

Statistical inference of assortative community structures

Lizhi Zhang*

*Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom
and The Alan Turing Institute, 96 Euston Road, London NW1 2DB, United Kingdom*

Tiago P. Peixoto[†]

*Department of Network and Data Science, Central European University, 1100 Vienna, Austria;
ISI Foundation, Via Chisola 5, 10126 Torino, Italy;
and Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom*



(Received 14 July 2020; accepted 23 September 2020; published 23 November 2020)

We develop a principled methodology to infer assortative communities in networks based on a nonparametric Bayesian formulation of the planted partition model. We show that this approach succeeds in finding statistically significant assortative modules in networks, unlike alternatives such as modularity maximization, which systematically overfits both in artificial as well as in empirical examples. In addition, we show that our method is not subject to an appreciable resolution limit, and can uncover an arbitrarily large number of communities, as long as there is statistical evidence for them. Our formulation is amenable to model selection procedures, which allow us to compare it to more general approaches based on the stochastic block model, and in this way reveal whether assortativity is in fact the dominating large-scale mixing pattern. We perform this comparison with several empirical networks and identify numerous cases where the network's assortativity is exaggerated by traditional community detection methods, and we show how a more faithful degree of assortativity can be identified.

DOI: [10.1103/PhysRevResearch.2.043271](https://doi.org/10.1103/PhysRevResearch.2.043271)

I. INTRODUCTION

Community detection is one of the most central methods in network science [1,2], and it consists in the algorithmic partition of the nodes of a network into cohesive groups, according to a mathematical definition of this concept (for which there are many). Historically, most community detection methods proposed have focused on the detection of *assortative* communities, i.e., groups of nodes that tend to be more connected to themselves than to other nodes in the network. However, there are also community detection methods that are more general, and attempt to cluster together nodes that have similar patterns of connection, regardless if they are assortative or not [3–5]. The widespread use of assortative community detection methods has led to the belief that the presence of communities is a pervasive feature of many different kinds of real networks [6]. Although the concept of assortativity is a central one in the study of social networks (known as “homophily” in that context) [7], and is also an appealing construct in biology [8–10], it is to some extent unclear if the

perceived assortativity of many networks is a byproduct of using algorithms that can only find this kind of structure. This is particularly problematic since many popular methods do not take into account the statistical significance of the patterns they uncover, and find seemingly strong community structure in completely random graphs [11,12], as well as in trees [13] and other manifestly non-modular networks [14]. More recently, these shortcomings have been addressed by employing Bayesian inference of generative network models [15], which accounts for statistical significance with a built-in Occam's razor, that decides to partition the network into groups only if this is necessary to explain its structure, beyond what can be done by a uniformly random placement of the edges. These approaches, however, are based on general mixing patterns, which include assortativity only as a special case. In many ways this is useful, and in fact arguably advantageous, since if assortativity happens to be the dominating pattern, then the general approach will capture it, otherwise it will reveal a different structure. However, having only a more general method at our disposal also has its shortcomings. First, if it is true that assortativity is the main pattern for a class of networks, then the more general representation is needlessly wasteful for them, as it forces us to represent this simple pattern as a special case of a more complex one. In so doing, we run the risk of marginally overfitting the data, and not focusing on the more central features, at the cost of algorithmic precision. Second, with a more general method alone, it can be difficult to quantify precisely how much has been wasted in the representation, and what is indeed the simpler pattern hiding inside it.

*lz848@bath.ac.uk

†peixotot@ceu.edu

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

In this work we develop a Bayesian inference approach designed to uncover assortative communities in networks, based on the planted partition (PP) model [16–18], which is itself a special case of the more general stochastic block model (SBM) [3,5]. Our approach is nonparametric, and can uncover communities even when their number is unknown, without overfitting. Furthermore we show that it does not suffer from the resolution limit present in other approaches, such as modularity maximization [19], and it can find an arbitrarily large number of communities, provided they are statistically significant. We also revisit an existing equivalence between the inference of the PP model and modularity maximization [20,21], and dispel the notion that both methods are interchangeable in practice, by showing that the equivalence is in general inconsistent with maximum likelihood estimation, and discuss the fact that even if this were not the case, the parametric nature of that approach would not address the overfitting problem. Our approach is not more complicated to employ than modularity maximization, and can be used as a drop-in replacement for it and other quality functions in popular community detection heuristics [22], and we describe how it can be used with an unbiased merge-split MCMC algorithm [23] that can explore the entire posterior distribution of partitions.

Furthermore, we perform a comparison of the PP model with the more general SBM for a variety of empirical networks, allowing us to determine if and to what extent is assortativity the most salient characteristic of the large-scale network structure. We find a variety of outcomes, ranging from very similar to very different results obtained with both models, demonstrating that there are indeed many cases where searching exclusively for assortative structures can give a very misleading representation of the network.

This work is organized as follows. In Sec. II we present the planted partition model, and we revisit its equivalence with modularity. In Sec. III we describe our Bayesian approach at inferring the PP model, and introduce a more realistic nonuniform version of the model. In Sec. IV we analyze the method when applied to artificial networks with known community structures, and compare it to variations of modularity-based approaches, demonstrating that our method prevents overfitting. We also show that the Bayesian method does not suffer from the resolution limit of modularity maximization, and hence it does not generically underfit as well. In Sec. V we employ our method in a variety of empirical networks, and we compare the results with those obtained by the more general SBM, as well as modularity maximization. We demonstrate once more that modularity tends to either massively overfit or underfit, and when comparing to the SBM we can determine the true nature of the assortativity of empirical networks. We finalize in Sec. VI with a conclusion.

II. THE PLANTED PARTITION (PP) MODEL

The statistical inference approach to community detection is based on the definition of generative models that contain communities as part of their parameters. Before we consider the particular case of inferring assortative communities, it is useful at first to review the more general case of arbitrary mixing patterns between groups, as characterized by the

Poisson degree-corrected stochastic block model (DC-SBM) [5]. This formulation describes a network of N nodes that are divided into B groups, amounting to a partition $\mathbf{b} = \{b_i\}$, where $b_i \in [1, B]$ is the group membership of node i , and a specific multigraph A is generated with probability

$$P(A|\lambda, \theta, \mathbf{b}) = \prod_{i < j} e^{-\theta_i \theta_j \lambda_{b_i b_j}} \frac{(\theta_i \theta_j \lambda_{b_i b_j})^{A_{ij}}}{A_{ij}!} \times \prod_i e^{-\theta_i^2 \lambda_{b_i b_i} / 2} \frac{(\theta_i^2 \lambda_{b_i b_i} / 2)^{A_{ii} / 2}}{(A_{ii} / 2)!}, \quad (1)$$

where A_{ij} determines the number of edges between nodes i and j , and by convention A_{ii} corresponds to twice the number of self-loops incident on node i . Note that the parameters λ and θ always appear multiplying each other in the likelihood, so they are not uniquely identifiable, i.e., there are many choices that yield the same model. To uniquely specify the model, it is useful to introduce the quantity

$$\hat{\theta}_r = \sum_i \theta_i \delta_{b_i r}, \quad (2)$$

which in the model above can be set to arbitrary values, without sacrificing its generality. For example, if we set $\hat{\theta}_r = 1$, then we can interpret λ_{rs} as determining the expected number of edges between groups r and s (or twice this value for $r = s$), and θ_i is the relative probability with which a node i is selected to form an edge among those that belong to the same group. However, any other choice for $\hat{\theta}_r$ would be equally valid, with the only immaterial consequence being a different interpretation of the parameters. The maximum likelihood estimate of the above parameters is given by

$$\lambda_{rs}^* = \frac{e_{rs}}{\hat{\theta}_r \hat{\theta}_s}, \quad \theta_i^* = \frac{k_i}{e_{b_i}} \hat{\theta}_{b_i}, \quad (3)$$

where e_{rs} is the number of edges that go between groups r and s (or twice that for $r = s$), and $e_r = \sum_s e_{rs} = \sum_i k_i \delta_{b_i r}$ is the sum of degrees in group r . Indeed, the values of $\hat{\theta}_r$ cannot be uniquely obtained from these equations, since any value $\hat{\theta}_r > 0$ offers a valid solution, and more importantly, any choice disappears when we compute the probabilities $\lambda_{b_i b_j}^* \theta_i^* \theta_j^*$. This means we can choose these values independently of the inference procedure, with any particular choice functioning as a mere technical convention.¹

The degree-corrected planted partition (PP) model corresponds to the special case of the DC-SBM given by

$$\lambda_{rs} = \lambda_{\text{in}} \delta_{rs} + \lambda_{\text{out}} (1 - \delta_{rs}). \quad (4)$$

In this situation there are only two parameters that determine the placement of edges between groups, λ_{in} and λ_{out} , that set the expected number of edges inside and outside groups. A choice $\lambda_{\text{in}} \sum_r \hat{\theta}_r^2 / 2 > \lambda_{\text{out}} \sum_{r < s} \hat{\theta}_r \hat{\theta}_s$ corresponds to the assortative case, where edges connect mostly nodes of the same

¹Strictly speaking, this is no longer true in a Bayesian setting, where we are required to integrate the likelihood over the set of parameters. As shown in Ref. [24], choosing $\hat{\theta}_r$ has an effect on the parameter space and choice of priors, and ultimately changes the integrated marginal likelihood.

group. Therefore, this model captures what is more typically known as community structure in the proper sense, at least if the condition just mentioned is met. With this parametrization, the model likelihood of Eq. (1) becomes

$$P(\mathbf{A}|\lambda_{\text{in}}, \lambda_{\text{out}}, \boldsymbol{\theta}, \mathbf{b}) = \frac{e^{-\lambda_{\text{out}} \sum_{r<s} \hat{\theta}_r \hat{\theta}_s} \lambda_{\text{out}}^{e_{\text{out}}} e^{-\lambda_{\text{in}} \sum_r \hat{\theta}_r^2 / 2} \lambda_{\text{in}}^{e_{\text{in}}} \prod_i \theta_i^{k_i}}{\prod_{i<j} A_{ij}! \prod_i A_{ii}!!}, \quad (5)$$

where,

$$e_{\text{in}} = \frac{1}{2} \sum_{ij} A_{ij} \delta_{b_i, b_j}, \quad e_{\text{out}} = \sum_{i<j} A_{ij} (1 - \delta_{b_i, b_j}), \quad (6)$$

are the number of edges inside and outside groups, respectively. The maximum likelihood estimate of the parameters of the PP model are then given by

$$\lambda_{\text{in}}^* = \frac{2e_{\text{in}}}{\sum_r \hat{\theta}_r^2}, \quad (7)$$

$$\lambda_{\text{out}}^* = \frac{e_{\text{out}}}{\sum_{r<s} \hat{\theta}_r \hat{\theta}_s}, \quad (8)$$

$$\theta_i^* = k_i \left[\frac{2e_{\text{in}} \hat{\theta}_{b_i}}{\sum_r \hat{\theta}_r^2} + \frac{e_{\text{out}} \sum_{r \neq b_i} \hat{\theta}_r}{\sum_{r<s} \hat{\theta}_r \hat{\theta}_s} \right]^{-1}. \quad (9)$$

