

Mapping flows on sparse networks with missing links

Jelena Smiljanić ^{1,2,*}, Daniel Edler ^{1,3,4} and Martin Rosvall ¹

¹*Integrated Science Lab, Department of Physics, Umeå University, SE-901 87 Umeå, Sweden*

²*Scientific Computing Laboratory, Center for the Study of Complex Systems, Institute of Physics Belgrade, University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia*

³*Gothenburg Global Biodiversity Centre, Box 461, SE-405 30 Gothenburg, Sweden*

⁴*Department of Biological and Environmental Sciences, University of Gothenburg, Carl Skottsbergs gata 22B, Gothenburg 41319, Sweden*



(Received 12 December 2019; revised 12 May 2020; accepted 9 June 2020; published 6 July 2020)

Unreliable network data can cause community-detection methods to overfit and highlight spurious structures with misleading information about the organization and function of complex systems. Here we show how to detect significant flow-based communities in sparse networks with missing links using the map equation. Since the map equation builds on Shannon entropy estimation, it assumes complete data such that analyzing undersampled networks can lead to overfitting. To overcome this problem, we incorporate a Bayesian approach with assumptions about network uncertainties into the map equation framework. Results in both synthetic and real-world networks show that the Bayesian estimate of the map equation provides a principled approach to revealing significant structures in undersampled networks.

DOI: [10.1103/PhysRevE.102.012302](https://doi.org/10.1103/PhysRevE.102.012302)

I. INTRODUCTION

Unraveling the modular organization of social and biological systems with interactions comprising measured movements of some entity such as people, money, or information requires reliable maps of network flows [1–5]. To find modular regularities in network flows, the map equation estimates a modular description length of the flows with information-theoretic measures. Optimizing the map equation with the search algorithm Infomap maximally compresses the modular description and detects significant flow-based communities when enough links are observed [2,6]. However, if too many links are missing, then the map equation may highlight spurious communities resulting from mere noise. While there are generative methods that can deal with uncertain network structures, including link-prediction algorithms [7–9] and network reconstruction approaches that often build on the stochastic block model [10–14], no method can reliably identify flow-based communities in networks with missing links.

The map equation estimates the modular description length of network flows with the Shannon entropy [15]. With missing data, the Shannon entropy underestimates the actual entropy of the complete data [16]. Consequently, when a network has many missing links, the map equation underestimates the actual description length of the complete network, capitalizes on

details in the observed network, and favors network partitions with many small communities. While higher model complexity can further compress the description length, the resulting communities become sensitive to network perturbations. Having more missing links further obscures the community structure and leads to higher sensitivity. Overfitting happens when the communities poorly compress the description length of the complete network or other samples of the complete network [17,18].

Underestimating the entropy in networks with missing links also causes problems for standard procedures that evaluate model-prediction performance, including cross-validation: When the modular description length depends on the number of observed links, it also depends on the number of cross-validation folds such that only balanced but wasteful equal-sized splits of a network into training and test networks give useful results.

To overcome these problems, we present two regularization methods based on entropy estimation for undersampled discrete data. First, we incorporate a Bayesian approach in the map equation framework [19] and derive a closed-form formula for the posterior mean of the map equation under the Dirichlet prior distribution of network flows. Second, to enable more effective cross-validation, we measure the modular description length of the training and test networks for a given partition using Grassberger entropy estimation [20].

We show that the Bayesian estimate of the map equation does not detect spurious communities in the undersampled regime in either synthetic or real-world networks. Also, compared with the degree-corrected stochastic block model [21,22], this approach gives solutions that are more robust to missing links in the analyzed networks. Moreover, with Grassberger entropy estimation, the modular description length becomes nearly independent of the amount of data: Instead of wasteful equal-sized splits, we can use most links in

*jelena.smiljanic@umu.se

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. Funded by Bibsam.

the training network to detect communities with Infomap and validate them using the remaining links in the test network. These two complementary solutions help us reduce overfitting and allow us to detect significant flow-based communities in networks with missing links.

II. MAPPING FLOWS ON COMPLETE NETWORKS

The map equation is an information-theoretic objective function for community detection based on the equivalence between data compression and identifying regularities in data. Building on this minimum description length principle, the map equation estimates the per-step theoretical lower limit of the average code word length needed to describe network flows with a modular description [2,6]. When the links themselves do not represent flows, we can model the network flows with a random walker traversing the network. The goal is to identify the network partition that maximally compresses the modular description, which, at the same time, best captures the modular regularities of the network flows.

For simplicity, here we consider modular descriptions with a two-level community hierarchy (for the multilevel map equation, see Appendix B). In a network with a well-defined community structure, the network flows stay for a relatively long time within communities. Therefore, to encode movements of the random walker between nodes with better compression, the map equation reuses short code words in modular codebooks instead of using unique code words for each node. For a uniquely decodable description, this approach requires an additional index codebook to encode transitions between communities.

The map equation measures the theoretical lower limit of the code length using the Shannon entropy [15]. For partition \mathbf{M} of nodes $\alpha = 1 \dots V$ in communities $i = 1 \dots m$, the map equation takes as input the probability that the random walker enters community i , $q_{i\curvearrowright}$, the probability to visit node α , p_α , and the probability to exit community i , $q_{i\curvearrowleft}$. With $p_i^\circ = q_{i\curvearrowleft} + \sum_{\alpha \in i} p_\alpha$ for the total use rate of module codebook i , the average per-step code length needed to describe random walker movements within community i is

$$H(\mathcal{P}_i) = -\frac{q_{i\curvearrowleft}}{p_i^\circ} \log_2 \frac{q_{i\curvearrowleft}}{p_i^\circ} - \sum_{\alpha \in i} \frac{p_\alpha}{p_i^\circ} \log_2 \frac{p_\alpha}{p_i^\circ}. \quad (1)$$

Similarly, the average per-step code length needed to describe random walker transitions between communities is

$$H(\mathcal{Q}) = -\sum_{i=1}^m \frac{q_{i\curvearrowleft}}{q_{\curvearrowleft}} \log_2 \frac{q_{i\curvearrowleft}}{q_{\curvearrowleft}}, \quad (2)$$

where $q_{\curvearrowleft} = \sum_{i=1}^m q_{i\curvearrowleft}$ is the total use rate of the index codebook. Therefore, we can express the map equation as the sum of the average code length of all codebooks weighted by their use rate:

$$L(\mathbf{M}) = q_{\curvearrowleft} H(\mathcal{Q}) + \sum_{i=1}^m p_i^\circ H(\mathcal{P}_i). \quad (3)$$

To identify the partition that minimizes the map equation, Infomap explores the space of possible solutions in a stochastic and greedy fashion.

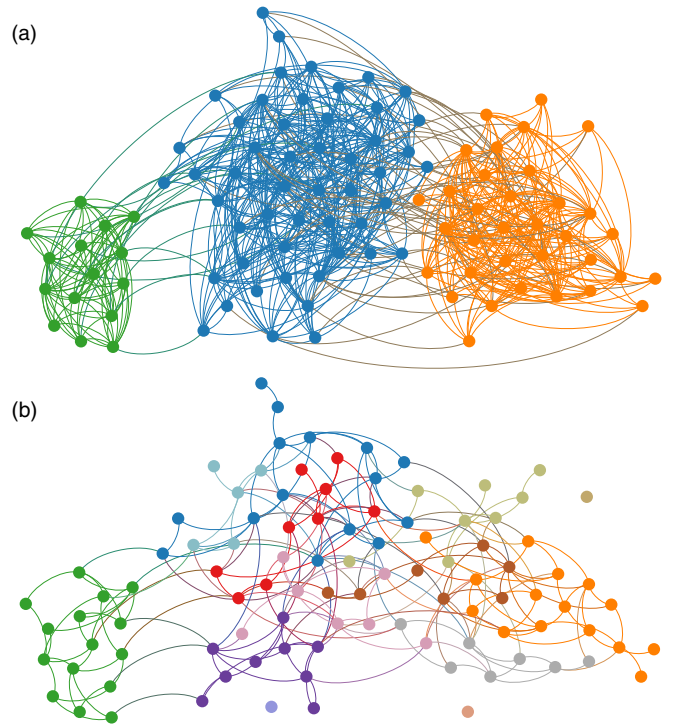


FIG. 1. Illustration of the overfitting problem in a small modular network. (a) The network has three communities. (b) When observing only a fraction of the links, the identified thirteen communities misrepresent the underlying network structure.

