{ttr} . c) Statistical comparison between θ_3, θ_4 for stroke patients (red shape) and healthy controls (white shape). Violin plots show the distribution of the values, while box-plots represent median and quartiles. Significant increases in the stroke group are reported for both the parameters considered ($p = 0.045$). d) Correlation of the TERGM coefficients and future language outcome of stroke patients. Both temporal edges and triangles in the subacute phase (2 weeks, 3 months) are significantly associated with the future language score in the chronic phase (1 year) ($p < 0.03$). Illustration adapted from [96].

score (1 year) of the patients, suggesting their potential use as clinical biomarkers (**Fig.8d**).

An alternative approach was adopted in [93] to characterize the recovery from coma after traumatic brain injuries. Authors evaluated longitudinal fMRI brain networks obtained at 11, 18, 25 and 181 days after the trauma. In particular, they used separable temporal ERGMs (sTERGMs) [72], which model separately and independently the tendencies to form and dissolve instances of different graph statistics.

The network at the time-step t is therefore given by the following:

$$G^t = G^+ - (G^{t-1} - G^-) \quad (28)$$

where G^{t-1} is the graph time $t - 1$, while G^+ and G^- are the so-called *formation* and *dissolution* graphs. More in detail, the formation graph is drawn from $\sim P_+(G^t|G^{t-1}, \theta^+)$, where P_+ is the same as in Eq.26 but with the formation statistics \mathbf{x}^+ and parameters θ^+ . In this formation process, ties can only be added to G^{t-1} . Analogously, the dissolution graph is obtained by removing edges from G^{t-1} and the

graph is drawn from $\sim P_-(G^t|G^{t-1}, \boldsymbol{\theta}^-)$, with dissolution statistics \boldsymbol{x}^- and parameters $\boldsymbol{\theta}^-$. The goal of a sTERMG is to estimate both sets of parameters $\boldsymbol{\theta}^+, \boldsymbol{\theta}^-$.

The approach in this study was to include in the model a huge number of potentially relevant statistics and let sTERMGs select the most relevant ones. Results showed indeed that only a subset of the estimated parameters turned out to be statistically significant. The recovery process was characterized by a tendency to preserve and strengthen connections within pairs of subsystems (*e.g.*, default, frontoparietal, thalamus, visual) and between some of them (eg, default-visual, somatomotor-frontoparietal, ventral-visual). Similarly to the results found in [96], a significant propensity to close triangles over time also emerged from this study, thus confirming the importance of the local clustering connections in shaping the brain reorganization process from brain damage-induced coma.

5. On the interpretation of ERGMs

Up to now, we have not lingered on the foundation and subtleties of ERGMs, albeit they are crucial to understand and interpret the related results. Due to the number and variety of scientists and practitioners that have developed or used these models, there is sometimes confusion in the scientific literature of ERGMs, if not in the notions at least in the lexicon. This might lead to unsupported claims or reckless interpretations of the results. We therefore aim at clarifying in this section the philosophy of the method and its theoretical underpinnings.

Our discussion lies on three fundamental concepts: *i*) description, *ii*) prediction, and *iii*) explanation. We briefly pin-down for clarity the corresponding definitions. *Description* means a characterization, or a summary, of the state of the system under investigation, based on the choice of suitable variables and observables to be measured. By definition, a description refers to the past or present state of the system. *Prediction* is the anticipation (forecast) of the future state of a system. *Explanation* aims at reaching a causal or mechanistic understanding of why a particular phenomenon or event has occurred, based on a theoretical framework or empirical evidence [122, 123, 124].

We argue that in general ERGMs as those described in (Eq.1) and their derivatives, are certainly descriptive, might be predictive, but they are not explicative. Note that a method can be tremendously efficient in predicting outcomes of an experiment while at the same time providing no explanation on the underlying physics. Analogously, a description of the phenomena under investigation is radically different from a scientific explanation [122].

5.1. ERGMs and the maximum entropy principle (MEP)

First, let us explicit the origin of the ERGM formulation from basic information theoretic principles. ERGMs can be indeed seen as one of the many particular instances belonging to the broad class of Maximum Entropy (Max.Ent.) models [14, 15, 125]. Similar models have been applied to neural populations [126], antibody diversity [127], protein sequences [62], economics [128], population genetics [129] and many more. Let us then derive the ERGMs from the MEP.||

|| Note, however, that this is not the way the ERGMs were originally introduced, which instead was based on the Hammersley-Clifford theorem for Markov graphs [47], building on a previous work of J. Besag in the context of spatial models of lattice systems [49].

Consider an observed graph $G^* \in \mathcal{G}$ and assume that all the *relevant* information about G^* is encoded in a set of sufficient statistics $\mathbf{x}(G^*)$. According to the MEP, the most unbiased probability density function $P(G)$ consistent with available knowledge is obtained by maximizing the Shannon information entropy \mathcal{S}

$$\mathcal{S}[P] = - \sum_{G \in \mathcal{G}} P(G) \log P(G) \quad (29)$$

under the constraints

$$\sum_{G \in \mathcal{G}} P(G) = 1, \quad (30)$$

$$\sum_{G \in \mathcal{G}} \mathbf{x}(G) P(G) = \mathbf{x}(G^*) : \quad (31)$$

the first one is just a normalization, the others fix the expected values of the sufficient statistics over the ensemble \mathcal{G} to the correspondent values computed for G^* .

The rationale at the basis of the Max.Ent. recipe is very pragmatic, as already stressed in [130], and dates back in similar forms to Bernoulli and Laplace. Let us consider the total number of ways in which the probability distribution of a system can be realized, given that all its possible states are equivalently likely. By a simple combinatorial theorem [130], it can be shown that the Max.Ent. distribution is the most likely and also that any other distribution becomes highly atypical as the number of degrees of freedom - in our case, the number of possible edges - becomes large.