Looking at this result, we see that, unlike in the general DC-SBM, we no longer have full freedom to choose $\hat{\theta}_r$, since its maximum likelihood value must be a solution of the following system of nonlinear equations:

$$\hat{\theta}_r^* = e_r \left[\frac{2e_{\text{in}} \hat{\theta}_r^*}{\sum_s \hat{\theta}_s^{*2}} + \frac{e_{\text{out}} \sum_{s \neq r} \hat{\theta}_s^*}{\sum_{s<t} \hat{\theta}_s^* \hat{\theta}_t^*} \right]^{-1}. \quad (10)$$

We recover partial freedom to choose $\hat{\theta}_r$ in the special situation where all groups are uniform, with $e_r = 2E/B$, in which case $\hat{\theta}_r$ can take any value, as long as it is the same for every group, i.e., $\hat{\theta}_r = \hat{\theta}$. Note that we are, in fact, allowed to make an arbitrary prior assumption for the values of $\hat{\theta}_r$ before doing inference, making them imposed constraints that are part of our model specification. In this case Eq. (9) becomes simply

$$\theta_i^* = \frac{k_i \hat{\theta}_r}{e_{b_i}}. \quad (11)$$

We emphasize, however, that imposing this sort of constraint does jeopardize the degree correction of the model, since the expected degree of node i is given by

$$\langle k_i \rangle = \sum_j \theta_i \theta_j [\lambda_{\text{in}} \delta_{b_i, b_j} + \lambda_{\text{out}} (1 - \delta_{b_i, b_j})] \quad (12)$$

$$= \theta_i \left[\hat{\theta}_{b_i} \lambda_{\text{in}} + \lambda_{\text{out}} \sum_{r \neq b_i} \hat{\theta}_r \right]. \quad (13)$$

If we now substitute the maximum likelihood estimates of Eqs. (11), (7), and (8) in the above, then we obtain

$$\langle k_i \rangle = \frac{k_i \hat{\theta}_{b_i}}{e_r} \left[\frac{2e_{\text{in}} \hat{\theta}_{b_i}}{\sum_r \hat{\theta}_r^2} + \frac{e_{\text{out}} \sum_{r \neq b_i} \hat{\theta}_r}{\sum_{r<s} \hat{\theta}_r \hat{\theta}_s} \right]. \quad (14)$$

Therefore, the inferred model generates the observed degrees in expectation, i.e., $\langle k_i \rangle = k_i$, only if all groups have the

same sum of degrees $e_r = 2E/B$ and all imposed $\hat{\theta}_r$ are the same, or if we do not impose any constraints on $\hat{\theta}_r$, and use Eq. (9) instead. This means that we face a trade-off between consistent degree correction and ease of inference with the planted partition model, which is important to keep in mind as we consider the connection between statistical inference and modularity maximization, which we address in the following.

On the consistency between statistical inference and modularity maximization

As was shown in Refs. [20,21], it is possible to manipulate the likelihood of the PP model to expose a connection with modularity maximization [25]. We can rewrite the likelihood of Eq. (5) as

$$\begin{aligned} \ln P(\mathbf{A}|\lambda_{\text{in}}, \lambda_{\text{out}}, \boldsymbol{\theta}, \mathbf{b}) &= \frac{\mu}{2} \sum_{ij} (A_{ij} - \gamma \theta_i \theta_j) \delta_{b_i, b_j} + E \ln \lambda_{\text{out}} \\ &\quad - \frac{\lambda_{\text{out}}}{2} \left(\sum_i \theta_i \right)^2 + \sum_i k_i \ln \theta_i, \end{aligned} \quad (15)$$

up to unimportant additive constants, and where

$$\mu = \ln \lambda_{\text{in}} - \ln \lambda_{\text{out}}, \quad \gamma = \frac{\lambda_{\text{in}} - \lambda_{\text{out}}}{\ln \lambda_{\text{in}} - \ln \lambda_{\text{out}}}. \quad (16)$$

If we now enforce the following constraint as part of our model specification,

$$\hat{\theta}_r = \frac{e_r}{\sqrt{2E}}, \quad (17)$$

and replace the maximum likelihood estimate for $\theta_i^* = k_i / \sqrt{2E}$ obtained from Eq. (11) in the above, then we obtain

$$\begin{aligned} \ln P(\mathbf{A}|\lambda_{\text{in}}, \lambda_{\text{out}}, \boldsymbol{\theta} = \boldsymbol{\theta}^*, \mathbf{b}) &= \frac{\mu}{2} \sum_{ij} \left(A_{ij} - \gamma \frac{k_i k_j}{2E} \right) \delta_{b_i, b_j} + E (\ln \lambda_{\text{out}} - \lambda_{\text{out}}), \end{aligned} \quad (18)$$

again up to unimportant additive constants. Therefore, maximizing the above likelihood with respect to the partition \mathbf{b} alone, while keeping λ_{in} and λ_{out} constant, is equivalent to maximizing the generalized modularity [26],

$$Q(\mathbf{A}, \mathbf{b}) = \frac{1}{2E} \sum_{ij} \left(A_{ij} - \gamma \frac{k_i k_j}{2E} \right), \quad (19)$$

with γ playing the role of the resolution parameter. However, before concluding that modularity maximization and the inference of the PP model amount to the same task, we need to make the following crucial observations:

(1) The imposed constraint of Eq. (17) involves the knowledge of the sum of all *observed* degrees in each group $e_r = \sum_i k_i \delta_{b_i, r}$, which cannot be known before doing inference, and thus cannot be part of our model specification.² However,

²The same is true for the total number of edges E , but to a lesser degree since this is a global value that does not depend on the network partition (and hence amounts to a more innocuous

any other choice of $\hat{\theta}_r$ will not yield $\theta_i^* = k_i/\sqrt{2E}$ via maximum likelihood, which is required to recover modularity. Not imposing any prior constraint on $\hat{\theta}_r$ also does not yield the appropriate value via Eq. (9) in general, and will result in the necessary value of θ_i^* only for a particular uniform partition of the network where $e_r = 2E/B$ [in which case Eq. (17) holds without being imposed]. Therefore, the modularity of Eq. (19) is consistent with maximum likelihood of the PP model only in the very narrow case where all groups have the same sum of degrees.

(2) In addition, we must keep in mind that the values of λ_{in} and λ_{out} are never known *a priori* in empirically relevant settings. Therefore, we are required to infer them as well, together with the network partition. When employing maximum likelihood, the resulting values of μ and γ , as well as the second term of Eq. (18) all depend on the network partition, and are no longer just constants. In this situation, the partial equivalence with modularity maximization breaks down (even if $e_r = 2E/B$ as per point 1 above), as the functional form resulting from substituting Eqs. (7)–(9) and Eq. (16) into Eq. (18) makes the latter very different from Eq. (19). We emphasize that the scheme suggested in Ref. [21] of updating the value of γ according to Eq. (16) is insufficient to restore consistency since the contribution of the nonnegligible terms μ and $E(\ln \lambda_{\text{out}} - \lambda_{\text{out}})$ remain unaccounted for.

Based on the above, we see that the overall connection between the inference of the PP model and modularity maximization is in fact rather tenuous, and we should not expect in general to obtain the same results with both approaches. As explained in Ref. [20], the only statement that can be made is that there exists a particular choice of parameters λ_{in} , λ_{out} and θ such that maximizing modularity with the appropriate choice of γ and the PP likelihood conditioned on these parameters will yield the same partition. But since these parameters are unknown in practice, and are in general inconsistent with maximum likelihood estimation, the relevance of this equivalence is arguably limited.

Furthermore, as we discuss in Appendix B, it is easy to establish a formal equivalence between any community detection method and the statistical inference of a suitably chosen generative model. Therefore, the central issue is not whether this mapping exists, but if the procedure itself is consistent and behaves well. In fact, neither approach considered above, i.e., maximum likelihood inference of the PP model and modularity maximization, actually offers a robust method to uncover community structure in networks. As is well known, modularity maximization suffers from severe shortcomings, such as a strong tendency to identify spurious communities in fully random [11] and nonmodular [12–14] networks, a systematic failure to identify relatively small communities in large networks [19], it exhibits extreme degeneracy in key empirically relevant cases [27], and has been recently shown to systematically overfit on a broad range of empirical networks [28]. Any equivalence with the statistical inference of a parametric model would just mean that the latter also inherits these same

limitations. However, the full maximum likelihood inference approach of the PP model outlined above (which is not equivalent to modularity optimization) is not substantially superior. Even though it has a better justification, it does not really address any of the core problems of modularity. Most prominently, the inference approach is still prone to overfitting, with the uncontrolled detection of an ever increasing number of meaningless communities in fully random networks, as long as those increase the likelihood of the model. This happens in the same manner as fitting a polynomial to a set of points will also overfit, even if we use maximum likelihood, as long as we are allowed to increase the polynomial order without any constraint. We will demonstrate this problem with some simple examples, but before we do so we turn instead to a Bayesian approach, which includes the correct penalization of model complexity, and hence addresses the overfitting problem at its root, in a manner analogous to what has been done for the general SBM [15], as we describe in the next session.

