

RELIABLE PREDICTION IN THE MARKOV STOCHASTIC BLOCK MODEL

QUENTIN DUCHEMIN*

Abstract. We introduce the Markov Stochastic Block Model (MSBM): a growth model for community based networks where node attributes are assigned through a Markovian dynamic. We rely on HMMs' literature to design prediction methods that are robust to local clustering errors. We focus specifically on the link prediction and collaborative filtering problems and we introduce a new model selection procedure to infer the number of hidden clusters in the network. Our approaches for reliable prediction in MSBMs are not algorithm-dependent in the sense that they can be applied using your favourite clustering tool. In this paper, we use a recent SDP method to infer the hidden communities and we provide theoretical guarantees. In particular, we identify the relevant signal-to-noise ratio (SNR) in our framework and we prove that the misclassification error decays exponentially fast with respect to this SNR.

Mathematics Subject Classification. 62F35, 62H30, 05C80, 60J10.

Received June 3, 2022. Accepted November 28, 2022.

1. INTRODUCTION

Large random graphs have been very popular in the last decade since they are powerful tools to model complex phenomena like interactions on social networks [66] or the spread of a disease [5]. The relevance of random graph models can be quantified by understanding their ability (i) to reproduce properties observed in most real world networks and (ii) to make predictions regarding missing information or future evolution of the system.

1.1. Random graph models and characteristics of real world networks

Among important features found in most of real world networks, one can mention the so-called scale-free property which states that degree distribution follows a power-law, or the small-world phenomenon in social networks. A large span of random graph models that exhibit such behaviours have been proposed such as the scale-free network model of Barabasi and Albert [7] or the small-world networks of Watts and Strogatz [61]. Another important characteristic of networks in applications is the existence of groups of nodes that are more densely connected internally than with the rest of the network. To deal with such cases, latent space models for network data emerged (see [59]). One of the most studied latent model is the Stochastic Block Model (SBM): a community based random graph model where each node is assumed to belong to one community while the

Keywords and phrases: Random graphs, growth model, Markov chains, clustering, collaborative filtering, link prediction.

LAMA, Univ Gustave Eiffel, Univ Paris Est Créteil, CNRS, 77447 Marne-la-Vallée, France.

* Corresponding author: quentin.duchemin@univ-eiffel.fr

connection probabilities between two nodes in the graph depend on their respective membership. The SBM gives a relevant framework to detect communities of well connected nodes in a graph and a large number of algorithms have been proposed to recover the hidden groups in SBMs from the observation of the edges. One can mention belief propagation algorithms [3], spectral methods [18], neural networks [58], Bayesian approaches [66] or Maximum Likelihood estimation [13]. Another powerful and popular tool is Semi-Definite Programming (SDP) which is known to have interesting robustness features [26, 55]. Recently, [29] proposed a SDP method to address community detection by solving a relaxed version of K -means. They prove that the proportion of clustering errors made by their algorithm decays exponentially fast with respect to a well-chosen Signal to Noise Ratio (SNR). Such result is known as a *partial recovery bound* in the literature and we refer to the survey [1] for further details regarding recent developments for community detection in the SBMs.

1.2. Link prediction in random graphs

Random graph models are not only used to answer questions about the properties of the studied network, but their goal is also to make predictions about missing information or future events (*cf.* [6]). In the last decade, a lot of work has been done on the link prediction problem and several review articles have been written to synthesize this abundant literature (*cf.* [60] or [40]). Papers tackling this question can be classified into three groups. In the first category, a network is partially observed and one aims at predicting the missing links. The simplest approaches are based on similarity-metrics [4, 51, 69] where for each pair of nodes a similarity score is calculated. Among the non-observed edges, the pair of nodes having a higher score represents the predicted link. Another important line of research is based on probabilistic models. One can mention for example [9] where the authors consider that covariates are observed for each node and propose a generative model formulated as a matrix logistic regression. Other approaches include matrix factorization techniques (*cf.* [63]) or the use of kernel methods (*cf.* [67]).