III. MAPPING FLOWS ON SPARSE NETWORKS WITH MISSING LINKS

Combined with Infomap, the map equation is an accurate method for community detection when complete network data are available [23]. However, empirical network data can lack data or contain measurement errors that cause missing or spurious links. When the map equation is applied to such unreliable network data, it may identify spurious communities with misleading information about the underlying network structure and function (Fig. 1).

We focus on missing links, a common problem in social and biological networks, that causes the sample estimates of the random walker's transition probabilities to lose precision. When plugging the estimates into the Shannon entropy, the obtained entropy estimator suffers from a negative bias and underestimates the entropy terms of the map equation [16]. Consequently, for the same partition \mathbf{M} , the description length decreases and the relative code length savings over the one-module solution, $l = 1 - L(\mathbf{M})/L(1)$, increases with the number of missing links (Fig. 2).

Worse yet, underestimating the index and module codebooks distorts their balance and shifts the optimal solution. The index codebook underrates the increase in between-module description length when using more communities, and the module codebooks overrate the within-module compression gain when using smaller communities. Also, stochastic fluctuations in missing links can lead the search algorithm off track because more undersampled regions attract community boundaries. Capitalizing on noise in this way underestimates




FIG. 2. Modular compression in sparse networks. (a) Modular code length for planted partitions after link removal. (b) Relative code length savings in networks for planted partitions. With standard entropy estimation, the average description length decreases as the number of missing links increases, and we cannot compare relative code length savings in networks with different densities. In contrast, Grassberger entropy estimation almost eliminates the code length's density dependency. For $r > 0.7$, the code length savings are negative for the Bayesian estimate of the map equation with prior $a_\alpha = \ln(V)$. By preferring the one-module solution over the planted partition in severely undersampled networks, the Bayesian estimate of the map equation avoids overfitting. For each r , we plot averages and variances over 100 network samplings of the synthetic network described in Sec. IV.

not only the codebooks but also, primarily, the transition rates between communities. As a result, the map equation favors more and smaller communities in sparse networks with missing links [9] (Fig. 1). This effect is evident when so many links are missing that actual communities become sparse or even form disconnected components. Then the map equation cannot detect the actual communities; instead it overfits and identifies spurious communities from mere noise in the network. To overcome overfitting, we incorporate a Bayesian estimate of the map equation.

A. Bayesian estimate of the map equation

Different methods have been proposed to address the problem of entropy underestimation [19,20,24–27]. Methods based on bias reduction cannot prevent overfitting of the map equation because they have a high variance in the undersampled regime [20,24,25] and cannot deal with the underestimation of the transition rates between communities. Instead, we use a Bayesian approach proposed by Wolpert and Wolf to estimate the function of probability distributions [19]. This method not only prevents overfitting to noisy structures better than other Bayesian estimators [26,27]; it also enables an analytical estimation of the map equation and a computationally efficient implementation in Infomap.

In general, we seek the Bayesian estimator \hat{f}_B of a function $f(\boldsymbol{\rho})$ that takes a discrete probability distribution $\boldsymbol{\rho} = (\rho_1, \rho_2, \dots, \rho_m)$ as input. When $\boldsymbol{\rho}$ is not given and we have only observations $\mathbf{n} = (n_1, n_2, \dots, n_m)$, with $\sum_{i=1}^m n_i = N$ sampled according to the distribution $\boldsymbol{\rho}$ ($E(n_i) = \rho_i N$), we must estimate $f(\boldsymbol{\rho})$ using the observed data \mathbf{n} . The Bayesian estimator for $f(\boldsymbol{\rho})$ is the posterior average,

$$\hat{f}_B(\mathbf{n}) = E[f|\mathbf{n}] = \int f(\boldsymbol{\rho})P(\boldsymbol{\rho}|\mathbf{n})d\boldsymbol{\rho}, \quad (4)$$

where $P(\boldsymbol{\rho}|\mathbf{n})$ is the posterior over the unknown distribution $\boldsymbol{\rho}$ given by Bayes' rule,

$$P(\boldsymbol{\rho}|\mathbf{n}) = \frac{P(\mathbf{n}|\boldsymbol{\rho})P(\boldsymbol{\rho})}{P(\mathbf{n})}. \quad (5)$$

To obtain $P(\boldsymbol{\rho}|\mathbf{n})$, we choose an appropriate prior probability distribution $P(\boldsymbol{\rho})$ and use the fact that the likelihood

$$P(\mathbf{n}|\boldsymbol{\rho}) = N! \prod_{i=1}^m \frac{\rho_i^{n_i}}{n_i!} \quad (6)$$

and the total probability of the data

$$P(\mathbf{n}) = \int d\boldsymbol{\rho}P(\mathbf{n}|\boldsymbol{\rho})P(\boldsymbol{\rho}). \quad (7)$$

Applied to the map equation, we seek the Bayesian estimator of $f(\boldsymbol{\rho}) = L(\mathbf{M})$. Assuming undirected and unweighted links, the transition rate estimates are [28]

$$p_\alpha = \frac{k_\alpha}{\sum_{\alpha=1}^V k_\alpha}, \quad (8)$$

$$q_{i\curvearrowright} = \frac{k_{i\curvearrowright}}{\sum_{\alpha=1}^V k_\alpha}, \quad (9)$$

$$q_{i\curvearrowleft} = \frac{k_{i\curvearrowleft}}{\sum_{\alpha=1}^V k_\alpha}, \quad (10)$$

where k_α is the degree of node α and $k_{i\curvearrowright} = k_{i\curvearrowleft}$ is the degree of module i , the number of links that connect nodes of module i with nodes of other modules j , $j \neq i$. However, when the information about links is incomplete, the actual values of node and module degrees can deviate from these estimates. Therefore, we must apply a probabilistic approach, or the map equation will overfit and exploit spurious network structures.

To develop a Bayesian treatment of the map equation, for a given partition M , we specify a prior distribution $P(p_\alpha, q_{i\curvearrowright}, q_{i\curvearrowleft})$ over the transition rates p_α , $q_{i\curvearrowright}$, and $q_{i\curvearrowleft}$. A convenient choice is the Dirichlet distribution, which has simple analytical properties and can be interpreted as a probability distribution over the multinomial distribution of the transition rates,

$$P(p_\alpha, q_{i\curvearrowright}, q_{i\curvearrowleft} | a_\alpha, a_{i\curvearrowright}, a_{i\curvearrowleft}) = \frac{\Gamma(a_1 + \dots + a_{m\curvearrowright})}{\Gamma(a_1) \dots \Gamma(a_{m\curvearrowright})} \prod_{\alpha=1}^V p_\alpha^{a_\alpha - 1} \prod_{i=1}^m q_{i\curvearrowright}^{a_{i\curvearrowright} - 1} \prod_{i=1}^m q_{i\curvearrowleft}^{a_{i\curvearrowleft} - 1}. \quad (11)$$

Here $\Gamma(x)$ is the gamma function and a_1, \dots, a_V , $a_{1\curvearrowright}, \dots, a_{m\curvearrowright}$, and $a_{1\curvearrowleft}, \dots, a_{m\curvearrowleft}$ are the parameters of the distribution. While $\sum_{\alpha=1}^V p_\alpha + \sum_{i=1}^m q_{i\curvearrowright} + \sum_{i=1}^m q_{i\curvearrowleft} \neq 1$,

we can use normalized transition rates because the map equation is scale invariant (see Appendix A).