III. BAYESIAN INFERENCE OF THE PLANTED PARTITION MODEL

Instead of maximum likelihood, a more formally correct approach to statistical inference is to sample or maximize from the posterior distribution of partitions [15],

$$P(\mathbf{b}|\mathbf{A}) = \frac{P(\mathbf{A}|\mathbf{b})P(\mathbf{b})}{P(\mathbf{A})}, \quad (20)$$

where

$$P(\mathbf{A}|\mathbf{b}) = \int P(\mathbf{A}|\lambda_{\text{in}}, \lambda_{\text{out}}, \theta, \mathbf{b}) \times P(\lambda_{\text{in}})P(\lambda_{\text{out}})P(\theta|\mathbf{b}) d\lambda_{\text{in}}d\lambda_{\text{out}}d\theta \quad (21)$$

is the marginal likelihood integrated over all model parameters, weighted according to their prior probabilities, and

$$P(\mathbf{b}) = \frac{\prod_r n_r!}{N!} \binom{N-1}{B(\mathbf{b})-1}^{-1} \frac{1}{N} \quad (22)$$

is the prior probability for partition \mathbf{b} , with $B(\mathbf{b})$ denoting the number of groups of \mathbf{b} (see Ref. [24] for a derivation). The remaining term $P(\mathbf{A}) = \sum_{\mathbf{b}} P(\mathbf{A}|\mathbf{b})P(\mathbf{b})$ is called the evidence, and it has the role of a normalization constant, and therefore will not play an important role in our calculations. To compute the integral of Eq. (21), we must specify our priors, which involves us also dealing with the model specification problem exposed earlier, with respect to the parameters θ . Here we will make the simple choice

$$\hat{\theta}_r = 1, \quad (23)$$

which allows the model parameters to have a straightforward interpretation, namely, the θ_i are the relative probabilities of selecting a node randomly from the group it belongs, and $B\lambda_{\text{in}}$ will determine twice the expected total number of edges inside communities, and $\binom{B}{2}\lambda_{\text{in}}$ the number of edges outside communities. (Remember that we are allowed to make any choice of $\hat{\theta}_r$ as part of our model specification, as long as the choice is made *a priori*, and does not depend on the data being modelled. As discussed previously, this choice does limit the accuracy of degree correction of the PP model

inconsistency), and replacing it by any other constant would only amount to a different effective value of γ in Eq. (19).

when performing maximum likelihood estimation. However, as we will see in a moment, this will not be a problem in the Bayesian formulation.) Now we can proceed in a manner that reflects our *a priori* indifference to any kind of model pattern, namely, we select a uniform prior for θ ,

$$P(\theta|\mathbf{b}) = \prod_r (n_r - 1)! \delta(\sum_i \theta_i \delta_{b_i, r} - 1), \quad (24)$$

and we choose maximum-entropy priors for the remaining parameters,

$$P(\lambda_{\text{in}}|\bar{\lambda}) = e^{-\lambda_{\text{in}}/(2\bar{\lambda})}/(2\bar{\lambda}), \quad (25)$$

$$P(\lambda_{\text{out}}|\bar{\lambda}) = e^{-\lambda_{\text{out}}/\bar{\lambda}}/\bar{\lambda}, \quad (26)$$

where $\bar{\lambda}$ is a hyperparameter that determines the expected total number of edges, with $\bar{\lambda} = 2\langle E \rangle/B^2$. Performing the integral of Eq. (21), we obtain

$$P(\mathbf{A}|\bar{\lambda}, \mathbf{b}) = \frac{e_{\text{in}}! e_{\text{out}}!}{2\bar{\lambda}^2 \left[\frac{B}{2} + \frac{1}{2\bar{\lambda}} \right]^{e_{\text{in}}+1} \left[\left(\frac{B}{2} \right) + \frac{1}{\bar{\lambda}} \right]^{e_{\text{out}}+1}} \times \prod_r \frac{(n_r - 1)!}{(e_r + n_r - 1)!} \times \frac{\prod_i k_i!}{\prod_{i<j} A_{ij}! \prod_i A_{ii}!}. \quad (27)$$

This marginal likelihood still depends on global hyperparameter $\bar{\lambda}$, which we can infer together with the other model parameters. However, there is an alternative that allows us to remove it altogether. We can reinterpret this marginal likelihood as an entirely equivalent model formulation given by

$$P(\mathbf{A}|\bar{\lambda}, \mathbf{b}) = P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b}) P(\mathbf{k}|\mathbf{e}, \mathbf{b}) P(\mathbf{e}|e_{\text{in}}, e_{\text{out}}, \mathbf{b}) \times P(e_{\text{in}}|\bar{\lambda}) P(e_{\text{out}}|\bar{\lambda}), \quad (28)$$

where

$$P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b}) = \frac{\prod_{r<s} e_{rs}! \prod_r e_{rr}!! \prod_i k_i!}{\prod_r e_r!! \prod_{i<j} A_{ij}! \prod_i A_{ii}!} \quad (29)$$

is the likelihood of the microcanonical DC-SBM [24], where e_{rs} specifies the exact number of edges between groups r and s (or twice that for $r = s$) and k_i is the exact degree of node i . We can recover Eq. (27) by making the following choice of priors,

$$P(\mathbf{e}|e_{\text{in}}, e_{\text{out}}, \mathbf{b}) = \frac{e_{\text{in}}!}{B^{e_{\text{in}}} \prod_r (e_{rr}/2)!} \times \frac{e_{\text{out}}!}{\left(\frac{B}{2} \right)^{e_{\text{out}}} \prod_{r<s} e_{rs}!}, \quad (30)$$

which is a product of uniform multinomial distributions for the diagonal and off-diagonal entries of the matrix e_{rs} , conditioned on the total sums e_{in} and e_{out} , respectively. For e_{in} and e_{out} themselves we use geometric distributions,

$$P(e_{\text{in}}|\bar{\lambda}, \mathbf{b}) = \frac{(B\bar{\lambda})^{e_{\text{in}}}}{(B\bar{\lambda} + 1)^{e_{\text{in}}+1}}, \quad (31)$$

$$P(e_{\text{out}}|\bar{\lambda}, \mathbf{b}) = \frac{\left(\left(\frac{B}{2} \right) \bar{\lambda} \right)^{e_{\text{out}}}}{\left(\left(\frac{B}{2} \right) \bar{\lambda} + 1 \right)^{e_{\text{out}}+1}}, \quad (32)$$

and finally for the degrees we choose uniform distributions inside each group [24],

$$P(\mathbf{k}|\mathbf{e}, \mathbf{b}) = \prod_r \frac{e_r!(n_r - 1)!}{(e_r + n_r - 1)!}. \quad (33)$$

Inserting these priors in Eq. (28) and rearranging leads to Eq. (27). Interestingly, and somewhat surprisingly, since this microcanonical model generates the exact degrees \mathbf{k} that are observed, we no longer have the same inconsistency as in the ‘‘canonical’’ model under maximum likelihood that we discussed earlier, where the inferred degrees were different from the observed, even though we have made use of the constraint $\hat{\theta}_r = 1$ in its derivation. We can therefore rest assured this model can accommodate arbitrary degree sequences. This equivalence also allows us to replace some of the priors of the microcanonical formulation by more convenient choices that make the approach fully nonparametric. In particular, for e_{in} and e_{out} we can use instead the following:

$$P(e_{\text{in}}, e_{\text{out}}|\mathbf{b}) = P(e_{\text{in}}, e_{\text{out}}|E, \mathbf{b}) P(E), \quad (34)$$

where

$$P(e_{\text{in}}, e_{\text{out}}|E, \mathbf{b}) = \left(\frac{1}{E + 1} \right)^{1-\delta_{B,1}} \quad (35)$$

is a uniform distribution of the E edges into two values (unless $B = 1$, where we must have $e_{\text{in}} = E$). The prior for the total number of edges $P(E)$ can now be chosen arbitrarily, as it will only amount to a unimportant constant in the marginal distribution, and hence vanish from the posterior. With this, we have a fully nonparametric marginal distribution

$$P(\mathbf{A}|\mathbf{b}) = P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b}) P(\mathbf{k}|\mathbf{e}, \mathbf{b}) \times P(\mathbf{e}|e_{\text{in}}, e_{\text{out}}|\mathbf{b}) P(e_{\text{in}}, e_{\text{out}}|E, \mathbf{b}) P(E) \quad (36)$$

that reads

$$P(\mathbf{A}|\mathbf{b}) = \frac{e_{\text{in}}! e_{\text{out}}!}{\left(\frac{B}{2} \right)^{e_{\text{in}}} \left(\frac{B}{2} \right)^{e_{\text{out}}} (E + 1)^{1-\delta_{B,1}}} \times \prod_r \frac{(n_r - 1)!}{(e_r + n_r - 1)!} \times \frac{\prod_i k_i!}{\prod_{i<j} A_{ij}! \prod_i A_{ii}!}. \quad (37)$$

This expression, together with the partition prior $P(\mathbf{b})$ of Eq. (22), are not much more difficult to compute than the modularity of Eq. (19). In fact, it is easy to see that if we consider the change in the posterior probability that is incurred if we move a node i from group r to group s , we need to compute only a few terms that depend on e_{in} , e_{out} , e_r , e_s , n_r , n_s , and B . To compute the change, we need only to inspect the neighborhood of the node, which takes time $O(k_i)$, independently of any other quantity, such as the number of groups. This is the same algorithmic complexity of computing changes in modularity, so the quantity $\ln P(\mathbf{A}, \mathbf{b})$ can be used as a drop-in replacement of the quality function in any modularity maximization algorithm,³ thereby addressing many existing fundamental limitations. In fact, we can understand in more detail why this approach prevents overfitting by exploiting a direct connection between Bayesian inference

³For reasons of numerical stability, it is better to work with the log-probability $\ln P(\mathbf{A}, \mathbf{b})$, such that products become sums, and we can also use tables of pre-computed log-factorial values to improve the speed of the computation.

and information theory. Namely we can write the negative joint log-likelihood as follows:

$$\Sigma = -\ln P(\mathbf{A}, \mathbf{b}) = -\ln P(\mathbf{A}, \mathbf{k}, \mathbf{e}, \mathbf{b}) \quad (38)$$

$$= -\ln P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b}) - \ln P(\mathbf{e}, \mathbf{k}, \mathbf{b}). \quad (39)$$

The quantity Σ is called the description length of the data [29], as it measures the amount of information required to describe the network \mathbf{A} when the parameters \mathbf{e} , \mathbf{k} , and \mathbf{b} are known, together with the information necessary to describe the parameters themselves. Therefore, the most likely partition of the network is the one that allow us to compress it the most. This means that this approach amounts to a formal implementation of Occam's razor, that favors the most parsimonious explanation for the data: As we increase the complexity of the model, by considering a larger number of communities, the first term $-\ln P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b})$ will tend to decrease, as the model becomes more constrained; however, the second term $-\ln P(\mathbf{e}, \mathbf{k}, \mathbf{b})$ will tend to increase, functioning as a penalty for more complex models. Since it is not possible to compress fully random data using any method, this approach cannot find communities in fully random networks. (This also explains why maximum likelihood overfits, since it omits the contribution of the second term of the description length, and hence there is no penalization for model complexity.) We will also show in Sec. IV A that this approach also does not suffer from the "resolution limit" underfitting problem present with modularity maximization. However, before we do so, we will first consider a small variation of the PP model that is slightly more realistic, and allows for a larger amount of heterogeneity in the community structure.