In order to solve the constrained maximization problem defined by Eqs.29-31 we introduce a set of Lagrange multipliers $\boldsymbol{\lambda} \in \mathbb{R}^{r+1}$ and maximizing over P the expression:

$$- \sum_{G \in \mathcal{G}} P(G) \log P(G) + \lambda_0 \left(\sum_{G \in \mathcal{G}} P(G) - 1 \right) + \sum_{\iota=1}^R \lambda_{\iota} \left(\sum_{G \in \mathcal{G}} x_{\iota}(G) P(G) - x_{\iota}(G^*) \right), \quad (32)$$

which gives $P(G|\boldsymbol{\lambda}) = e^{\lambda_0 - 1 + \sum_{\iota=1}^R \lambda_{\iota} x_{\iota}(G)}$. By imposing (Eq.30-31), we find $\lambda_0 = 1 - \log \mathcal{Z}$ and $\lambda_{\iota} = \theta_{\iota}$ for $\iota = 1, \dots, R$, with the $\boldsymbol{\theta}$ being *defined* as the choice of the corresponding multipliers which satisfy Eq.31. The result is precisely Eq.1.

The same argument can be used for a temporal series of graphs $\mathbf{G} = G^1, \dots, G^T$. If the information is available in the form of graph statistics $\mathbf{x}(\mathbf{G}^*)$ evaluated on the observed sequence \mathbf{G}^* , the path entropy reads as

$$\mathcal{S}[P] = \sum_{G^1 \in \mathcal{G}} \dots \sum_{G^T \in \mathcal{G}} P(\mathbf{G}) \log P(\mathbf{G}), \quad (33)$$

and the mass probability function becomes

$$P(\mathbf{G}|\boldsymbol{\theta}) = \frac{e^{\boldsymbol{\theta} \cdot \mathbf{x}(\mathbf{G})}}{\sum_{\tilde{G}^1 \in \mathcal{G}} \dots \sum_{\tilde{G}^T \in \mathcal{G}} e^{\boldsymbol{\theta} \cdot \mathbf{x}(\tilde{\mathbf{G}})}}, \quad (34)$$

which is the more general case of the TERGMs presented before (Eq.25-26). This generalized variational principle is sometimes referred as maximum caliber (Max.Cal.) [131].

5.2. Description and prediction: the choice of graph statistics

As discussed in (Sec.3.3.1), when formulating ERGMs it is crucial to avoid degeneracy issues which might occur especially with too simplistic model formulations. The possibility of including several potentially relevant effects (model terms) has two advantages. On the one hand, it helps alleviating degeneracy issues, on the other hand it allows to statistically test the joint effect of multiple network properties. Once the model is specified and as long as the estimation process converges properly, an ERGM is always descriptive, the estimated parameters convey information on the relevance of the associated statistics. Whether an ERGM is also predictive is a less trivial question that we address in the following.

A delicate implicit assumption of ERGMs is that the selected statistics $\mathbf{x}(G^*)$ encode all the *relevant* information of the system, *i.e.* they are *sufficient* statistics. This assumption raised major criticism around Max.Ent. and related methods, see *e.g.* [132, 133, 129]. Indeed, what is deemed to be *relevant* entirely relies on the modeler’s hypothesis and is consequently affected by some arbitrariness. This is particularly true in complex systems science where, almost by definition, we rarely have control on the hidden degrees of freedom [6, 134]. It might be that what we aim to measure, or what we are able to measure, simply is not a relevant feature of the system.

In the case of TERGMs, such as those in Eq.25-26, there are two more subtleties. First, the Markov hypothesis in Eq.25 is in general hard to justify a-priori [135, 133]. Second, when writing Eq.26 we implicitly assume that the selected statistics are relevant throughout the observed time window. This is true if we can assume the stationarity of the process [136], whose biological plausibility is however disputable [137]. Both these two assumptions are further complicated by the fact that in many real applications the measurements of networks are not equally spaced in time [138].

To address these issues, *a-posteriori* model-selection approaches like goodness of fit (GoF) or out-of-sample assessment can be adopted to rule out models that are markedly wrong. These methods evaluate the ability of the graphs simulated with a fitted ERGM to reproduce properties of the system other than those explicitly included. When this happens then the estimated ERGM is not only *descriptive* but also *predictive*. Thus, when using ERGMs with GOF assessments, it is not only important to use appropriate *relevant* graph statistics but also to have a precise hypothesis on the network features against which to test the predictive performance.

5.3. Explanation: ERGMs and Stat.Mech.

In the previous sections, we have insisted on the role of the modeler in shaping the mass probability distribution by selecting a set of graph statistics. This is consistent with the E.T. Jaynes’ subjective interpretation of probability [139, 131], according to which probabilities are epistemic statements about physical systems. Put differently, they reflect our state of knowledge or, equivalently, of ignorance, about the world.

Such subjective probability is enough for purely inferential purposes. For example, even if Eq.1 is not the *true* probability density function of the system under investigation, it may nevertheless have good predictive performance (GoF) for a given set of test statistics. Indeed, the ERGM inferential scheme is agnostic about the data-generating process.

This property is somewhat unsatisfactory if the ultimate goal is to explain the physics underlying the observed data. In that case, probabilities should not reflect a

subjective state of knowledge but they should have an objective *physical* meaning.