We obtain the posterior distribution of the transition rates in Eq. (5) by multiplying the Dirichlet prior by the likelihood function and normalizing:

$$P(p_\alpha, q_{i\curvearrowright}, q_{i\curvearrowleft} | k_\alpha, k_{i\curvearrowright}, k_{i\curvearrowleft}, a_\alpha, a_{i\curvearrowright}, a_{i\curvearrowleft}) \propto \prod_{\alpha=1}^V p_\alpha^{k_\alpha + a_\alpha - 1} \prod_{i=1}^m q_{i\curvearrowright}^{k_{i\curvearrowright} + a_{i\curvearrowright} - 1} \prod_{i=1}^m q_{i\curvearrowleft}^{k_{i\curvearrowleft} + a_{i\curvearrowleft} - 1}. \quad (12)$$

By combining this distribution and the expanded form of the map equation,

$$\begin{aligned} L(\mathbf{M}) = & - \sum_{\alpha=1}^V p_\alpha \log_2(p_\alpha) - \sum_{i=1}^m q_{i\curvearrowright} \log_2(q_{i\curvearrowright}) \\ & + \sum_{i=1}^m \left(q_{i\curvearrowright} + \sum_{\alpha \in i} p_\alpha \right) \log_2 \left(q_{i\curvearrowright} + \sum_{\alpha \in i} p_\alpha \right) \\ & - \sum_{i=1}^m q_{i\curvearrowleft} \log_2(q_{i\curvearrowleft}) \\ & + \left(\sum_{i=1}^m q_{i\curvearrowleft} \right) \log_2 \left(\sum_{i=1}^m q_{i\curvearrowleft} \right), \end{aligned} \quad (13)$$

in Eq. (4), and integrating, we obtain a closed formula for the posterior average of the map equation,

$$\begin{aligned} \hat{L}_B(\mathbf{M}) = & \frac{1}{\ln(2)} \frac{1}{\sum_{\alpha=1}^V u_\alpha} \\ & \times \left[- \sum_{\alpha=1}^V u_\alpha \psi(u_\alpha + 1) - \sum_{i=1}^m u_{i\curvearrowright} \psi(u_{i\curvearrowright} + 1) \right. \\ & + \sum_{i=1}^m \left(u_{i\curvearrowright} + \sum_{\alpha \in i} u_\alpha \right) \psi \left(u_{i\curvearrowright} + \sum_{\alpha \in i} u_\alpha + 1 \right) \\ & - \sum_{i=1}^m u_{i\curvearrowleft} \psi(u_{i\curvearrowleft} + 1) \\ & \left. + \left(\sum_{i=1}^m u_{i\curvearrowleft} \right) \psi \left(\sum_{i=1}^m u_{i\curvearrowleft} + 1 \right) \right], \end{aligned} \quad (14)$$

where $u_x = k_x + a_x$ and $\psi(x)$ is the digamma function.

The parameters \mathbf{a} reflect our prior assumption of the link distribution in the network before we observed the network data. After seeing the data, we update our assumption by increasing the value of a_x by k_x and obtain the posterior distribution. For a sparse, undersampled network, therefore, the prior parameters \mathbf{a} dominate the posterior link distribution. Conversely, as the network density increases, the posterior distribution becomes sharply peaked and the network data dominate the posterior link distribution. Proper selection of prior parameters \mathbf{a} is important for good performance.

We consider as an uninformative prior an Erdős-Rényi network with V nodes, where each pair of nodes is connected with some constant probability p [29]. The average degree is $\langle k \rangle = pV$ and sets the prior parameters to $a_\alpha = \langle k \rangle$ and $a_{i\curvearrowright} = a_{i\curvearrowleft} = V_i(V - V_i) \binom{\langle k \rangle}{V-1}$, where V_i is the number of nodes

in module i . We aim to choose the average degree $\langle k \rangle$ such that the prior prevents the map equation from overfitting in the undersampled network, but also enables the map equation to detect well-formed communities. Since the random network experiences a phase transition from disconnected to connected at $\langle k \rangle = \ln(V)$ [29], for $\langle k \rangle \ll \ln(V)$ the random network has isolated components and the prior cannot prevent overfitting, while for $\langle k \rangle \gg \ln(V)$ well-formed communities can merge such that the map equation underfits. At the phase transition between these extremes, $a \sim \ln(V)$ forms a principled prior.

Because there are no modular regularities in an Erdős-Rényi network, this choice of prior induces positive bias in the code length estimation [Fig. 2(a)]. When observing fewer links in a network, the prior network influences the posterior link distribution more such that the code length increases for the planted partition. Eventually, for severely undersampled networks, the Bayesian estimate of the map equation prefers the one-module solution and thereby avoids overfitting [Fig. 2(b)].

This Bayesian estimate of the map equation extends to weighted networks where complete information about link weights is missing. If the link weights represent flows such that no flow modeling is necessary, then the method also works for directed networks.

We have implemented the Bayesian estimate of the map equation in Infomap, available for anyone to use [30]. While we restrict our paper to the two-level formulation of the map equation for the sake of simplicity, the code also handles the Bayesian estimate of the multilevel map equation (see Appendix B).

B. The map equation with Grassberger entropy estimation

An informative comparison between the standard map equation and a map equation with corrected entropy terms must take into account the structural properties of the detected communities. When possible, we can compare detected communities with planted communities; however, this approach does not work for real networks without known communities. To test for under- or overfitting in any network, we use cross-validation.

We first split the network data into training and test sets and apply Infomap to identify the partition that maximally compresses the description length of the training network. If Infomap successfully recovers a significant partition of the training network, then the partition with maximal modular code length savings over the one-level code length will also successfully compress the description length of the test network. The opposite happens when there is not enough evidence in the data. Then Infomap overfits and detects a partition in the training network without code length savings in the test network. Thus, if Infomap detects a significant partition \mathbf{M} without overfitting, the relative code length savings in the test network should be positive, $l^{\text{test}} = 1 - L_{\text{test}}(\mathbf{M})/L_{\text{test}}(1) > 0$ and close to the relative code length savings of the training network, $l^{\text{test}} \sim l^{\text{train}}$. Conversely, if Infomap overfits we expect $l^{\text{test}} < 0$.

However, the fact that the description length and the relative code length savings vary with the fraction of

observed links limits the choice of training and test networks (Fig. 2). Only with equal-sized training and test networks will the standard map equation underestimate their true description lengths to the same degree. But since equal splits waste half of the links on the test network, the training network of already sparse networks will be severely undersampled and possibly below the detectability limit. To reduce the description length's dependency on the fraction of observed links and enable effective cross-validation, we incorporate Grassberger entropy estimation [20] into the map equation.

For effective cross-validation, Grassberger entropy estimation enables the use of most of the links in the training network. We construct a test network by randomly removing a fraction r of links from the network. The remaining links form a training network. With E for the total number of links in the network and k_α for the degree of node α , the probability that k'_α links of node α remain in the training network after removing $E - E' = rE$ links follows the hypergeometric distribution:

$$P(k'_\alpha) = \frac{\binom{k_\alpha}{k'_\alpha} \binom{E-k_\alpha}{E-k'_\alpha}}{\binom{E}{E'}}. \quad (15)$$

If E , E' , and k_α are sufficiently large, then the hypergeometric distribution converges toward the Poisson distribution,

$$P(k'_\alpha) = \frac{\lambda^{k'_\alpha}}{k'_\alpha!} e^{-\lambda}, \quad (16)$$

where the parameter $\lambda = \frac{E'k_\alpha}{E} = (1-r)k_\alpha$ such that $\langle k'_\alpha \rangle = (1-r)k_\alpha$.