A. The nonuniform PP model

Even if we commit ourselves to search exclusively for assortative community structures, the particular formulation of the PP model considered previously seems needlessly restrictive. This is because it assumes that, if all θ_i are the same, then the expected number of edges inside communities is the same for every community, which is likely to be an inadequate assumption in a variety of empirical scenarios. We can relax this constraint by formulating instead a nonuniform version of the PP model, with

$$\lambda_{rs} = \lambda_r \delta_{rs} + \omega(1 - \delta_{rs}). \quad (40)$$

This parametrization allows for the expected number of edges inside communities to vary arbitrarily, via the parameters $\lambda = \{\lambda_r\}$ that can be different for every community. Given this formulation, we can essentially repeat the same calculations as before, as we show in Appendix A. In the end, we obtain a marginal likelihood given by

$$P(\mathbf{A}|\mathbf{b}) = \frac{e_{\text{out}}! \prod_r e_{rr}!!}{\binom{B}{2}^{e_{\text{out}}} (E+1)^{1-\delta_{B,1}}} \times \binom{B + e_{\text{in}} - 1}{e_{\text{in}}}^{-1} \times \prod_r \frac{(n_r - 1)!}{(e_r + n_r - 1)!} \times \frac{\prod_i k_i!}{\prod_{i < j} A_{ij}! \prod_i A_{ii}!!}. \quad (41)$$

This likelihood is very similar to the uniform planted partition model and is just as easy to compute, but it should work better when the communities are sufficiently heterogeneous.

B. Inference algorithm

The posterior distribution of Eq. (20) is not simple enough to allow us to sample directly from it, so we have to perform this indirectly using Markov chain Monte Carlo (MCMC). This is done by defining move proposals that are conditioned on the current partition \mathbf{b} ,

$$P(\mathbf{b}'|\mathbf{b}), \quad (42)$$

and accepting a new partition \mathbf{b}' sampled from this distribution according to the Metropolis-Hastings probability [30,31]

$$\min \left[1, \frac{P(\mathbf{b}'|\mathbf{A})P(\mathbf{b}|\mathbf{b}')}{P(\mathbf{b}|\mathbf{A})P(\mathbf{b}'|\mathbf{b})} \right], \quad (43)$$

otherwise we reject the move and remain at the previous partition \mathbf{b} . Note that the computation of the above ratio does not depend on the intractable normalization constant $P(\mathbf{A})$ of Eq. (20), since it cancels out in the computation. By iterating the above procedure sufficiently often, we are guaranteed to sample from the target distribution $P(\mathbf{b}|\mathbf{A})$ asymptotically, provided our proposals $P(\mathbf{b}'|\mathbf{b})$ are ergodic and aperiodic. However, the time required to reach the target distribution will depend on the quality of our proposals, which will determine the practical feasibility of the algorithm. In this work we use the merge-split proposals described in Ref. [23], which have been shown to work well in many cases, in particular when the number of groups tends to vary. The only modification we make of that algorithm is that when proposing the move of a single node i from its current group to group r , we do it according to the following probability,

$$P(r|i, \epsilon) = (1 - \epsilon) \frac{\sum_j A_{ij} \delta_{b_j, r}}{k_i} + \frac{\epsilon}{B}, \quad (44)$$

where B is the number of occupied groups. The parameter ϵ determines the probability with which we look at a random neighbor of node i to copy its group membership, otherwise we select a group at random. We require a value $\epsilon > 0$ to guarantee ergodicity, but otherwise any other value yields a valid algorithm (we have used $\epsilon = 1/2$ in our analysis, which provided good acceptance rates).

As we mentioned before, when using our model formulation, the likelihood ratio when changing the membership of node i , as well as the move proposal probability, can be computed in time $O(k_i)$, where k_i is its degree. This means that a single MCMC "sweep," where every node had a chance to be moved once, takes a time $O(N + E)$, where E is the number of edges, which is the best we can hope for this kind of problem. Therefore, we can use this algorithm to approach networks with a very large size. A reference implementation of this algorithm is freely available as part of the `graph-tool` library [32].

In some cases, we may seek to maximize from the posterior distributions instead of sampling from it. This is achieved via a simple modification of the above algorithm, where we replace the target distribution with $P(\mathbf{b}|\mathbf{A}) \rightarrow P(\mathbf{b}|\mathbf{A})^\beta$, where β is an inverse temperature parameter. If we increase $\beta \rightarrow \infty$

(preferably slowly, to avoid getting trapped in local optima), then we obtain a maximization algorithm. The merge-split MCMC often shows a good behavior when employed in this manner, as it can more easily escape local maxima that would trap alternative schemes, such as those based on the change of a single node at a time.

IV. ARTIFICIAL NETWORKS

Here we show how our approach behaves for artificial networks that have imposed community structure. We compare the inference of the PP model with the DC-SBM [24], as well as with variations of modularity. We focus on the overfitting problem, and the potential identification of nonexistent communities. We do so by sampling networks with $N = 10^5$ nodes and average degree $\langle k \rangle = 5$, and a specific number of equal-sized groups B , from the PP model defined above, with a choice of parameters given by $\theta_i = 1/B$, $\lambda_{\text{in}} = c\langle k \rangle N/B$, and $\lambda_{\text{out}} = (1 - c)\langle k \rangle N/[B(B - 1)]$, with $c = 1/B + (1 - 1/B)\epsilon$, such that if $\epsilon = 0$ we have fully random networks, and $\epsilon = 1$ we have perfectly assortative communities. For the inference of the PP model and the DC-SBM, we sample from the posterior distribution of Eq. (20), using the algorithm above. When using the modularity function, we sample from the target distribution

$$P(\mathbf{b}|\mathbf{A}) = \frac{e^{\beta Q(\mathbf{b}, \mathbf{A})}}{Z(\mathbf{A})}, \quad (45)$$

where $Z(\mathbf{A}) = \sum_{\mathbf{b}} e^{\beta Q(\mathbf{b}, \mathbf{A})}$, and $Q(\mathbf{b}, \mathbf{A})$ is given by Eq. (19). We choose $\beta = 2E\mu = 2E \ln(\lambda_{\text{in}}/\lambda_{\text{out}})$, such that if

$$\gamma = \gamma_{\text{true}} = \frac{\lambda_{\text{in}} - \lambda_{\text{out}}}{\ln \lambda_{\text{in}} - \ln \lambda_{\text{out}}}, \quad (46)$$

then the posterior will be proportional to the likelihood of the true underlying model, i.e., $P(\mathbf{b}|\mathbf{A}) \propto P(\mathbf{A}|\lambda_{\text{in}}, \lambda_{\text{out}}, \theta^*, \mathbf{b})$. We also compare with the results obtained with the maximum likelihood choice for γ ,

$$\gamma = \gamma_{\text{fit}} = \frac{\lambda_{\text{in}}^* - \lambda_{\text{out}}^*}{\ln \lambda_{\text{in}}^* - \ln \lambda_{\text{out}}^*}, \quad (47)$$

where $\lambda_{\text{in}}^* = Be_{\text{in}}/E$ and $\lambda_{\text{out}}^* = Be_{\text{out}}/[E(B - 1)]$ [assuming Eq. (17) holds]. Finally, we also compare with $\gamma = 1$, which corresponds to the original definition of modularity, still widely used in practice.

The results for the inferred number of groups can be seen in Fig. 1. The Bayesian inference of both versions of the PP model (uniform and nonuniform) as well as the DC-SBM yield identical results, always returning the true number of groups. All versions of the modularity-based approach overfit systematically, often finding a number of groups which is orders of magnitude wrong. The bad behavior of the case $\gamma = \gamma_{\text{true}}$ may seem surprising, since it corresponds to the true likelihood of the model, which one could expect to be ‘‘Bayes optimal,’’ in the sense that since it already includes the correct model parameters other than the partition itself, then any other approach would need to yield a strictly worse performance. However, this would only be true if the number of groups would also be set to its true value (rendering its inference moot), otherwise this choice of parameter is no longer optimal. The behavior with $\gamma = \gamma_{\text{fit}}$ is considerably

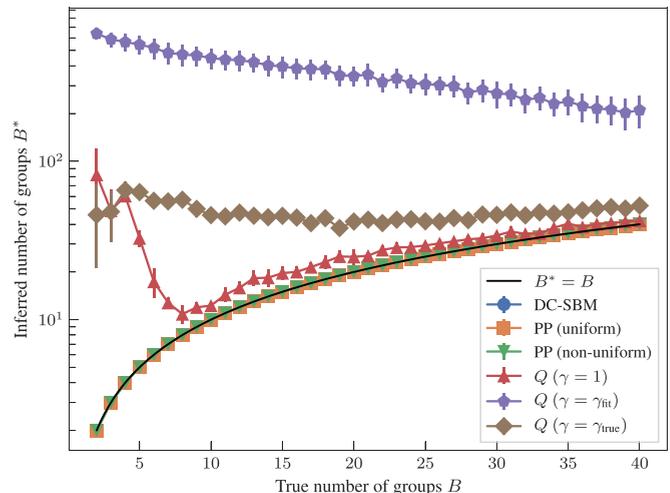


FIG. 1. Inferred number of groups B^* , averaged from the posterior distribution, as a function of the true number of groups B , according to the procedures shown in the legend, for networks sampled from the PP model with $N = 10^5$ nodes and average degree $\langle k \rangle = 5$, and mixing $\epsilon = 0.8$, as described in the text. The error bars show the standard deviation of the distribution (not its mean). The solid line shows the identity curve $B^* = B$. All results for the PP models (uniform and not) and DC-SBM are identical.

worse than all others, showing how maximum likelihood is inadequate for models with unconstrained degrees of freedom, as it trivially overfits. Interestingly, the choice $\gamma = 1$ seem to yield a better regularization than the alternatives, although the approach still systematically overfits, specially for a small number of planted communities. Results like this should give us pause when employing modularity to uncover communities in networks. Our Bayesian approach, on the other hand, behaves robustly, without requiring us to tune any parameter.

Bayesian inference of the PP model has no resolution limit

As was shown by Fortunato and Barthélemy [19], the method of modularity maximization possesses an intrinsic preferred scale for the size of the communities, which results in the so-called ‘‘resolution limit’’ that prevents relatively small modules to be uncovered, even if they have a very clear structure. Here we show that our Bayesian method does not suffer from the same problem.

We begin by briefly revisiting the result of Ref. [19], and we consider the structure of a maximally modular network, i.e., one that is constructed to maximize modularity. Following Ref. [19] we consider, without loss of generality, a network of N nodes and E edges that are divided into B equal-sized groups, each with $(E - B)/B$ internal edges, connecting nodes of the same group, and in total B edges connecting nodes of different communities, forming a circular ring between communities (the ring construction simply enforces that the network can in principle be connected, but plays no other role in the results). With this parametrization we have $e_r = 2E/B$, $e_{\text{in}} = E - B$, $e_{\text{out}} = B$. The number of groups itself is a free parameter, and it determines the overall

modularity, which from Eq. (19) we obtain

$$Q(\mathbf{A}, \mathbf{b}) = 1 - \frac{B}{2E} - \frac{\gamma}{B}. \quad (48)$$

We now seek to find the value $B = B^*$ that maximizes the above equation. Treating B as a continuous value for this purpose, and taking the derivative and setting it to zero, $dQ/dB = 0$, we obtain

$$B^* = \sqrt{2\gamma E}. \quad (49)$$

This result tells us that if we construct a network in the above way but with $B > B^*$, even if the groups themselves happen to be obvious assortative communities, e.g., large cliques connected by single edges, then these communities will be unintuitively merged together to achieve a larger modularity. The above result also reveals the role of the resolution parameter γ , which serves as the base of methods that attempt to determine its most appropriate value to counteract the limit in resolution [26,33–36].