This is, for instance, the case of statistical mechanics where probabilities have an objective irreducible meaning, they do not depend on the observer. The role of the theory is then to explain how the macroscopic properties of bodies, mainly thermodynamic properties, statistically arise from their microscopic structure [140]. The formal analogy between the ERGM in (Eq.1) and the Gibbs-Boltzmann distribution for physical systems at equilibrium has instilled hopes for a similar micro-level objective interpretation of the ERGM probability density function. However, on a closer inspection we immediately realize that this equivalence is problematic.

These two exponential distributions have in fact very different theoretical underpinnings. Just as an example: at the very foundation of statistical mechanics, and hence of Gibbs-Boltzmann exponential distributions, there is the assumption that the system dynamics obeys a *detailed balance* [141]. In its simplest form, detailed balance means that for any two accessible states r, s the transition rate $W_{r \rightarrow s}$ is equal to its inverse $W_{s \rightarrow r}$. This assumption becomes highly disputable when it comes to biological processes. Detailed balance is indeed violated at the molecular level and apparently so also for large-scale brain dynamics even during intrinsic spontaneous activity [142]. Therefore, in general, that between the ERGM and the Gibbs-Boltzmann distributions should not be considered more than a formal analogy, and their respective interpretations should not be confused.

6. Conclusion and perspectives

Statistical models are crucial to gain intuition on the connection rules and generative mechanisms of brain networks across multiple time and spatial scales [143]. Furthermore, they can be used to improve the detection of the underlying network properties in the presence of statistical noise in the data [144, 145, 146].

Here, we have discussed a family of models - exponential random graph models (ERGMs) - which allows for simultaneously taking into account different local connection properties, disentangle their interactions and quantify their statistical influence on the global observed network. The number of studies using ERGMs to model brain connectivity networks has rapidly increased in the last decade. We have discussed how they have been used to characterize intrinsic brain connectivity, build group representative networks, discriminate different brain states, as well as model temporal patterns, with focused applications to brain diseases.

In addition, we have placed the ERGM framework in the broader context of maximum entropy models and discussed in detail the subtleties of network analyses based on the Max.Ent. recipe, with particular focus on the meaning of the inferred model parameters. We have discussed the descriptive and predictive power of ERGMs, and warned about its interpretation as an equilibrium *theory* of the brains' biological processes, for which a non-equilibrium picture is more appropriate.

Despite the increasing number of works using or developing ERGMs to address neuroscience-related questions, the field is still in its infancy and we hope this review will encourage curious physicists, network scientists, and neuroscientists to pursue the research direction we have laid out here. In the years to come, we anticipate a boost in the field fueled by an increasing availability of high-quality experimental data, which is crucial to validate the hypothesized mechanisms. This includes the use of ERGMs to model brain diseases not only from a recovery perspective (*e.g.*, stroke or coma), but also in terms of degeneration (*e.g.* Parkinson, Alzheimer's), as well as to evaluate the

efficacy of treatments for psychiatric conditions (*e.g.* Schizophrenia) and developmental disorders (*e.g.* autism).

On a more methodological note, only a small fraction of the ERGM-like galaxy of statistical models has been employed so far in network neuroscience. Among the unexplored alternatives, we spot out the potential use of multilayer ERGMs [147] to capture interactions between different scales of neural organization [143]. As for temporal analyses, a promising family of models beyond the (T)ERGM paradigm is the related stochastic actor oriented model (SAOM), which simulates time-varying network changes through a microscale process defined at the level of the nodes, *i.e.* the actors [148]. Critically, both TERGMS and SAOM only take into account monotonic network changes increasing or decreasing over time. However, modeling brain reorganization might require capturing non-monotonic mechanisms, too. This is particularly true for intrinsic brain functioning and cognitive/motor tasks, which typically exhibit oscillatory time-varying dynamics [113]. In this direction, the development of models that integrate this dynamic component, such as the varying-coefficient ERGM [149], appear particularly important to overcome such limitations. A distinct mention is for stochastic block models (SBMs), which allows flexible and scalable analysis of complex networks [150]. In addition to handling community structures, recent SBM extensions also include several types of structures, such as node degree sequences [151], hierarchical structures [152], edge weights [153], and triadic closures [154]. SBM applications to brain networks are still in their infancy (see [13] for a recent example), but they surely represent a future promising research avenue. Finally, a long-term and ambitious goal would be to enrich statistical models by including higher-order interactions, which might reveal previously unappreciated brain organization principles [155].

We believe that combining physics-inspired concepts, biological knowledge, and statistical methods is essential for unraveling the intricacies of the human brain's structure and dynamics. This synergy holds the potential to provide crucial insights into the complexity of the brain.

Acknowledgements

We would like to thank the anonymous reviewers, whose comments helped us improving the initial version of the review. We thank Erik Aurell and Mario Chavez for many insightful discussions and for providing useful remarks on the manuscript. We are also thankful to Thibault Rolland and Remy Ben Messaoud for their contribution in the preparation of the figures. FDVF acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 864729)

References

- [1] Herbert A Simon. The architecture of complexity. In *Facets of systems science*, pages 457–476. Springer, 1991.
- [2] Christopher W Lynn and Danielle S Bassett. The physics of brain network structure, function and control. *Nature Reviews Physics*, 1(5):318–332, 2019.
- [3] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, 10(3):186–198, 2009.
- [4] Danielle S Bassett and Olaf Sporns. Network neuroscience. *Nature neuroscience*, 20(3):353–364, 2017.