For a given incomplete set of observations (n_1, n_2, \dots, n_m) , Grassberger entropy estimation assumes that they come from Poisson distributions with mean values (z_1, z_2, \dots, z_m) and aims to construct a function $\phi(n)$ that minimizes the error $|z_i \ln(z_i) - E[n_i \phi(n_i)]|$ across all values of z_i [20]. The solution that minimizes the error is a recursive function $\phi(n) = G_n$ defined as

$$\begin{aligned} G_1 &= -\gamma - \ln(2), & G_2 &= 2 - \gamma - \ln(2), \\ G_{2n+1} &= G_{2n}, & G_{2n+2} &= G_{2n} + \frac{2}{2n+1}, \end{aligned} \quad (17)$$

where γ is Euler's constant [20].

While we cannot use Grassberger entropy estimation for weighted or directed networks, where visit rates correspond to the PageRank of the nodes [6], it does work for unweighted and undirected networks, where node visit and module transition rate estimates are given by link counts, Eqs. (8)–(10). Assuming incomplete observations, we can incorporate Grassberger entropy estimation into the map equation such that Eq. (13) takes the form

$$\begin{aligned} \hat{L}_G(M) &= \frac{1}{\ln(2)} \frac{1}{\sum_{\alpha=1}^V k_\alpha} \\ &\times \left[-\sum_{\alpha=1}^V k_\alpha G_{k_\alpha} - \sum_{i=1}^m k_{i \curvearrowright} G_{k_{i \curvearrowright}} \right] \end{aligned}$$

$$\begin{aligned} &+ \sum_{i=1}^m \left(k_{i \curvearrowright} + \sum_{\alpha \in i} k_\alpha \right) G_{k_{i \curvearrowright} + \sum_{\alpha \in i} k_\alpha} \\ &- \sum_{i=1}^m k_{i \curvearrowright} G_{k_{i \curvearrowright}} + \left(\sum_{i=1}^m k_{i \curvearrowright} \right) G_{\sum_{i=1}^m k_{i \curvearrowright}} \end{aligned} \quad (18)$$

Grassberger entropy estimation also works for the multilevel formulation of the map equation [31].

Grassberger entropy estimation has high variance and low bias [32]. Due to its high variance in the undersampled regime (Fig. 2) and its lack of prior that can deal with underestimating the transition rates between communities, the map equation with Grassberger entropy estimation paired with Infomap does not perform better than the standard map equation on sparse networks with missing links. However, thanks to its low bias, the map equation with Grassberger entropy estimation applied to cross-validation with averaged code length over several network samplings can dramatically reduce the code length dependency on network density [Fig. 2(a)]. Also, for planted partitions, the average relative code length savings is practically independent of network density [Fig. 2(b)]. Consequently, we can use most links in the training network to reliably detect communities with Infomap.

IV. RESULTS AND DISCUSSION

We first analyze a synthetic network with planted community structure and a real-world Jazz collaboration network [33]. We generate the synthetic network with the Lancichinetti-Fortunato-Radicchi (LFR) method [34]. It has $V = 1000$ nodes, average node degree $\langle k \rangle = 16$, and nodes partitioned into $M = 35$ communities. The mixing parameter $\mu = 0.3$ is the probability that a randomly chosen link will connect nodes from different communities. In the Jazz collaboration network, each node represents a band and two nodes are connected if there is at least one musician who has played in both bands. For this network with 198 nodes and 2742 links, there is no information about ground-truth communities and no consensus about an optimal community partition [35,36]. To generate sparse networks with missing links, we randomly remove a fraction r of links from the networks, and average the results for each value of r over 100 samplings.

Using these two networks, we compare the performance of the standard map equation, the Bayesian estimate of the map equation with different values of Dirichlet prior parameter a_α , and the degree-corrected stochastic block model [21,22]. We are interested in the number of communities, the partition similarities measured with the adjusted mutual information (AMI), and the predictive accuracy with cross-validation. Since the map equation and the degree-corrected stochastic block model use stochastic search algorithms to detect communities, we average the results over ten searches for each of the 100 network samplings.

We analyze the Bayesian approach for prior $a \sim \ln(V)$. For the node degree, therefore, we use $a_\alpha = C \ln(V)$, where $\alpha = 1 \dots V$ and C is a constant that we need to specify. For the module degree, we use $a_{i \curvearrowright} = a_{i \curvearrowleft} = v_i C \ln(V)$, where $v_i = V_i \frac{V-V_i}{V-1}$ for $i = 1 \dots M$ and V_i is the number of nodes in module i .




FIG. 3. Mean number of communities obtained by the standard map equation, the Bayesian estimate of the map equation with different values of Dirichlet prior parameter a , and the degree-corrected stochastic block model (DC-SBM). The Bayesian estimate of the map equation with prior $a_\alpha = \ln(V)$ provides the best solution: when sufficient network data are available it distinguishes significant communities from mere noise, while in the undersampled regime it detects no community structure. Results are averaged over 100 network samplings and ten algorithm searches. The standard error of the mean is never higher than 0.58.

A. Number of communities

Applied to the synthetic network, the standard map equation favors the planted partition until we remove more than approximately 55% of the links [Fig. 3(a)]. As we remove more links, the network also becomes sparse within communities. In the undersampled regime below the detectability limit where it is not possible to recover the planted partition, the map equation overfits to random fluctuations and favors more, smaller communities. The Bayesian estimate of the map equation behaves differently. For $C = 0.5$, the random prior network is weakly connected and cannot prevent overfitting when we remove 70–95% of the links. In contrast, for $C = 2$, the random prior network is densely connected and hides the communities in the noise induced by the prior such that the Bayesian estimate of the map equation underfits even when sufficient network data are available. In between, at the critical point where the random prior network becomes connected, the prior constant $C = 1$ balances over- and underfitting and prevents the detection of spurious communities. Moreover, the amount of noise that this prior network induces in the original network is so low that it does not wash out any significant community structure. While prior parameter C between 0.5 and 1 performs best for some analyzed networks, $C = 1$ remains a robust choice in general (Appendix C).

The degree-corrected stochastic block model detects the planted partition until we remove more than 40% of the links from the synthetic network. Compared to the Bayesian estimate of the map equation with the prior constant $C = 1$, the degree-corrected stochastic block model starts to underfit the planted partition earlier. For $r > 40\%$, the number of communities decreases continuously and when $r > 80\%$, the degree-corrected stochastic block model detects no community structure.

Similar behaviors appear accentuated when we apply the methods to the real-world Jazz collaboration network [Fig. 3(b)]. For the standard map equation, the number of detected communities increases with the number of missing links, whereas the degree-corrected stochastic block model

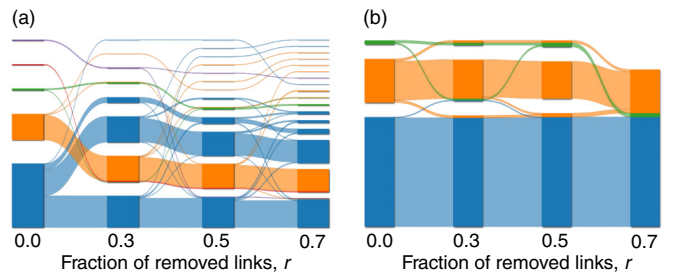


FIG. 4. Alluvial diagrams of the Jazz collaboration network show changes in community structure with missing links for (a) the standard map equation and (b) the Bayesian estimate of the map equation with prior $a_\alpha = \ln(V)$. Compared to the standard map equation, the communities detected using the Bayesian estimate of the map equation are more robust to missing links.

shows the opposite trend. Unlike when applied to the synthetic network, the various map equation variants already favor different partitions before removing any links. The Bayesian estimate of the map equation detects fewer communities than the standard map equation, and its performance depends on the choice of the prior. For $C = 0.5$, the average number of communities is relatively stable when more than 50% of the links remain. However, if we remove more than 50% of the links, the number of communities increases because the prior parameter is too low. As for the synthetic network, the prior parameter $C = 2$ is too high and causes underfit: the method detects no community structure when we remove more than 10% of the links. Again, $C = 1$ offers a good tradeoff. The number of communities is approximately constant as long as at least 50% of the links remain and then decreases to 1 when fewer than 40% of the links remain, where the method deduces that there no longer exists any significant community structure.