We now turn to the PP model, to determine if the same natural scale emerges. We need to consider the value of $\ln P(\mathbf{A}, \mathbf{b})$ for the same construction above, and determine the value of $B = B^*$ that maximizes it. We will use the nonuniform PP model with Eqs. (41) and (22), although the final result is the same with the uniform version. We can obtain a simpler expression for the joint log-likelihood by assuming a large network with $N \gg 1$ and $B \gg 1$ (although we make no assumption on the value of B relative to N or E), so that we can use Stirling's formula $\ln x! = x \ln x - x + O(\log x)$, which yields

$$\begin{aligned} \ln P(\mathbf{A}, \mathbf{b}) &= (E - N) \ln B + 2(E - B) \ln(E - B) \\ &+ 2B \ln B + 2(N - B) \ln(N - B) + B \\ &- (2E + N - B) \ln(2E + N - B) + O(\ln B), \end{aligned} \quad (50)$$

up to unimportant additive constants. Taking the derivative and setting it to zero, we obtain the equation

$$\frac{d}{dB} \ln P(\mathbf{A}, \mathbf{b}) = \frac{E - N}{B} + \ln \frac{(N + 2E - B)B^2}{(E - B)^2(N - B)^2} = 0. \quad (51)$$

If we now assume we have a sparse graph with $E = \langle k \rangle N/2$, with $\langle k \rangle > 2$ being the average degree independent of N , then for $N \gg \langle k \rangle$ the solution of the above equation is

$$B^* = \frac{\langle k \rangle - 2}{2} \frac{N}{\ln N}. \quad (52)$$

This means that the Bayesian approach has a natural scale which prefers group sizes $N/B^* = O(\ln N)$, which is significantly smaller than the modularity scale $N/B^* = O(\sqrt{N})$. The scale of the Bayesian approach arises mostly due to the requirements of statistical evidence—we should partition a network only if its structure cannot be explained by a uniformly random placement of the edges. This explains also why for $\langle k \rangle < 2$ we obtain a value $B^* = 1$, since sparser networks inherently contain less information about the existing

community structure.⁴ As the size of the network increases, so does the possible ways of partitioning it and the number of parameters involved in the inference of the model, and as a consequence the required statistical evidence to support it also increases, and hence it becomes impossible to uncover groups smaller than $O(\ln N)$. However, this threshold grows so slowly that it can barely be compared with what exists for modularity maximization. We emphasize that this approach virtually eliminates the resolution limit without the introduction of a single parameter that needs to be tuned.

It is interesting to compare the value of B^* for the PP model with the same value obtained for the general SBM. As shown in Refs. [24,38], when using noninformative priors, the SBM has a resolution limit $B^* = O(\sqrt{N})$, which is similar to modularity, although it occurs for a completely different reason, namely the model depends on a matrix of parameters of size $O(B^2)$, which results in a penalty in the joint log-likelihood in the order of $O(B^2 \ln E)$, which becomes comparable to the likelihood when $B \sim \sqrt{N}$ for sparse networks. This limitation is lifted when the noninformative priors are replaced by a sequence of nested priors and hyperpriors, resulting in the nested SBM [24,39], which exhibits the natural scale $B^* = O(N/\ln N)$, similar to the PP model. However, the PP model achieves this high resolution already with simple noninformative priors, since it depends on a set of parameters which has total size $O(N + B)$ in the case of the nonuniform model, and $O(N)$ in the case of the uniform variant. This illustrates the usefulness of simpler models, which can achieve a higher performance than more general ones, if they happen to be a good description of the data.

V. EMPIRICAL NETWORKS

The existence of assortative community structure is often assumed to be a ubiquitous property of many kinds of networks across different domains. However, this kind of latent structure is not something that can be readily obtained from network data, and most methods that are used to detect it search exclusively for assortative structure, ignoring other patterns. Therefore, they cannot be used to rule out the existence of more fundamental nonassortative mixing patterns that are qualitatively different. A comparison between the assortative PP models that we consider here, together with more general SBM formulations allow us to address this comparison in a principled way, to understand how pervasive assortativity really is.

Here we compare the results obtained with the inference of PP model (both uniform and nonuniform versions) for a variety of empirical networks, together with those obtained using a Bayesian version of the DC-SBM [24], using both noninformative priors as well as nested hierarchical priors [39]. A powerful feature of the Bayesian inference approach is that it permits principled model selection, in the following way. Suppose we want to compare the community structure

⁴This does not mean that if $\langle k \rangle < 2$ the modules are completely undetectable, only that the posterior distribution has a maximum at $B = 1$. The actual detectability threshold is determined by averaging over the posterior distribution with the true parameters imposed [37].

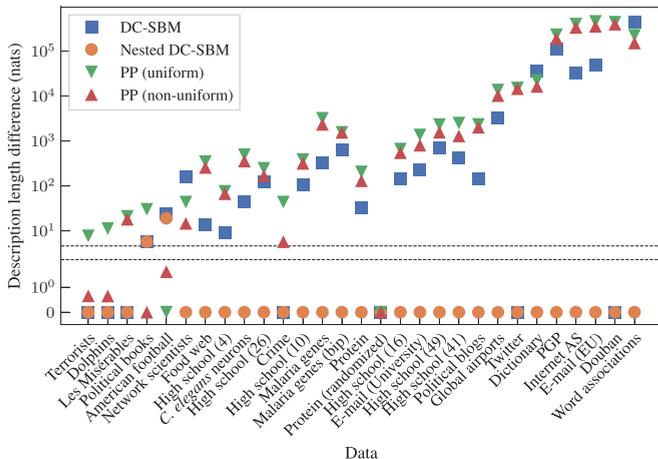


FIG. 2. Difference in the description length between the best fitting and the remaining models, as specified in the legend, for a selection of networks obtained from the KONECT repository [40]. The best fitting model always appears in the bottom. For reference, the values $\ln 10$ and $\ln 100$ are shown as dashed lines.

\mathbf{b}_1 found with model \mathcal{M}_1 with structure \mathbf{b}_2 found with model \mathcal{M}_2 , both for the same network \mathbf{A} . We can do this by comparing their posterior probability ratio

$$\Lambda = \frac{P(\mathbf{b}_1, \mathcal{M}_1 | \mathbf{A})}{P(\mathbf{b}_2, \mathcal{M}_2 | \mathbf{A})} \quad (53)$$

$$= \frac{P(\mathbf{A}, \mathbf{b}_1 | \mathcal{M}_1) P(\mathcal{M}_1)}{P(\mathbf{A}, \mathbf{b}_2 | \mathcal{M}_2) P(\mathcal{M}_2)}. \quad (54)$$

Therefore, if we have no prior preference toward any model, i.e., $P(\mathcal{M}_1) = P(\mathcal{M}_2)$, then this ratio will be given by the difference in the description length obtained with both models,

$$\Lambda = \exp(\Sigma_2 - \Sigma_1), \quad (55)$$

where $\Sigma_1 = -\ln P(\mathbf{A}, \mathbf{b}_1 | \mathcal{M}_1)$ and $\Sigma_2 = -\ln P(\mathbf{A}, \mathbf{b}_2 | \mathcal{M}_2)$. Hence, the most likely model is the one that offers the best compression for the data, and the difference in the compression itself yields the statistical significance of the preference toward the best model.

We performed the inference of the four models on a selection of 29 networks, representing different scientific domains, obtained from the KONECT repository [40]. In Fig. 2 we show the difference in description length obtained from the best fitting to the other models. Perhaps unsurprisingly, we find that the general DC-SBM provides a better fit for most networks, indicating that the strictly assortative structure of the PP model is insufficient to account for the observed networks. However, the PP model is selected as the best fitting model in a minority of the cases, and it is instructive to inspect those more closely. In Fig. 3 we show the communities uncovered using the PP model and the nested DC-SBM, for a network of games between American college football teams [41] and a network of co-purchases of books about American politics [42]. In both cases, PP and the nested DC-SBM find very similar partitions, but the PP model finds a slightly larger number of communities. The model selection criterion outlined above selects the PP model as the more

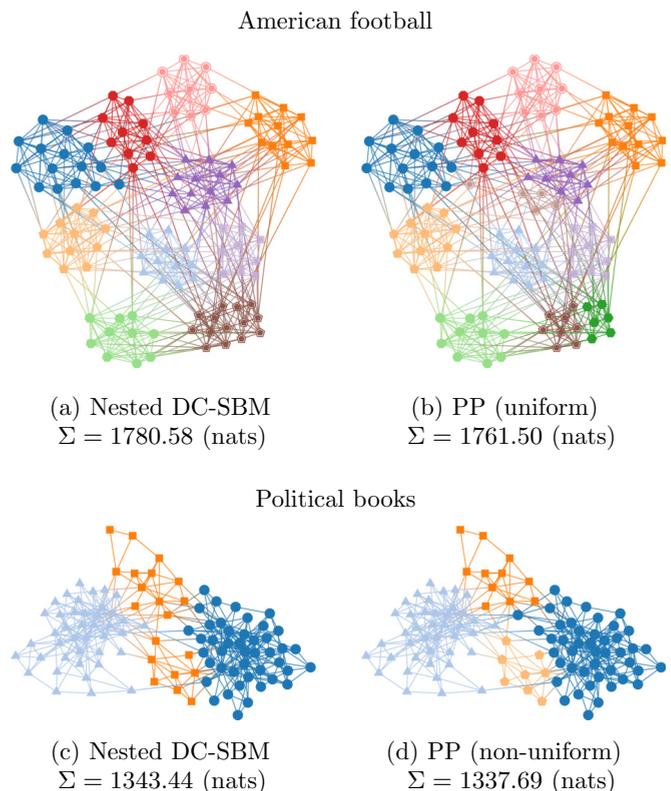


FIG. 3. Inferred community structure of a network of games between American college football teams [41] and a network of co-purchases of books about American politics [42], using both the PP model and the nested DC-SBM, as indicated in the legend, which also shows the description length of each fit.

plausible alternative due to the strong assortativity observed. The result found for the football network is particularly interesting, since it is a rare case where the *uniform* PP model is the one that gets selected. This is because the number of edges inside each community is indeed very similar for all of them, and the connections between the communities seem fairly random, exactly how the PP model prescribes. This highlights the robust character of our approach, which will not favor a more complicated model when it is unnecessary, and gives us confidence that when the PP model is not selected, it is indeed because it does not fully account for the actual structure observed in the network.