- [5] Cornelis J Stam. Modern network science of neurological disorders. *Nature Reviews Neuroscience*, 15(10):683–695, 2014.
- [6] Danielle S Bassett, Perry Zurn, and Joshua I Gold. On the nature and use of models in network neuroscience. *Nature Reviews Neuroscience*, 19(9):566–578, 2018.
- [7] Olaf Sporns, Dante R Chialvo, Marcus Kaiser, and Claus C Hilgetag. Complex networks: small-world and scale-free architectures. *Trends in Cognitive Sciences*, 9(8):418–425, 2004.
- [8] Richard F Betzel, Andrea Avena-Koenigsberger, Joaquín Goñi, Ye He, Marcel A De Reus, Alessandra Griffa, Petra E Vértes, Bratislav Mišić, Jean-Philippe Thiran, Patric Hagmann, et al. Generative models of the human connectome. *Neuroimage*, 124:1054–1064, 2016.
- [9] Ed Bullmore and Olaf Sporns. The economy of brain network organization. *Nature reviews neuroscience*, 13(5):336–349, 2012.
- [10] Fabrizio de Vico Fallani, Jonas Richiardi, Mario Chavez, and Sophie Achard. Graph analysis of functional brain networks: practical issues in translational neuroscience. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1653):20130521, 2014.
- [11] Sergei Maslov and Kim Sneppen. Specificity and stability in topology of protein networks. *Science*, 296(5569):910–913, 2002.
- [12] Richard F Betzel, John D Medaglia, and Danielle S Bassett. Diversity of meso-scale architecture in human and non-human connectomes. *Nature communications*, 9(1):1–14, 2018.
- [13] Joshua Faskowitz, Xiaoran Yan, Xi-Nian Zuo, and Olaf Sporns. Weighted stochastic block models of the human connectome across the life span. *Scientific reports*, 8(1):1–16, 2018.
- [14] Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.
- [15] Giulio Cimini, Tiziano Squartini, Fabio Saracco, Diego Garlaschelli, Andrea Gabrielli, and Guido Caldarelli. The statistical physics of real-world networks. *Nature Reviews Physics*, 1(1):58–71, 2019-01.
- [16] David R Hunter and Mark S Handcock. Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics*, 15(3):565–583, 2006.
- [17] Stefano Boccaletti, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. Complex networks: Structure and dynamics. *Physics reports*, 424(4-5):175–308, 2006.
- [18] Olaf Sporns and Richard F Betzel. Modular brain networks. *Annual review of psychology*, 67:613–640, 2016.
- [19] Danielle Smith Bassett and ED Bullmore. Small-world brain networks. *The neuroscientist*, 12(6):512–523, 2006.
- [20] Danielle S Bassett and Edward T Bullmore. Small-world brain networks revisited. *The Neuroscientist*, 23(5):499–516, 2017.
- [21] Gaolang Gong, Yong He, Luis Concha, Catherine Lebel, Donald W Gross, Alan C Evans, and Christian Beaulieu. Mapping anatomical connectivity patterns of human cerebral cortex using in vivo diffusion tensor imaging tractography. *Cerebral cortex*, 19(3):524–536, 2009.
- [22] Sophie Achard, Raymond Salvador, Brandon Whitccher, John Suckling, and ED Bullmore. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *Journal of Neuroscience*, 26(1):63–72, 2006.
- [23] Robert Rosenbaum, Matthew A Smith, Adam Kohn, Jonathan E Rubin, and Brent Doiron. The spatial structure of correlated neuronal variability. *Nature neuroscience*, 20(1):107–114, 2017.
- [24] Richard F Betzel and Danielle S Bassett. Multi-scale brain networks. *Neuroimage*, 160:73–83, 2017.
- [25] Alex Fornito, Andrew Zalesky, and Michael Breakspear. The connectomics of brain disorders. *Nature Reviews Neuroscience*, 16(3):159–172, 2015.
- [26] Derek K Jones, Thomas R Knösche, and Robert Turner. White matter integrity, fiber count, and other fallacies: the do’s and don’ts of diffusion mri. *Neuroimage*, 73:239–254, 2013.
- [27] Onerva Korhonen, Massimiliano Zanin, and David Papo. Principles and open questions in functional brain network reconstruction. *Human Brain Mapping*, 42(11):3680–3711, 2021.
- [28] Alex Fornito, Andrew Zalesky, and Edward Bullmore. *Fundamentals of brain network analysis*. Academic Press, 2016.
- [29] Peter J Basser, Sinisa Pajevic, Carlo Pierpaoli, Jeffrey Duda, and Akram Aldroubi. In vivo fiber tractography using dt-mri data. *Magnetic resonance in medicine*, 44(4):625–632, 2000.
- [30] Juliana Gonzalez-Astudillo, Tiziana Cattai, Giulia Bassignana, Marie-Constance Corsi, and Fabrizio De Vico Fallani. Network-based brain-computer interfaces: principles and applications. *Journal of Neural Engineering*, 18(1):011001, 2021.
- [31] J Craig Henry. Electroencephalography: basic principles, clinical applications, and related fields. *Neurology*, 67(11):2092–2092, 2006.