We illustrate differences in the community structure of the Jazz collaboration network induced by missing links for the standard and Bayesian map equation with alluvial diagrams [37]. The standard map equation identifies more and smaller communities with sparser networks, whereas its Bayesian estimate keeps similar communities with few changes before collapsing into one community when only 30% of the links remain. The Bayesian estimate's prior assumption of missing links prevents the map equation from splitting communities when the networks lose links (Fig. 4).

B. Adjusted mutual information

AMI is a standard measure used to compare two different partitions [38]. For the synthetic network, we compare identified partitions with the planted partition. The standard map equation successfully recovers the planted partition when more than 60% of the links are available ($\text{AMI} = 1$). When we remove more links, the accuracy decreases [Fig. 5(a)]. The Bayesian estimate of the map equation with prior constant $C = 0.5$ has almost the same accuracy. If we use $C = 1$ instead, then the method performs slightly better when we remove 40–60% of the links. Again, when we remove more than 65% of the links, the Bayesian estimate of the map equation with prior constant $C = 1$ deduces that there no longer exists any significant community structure and $\text{AMI} = 0$.




FIG. 5. Performance tests of the community-detection algorithms using AMI. (a) AMI scores with the planted partition of the synthetic network as reference. (b) AMI scores with a partition obtained for the complete Jazz collaboration network as reference. The Bayesian estimate of the map equation with prior $a_\alpha = \ln(V)$ gives the most robust results when it is possible to detect significant communities. Results are averaged over 100 network samplings and ten algorithm searches. The standard error of the mean is never higher than 0.01.

To measure the AMI for the Jazz collaboration network, which has no planted partition, we compare the partitions that the community detection methods return for networks with different fractions of missing links to the partitions they return for the complete network. For the complete network, we measure the average AMI over ten searches. The Bayesian estimate of the map equation with prior $a_\alpha = \ln(V)$ is the most consistent method when it is possible to detect significant communities [Fig. 5(b)].

In both synthetic and real-world networks when $\langle k \rangle > \ln(V)$, the Bayesian estimate of the map equation with prior constant $C = 1$ shows robust performance. However, when $C = 2$ it can fail to detect their community structure due to the high level of noise induced by the prior. To understand how the noise induced by the prior in the Bayesian estimate of the map equation affects community detection in sparse networks with $\langle k \rangle \sim \ln(V)$ and weak community structure, we test the performance on a range of different networks. We generate LFR networks with various values of average degree and mixing parameter, randomly remove a fraction of links, detect communities using the standard map equation and its Bayesian estimate with prior $a_\alpha = 0.5 \ln(V)$ and $\ln(V)$, and classify the community detection as successful when the AMI between the planted partition and the identified partition is 0.9 or higher. Even if the random prior network has higher density than the original network, the Bayesian estimate of the map equation achieves the same performance as the standard map equation when the community structure is well defined ($\mu < 0.5$). However, if the community structure is weak ($\mu = 0.5$), then the prior $a_\alpha = \ln(V)$ can cause underfit before the standard map equation starts to overfit to noise induced by missing links (Fig. 6). These results rely on the cost of overfitting and underfitting implied by the AMI. Specific networks or research questions may require other penalties for many or few communities.

C. Cross-validation

Cross-validation allows us to compare model-selection performance without planted or known partitions. We validate

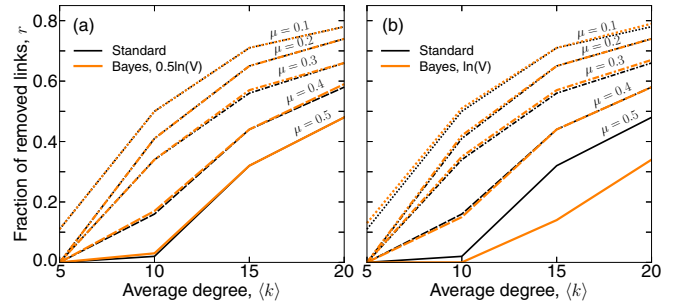


FIG. 6. Impact of network structure on the performance of the standard map equation and its Bayesian estimate. Prior parameter $C = 0.5$ in (a) and $C = 1$ in (b). For LFR networks with $V = 1000$ nodes and various densities $\langle k \rangle$ and mixing parameters μ , we show the critical fraction of removed links $r(\langle k \rangle, \mu)$ where the AMI between the planted partition and the identified partition falls below 0.9. Except for weak community structures ($\mu = 0.5$), where the Bayesian estimate with prior constant $C = 1$ underfits for lower fraction of removed links than the standard map equation overfits, the methods are on par. Results are averaged over ten network samplings and ten algorithm searches.

the significance of network partitions returned by Infomap for training networks with a fraction $1 - r$ of links using the standard map equation and its Bayesian estimate (Fig. 7).

As the link density of the training network decreases below the detectability limit, the standard map equation mistakes noisy substructures in the sparse training networks for actual communities. As a result, the relative code length savings in the training and test networks diverge, and partitions obtained with the standard map equation give negative code length savings in the test network. In contrast, the Bayesian estimate of the map equation with prior constant $C = 1$ prevents overfitting in the sparse training network, implying that there is no significant community structure.

To complement with results for other networks, we provide summary statistics for six real-world networks often used to evaluate the performance of community detection algorithms (Table I). The networks include a collaboration network in Astrophysics extracted from the arXiv (AstroPh)




FIG. 7. Performance tests of the map equation with and without Bayesian estimates using cross-validation. The Bayesian estimate of the map equation with prior $a_\alpha = \ln(V)$ prevents overfitting in the undersampled regime. Results are averaged over 100 network samplings and ten algorithm searches. The code length is measured with Grassberger entropy estimation. The standard error of the mean is never higher than 0.38.

TABLE I. Comparison between partitions detected by the standard map equation and the Bayesian estimate of the map equation for six real-world networks. The notations $m_{0.25}$ and $m_{1.0}$ refer to the number of communities in the network with 25% removed links and the complete network, respectively. The last two columns report the code length savings of test and training networks for partitions detected in the training networks with 25% removed links

Network	Nodes	Links	Method	m_0	$m_{0.25}$	$l_{0.25}^{\text{train}}(\%)$	$l_{0.25}^{\text{test}}(\%)$
AstroPh	17,903	197,031	Bayes	707	771	24	18
			Standard	663	1,080	24	18
Email	1,133	5,451	Bayes	34	1	0	0
			Standard	50	104	16	2
Erdős N1	466	1,600	Bayes	1	1	0	0
			Standard	38	67	17	-9
Football	115	613	Bayes	9	9	18	15
			Standard	10	11	20	16
PGP	10,680	24,316	Bayes	956	1,057	49	19
			Standard	897	2,210	49	16
Polblogs	1,222	16,717	Bayes	24	23	6	5
			Standard	33	80	6	5

[39], the network of e-mails exchanged between members of the University Rovira i Virgili (Email) [40], a collaboration network of authors with Erdős number 1 (Erdős N1) [41], the American College Football network (Football) [42], the PGP social network of trust (PGP) [43], and the network of political weblogs (Polblogs) [44]. In all networks, the standard map equation returns partitions with a higher number of communities when links are missing. Except for the Football network, the number of detected communities increases by 60% or more compared with the number of communities detected in the complete network. In contrast, except for the AstroPh and PGP networks, the Bayesian estimate of the map equation with prior constant $C = 1$ identifies partitions with fewer communities. Nevertheless, the different community structures detected by the two methods result in similar relative code length savings in all networks but the Email and Erdős N1 networks. They are sparse with $\langle k \rangle < \ln(V)$. In the complete Email network, the Bayesian estimate of the map equation detects 34 communities but underfits and detects no community structure after removing 25% of the links. After removing links in the Erdős N1 network, the standard map equation overfits and detects communities that, when applied to the test network, gives worse compression than the one-module solution. The Bayesian estimate of the map equation prevents this overfitting by preferring the one-module solution over any non-trivial solution.