For other networks such as the associations between terrorists [45] and the social network between dolphins [46], even though the DC-SBM is strictly preferred, the difference between the nonuniform PP model is negligible, and therefore there is no sufficient evidence in the data to reliably distinguish between both models. For all other data, however, we find substantial evidence in favor of the more general DC-SBM. What is particularly interesting is that the DC-SBM is often preferred even when the uncovered structures are in fact very assortative. We give an example of this in Fig. 4, which shows the communities found with the nonuniform PP model and the nested DC-SBM for a social network of high school students [43]. Even though all communities found have a larger probability of forming internal than

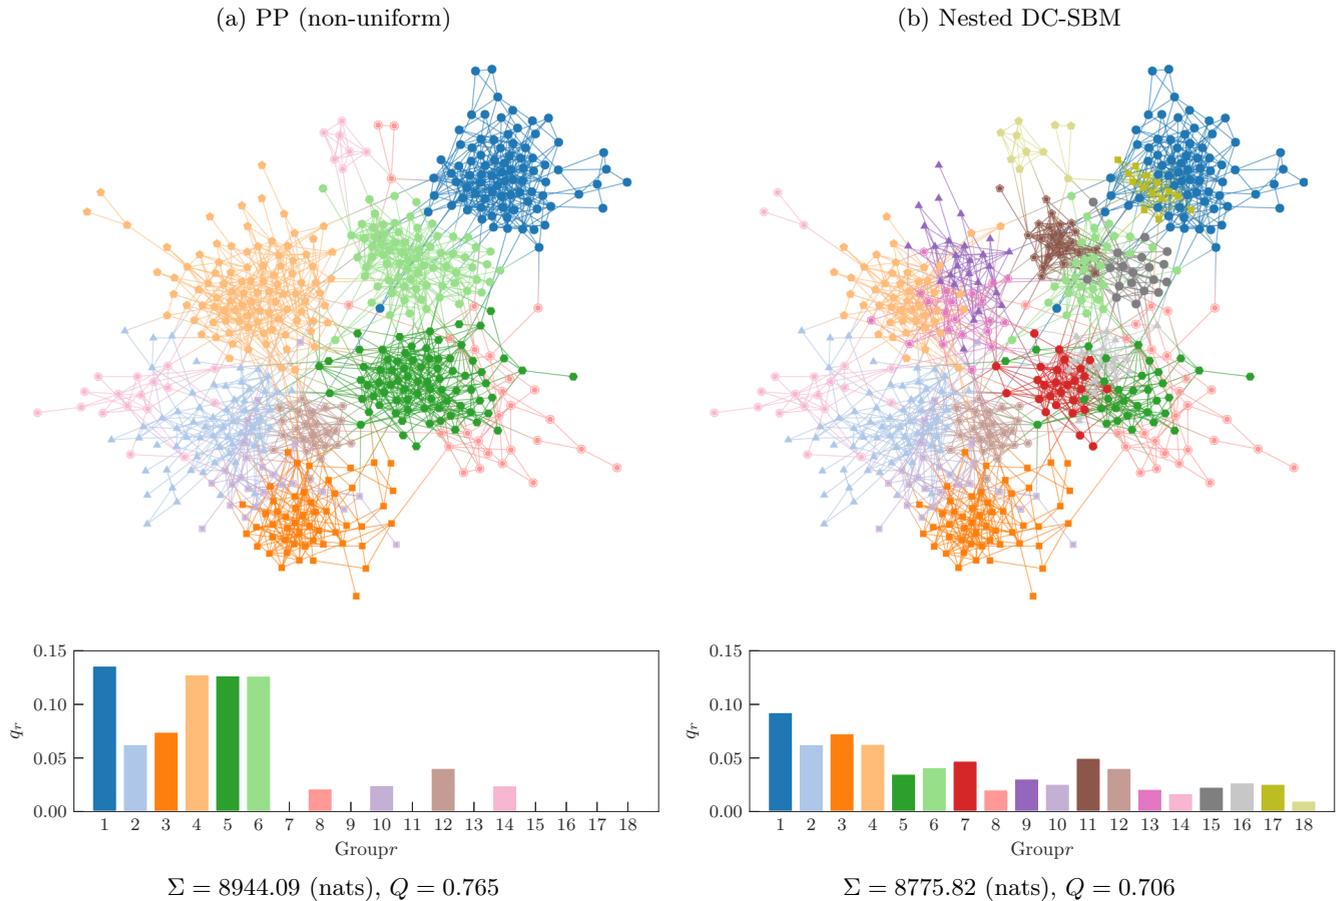


FIG. 4. Inferred community structure of a social network of high school students [43], using both the nonuniform PP model and the nested DC-SBM. The bottom panels show the community-wise modularity values $q_r = (e_{rr} - e_r^2/2E)/2E$, such that $Q = \sum_r q_r$. A value of $q_r > 0$ indicates that group r has a predominantly assortative contribution. The bottom legend shows the description length obtained with both models, as well as the value of modularity of the partitions. Both divisions have a normalized maximum overlap distance [44] of $d = 0.299$. The group colors are chosen to maximize the matching between both partitions, as described in Ref. [44], and the same colors are used in the bottom panels.

external connections, the ones found by the DC-SBM yield a larger plausibility. If we inspect it more closely, then we see that the divisions found by the DC-SBM amount largely to subdivisions of the ones found by the PP model. This can be explained by the DC-SBM using the preference of connections *between* the different communities as additional evidence for their existence, instead of merely their assortativity strength. This illustrates how a more general model like the DC-SBM can be more useful even when assortativity is a dominant but not unique pattern. (We note that the result found by the PP model has a slightly larger modularity, but this is not a very significant fact, given that modularity is in general largely decoupled from statistical significance.)

In Fig. 5 we show more details about the inferences obtained for all the networks, including the number of communities found, the modularity of the partitions, and the normalized maximum overlap distance [44] between the best fitting partition \mathbf{b} and the remaining partitions \mathbf{c} , defined as

$$d(\mathbf{b}, \mathbf{c}) = 1 - \frac{1}{N} \max_{\mu} \sum_i \delta_{b_i, \mu(c_i)}, \quad (56)$$

where $\mu(r)$ is a bijection between the group labels of \mathbf{b} and \mathbf{c} .

Overall, we observe a fair amount of variability in the comparisons between models for the different networks. Very often, the PP models yield a more conservative view of the networks, uncovering a smaller number of groups when compared to the DC-SBM, but there are also cases where the opposite is true. We also observe that although there are many cases where both the DC-SBM and PP models yield partitions with similar modularity, the overlap distance between partitions is very high, indicating that these networks admit a variety of divisions that have a similar overall level of assortativity (a good example of this is the high school network we considered in Fig. 4). Therefore, despite similar values of modularity found with the DC-SBM, the more general model rarely yields partitions that are very similar to the ones returned by any of the PP variants.

The values of modularity obtained with the best fitting model (which is most often the nested DC-SBM) are in some situations similar to what is found with the PP models, like for the high school social networks. However, for networks like the douban.com online social recommendation network [47], political blogs [48], Internet at the autonomous system level [39], and others, the modularity obtained with the best-fitting

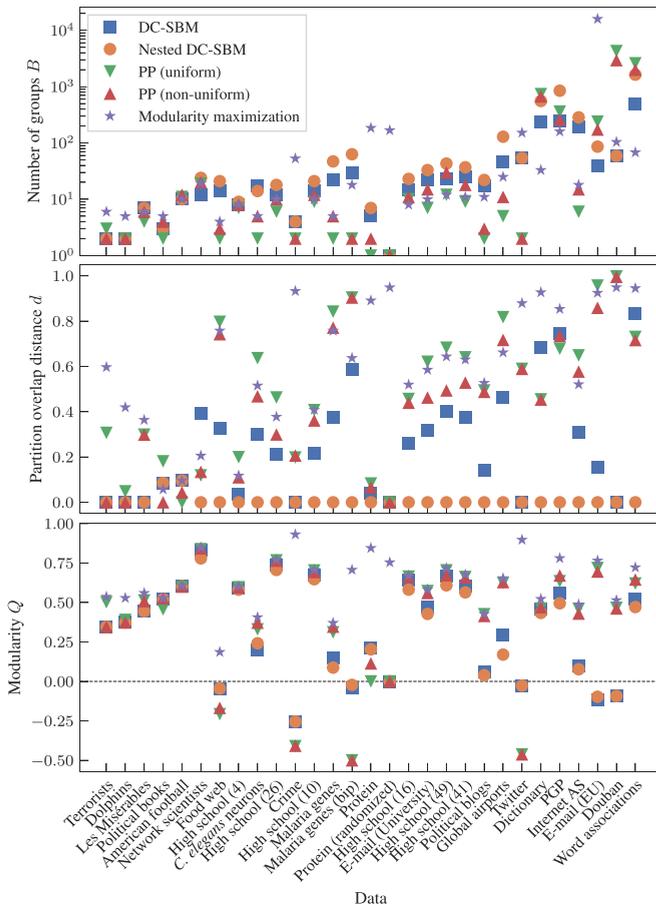


FIG. 5. More details about the results of Fig. 2, including the number of communities found with each method (top panel), the normalized maximum overlap distance [44] between the best fitting and the remaining partitions (middle panel), and the modularity of the partitions found (bottom panel).

DC-SBM is significantly smaller than with the PP models, indicating that assortativity is not the most fundamental pattern in these networks, and using a community detection method that searches exclusively for these patterns gives us a significantly biased view.