- [32] Marcus E Raichle. Behind the scenes of functional brain imaging: a historical and physiological perspective. *Proceedings of the National Academy of Sciences*, 95(3):765–772, 1998.
- [33] Raymond Salvador, John Suckling, Martin R Coleman, John D Pickard, David Menon, and ED Bullmore. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral cortex*, 15(9):1332–1342, 2005.
- [34] Bernadette CM Van Wijk, Cornelis J Stam, and Andreas Daffertshofer. Comparing brain networks of different size and connectivity density using graph theory. *PloS one*, 5(10):e13701, 2010.
- [35] Fabrizio De Vico Fallani, Vito Latora, and Mario Chavez. A topological criterion for filtering information in complex brain networks. *PLoS computational biology*, 13(1):e1005305, 2017.
- [36] Kathleen A Garrison, Dustin Scheinost, Emily S Finn, Xilin Shen, and R Todd Constable. The (in) stability of functional brain network measures across thresholds. *Neuroimage*, 118:651–661, 2015.
- [37] František Váša and Bratislav Mišić. Null models in network neuroscience. *Nature Reviews Neuroscience*, pages 1–12, 2022.
- [38] Fabrizio De Vico Fallani, Floriana Pichiorri, Giovanni Morone, Marco Molinari, Fabio Babiloni, Febo Cincotti, and Donatella Mattia. Multiscale topological properties of functional brain networks during motor imagery after stroke. *Neuroimage*, 83:438–449, 2013.
- [39] M Angeles Serrano, Marián Boguná, and Alessandro Vespignani. Extracting the multiscale backbone of complex weighted networks. *Proceedings of the national academy of sciences*, 106(16):6483–6488, 2009.
- [40] Bernard O Koopman. On distributions admitting a sufficient statistic. *Transactions of the American Mathematical society*, 39(3):399–409, 1936.
- [41] Edwin James George Pitman. Sufficient statistics and intrinsic accuracy. In *Mathematical Proceedings of the cambridge Philosophical society*, volume 32, pages 567–579. Cambridge University Press, 1936.
- [42] Ronald A Fisher. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 222(594-604):309–368, 1922.
- [43] Ronald A Fisher. Theory of statistical estimation. In *Mathematical proceedings of the Cambridge philosophical society*, volume 22, pages 700–725. Cambridge University Press, 1925.
- [44] Ole Barndorff-Nielsen. *Information and exponential families: in statistical theory*. John Wiley & Sons, 2014.
- [45] Lawrence D Brown. *Fundamentals of statistical exponential families: with applications in statistical decision theory*. Ims, 1978.
- [46] Paul W Holland and Samuel Leinhardt. An exponential family of probability distributions for directed graphs. *Journal of the american Statistical association*, 76(373):33–50, 1981.
- [47] Ove Frank and David Strauss. Markov Graphs. *Journal of the American Statistical Association*, 81(395):832–842, 1986.
- [48] David Strauss. On a general class of models for interaction. *SIAM review*, 28(4):513–527, 1986.
- [49] Julian Besag. Spatial Interaction and the Statistical Analysis of Lattice Systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236, 1974.
- [50] Stanley Wasserman and Philippa Pattison. Logit models and logistic regressions for social networks: I. an introduction to markov graphs andp. *Psychometrika*, 61(3):401–425, 1996.
- [51] Carolyn J Anderson, Stanley Wasserman, and Bradley Crouch. A p* primer: Logit models for social networks. *Social networks*, 21(1):37–66, 1999.
- [52] Ginestra Bianconi. Entropy of network ensembles. *Physical Review E*, 79(3):036114, 2009.
- [53] Filippo Radicchi, Dmitri Krioukov, Harrison Hartle, and Ginestra Bianconi. Classical information theory of networks. *Journal of Physics: Complexity*, 1(2):025001, 2020.
- [54] Matthew J Silk, Darren P Croft, Richard J Delahay, David J Hodgson, Nicola Weber, Mike Boots, and Robbie A McDonald. The application of statistical network models in disease research. *Methods in Ecology and Evolution*, 8(9):1026–1041, 2017.
- [55] Mark S Handcock. Statistical models for social networks: Inference and degeneracy. In *Dynamic Social Network Modelling and Analysis: Workshop Summary and Papers*, pages 229–240. National Academy Press, 2003.
- [56] Garry Robins, Pip Pattison, Yuval Kalish, and Dean Lusher. An introduction to exponential random graph (p*) models for social networks. *Social networks*, 29(2):173–191, 2007.
- [57] Dean Lusher, Johan Koskinen, and Garry Robins. *Exponential random graph models for social networks: Theory, methods, and applications*, volume 35. Cambridge University Press, 2013.