Overall, the model-accuracy results quantified by number of communities, AMI scores, and code length savings in cross-validation on synthetic and real-world networks suggest that the analyzed network and research question should determine whether to use the standard map equation or its Bayesian estimate. Choose the standard map equation when the network data are complete or when extra communities caused by missing links are not a problem. Choose its Bayesian estimate when spurious communities can harm the analysis.

V. CONCLUSION

We have derived a Bayesian approach of the map equation that imposes prior information about the network structure to reduce overfitting for sparse networks with missing links. Using an uninformative Dirichlet prior, we show that the Bayesian estimate of the map equation avoids finding spurious communities in sparse synthetic and real-world networks with missing links. With a properly chosen prior constant, the proposed method successfully balances the impact of the imposed prior against the observed network data: The Bayesian estimate of the map equation provides a principled approach to reducing overfitting and detecting significant communities in two or more levels. We also show how to assess whether communities are significant using more effective cross-validation with Grassberger entropy estimation, which enables larger training networks. The computational overhead of the methods compared with the standard map equation is low. We anticipate that more reliable flow-based community detection of undersampled networks will be useful in many applications, including better prediction of missing links.

ACKNOWLEDGMENTS

We thank Vincenzo Nicosia and Leto Peel for stimulating discussions and Christopher Blöcker, Anton Eriksson, and Alexis Rojas for useful comments on the manuscript. J.S., D.E., and M.R. were supported by the Swedish Research Council, Grant No. 2016-00796.

APPENDIX A: NORMALIZED TRANSITION RATES

Proposition. The map equation,

$$\begin{aligned}
 L(\mathbf{M}) = & - \sum_{\alpha=1}^V p_{\alpha} \log_2(p_{\alpha}) - \sum_{i=1}^m q_{i\curvearrowright} \log_2(q_{i\curvearrowright}) \\
 & + \sum_{i=1}^m \left(q_{i\curvearrowright} + \sum_{\alpha \in i} p_{\alpha} \right) \log_2 \left(q_{i\curvearrowright} + \sum_{\alpha \in i} p_{\alpha} \right) \\
 & - \sum_{i=1}^m q_{i\curvearrowright} \log_2(q_{i\curvearrowright}) + \left(\sum_{i=1}^m q_{i\curvearrowright} \right) \log_2 \left(\sum_{i=1}^m q_{i\curvearrowright} \right),
 \end{aligned} \tag{A1}$$

is a scale invariant function.

Proof. If we scale the transition rates p_{α} , $q_{i\curvearrowright}$ and $q_{i\curvearrowleft}$ by a constant K , where $K > 0$, and change $L(\mathbf{M})$ to

$$\begin{aligned}
 L'_{\mathbf{M}} = & - \sum_{\alpha=1}^V K p_{\alpha} \log_2(K p_{\alpha}) - \sum_{i=1}^m K q_{i\curvearrowright} \log_2(K q_{i\curvearrowright}) \\
 & + \sum_{i=1}^m \left(K q_{i\curvearrowright} + \sum_{\alpha \in i} K p_{\alpha} \right) \log_2 \left(K q_{i\curvearrowright} + \sum_{\alpha \in i} K p_{\alpha} \right) \\
 & - \sum_{i=1}^m K q_{i\curvearrowright} \log_2(K q_{i\curvearrowright}) \\
 & + \left(\sum_{i=1}^m K q_{i\curvearrowright} \right) \log_2 \left(\sum_{i=1}^m K q_{i\curvearrowright} \right),
 \end{aligned}$$

then

$$\begin{aligned}
 L'_M &= -K \sum_{\alpha=1}^V p_\alpha \log_2(K) - K \sum_{\alpha=1}^V p_\alpha \log_2(p_\alpha) \\
 &\quad - K \sum_{i=1}^m q_{i\curvearrowright} \log_2(K) - K \sum_{i=1}^m q_{i\curvearrowright} \log_2(q_{i\curvearrowright}) \\
 &\quad - K \sum_{i=1}^m q_{i\curvearrowleft} \log_2(K) - K \sum_{i=1}^m q_{i\curvearrowleft} \log_2(q_{i\curvearrowleft}) \\
 &\quad + K \sum_{i=1}^m q_{i\curvearrowleft} \log_2(K) + K \sum_{\alpha=1}^V p_\alpha \log_2(K) \\
 &\quad + K \sum_{i=1}^m \left(q_{i\curvearrowright} + \sum_{\alpha \in i} p_\alpha \right) \log_2 \left(q_{i\curvearrowright} + \sum_{\alpha \in i} p_\alpha \right) \\
 &\quad + K \left(\sum_{i=1}^m q_{i\curvearrowleft} \right) \log_2(K) \\
 &\quad + K \left(\sum_{i=1}^m q_{i\curvearrowright} \right) \log_2 \left(\sum_{i=1}^m q_{i\curvearrowright} \right) \\
 &= KL(M).
 \end{aligned}$$

If we choose

$$K = \frac{\sum_{\alpha=1}^V k_\alpha}{\sum_{\alpha=1}^V k_\alpha + \sum_{i=1}^m k_{i\curvearrowright} + \sum_{i=1}^m k_{i\curvearrowleft}}, \quad (\text{A2})$$

such that

$$p'_\alpha = K p_\alpha = \frac{k_\alpha}{\sum_{\alpha=1}^V k_\alpha + \sum_{i=1}^m k_{i\curvearrowright} + \sum_{i=1}^m k_{i\curvearrowleft}}, \quad (\text{A3})$$

$$q'_{i\curvearrowright} = K q_{i\curvearrowright} = \frac{k_{i\curvearrowright}}{\sum_{\alpha=1}^V k_\alpha + \sum_{i=1}^m k_{i\curvearrowright} + \sum_{i=1}^m k_{i\curvearrowleft}}, \quad (\text{A4})$$

$$q'_{i\curvearrowleft} = K q_{i\curvearrowleft} = \frac{k_{i\curvearrowleft}}{\sum_{\alpha=1}^V k_\alpha + \sum_{i=1}^m k_{i\curvearrowright} + \sum_{i=1}^m k_{i\curvearrowleft}}, \quad (\text{A5})$$

then we will have

$$\sum_{\alpha=1}^V p'_\alpha + \sum_{i=1}^m q'_{i\curvearrowright} + \sum_{i=1}^m q'_{i\curvearrowleft} = 1. \quad (\text{A6})$$

Now we can use

$$\begin{aligned}
 L(M) &= \frac{1}{K} \left[- \sum_{\alpha=1}^V p'_\alpha \log_2(p'_\alpha) - \sum_{i=1}^m q'_{i\curvearrowright} \log_2(q'_{i\curvearrowright}) \right. \\
 &\quad + \sum_{i=1}^m \left(q'_{i\curvearrowright} + \sum_{\alpha \in i} p'_\alpha \right) \log_2 \left(q'_{i\curvearrowright} + \sum_{\alpha \in i} p'_\alpha \right) \\
 &\quad \left. - \sum_{i=1}^m q'_{i\curvearrowleft} \log_2(q'_{i\curvearrowleft}) + \left(\sum_{i=1}^m q'_{i\curvearrowleft} \right) \log_2 \left(\sum_{i=1}^m q'_{i\curvearrowleft} \right) \right] \\
 & \quad (\text{A7})
 \end{aligned}$$

to calculate the posterior average of the map equation

$$\begin{aligned}
 \hat{L}_B(M) &= E[L(M)|k, a] \\
 &= \int L(M) P(p', q'_{i\curvearrowright}, q'_{i\curvearrowleft} | k, a) dp' dq'_{i\curvearrowright} dq'_{i\curvearrowleft}, \quad (\text{A8})
 \end{aligned}$$

where posterior probability distribution equals

$$\begin{aligned}
 P(p', q'_{i\curvearrowright}, q'_{i\curvearrowleft} | k, a) \\
 \propto \prod_{\alpha=1}^V (p'_\alpha)^{k_\alpha + a_\alpha - 1} \prod_{i=1}^m [(q'_{i\curvearrowright})^{k_{i\curvearrowright} + a_{i\curvearrowright} - 1} (q'_{i\curvearrowleft})^{k_{i\curvearrowleft} + a_{i\curvearrowleft} - 1}].
 \end{aligned} \quad (\text{A9})$$