In Fig. 5 we also include results obtained with the method of modularity maximization. As it must be the case, this approach yields the division of the network with the highest values of modularity among the alternative ones considered. When we compare the results obtained with more statistically grounded approaches, we observe a rather erratic behavior. For some networks, such as word associations [50], modularity maximization yields a seemingly more conservative result with fewer groups, which could be an underfit potentially due to the resolution limit [19]. In other instances, like the E-mail network of an undisclosed European institution [51], protein-protein interactions [49], and bipartite person-crime associations [52], modularity maximization finds a number of communities that is multiple orders of magnitude larger than what is obtained with the inference methods, strongly indicating a massive overfit of these datasets. We illustrate this point in more detail by focusing on the protein-protein interaction network in Fig. 6. There we see that while mod-

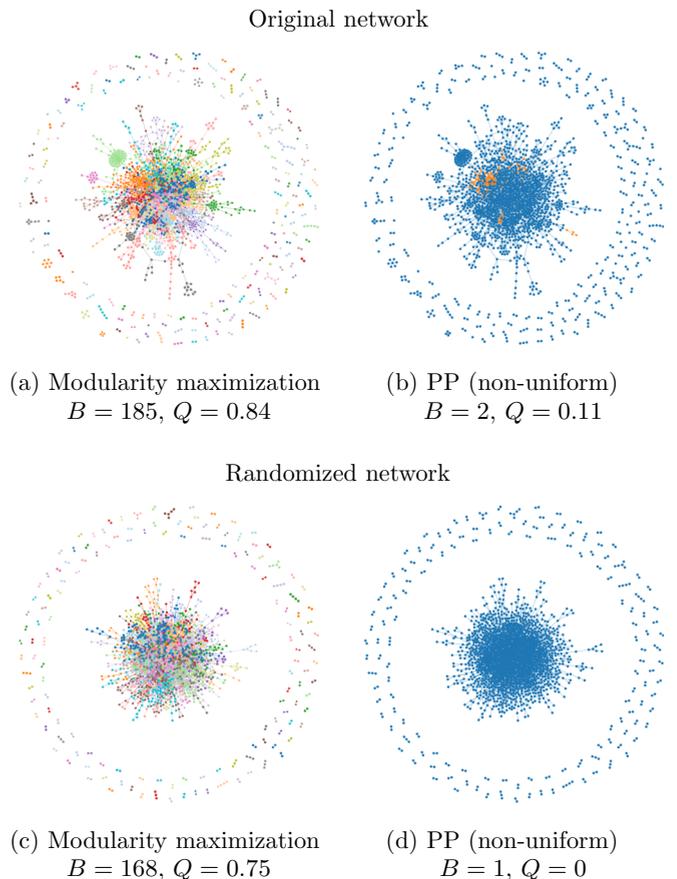


FIG. 6. Community structures found in a network of protein-protein interactions [49], using modularity maximization and Bayesian inference of the PP model, as indicated in the legends. The top panel shows the results for the original network, and in the bottom panel the results obtained for a randomized version with the same degree sequence. The legends show the number of groups and the value of modularity of the corresponding partitions.

ularity maximization finds over a hundred communities, the inference of the PP model finds only two, with one of them being relatively small. If we now consider a fully randomized version of the network, shown in Fig. 6 as well, then we see that modularity maximization still finds a very similar number of communities in it, with a high value of modularity, while the inference of the PP model finds, correctly, only a single group. This example clearly shows that while the structure of the original network is probably not completely random, most of it (including its disconnectedness) can be explained by its degree sequence alone, with no convincing evidence of community structure, and that the results obtained with modularity maximization are mostly spurious.

Our findings corroborate a recent analysis based on link prediction, carried over a large corpus of empirical networks, that showed that modularity maximization tends to systematically overfit [28]. Together with our results, this serves to illustrate that the tenuous connection with maximum likelihood of the PP model should not encourage practitioners to employ modularity maximization in the analysis of real networks, if they expect to be guided by statistical significance, or have any inherent guarantee against overfitting or underfitting.

VI. CONCLUSION

We have described how to perform a nonparametric Bayesian inference of the planted partition generative model, resulting in a principled community detection algorithm tailored for assortative structures. Our method separates structure from randomness, and does not find spurious communities in fully random networks. We also showed that it does not suffer from the resolution limit present in modularity-based methods and is capable of uncovering arbitrarily small communities, provided those are statistically significant, and without the tuning of any parameter. Our approach is based on the sampling or a maximization of a posterior probability function that is not much more complicated to implement than popular heuristics like modularity, and hence can be used as a drop-in replacement for it in a variety of algorithms, intrinsically providing better regularization for them.

We showed how our inference approach is amenable to statistical model selection, and we have compared our model variations, together with more general stochastic block models, on a variety of empirical networks. We discussed how this comparison allows us to determine the true assortativity of community structures, by removing a systemic bias that exists when only constrained methods are employed. We have shown that in many cases the assortativity of real networks is exaggerated when viewed through the lenses of community detection methods that search exclusively for assortative patterns, and how model selection can reveal more fundamental mixing patterns.

A recent investigation of a variety of community detection methods on a large corpus of empirical networks found that most of them tend to yield a number of communities that seems compatible with an overall $O(\sqrt{N})$ scaling [28], indicating that this might be a limitation that is present in a larger set of community detection algorithms. That analysis did not include statistical inference methods that are known not to have this particular limitation, like the nested SBM [24,39] and our Bayesian PP model, as we have demonstrated. Incorporating these methods into such large-scale comparisons would allow us to better understand what are the true fundamental limitations of the community detection task.

The inference approach is infinitely extensible, as it admits any conceivable generative model, and it provides a general platform for a meaningful comparison between them. It is easy to envision a more general comparison across network models that are tailored toward other kinds of specific mixing patterns, such as bipartiteness [53] and core-peripheries [54,55], as well as different classes of models such as those based on latent spaces [56,57]. A systematic comparison under such a framework would shed important light on the inherent trade-offs between more general and specific models, and how they relate to the various empirical domains.

ACKNOWLEDGMENTS

L.Z. is supported by the EPSRC Centre for Doctoral Training in Statistical Applied Mathematics at Bath (Grant No. EP/L015684/1) and by a scholarship from the Alan Turing Institute. This research made use of the Balena High Performance Computing (HPC) Service at the University of Bath.

APPENDIX A: MARGINAL LIKELIHOOD OF THE NONUNIFORM PP MODEL

The model likelihood of the nonuniform PP model described in the text can be written as

$$P(A|\lambda, \omega, \theta, \mathbf{b}) = \frac{e^{-\omega \sum_{r<s} \hat{\theta}_r \hat{\theta}_s} \omega^{e_{\text{out}}} \prod_r e^{-\lambda_r \hat{\theta}_r^2 / 2} \lambda_r^{e_{rr}/2} \prod_i \theta_i^{k_i}}{\prod_{i<j} A_{ij}! \prod_i A_{ii}!}. \quad (\text{A1})$$

Enforcing the constraint $\hat{\theta}_r = 1$, and using the noninformative priors

$$P(\lambda_r|\bar{\lambda}) = e^{-\lambda_r/(2\bar{\lambda})}/(2\bar{\lambda}),$$

$$P(\omega|\bar{\lambda}) = e^{-\omega/\bar{\lambda}}/\bar{\lambda},$$

$$P(\theta|\mathbf{b}) = \prod_r (n_r - 1)! \delta(\sum_r \theta_i \delta_{b_i, r} - 1),$$

we obtain the following marginal likelihood, after integrating over all parameters:

$$P(A|\lambda, \mathbf{b}) = \frac{e_{\text{out}}! \prod_r (e_{rr}/2)! \prod_i k_i!}{\left(\frac{1}{2} + \frac{1}{2\bar{\lambda}}\right)^{e_{\text{in}}+B} \left[\binom{B}{2} + \frac{1}{\bar{\lambda}}\right]^{e_{\text{out}}+1} \prod_{i<j} A_{ij}! \prod_i A_{ii}!}. \quad (\text{A2})$$

This likelihood is once more identical to a microcanonical model,

$$P(A|\lambda, \mathbf{b}) = P(A|\mathbf{e}, \mathbf{k}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\{e_{rr}\}, e_{\text{out}}, \mathbf{b}) \times P(\{e_{rr}\}|\bar{\lambda}, \mathbf{b})P(e_{\text{out}}|\bar{\lambda}, \mathbf{b})P(E), \quad (\text{A3})$$

where we now have the priors

$$P(\mathbf{e}|\{e_{rr}\}, e_{\text{out}}, \mathbf{b}) = \frac{e_{\text{out}}!}{\binom{B}{2}^{e_{\text{out}}} \prod_{r<s} e_{rs}!},$$

$$P(\{e_{rr}\}|\bar{\lambda}, \mathbf{b}) = \prod_r \frac{\bar{\lambda}^{e_{rr}}}{(\bar{\lambda} + 1)^{e_{rr}+1}},$$

$$P(e_{\text{out}}|\bar{\lambda}, \mathbf{b}) = \frac{[\bar{\lambda} \binom{B}{2}]^{e_{\text{out}}}}{[\bar{\lambda} \binom{B}{2} + 1]^{e_{\text{out}}+1}}.$$

We can improve this by replacing the last two equations with the following choice:

$$P(\{e_{rr}\}, e_{\text{out}}|\mathbf{b}, E) = P(\{e_{rr}\}|e_{\text{in}}, \mathbf{b})P(e_{\text{in}}|E, \mathbf{b}) \quad (\text{A4})$$

$$= \binom{B + e_{\text{in}} - 1}{e_{\text{in}}}^{-1} \left(\frac{1}{E + 1}\right)^{1 - \delta_{B,1}}, \quad (\text{A5})$$

which amounts to first choosing the value of e_{in} uniformly at random in the range $[0, E]$, and then likewise for the distribution of the diagonal values $\{e_{rr}\}$. Multiplying the above equations as

$$P(A|\mathbf{b}) = P(A|\mathbf{e}, \mathbf{k}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\{e_{rr}\}, e_{\text{out}}, \mathbf{b}) \times P(\{e_{rr}\}, e_{\text{out}}|E, \mathbf{b})P(E), \quad (\text{A6})$$

we arrive at Eq. (41) in the main text.

APPENDIX B: GENERAL EQUIVALENCES WITH STATISTICAL INFERENCE

Here we show that it is not difficult to establish a formal connection between any community detection method and statistical inference. Let us consider an arbitrary quality function,

$$W(\mathbf{A}, \mathbf{b}) \in \mathbb{R}, \quad (\text{B1})$$

which is used to perform community detection via the optimization

$$\mathbf{b}^* = \operatorname{argmax}_{\mathbf{b}} W(\mathbf{A}, \mathbf{b}). \quad (\text{B2})$$

We can retrofit any such method, and transform it into a statistical inference procedure by using $W(\mathbf{A}, \mathbf{b})$ as the Hamiltonian of an *ad hoc* generative model given by

$$P(\mathbf{A}|\mathbf{b}) = \frac{e^{W(\mathbf{A}, \mathbf{b})}}{Z(\mathbf{b})}, \quad (\text{B3})$$

with normalization given by

$$Z(\mathbf{b}) = \sum_{\mathbf{A}} e^{W(\mathbf{A}, \mathbf{b})}. \quad (\text{B4})$$

In general, performing a maximum likelihood estimation of this model will not be equivalent to the original optimization problem, due to the role of the normalization constant $Z(\mathbf{b})$. However, we can cast it as a Bayesian procedure to achieve a trivial equivalence, via the posterior distribution

$$P(\mathbf{b}|\mathbf{A}) = \frac{P(\mathbf{A}|\mathbf{b})P(\mathbf{b})}{P(\mathbf{A})}, \quad (\text{B5})$$

and by choosing the prior

$$P(\mathbf{b}) = \frac{Z(\mathbf{b})}{\Omega}, \quad (\text{B6})$$

with $\Omega = \sum_{\mathbf{b}} Z(\mathbf{b})$, and $P(\mathbf{A}) = \sum_{\mathbf{b}} e^{W(\mathbf{A}, \mathbf{b})} / \Omega$. Based on this, we recover easily the original optimization by maximizing from the posterior distribution,

$$\mathbf{b}^* = \operatorname{argmax}_{\mathbf{b}} \ln P(\mathbf{b}|\mathbf{A}) \quad (\text{B7})$$

$$= \operatorname{argmax}_{\mathbf{b}} W(\mathbf{A}, \mathbf{b}). \quad (\text{B8})$$

Therefore, finding a mere equivalence between any given community detection method and statistical inference, by itself, is not a very insightful exercise, as it can amount to little more than tautology. This also shows that not every inference procedure is any more meaningful or principled than using an arbitrary quality function. Instead, these features are contingent on the actual generative models used, which need to be properly justified, together with the choice of priors, and care should be taken to verify the consistency of the whole approach, which is not granted automatically in every case.