- [58] Mark S Handcock, David R Hunter, Carter T Butts, Steven M Goodreau, and Martina Morris. statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of statistical software*, 24(1):1548, 2008.
- [59] Steven M Goodreau, Mark S Handcock, David R Hunter, Carter T Butts, and Martina Morris. A statnet tutorial. *Journal of statistical software*, 24(9):1, 2008.
- [60] David R Hunter, Mark S Handcock, Carter T Butts, Steven M Goodreau, and Martina Morris. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of statistical software*, 24(3):nihpa54860, 2008.
- [61] H Chau Nguyen, Riccardo Zecchina, and Johannes Berg. Inverse statistical problems: from the inverse ising problem to data science. *Advances in Physics*, 66(3):197–261, 2017.
- [62] Simona Cocco, Christoph Feinauer, Matteo Figliuzzi, Rémi Monasson, and Martin Weigt. Inverse statistical physics of protein sequences: a key issues review. *Reports on Progress in Physics*, 81(3):032601, 2018.
- [63] Charles J Geyer. Markov chain monte carlo maximum likelihood. 1991.
- [64] Charles J Geyer and Elizabeth A Thompson. Constrained monte carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 54(3):657–683, 1992.
- [65] Bruce A Desmarais and Skyler J Cranmer. Statistical inference for valued-edge networks: The generalized exponential random graph model. *PloS one*, 7(1):e30136, 2012.
- [66] Pavel N Krivitsky, David R Hunter, Martina Morris, and Chad Klumb. Ergm 4: Computational improvements. *arXiv preprint arXiv:2203.08198*, 2022.
- [67] Mark S. Handcock Pavel N. Krivitsky, David R. Hunter, Carter T. Butts, Chad Klumb, Steven M. Goodreau, and Martina Morris. Statnet: Tools for the statistical modeling of network data. <https://statnet.org>, 2003/2022.
- [68] Alberto Caimo and Nial Friel. Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41–55, 2011.
- [69] James D Wilson, Matthew J Denny, Shankar Bhamidi, Skyler J Cranmer, and Bruce A Desmarais. Stochastic weighted graphs: Flexible model specification and simulation. *Social Networks*, 49:37–47, 2017.
- [70] Michael Schweinberger and Pamela Luna. Hergm: Hierarchical exponential-family random graph models. *Journal of Statistical Software*, 85(1):1–39, 2018.
- [71] Pavel N Krivitsky and Martina Morris. Inference for social network models from egocentrically sampled data, with application to understanding persistent racial disparities in hiv prevalence in the us. *The annals of applied statistics*, 11(1):427, 2017.
- [72] Pavel N Krivitsky and Mark S Handcock. A separable model for dynamic networks. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, 76(1):29, 2014.
- [73] Philip Leifeld, Skyler J Cranmer, and Bruce A Desmarais. Temporal exponential random graph models with btergm: Estimation and bootstrap confidence intervals. *Journal of Statistical Software*, 83(6), 2018.
- [74] V Dichio. ergm_minimal. https://github.com/dichio/ergm_minimal, 2023.
- [75] Pavel N Krivitsky, David R Hunter, Martina Morris, and Chad Klumb. ergm 4: New features for analyzing exponential-family random graph models. *Journal of Statistical Software*, 105:1–44, 2023.
- [76] Juyong Park and M. E. J. Newman. Statistical mechanics of networks. *Physical Review E*, 70(6):066117, 2007.
- [77] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.
- [78] Paul Erdos, Alfréd Rényi, et al. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.
- [79] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.
- [80] Tom AB Snijders et al. Markov chain monte carlo estimation of exponential random graph models. *Journal of Social Structure*, 3(2):1–40, 2002.
- [81] Mark S Handcock, Garry Robins, Tom Snijders, Jim Moody, and Julian Besag. Assessing degeneracy in statistical models of social networks. Technical report, Working paper, 2003.
- [82] Alessandro Rinaldo, Stephen E Fienberg, and Yi Zhou. On the geometry of discrete exponential families with application to exponential random graph models. *Electronic Journal of Statistics*, 3:446–484, 2009.
- [83] Michael Schweinberger. Instability, sensitivity, and degeneracy of discrete exponential families. *Journal of the American Statistical Association*, 106(496):1361–1370, 2011.
- [84] Juyong Park and Mark EJ Newman. Solution of the two-star model of a network. *Physical*

- Review E*, 70(6):066146, 2004.
- [85] Juyong Park and Mark EJ Newman. Solution for the properties of a clustered network. *Physical Review E*, 72(2):026136, 2005.
- [86] Michael Schweinberger, Pavel N Krivitsky, Carter T Butts, and Jonathan R Stewart. Exponential-family models of random graphs: Inference in finite, super and infinite population scenarios. *Statistical Science*, 35(4):627–662, 2020.
- [87] Tom AB Snijders, Philippa E Pattison, Garry L Robins, and Mark S Handcock. New specifications for exponential random graph models. *Sociological methodology*, 36(1):99–153, 2006.
- [88] Sean L Simpson, Satoru Hayasaka, and Paul J Laurienti. Exponential random graph modeling for complex brain networks. *PloS one*, 6(5):e20039, 2011.
- [89] Sean L Simpson, Malaak N Moussa, and Paul J Laurienti. An exponential random graph modeling approach to creating group-based representative whole-brain connectivity networks. *Neuroimage*, 60(2):1117–1126, 2012.
- [90] Michel RT Sinke, Rick M Dijkhuizen, Alberto Caimo, Cornelis J Stam, and Willem M Otte. Bayesian exponential random graph modeling of whole-brain structural networks across lifespan. *NeuroImage*, 135:79–91, 2016.
- [91] Catalina Obando and Fabrizio De Vico Fallani. A statistical model for brain networks inferred from large-scale electrophysiological signals. *Journal of The Royal Society Interface*, 14(128):20160940, 2017.
- [92] Paul E Stillman, James D Wilson, Matthew J Denny, Bruce A Desmarais, Shankar Bhamidi, Skyler J Cranmer, and Zhong-Lin Lu. Statistical modeling of the default mode brain network reveals a segregated highway structure. *Scientific reports*, 7(1):1–14, 2017.
- [93] John Dell’Italia, Micah A Johnson, Paul M Vespa, and Martin M Monti. Network analysis in disorders of consciousness: four problems and one proposed solution (exponential random graph models). *Frontiers in neurology*, 9:439, 2018.
- [94] Paul E Stillman, James D Wilson, Matthew J Denny, Bruce A Desmarais, Skyler J Cranmer, and Zhong-Lin Lu. A consistent organizational structure across multiple functional subnetworks of the human brain. *NeuroImage*, 197:24–36, 2019.
- [95] BCL Lehmann, RN Henson, Linda Geerligs, SR White, et al. Characterising group-level brain connectivity: a framework using bayesian exponential random graph models. *NeuroImage*, 225:117480, 2021.
- [96] Catalina Obando, Charlotte Rosso, Joshua Siegel, Maurizio Corbetta, and Fabrizio De Vico Fallani. Temporal exponential random graph models of longitudinal brain networks after stroke. *Journal of the Royal Society Interface*, 19(188):20210850, 2022.
- [97] David R Hunter, Steven M Goodreau, and Mark S Handcock. Goodness of fit of social network models. *Journal of the american statistical association*, 103(481):248–258, 2008.
- [98] Jonathan D Power, Bradley L Schlaggar, Christina N Lessov-Schlaggar, and Steven E Petersen. Evidence for hubs in human functional brain networks. *Neuron*, 79(4):798–813, 2013.
- [99] Jonathan D Power, Alexander L Cohen, Steven M Nelson, Gagan S Wig, Kelly Anne Barnes, Jessica A Church, Alecia C Vogel, Timothy O Laumann, Fran M Miezin, Bradley L Schlaggar, et al. Functional network organization of the human brain. *Neuron*, 72(4):665–678, 2011.
- [100] Shi Gu, Fabio Pasqualetti, Matthew Cieslak, Qawi K Telesford, Alfred B Yu, Ari E Kahn, John D Medaglia, Jean M Vettel, Michael B Miller, Scott T Grafton, et al. Controllability of structural brain networks. *Nature communications*, 6(1):1–10, 2015.
- [101] Nikola T Markov, Mária Ercsey-Ravasz, David C Van Essen, Kenneth Knoblauch, Zoltán Toroczkai, and Henry Kennedy. Cortical high-density counterstream architectures. *Science*, 342(6158):1238406, 2013.
- [102] Gustavo Deco, Giulio Tononi, Melanie Boly, and Morten L Kringelbach. Rethinking segregation and integration: contributions of whole-brain modelling. *Nature Reviews Neuroscience*, 16(7):430–439, 2015.
- [103] Ming Song, Yong Liu, Yuan Zhou, Kun Wang, Chunshui Yu, and Tianzi Jiang. Default network and intelligence difference. *IEEE Transactions on autonomous mental development*, 1(2):101–109, 2009.
- [104] Andrew Zalesky, Alex Fornito, and Edward T Bullmore. Network-based statistic: identifying differences in brain networks. *Neuroimage*, 53(4):1197–1207, 2010.
- [105] Robert J Barry, Adam R Clarke, Stuart J Johnstone, Christopher A Magee, and Jacqueline A Rushby. Eeg differences between eyes-closed and eyes-open resting conditions. *Clinical neurophysiology*, 118(12):2765–2773, 2007.
- [106] Yu A Boytsova and SG Danko. Eeg differences between resting states with eyes open and