As a result, we obtain

$$\begin{aligned}
 \hat{L}_B(M) &= \frac{1}{\ln(2)} \frac{1}{\sum_{\alpha=1}^V u_\alpha} \\
 &\quad \times \left[- \sum_{\alpha=1}^V u_\alpha \psi(u_\alpha + 1) - \sum_{i=1}^m u_{i\curvearrowright} \psi(u_{i\curvearrowright} + 1) \right. \\
 &\quad + \sum_{i=1}^m \left(u_{i\curvearrowright} + \sum_{\alpha \in i} u_\alpha \right) \psi \left(u_{i\curvearrowright} + \sum_{\alpha \in i} u_\alpha + 1 \right) \\
 &\quad - \sum_{i=1}^m u_{i\curvearrowleft} \psi(u_{i\curvearrowleft} + 1) \\
 &\quad \left. + \left(\sum_{i=1}^m u_{i\curvearrowleft} \right) \psi \left(\sum_{i=1}^m u_{i\curvearrowleft} + 1 \right) \right], \quad (\text{A10})
 \end{aligned}$$

where $u_x = k_x + a_x$ and ψ is digamma function, $\psi(x) = \frac{d}{dx} \ln[\Gamma(x)]$.

APPENDIX B: THE BAYESIAN ESTIMATE OF THE MULTILEVEL MAP EQUATION

The multilevel formulation of the map equation [6,31] measures the minimum average description length given a multilevel map M of V nodes clustered into m communities, for which each community i has a submap M_i with m_i subcommunities, for which each subcommunity ij has a submap M_{ij} with m_{ij} subcommunities, and so on. It uses hierarchically nested code structures,

$$L(M) = q_{i\curvearrowright} H(Q) + \sum_{i=1}^m L(M_i), \quad (\text{B1})$$

where the average per-step code length needed to describe random walker transitions between communities at the coarsest level is the same as in the case of two-level clusterings,

$$H(Q) = - \sum_{i=1}^m \frac{q_{i\curvearrowright}}{q_{i\curvearrowright}} \log_2 \frac{q_{i\curvearrowright}}{q_{i\curvearrowright}}, \quad (\text{B2})$$

and the average per-step code word length of the module codebook i recursively takes into account contributions of the description lengths of communities at finer levels,

$$L(M_i) = q_i^\circ H(Q_i) + \sum_{i=1}^{m_i} L(M_{ij}). \quad (\text{B3})$$




FIG. 8. Mean number of communities obtained by the Bayesian estimate of the map equation with different values of the prior constant C . Smaller prior constants give more communities when many links are missing. Results are averaged over 100 network samplings and ten algorithm searches.

Here, the average per-step code length needed to describe the random walker at intermediate level i exiting to a coarser level or entering the m_i subcommunities \mathbf{M}_{ij} at a finer level is

$$H(\mathcal{Q}_i) = -\frac{q_{i\curvearrowright}}{q_i^{\circlearrowleft}} \log_2 \frac{q_{i\curvearrowright}}{q_i^{\circlearrowleft}} - \sum_{j=1}^{m_i} \frac{q_{ij\curvearrowright}}{q_i^{\circlearrowleft}} \log_2 \frac{q_{ij\curvearrowright}}{q_i^{\circlearrowleft}}, \quad (\text{B4})$$

$$\hat{L}_B(\mathbf{M}) = \frac{1}{\ln(2)} \frac{1}{\sum_{\alpha=1}^V u_{\alpha}} \left[-\sum_{i=1}^m u_{i\curvearrowright} \psi(u_{i\curvearrowright} + 1) + \left(\sum_{i=1}^m u_{i\curvearrowright} \right) \psi \left(\sum_{i=1}^m u_{i\curvearrowright} + 1 \right) \right] + \sum_{i=1}^m \hat{L}_B(\mathbf{M}_i), \quad (\text{B9})$$

where

$$\begin{aligned} \hat{L}_B(\mathbf{M}_i) = & \frac{1}{\ln(2)} \frac{1}{\sum_{\alpha=1}^V u_{\alpha}} \left[-u_{i\curvearrowright} \psi(u_{i\curvearrowright} + 1) - \sum_{j=1}^{m_i} u_{ij\curvearrowright} \psi(u_{ij\curvearrowright} + 1) \right. \\ & \left. + \left(u_{i\curvearrowright} + \sum_{j=1}^{m_i} u_{ij\curvearrowright} \right) \psi \left(u_{i\curvearrowright} + \sum_{j=1}^{m_i} u_{ij\curvearrowright} + 1 \right) \right] + \sum_{j=1}^{m_i} \hat{L}_B(\mathbf{M}_{ij}) \end{aligned} \quad (\text{B10})$$

and at the finest level

$$\begin{aligned} \hat{L}_B(\mathbf{M}_{ij\dots l}) = & \frac{1}{\ln(2)} \frac{1}{\sum_{\alpha=1}^V u_{\alpha}} \left[-u_{ij\dots l\curvearrowright} \psi(u_{ij\dots l\curvearrowright} + 1) - \sum_{\alpha \in \mathbf{M}_{ij\dots l}} u_{\alpha} \psi(u_{\alpha} + 1) \right. \\ & \left. + \left(u_{ij\dots l\curvearrowright} + \sum_{\alpha \in \mathbf{M}_{ij\dots l}} u_{\alpha} \right) \psi \left(u_{ij\dots l\curvearrowright} + \sum_{\alpha \in \mathbf{M}_{ij\dots l}} u_{\alpha} + 1 \right) \right]. \end{aligned} \quad (\text{B11})$$

APPENDIX C: RESULTS FOR DIFFERENT VALUES OF THE PRIOR PARAMETER

The number of communities obtained by the Bayesian estimate of the map equation varies for different values of the prior constant C between 0.5 and 1 (Fig. 8). For the synthetic network in the undersampled regime, $C < 0.8$ can lead to severe overfitting before removing so many links that it becomes evident that there is no significant community structure. For the Jazz collaboration network, the number of

where

$$q_i^{\circlearrowleft} = q_{i\curvearrowright} + \sum_{j=1}^{m_i} q_{ij\curvearrowright} \quad (\text{B5})$$

is the total code rate use in subcommunity i . We add the description lengths of codebooks for subcommunities at finer levels in a recursive fashion down to the finest level,

$$L(\mathbf{M}_{ij\dots l}) = p_{ij\dots l}^{\circlearrowleft} H(\mathcal{P}_{ij\dots l}), \quad (\text{B6})$$

where

$$\begin{aligned} H(\mathcal{P}_{ij\dots l}) = & -\frac{q_{ij\dots l\curvearrowright}}{p_{ij\dots l}^{\circlearrowleft}} \log_2 \frac{q_{ij\dots l\curvearrowright}}{p_{ij\dots l}^{\circlearrowleft}} \\ & - \sum_{\alpha \in \mathbf{M}_{ij\dots l}} \frac{\pi_{\alpha}}{p_{ij\dots l}^{\circlearrowleft}} \log_2 \frac{\pi_{\alpha}}{p_{ij\dots l}^{\circlearrowleft}} \end{aligned} \quad (\text{B7})$$

and

$$p_{ij\dots l}^{\circlearrowleft} = q_{ij\dots l\curvearrowright} + \sum_{\alpha \in \mathbf{M}_{ij\dots l}} \pi_{\alpha} \quad (\text{B8})$$

is the total code word use rate of module codebook $ij\dots l$.