Papers that we classify in the second category tackle the link prediction problem in *temporal networks*: they observe several snapshots of a network with a fixed number of nodes where edges can appear or disappear over time. Based on the sequence of snapshots, the goal is to predict the connections in the network in the future. The difference between various temporal link prediction algorithms lies in how they capture the temporal or dynamic nature of the networks and also how they define the network property to be preserved. Several probabilistic approaches has been deployed to solve link prediction in temporal networks. In [22], the authors present a stochastic Markov model over a time-varying graph. The temporal analysis in this model considers link (local structural) evolution over fine-grained time scale and cluster (semi-global structural) evolution over coarse grained time scale. Some of the probabilistic temporal link prediction algorithms follow non-parametric approaches such as in [56] where the authors model the out-edges of a node i at time t as a function of the local neighbourhood of i over a moving time window. We refer to [24] for a more detailed description of the different methods proposed to tackle link prediction in temporal networks.

In the third and last category, the goal is to tackle link prediction in *growth models*, *i.e.* in networks where at each time step new nodes are joining the graph. Given the observation of the graph up to time t , the goal is to predict how a node that joins the network at time $T > t$ will connect to nodes already existing in the network. Growth models aim at mimicking the way the nature generates heterogeneous networks. Each of these models has its own growth dynamics or generating process, proposed as an hypothesis explaining the emergence of a target feature. Famous growth models include the preferential attachment model (*cf.* [51]), the team-based Yule model (*cf.* [50]) or the copying mechanism from [41] where at each time step a new node enters the network and copies a number of links from a “prototype” node that is selected randomly from the existing nodes whereas choosing the remaining neighbors is random. In [25], the authors introduced a new growth model based on the Random Geometric Graph on the euclidean sphere \mathbb{S}^{d-1} . A latent attribute $X_i \in \mathbb{S}^{d-1}$ is associated to each node $i \in [n]$ and they assume that the process $(X_i)_{i \geq 1}$ is a Markov chain. Two nodes i and j are connected with a probability that is a function of the euclidean distance between the latent representations X_i and X_j . Using non-parametric methods, the authors prove that they can estimate the probability of connection between

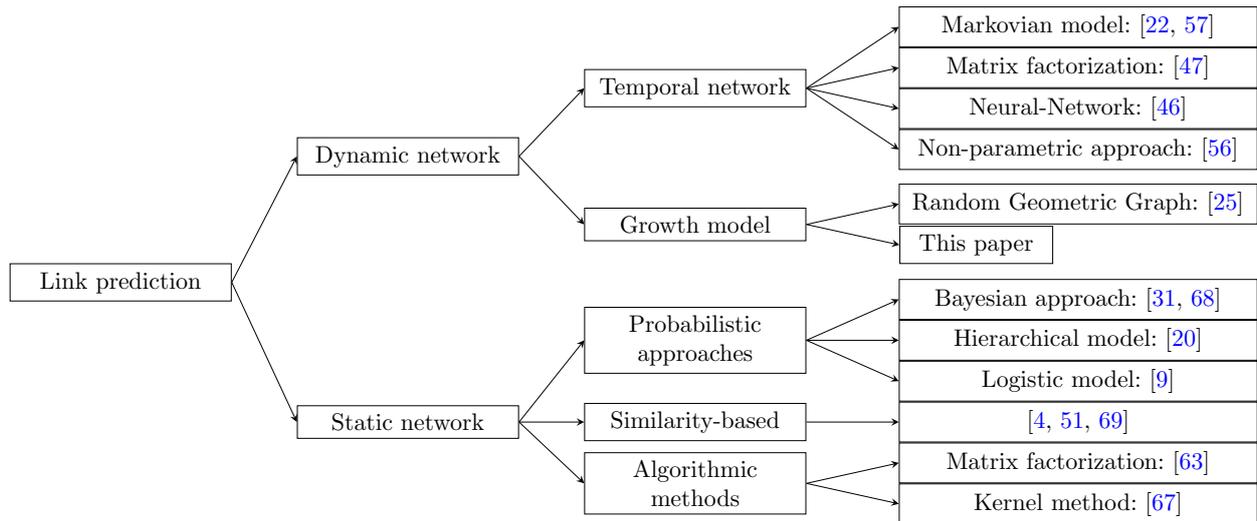


FIGURE 1. Classification of several methods proposed for link prediction in random graphs.

the upcoming node $n + 1$ and any node $i \in [n]$ that already exists in the graph. Figure 1 gives a synthetic presentation of the different link prediction approaches discussed so far.