- closed in darkness. *Human physiology*, 36(3):367–369, 2010.
- [107] Gaolang Gong, Pedro Rosa-Neto, Felix Carbonell, Zhang J Chen, Yong He, and Alan C Evans. Age-and gender-related differences in the cortical anatomical network. *Journal of Neuroscience*, 29(50):15684–15693, 2009.
- [108] Patric Hagmann, Olaf Sporns, Neel Madan, Leila Cammoun, Rudolph Pienaar, Van Jay Wedeen, Reto Meuli, J-P Thiran, and PE Grant. White matter maturation reshapes structural connectivity in the late developing human brain. *Proceedings of the National Academy of Sciences*, 107(44):19067–19072, 2010.
- [109] Willem M Otte, Eric van Diessen, Subhadip Paul, Rajiv Ramaswamy, VP Subramanyam Rallabandi, Cornelis J Stam, and Prasun K Roy. Aging alterations in whole-brain networks during adulthood mapped with the minimum spanning tree indices: the interplay of density, connectivity cost and life-time trajectory. *Neuroimage*, 109:171–189, 2015.
- [110] F De Vico Fallani, Vito Latora, Laura Astolfi, Febo Cincotti, Donatella Mattia, Maria Grazia Marciani, Serenella Salinari, Alfredo Colosimo, and Fabrizio Babiloni. Persistent patterns of interconnection in time-varying cortical networks estimated from high-resolution eeg recordings in humans during a simple motor act. *Journal of Physics A: Mathematical and Theoretical*, 41(22):224014, 2008.
- [111] Miguel Valencia, J Martinerie, Samuel Dupont, and M Chavez. Dynamic small-world behavior in functional brain networks unveiled by an event-related networks approach. *Physical Review E*, 77(5):050905, 2008.
- [112] Danielle S Bassett, Nicholas F Wymbs, Mason A Porter, Peter J Mucha, Jean M Carlson, and Scott T Grafton. Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, 108(18):7641–7646, 2011.
- [113] Maria Giulia Preti, Thomas AW Bolton, and Dimitri Van De Ville. The dynamic functional connectome: State-of-the-art and perspectives. *Neuroimage*, 160:41–54, 2017.
- [114] Kai Wu, Yasuyuki Taki, Kazunori Sato, Haochen Qi, Ryuta Kawashima, and Hiroshi Fukuda. A longitudinal study of structural brain network changes with normal aging. *Frontiers in human neuroscience*, 7:113, 2013.
- [115] Sophie Dautricourt, Robin de Flores, Brigitte Landeau, Géraldine Poinsnel, Matthieu Vanhoutte, Nicolas Delcroix, Francis Eustache, Denis Vivien, Vincent de la Sayette, and Gaël Chételat. Longitudinal changes in hippocampal network connectivity in alzheimer’s disease. *Annals of Neurology*, 90(3):391–406, 2021.
- [116] Joshua S Siegel, Benjamin A Seitzman, Lenny E Ramsey, Mario Ortega, Evan M Gordon, Nico UF Dosenbach, Steven E Petersen, Gordon L Shulman, and Maurizio Corbetta. Re-emergence of modular brain networks in stroke recovery. *Cortex*, 101:44–59, 2018.
- [117] Petter Holme and Jari Saramäki. Temporal networks. *Physics reports*, 519(3):97–125, 2012.
- [118] Naoki Masuda and Renaud Lambiotte. *A guide to temporal networks*. World Scientific, 2016.
- [119] Steve Hanneke, Wenjie Fu, and Eric P Xing. Discrete temporal models of social networks. *Electronic journal of statistics*, 4:585–605, 2010.
- [120] Joshua Sarfaty Siegel, Lenny E Ramsey, Abraham Z Snyder, Nicholas V Metcalf, Ravi V Chacko, Kilian Weinberger, Antonello Baldassarre, Carl D Hacker, Gordon L Shulman, and Maurizio Corbetta. Disruptions of network connectivity predict impairment in multiple behavioral domains after stroke. *Proceedings of the National Academy of Sciences*, 113(30):E4367–E4376, 2016.
- [121] Christian Grefkes and Gereon R Fink. Recovery from stroke: current concepts and future perspectives. *Neurological research and practice*, 2(1):1–10, 2020.
- [122] Carl G Hempel and Paul Oppenheim. Studies in the logic of explanation. *Philosophy of science*, 15(2):135–175, 1948.
- [123] Joseph F Hanna. Explanation, prediction, description, and information theory. *Synthese*, 20(3):308–334, 1969.
- [124] Galit Shmueli. To explain or to predict? *Statistical science*, 25(3):289–310, 2010.
- [125] Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- [126] Elad Schneidman, Michael J Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.
- [127] Thierry Mora, Aleksandra M Walczak, William Bialek, and Curtis G Callan Jr. Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences*, 107(12):5405–5410, 2010.
- [128] Sanjay G Reddy. What is an explanation? statistical physics and economics. *The European Physical Journal Special Topics*, 229(9):1645–1659, 2020.
- [129] Vito Dichio, Hong-Li Zeng, and Erik Aurell. Statistical genetics in and out of quasi-linkage