Despite the above, it should be mentioned that constructing a posterior distribution in the *ad hoc* way described above does have its uses. In particular, it allows us to formally define a distribution over all possible divisions of the network according to any given community detection method. As shown in Ref. [20], by characterizing this entire distribution, we have, to some extent, a mechanism to detect degeneracy and evaluate the statistical significance of the results, by seeking the consensus of a large fraction of the solutions. Nonetheless, this does not address the arbitrariness of the Hamiltonian chosen, and the ultimate interpretation of the results.

-
- [1] S. Fortunato, Community detection in graphs, *Phys. Rep.* **486**, 75 (2010).
 - [2] S. Fortunato and D. Hric, Community detection in networks: A user guide, *Phys. Rep.* **659**, 1 (2016).
 - [3] S. Wasserman and C. Anderson, Stochastic a posteriori block-models: Construction and assessment, *Soc. Netw.* **9**, 1 (1987).
 - [4] J. Reichardt and D. R. White, Role models for complex networks, *Eur. Phys. J. B* **60**, 217 (2007).
 - [5] B. Karrer and M. E. J. Newman, Stochastic blockmodels and community structure in networks, *Phys. Rev. E* **83**, 016107 (2011).
 - [6] M. E. J. Newman, Communities, modules and large-scale structure in networks, *Nat. Phys.* **8**, 25 (2012).
 - [7] M. A. Porter, J.-P. Onnela, and P. J. Mucha, *Communities in Networks*, Notices of the American Mathematical Society, Vol. 56, No. 9 (American Mathematical Society, Providence, RI, 2009), pp. 1082–1097, 1164–1166.
 - [8] R. Guimerá and L.-S. A. Nunes Amaral, Functional cartography of complex metabolic networks, *Nature* **433**, 895 (2005).
 - [9] E. Ravasz, A. L. Somera, D. A. Mongru, Z. N. Oltvai, and A.-L. Barabási, Hierarchical organization of modularity in metabolic networks, *Science* **297**, 1551 (2002).
 - [10] P. Holme, M. Huss, and H. Jeong, Subnetwork hierarchies of biochemical pathways, *Bioinformatics* **19**, 532 (2003).
 - [11] R. Guimerá, M. Sales-Pardo, and L.-S. A. N. Amaral, Modularity from fluctuations in random graphs and complex networks, *Phys. Rev. E* **70**, 025101(R) (2004).
 - [12] C. McDiarmid and F. Skerman, Modularity of regular and treelike graphs, *J. Complex Netw.* **6**, 596 (2018).
 - [13] James P. Bagrow, Communities and bottlenecks: Trees and tree-like networks have high modularity, *Phys. Rev. E* **85**, 066118 (2012).
 - [14] C. McDiarmid and F. Skerman, Modularity in random regular graphs and lattices, *Electronic Notes Discrete Math.* **43**, 431 (2013).
 - [15] Tiago P. Peixoto, Bayesian stochastic blockmodeling, in *Advances in Network Clustering and Blockmodeling* (John Wiley & Sons, New York, 2019), p. 289.
 - [16] T. N. Bui, S. Chaudhuri, F. T. Leighton, and M. Sipser, Graph bisection algorithms with good average case behavior, *Combinatorica* **7**, 171 (1987).
 - [17] M. E. Dyer and A. M. Frieze, The solution of some random NP-hard problems in polynomial expected time, *J. Algor.* **10**, 451 (1989).

- [18] A. Condon and R. M. Karp, Algorithms for graph partitioning on the planted partition model, *Random Struct. Algor.* **18**, 116 (2001).
- [19] S. Fortunato and M. Barthélemy, Resolution limit in community detection, *Proc. Natl. Acad. Sci. USA* **104**, 36 (2007).
- [20] P. Zhang and C. Moore, Scalable detection of statistically significant communities and hierarchies, using message passing for modularity, *Proc. Natl. Acad. Sci. USA* **111**, 18144 (2014).
- [21] M. E. J. Newman, Equivalence between modularity optimization and maximum likelihood methods for community detection, *Phys. Rev. E* **94**, 052315 (2016).
- [22] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, Fast unfolding of communities in large networks, *J. Stat. Mech.: Theory Exp.* (2008) P10008.
- [23] T. P. Peixoto, Merge-split Markov chain Monte Carlo for community detection, *Phys. Rev. E* **102**, 012305 (2020).
- [24] T. P. Peixoto, Nonparametric Bayesian inference of the microcanonical stochastic block model, *Phys. Rev. E* **95**, 012317 (2017).
- [25] M. E. J. Newman, Modularity and community structure in networks, *Proc. Natl. Acad. Sci. USA* **103**, 8577 (2006).
- [26] J. Reichardt and S. Bornholdt, Statistical mechanics of community detection, *Phys. Rev. E* **74**, 016110 (2006).
- [27] B. H. Good, Y.-A. de Montjoye, and A. Clauset, Performance of modularity maximization in practical contexts, *Phys. Rev. E* **81**, 046106 (2010).
- [28] A. Ghasemian, H. Hosseinmardi, and A. Clauset, Evaluating overfit and underfit in models of network community structure, *IEEE Trans. Knowl. Data Eng.* **32**, 1722 (2019).
- [29] P. D. Grünwald, *The Minimum Description Length Principle* (The MIT Press, Cambridge, MA, 2007).
- [30] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* **21**, 1087 (1953).
- [31] W. K. Hastings, Monte Carlo sampling methods using Markov chains and their applications, *Biometrika* **57**, 97 (1970).
- [32] T. P. Peixoto, The `graph-tool` python library, figshare (2014), available at <https://graph-tool.skewed.de>.
- [33] A. Arenas, A. Fernández, and S. Gómez, Analysis of the structure of complex networks at different resolution levels, *New J. Phys.* **10**, 053039 (2008).
- [34] P. Ronhovde and Z. Nussinov, Multiresolution community detection for megascale networks by information-based replica correlations, *Phys. Rev. E* **80**, 016109 (2009).
- [35] A. Lancichinetti and S. Fortunato, Limits of modularity maximization in community detection, *Phys. Rev. E* **84**, 066122 (2011).
- [36] P. Ronhovde and Z. Nussinov, Local multiresolution order in community detection, *J. Stat. Mech.: Theory Exp.* (2015) P01001.
- [37] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications, *Phys. Rev. E* **84**, 066106 (2011).
- [38] T. P. Peixoto, Parsimonious Module Inference in Large Networks, *Phys. Rev. Lett.* **110**, 148701 (2013).
- [39] T. P. Peixoto, Hierarchical Block Structures and High-Resolution Model Selection in Large Networks, *Phys. Rev. X* **4**, 011047 (2014).
- [40] J. Kunegis, KONECT: The Koblenz Network Collection, in *Proceedings of the 22nd International Conference on World Wide Web*, WWW '13 Companion (ACM, New York, NY, 2013) pp. 1343–1350.
- [41] M. Girvan and M. E. J. Newman, Community structure in social and biological networks, *Proc. Natl. Acad. Sci. USA* **99**, 7821 (2002).
- [42] V. Krebs, Political Books Network, unpublished, retrieved from <http://www-personal.umich.edu/~mejn/netdata/>.
- [43] K. M. Harris, C. T. Halpern, E. Whitsel, J. Hussey, J. Tabor, P. Entzel, and J. R. Udry, The national longitudinal study of adolescent to adult health: Research design, see <http://www.cpc.unc.edu/projects/addhealth/design> (accessed 9 April 2015) (2009).
- [44] T. P. Peixoto, Revealing consensus and dissensus between network partitions, [arXiv:2005.13977](https://arxiv.org/abs/2005.13977).
- [45] V. Krebs, Uncloaking terrorist networks, *First Monday* **7**, 4 (2002).
- [46] D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, and S. M. Dawson, The bottlenose dolphin community of Doubtful Sound features a large proportion of long-lasting associations, *Behav. Ecol. Sociobiol.* **54**, 396 (2003).
- [47] Z. Reza and L. Huan, Social computing data repository (2009), <http://socialcomputing.asu.edu/>.
- [48] L. A. Adamic and N. Glance, The political blogosphere and the 2004 U.S. election: Divided they blog, in *Proceedings of the 3rd International Workshop on Link Discovery*, LinkKDD'05 (ACM, New York, NY, 2005), pp. 36–43.
- [49] H. Jeong, S. P. Mason, A.-L. Barabási, and Z. N. Oltvai, Lethality and centrality in protein networks, *Nature* **411**, 41 (2001).
- [50] C. Fellbaum, *WordNet: An Electronic Lexical Database* (MIT Press, Cambridge, MA, 1998).
- [51] J. Leskovec, J. Kleinberg, and C. Faloutsos, Graph evolution: Densification and shrinking diameters, *ACM Trans. Knowl. Discov. Data* **1**, 2 (2007).
- [52] S. Decker, C. W. Kohfeld, R. Rosenfeld, and J. Sprague, *St. Louis Homicide Project: Local Responses to a National Problem* (University of Missouri-St. Louis, St. Louis, MO, 1991).
- [53] T.-C. Yen and D. B. Larremore, Community detection in bipartite networks with stochastic blockmodels, *Phys. Rev. E* **102**, 032309 (2020).
- [54] X. Zhang, T. Martin, and M. E. J. Newman, Identification of core-periphery structure in networks, *Phys. Rev. E* **91**, 032803 (2015).
- [55] R. J. Gallagher, J.-G. Young, and B. F. Welles, A clarified typology of core-periphery structure in networks, [arXiv:2005.10191](https://arxiv.org/abs/2005.10191).
- [56] M. Boguñá, F. Papadopoulos, and D. Krioukov, Sustaining the internet with hyperbolic mapping, *Nat. Commun.* **1**, 62 (2010).
- [57] M. E. J. Newman and T. P. Peixoto, Generalized Communities in Networks, *Phys. Rev. Lett.* **115**, 088701 (2015).