Among all the above mentioned link prediction techniques, only a small fraction of them discuss the reliability of the proposed algorithms in the presence of spurious links or when information is missing. Some works have investigated the reliability of link prediction for temporal networks using Bayesian approaches (*cf.* [27, 31]), but this question remains so far understudied.

In this paper, we propose a reliable link prediction method in a new growth model for community-based networks which is a dynamic extension of the standard SBM.

1.3. Dynamic community-based networks

Several time evolving SBMs have been recently introduced. In [49], a Stochastic Block Temporal Model is considered where the temporal evolution is modeled through a discrete hidden Markov chain on the nodes membership and where the connection probabilities also evolve through time. In [54], connection probabilities between nodes are functions of time, considering a maximum number of nodes that can switch their communities between two consecutive time steps. Following the work of [36, 43] study the Degree Corrected Stochastic Block Model where the degree of the nodes can vary within the same community. They show that for the relatively sparse case (*i.e.* when the maximum expected node degree is of order $\log(n)$ or higher), the proportion of misclassified nodes tends to 0 with a probability that goes to 1 when the number of nodes n increases using spectral clustering. This result inspired the recent paper [37] which considers a Dynamic Stochastic Block Model where the communities can change with time. They provide direct connection between the density of the graph and its smoothness (which measures how much the graph changes with time). Several other dynamic variants of the SBM have been proposed so far like in [64] where the presence of an edge at the time step $t + 1$ directly depends on its presence or absence at time t .

The above mentioned works are mainly considering temporal networks where membership of nodes or edges can evolve with time, but only few papers are interested in growth model for SBMs (meaning that the size of the graph increases over time) and we aim at filling this gap.

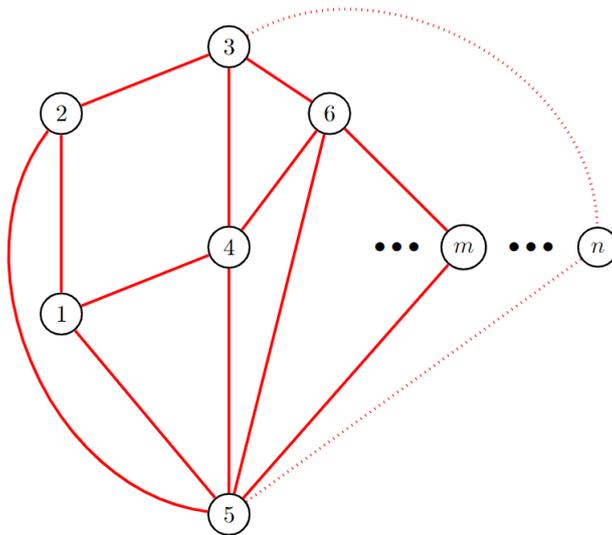


FIGURE 2. Collaborative filtering in a growth model: the graph is fully observed until time m and one aims at inferring the community of node n based on a partial information on the connections between node n and nodes in $[m]$.

1.4. Model and motivations

While previous works mainly consider a fixed number of nodes with an evolving graph where communities or connection probabilities can evolve, the Markov Stochastic Block Model (MSBM) is a growth model where a new node enters the graph at each time step and its community is drawn from a distribution depending only on the community of its predecessor. Our model could find interesting applications as in the study of bird migrations (see Sect. 6.3) where animals have regular seasonal movement between breeding and wintering grounds. Another possible application of our model is for recommendation systems or the analysis of tumor growth that we describe in Sections B.5 and B.6 of the Appendix. We provide a reliable link prediction method in the MSBM and we also propose an algorithm to solve *collaborative filtering problems*.