To obtain the Bayesian estimate of the multilevel map equation, we use Eq. (B1) to calculate the posterior average according to Eq. (4). Following the same procedure described in Sec. III A, we obtain a formula for the Bayesian estimate of the multilevel map equation,

detected communities is similar for prior constant $C > 0.6$ but is higher for all values of r when $C \leq 0.6$.

To compare the performance for different prior parameters, we also compute the AMI for C between 0.5 and 1 (Fig. 9). For the synthetic network, the AMI results confirm that the detected communities become sensitive to the choice of prior when we remove more than 65% of the links. For example, for $C \geq 0.8$, the detected communities have AMI down to 0.65 before dropping to 0. For $C < 0.8$, the method can detect communities in sparser networks but these communities have




FIG. 9. Performance tests of the Bayesian estimate of the map equation with different values of the prior constant C using AMI. (a) AMI scores with the planted partition of the synthetic network as reference. (b) AMI scores with a partition obtained for the complete Jazz collaboration network as reference. Smaller prior constants give communities with non-zero AMI scores when many links are missing at the cost of overall lower AMI-scores in the Jazz network. Results are averaged over 100 network samplings and ten algorithm searches.

AMI scores below 0.5. For the Jazz collaboration network, the AMI results confirm that the detected communities are more robust when $C > 0.6$.

Cross-validation further confirms these results for different prior parameters. For the synthetic network, the Bayesian estimate of the map equation is more robust to overfitting with prior constant $C \geq 0.8$ (Fig. 10). With $C < 0.8$ and more than 75% of the links removed, the communities detected in the training network applied to the test network give




FIG. 10. Performance tests of the Bayesian estimate of the map equation with different values of the prior constant C using cross-validation. Smaller prior constants give higher compression in a narrow range of missing links at the cost of lower compression for more missing links. We show relative code length savings for the test network compared to the one-community partition. The code length is measured with Grassberger entropy estimation. Results are averaged over 100 network samplings and ten algorithm searches.

worse compression than with a single community. For the Jazz collaboration network, a prior with $C \geq 0.6$ prevents the detection of communities in the training network that, when applied to the test network, give negative relative code length savings.

These results for different values of the prior parameter indicate that there is no single prior $C \ln(V)$ that achieves optimal performance for all networks. We suggest using $\ln(V)$ as a prior because it is robust to overfitting and has good overall performance. If desired for specific networks, then C can be optimized between 0.5 and 1 with cross-validation.

[1] P. Pons and M. Latapy, *J. Graph Algor. Appl.* **10**, 191 (2006).
 [2] M. Rosvall and C. T. Bergstrom, *Proc. Natl. Acad. Sci. USA* **105**, 1118 (2008).
 [3] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, *Proc. Natl. Acad. Sci. USA* **107**, 12755 (2010).
 [4] M. T. Schaub, J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, *PLoS One* **7**, e32210 (2012).
 [5] R. Lambiotte, J. Delvenne, and M. Barahona, *IEEE Trans. Network Sci. Eng.* **1**, 76 (2014).
 [6] D. Edler, L. Bohlin, and M. Rosvall, *Algorithms* **10**, 112 (2017).
 [7] R. Guimerà and M. Sales-Pardo, *Proc. Natl. Acad. Sci. USA* **106**, 22073 (2009).
 [8] L. Lu and T. Zhou, *Physica A: Stat. Mech. Appl.* **390**, 1150 (2011).
 [9] A. Ghasemian, H. Hosseinmardi, and A. Clauset, *IEEE Trans. Knowl. Data Eng.* **1**, 1 (2019).
 [10] T. Martin, B. Ball, and M. E. J. Newman, *Phys. Rev. E* **93**, 012306 (2016).
 [11] M. E. J. Newman, *Nat. Phys.* **14**, 542 (2018).
 [12] M. E. J. Newman, *Phys. Rev. E* **98**, 062321 (2018).
 [13] T. P. Peixoto, *Phys. Rev. X* **8**, 041011 (2018).
 [14] T. Squartini, G. Caldarelli, G. Cimini, A. Gabrielli, and D. Garlaschelli, *Phys. Rep.* **757**, 1 (2018).
 [15] C. E. Shannon, *Bell Syst. Tech. J.* **27**, 379 (1948).
 [16] G. P. Basharin, *Theory Probab. Appl.* **4**, 333 (1959).
 [17] T. P. Peixoto, *Phys. Rev. X* **4**, 011047 (2014).
 [18] T. Vallès-Català, T. P. Peixoto, M. Sales-Pardo, and R. Guimerà, *Phys. Rev. E* **97**, 062316 (2018).
 [19] D. H. Wolpert and D. R. Wolf, *Phys. Rev. E* **52**, 6841 (1995).
 [20] P. Grassberger, *arXiv:physics/0307138*.
 [21] T. P. Peixoto, *Phys. Rev. E* **95**, 012317 (2017).
 [22] T. P. Peixoto, *arXiv:2003.07070*.
 [23] A. Lancichinetti and S. Fortunato, *Phys. Rev. E* **80**, 056117 (2009).
 [24] G. Miller, in *Information Theory in Psychology; Problems and Methods*, edited by H. Quastler (Free Press, Glencoe, IL, 1955).
 [25] S. Zahl, *Ecology* **58**, 907 (1977).
 [26] I. Nemenman, F. Shafee, and W. Bialek, in *Advances in Neural Information Processing Systems 14* (MIT Press, Cambridge, MA, 2002).
 [27] E. Archer, I. M. Park, and J. W. Pillow, *J. Mach. Learn. Res.* **15**, 2833 (2014).
 [28] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis* (Cambridge University Press, New York, NY, 2005).
 [29] P. Erdős and A. Rényi, *Publ. Math. Debrecen* **6**, 290 (1959).
 [30] D. Edler, A. Eriksson, and M. Rosvall, The Infomap Software Package (2020), <https://www.mapequation.org>.
 [31] M. Rosvall and C. T. Bergstrom, *PLoS One* **6**, e18209 (2011).
 [32] T. Schürmann, *J. Phys. A* **37**, L295 (2004).
 [33] P. M. Gleiser and L. Danon, *Adv. Complex Syst.* **06**, 565 (2003).

- [34] A. Lancichinetti, S. Fortunato, and F. Radicchi, *Phys. Rev. E* **78**, 046110 (2008).
- [35] M. E. J. Newman, *Phys. Rev. E* **94**, 052315 (2016).
- [36] L. Peel, D. B. Larremore, and A. Clauset, *Sci. Adv.* **3**, e1602548 (2017).
- [37] M. Rosvall and C. T. Bergstrom, *PLoS One* **5**, e8694 (2010).
- [38] N. X. Vinh, J. Epps, and J. Bailey, *J. Mach. Learn. Res.* **11**, 2837 (2010).
- [39] J. Leskovec, J. Kleinberg, and C. Faloutsos, *ACM Trans. Knowl. Discovery Data* **1**, 2 (2007).
- [40] H. Ebel, L.-I. Mielsch, and S. Bornholdt, *Phys. Rev. E* **66**, 035103(R) (2002).
- [41] V. Batagelj and A. Mrvar, *Soc. Networks* **22**, 173 (2000).
- [42] M. E. J. Newman and M. Girvan, *Phys. Rev. E* **69**, 026113 (2004).
- [43] M. Boguñá, R. Pastor-Satorras, A. Díaz-Guilera, and A. Arenas, *Phys. Rev. E* **70**, 056122 (2004).
- [44] L. A. Adamic and N. Glance, in *Proceedings of the Workshop on the Weblogging Ecosystem (WWW'05)* (ACM, New York, 2005).