equilibrium. *Reports on Progress in Physics*, 86(5):052601, 2023.

[130] Edwin T Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9):939–952, 1982.

[131] Steve Pressé, Kingshuk Ghosh, Julian Lee, and Ken A Dill. Principles of maximum entropy and maximum caliber in statistical physics. *Reviews of Modern Physics*, 85(3):1115, 2013.

[132] Erik Aurell. The maximum entropy fallacy redux? *PLoS computational biology*, 12(5):e1004777, 2016.

[133] Gennaro Auletta, Lamberto Rondoni, and Angelo Vulpiani. On the relevance of the maximum entropy principle in non-equilibrium statistical mechanics. *The European Physical Journal Special Topics*, 226(10):2327–2343, 2017.

[134] James Ladyman, James Lambert, and Karoline Wiesner. What is a complex system? *European Journal for Philosophy of Science*, 3(1):33–67, 2013.

[135] Lars Onsager and Stefan Machlup. Fluctuations and irreversible processes. *Physical Review*, 91(6):1505, 1953.

[136] C W Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, volume 13 of *Springer Series in Synergetics*. Springer-Verlag, third edition, 2004.

[137] Xiaona Fang, Karsten Kruse, Ting Lu, and Jin Wang. Nonequilibrium physics in biology. *Reviews of Modern Physics*, 91(4):045004, 2019.

[138] Per Block, Johan Koskinen, James Hollway, Christian Steglich, and Christoph Stadtfeld. Change we can believe in: Comparing longitudinal network models on consistency, interpretability and predictive power. *Social Networks*, 52:180–191, 2018.

[139] Edwin T Jaynes. *Probability theory: The logic of science*. Cambridge university press, 2003.

[140] Luca Peliti. *Statistical mechanics in a nutshell*. Princeton University Press, 2011.

[141] Ludwig Boltzmann. Weitere studien über das wärmegleichgewicht unter gasmolekülen. In *Kinetische Theorie II*, pages 115–225. Springer, 1970.

[142] Christopher W Lynn, Eli J Cornblath, Lia Papadopoulos, Maxwell A Bertolero, and Danielle S Bassett. Broken detailed balance and entropy production in the human brain. *Proceedings of the National Academy of Sciences*, 118(47):e2109889118, 2021.

[143] Charley Presigny and Fabrizio De Vico Fallani. Colloquium: Multiscale modeling of brain network organization. *Reviews of Modern Physics*, 94(3):031002, 2022.

[144] Bradley Efron and Robert Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical science*, pages 54–75, 1986.

[145] David Gfeller, Jean-Cédric Chappelier, and Paolo De Los Rios. Finding instabilities in the community structure of complex networks. *Physical Review E*, 72(5):056135, 2005.

[146] Brian Karrer, Elizaveta Levina, and Mark EJ Newman. Robustness of community structure in networks. *Physical review E*, 77(4):046119, 2008.

[147] Pavel N Krivitsky, Laura M Koehly, and Christopher Steven Marcum. Exponential-family random graph models for multi-layer networks. *Psychometrika*, 85(3):630–659, 2020.

[148] Tom AB Snijders, Gerhard G Van de Bunt, and Christian EG Steglich. Introduction to stochastic actor-based models for network dynamics. *Social networks*, 32(1):44–60, 2010.

[149] Jihui Lee, Gen Li, and James D. Wilson. Varying-coefficient models for dynamic networks. *Computational Statistics & Data Analysis*, 152:107052, 2020.

[150] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.

[151] Brian Karrer and Mark EJ Newman. Stochastic blockmodels and community structure in networks. *Physical review E*, 83(1):016107, 2011.

[152] Tiago P Peixoto. Hierarchical block structures and high-resolution model selection in large networks. *Physical Review X*, 4(1):011047, 2014.

[153] Tiago P Peixoto. Nonparametric weighted stochastic block models. *Physical Review E*, 97(1):012306, 2018.

[154] Tiago P Peixoto. Disentangling homophily, community structure, and triadic closure in networks. *Physical Review X*, 12(1):011004, 2022.

[155] Federico Battiston, Giulia Cencetti, Iacopo Iacopini, Vito Latora, Maxime Lucas, Alice Patania, Jean-Gabriel Young, and Giovanni Petri. Networks beyond pairwise interactions: structure and dynamics. *Physics Reports*, 874:1–92, 2020.