Collaborative filtering is mainly studied in recommender systems and refers to the ability to exploit the relationships between users to recommend items to the active user according to the ratings of his/her neighbors (*cf.* [17]). In this paper, we consider the more general sense of collaborative filtering defined as the process of searching for information using strategies that involve several agents. More precisely, our goal is to infer the community of some node n when we have only partial information about how node n is connected to the other nodes of the graph. Typically, we consider a growth model where we fully observe the graph at time m while a poor information transmission occurs from time m making available only a small number of edges between node n and the nodes in $[m]$. The goal is to infer the hidden community of node n . Figure 2 gives a visualization of the collaborative filtering problem tackled in this paper.

1.5. Contributions

We show that the MSBM gives a convenient framework to study community-based growing networks with a Markovian underlying dynamic. Our goal is to predict future information based on historical data and we tackle specifically the problems of link prediction and collaborative filtering. We show that the standard plug-in method is highly sensitive to clustering errors. This is the reason why we propose a new approach that is more robust to possible errors in the estimated communities. To do so, we borrow ideas from mean field approximation (see *e.g.* [52]) by considering that the joint distribution of the communities of nodes (*hidden states*) and the communities returned by the clustering algorithm (*observed states*) factorizes according to the graph of a homogeneous Hidden

Markov Model (HMM). Using the Baum Welch algorithm, we learn the emission probabilities of this HMM, namely the probability that the clustering algorithm predicts community l for a node belonging to community k for any $k, l \in [K]$. Based on these quantities, we also propose a model selection procedure. In our simulations, we use the algorithm from [29] to recover communities and as far as we know, we are the first to provide an implementation of their algorithm. From a theoretical point of view, we show that the misclassification error decays exponentially fast with respect to the SNR and we provide regimes where we can estimate consistently the parameters of our model.

1.6. Outline

In Section 2, we formally define SBMs and we introduce MSBMs. In Section 3, we establish a partial recovery bound and we show that we can consistently estimate the parameters of our model. Sections 4 and 5 are dedicated to our main contributions: we address the link prediction (*cf.* Sect. 4) and collaborative filtering problems (*cf.* Sect. 5), and we give heuristics to be robust to potential local clustering errors of the algorithm. In Section 6, we propose a model selection procedure to estimate the number of communities in our model and we apply our methods on real data.

In the Appendix, we provide the proofs of the theoretical results presented in Section 3 and we present additional experiments. Let us mention that Appendix F contains a list of the different notations used in this paper.

1.7. Notations

For any $n \in \mathbb{N}$ with $n \geq 1$, we denote by $[n]$ the set $\{1, \dots, n\}$. Given any matrix $A = (A_{i,j})_{i \in [n], j \in [d]} \in \mathbb{R}^{n \times d}$, we denote by $\|A\|_\infty := \max_{i,j} |A_{i,j}|$ the max norm of the matrix A , while $\|A\|_2 := \sqrt{\sum_{i \in [n], j \in [d]} A_{i,j}^2}$ is the Frobenius norm of the matrix A . Given two sequences $(a_n)_{n \in \mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}}$ of real numbers such that for some $N \in \mathbb{N}$, $b_n \neq 0$ for all $n \geq N$, we write $b_n \underset{n \rightarrow +\infty}{=} \omega(a_n)$ if the sequence $(a_n/b_n)_{n \geq N}$ goes to zero as n tends to $+\infty$.

2. MODEL

An undirected graph is defined by a set of nodes V and a set of edges $E \subset V \times V$. For an undirected graph with n nodes, we define the adjacency matrix of this graph $X \in \{0, 1\}^{n \times n}$ such that for all $i, j \in [n]$,

$$X_{i,j} = \begin{cases} 1 & \text{if } \{i, j\} \in E \\ 0 & \text{otherwise.} \end{cases}$$

2.1. Stochastic Block model

Let us consider $K \geq 2$ communities and a set of n nodes $V = [n]$. The communities $(c_i)_{i \in [n]} \in K^n$ are assigned independently to each node according to a probability distribution $\nu \in [0, 1]^K$, $\sum_{k \in [K]} \nu_k = 1$. Stated otherwise, the community c_i of node $i \in [n]$ is randomly sampled from the distribution ν . Considering the symmetric connectivity matrix $Q \in [0, 1]^{K \times K}$, the adjacency matrix of the graph $X \in \{0, 1\}^{n \times n}$ related to the assignment of the communities $(c_i)_{i \in [n]}$ is defined by

$$X_{i,j} \sim \text{Ber}(Q_{c_i, c_j}),$$

where $\text{Ber}(p)$ indicates a Bernoulli random variable with parameter $p \in [0, 1]$. In the standard SBM, X is observed while the latent variables $(c_i)_{i \in [n]}$ are unknown.

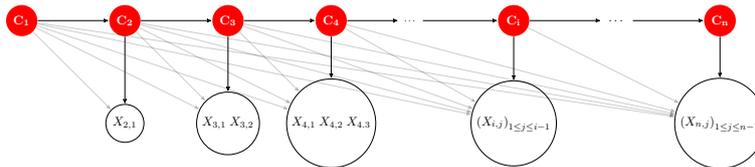


FIGURE 3. Graphical model presenting the SBM with Markovian assignment of the communities.

For a parameter $\alpha_n \in (0, 1)$ varying with the number of nodes n , we will be focused on connectivity matrix of the form

$$Q := \alpha_n Q_0,$$

where $Q_0 \in [0, 1]^{K \times K}$ is a matrix independent of n . As highlighted for example in [2], the rate of α_n as $n \rightarrow \infty$ is a key property to study random graphs sampled from SBMs. Typical regimes are $\alpha_n \sim 1$ (dense regime), $\alpha_n \sim \frac{\log(n)}{n}$ (relatively sparse regime) and $\alpha_n \sim \frac{1}{n}$ (sparse regime).

2.2. Markovian assignment of communities in the SBM

We introduce in this paper the Markov Stochastic Block Model (MSBM) which assigns a community to each node using a Markovian dynamic. We start by ordering the n nodes in V and without loss of generality, we consider the increasing order of the integers $1, 2, \dots, n$. For all $i \in [n]$, we denote $C_i \in [K]$ the random variable representing the community of the node i and we consider that they satisfy the following assumption.

Assumption A1. $(C_i)_{i \in [n]}$ is a positive recurrent Markov chain on the finite space $[K]$ with stationary measure π , with transition matrix $P \in \mathbb{R}^{K \times K}$ and initial distribution π . K is independent of n .

The community of the first node C_1 is drawn from the probability distribution π . Then for any $i \geq 2$, the community of the i -th node C_i is sampled from the probability distribution $P(C_{i-1}, \cdot)$. In the following, we will denote by

$$G_k := \{i \in [n] \mid C_i = k\},$$

the set of nodes belonging to some community $k \in [K]$. Once the community of each node is assigned, we draw an edge between the nodes i and j with probability Q_{C_i, C_j} ,

$$X_{i,j} \sim \text{Ber}(Q_{C_i, C_j}) \quad \text{with} \quad Q := \alpha_n Q_0.$$

Here, $Q_0 \in [0, 1]^{K \times K}$ is independent of n and $\alpha_n \in (0, 1)$ is varying with n . Figure 3 presents a graphical representation of our model. We observe the adjacency matrix X but the latent variables $(C_i)_{i \in [n]}$ are unknown.

The following quantities are independent of n and will be crucial in the definition of the SNR

$$L := \|Q_0\|_\infty, \quad \pi_m := \min_{c \in [K]} \pi(c), \quad D^2 := \min_{l \neq k} \|(Q_0)_{:,k} - (Q_0)_{:,l}\|_2^2.$$

2.3. Identifiability

Let us consider some integer $K \geq 1$ and \mathcal{P} (resp. \mathcal{Q}) a subset of the set of Markov kernels (resp. of connectivity matrices) of dimension $K \times K$. For any $(P, Q) \in \mathcal{P} \times \mathcal{Q}$, let us denote $\mathbb{P}_{n, (P, Q)}$ the distribution of the adjacency matrix of a graph of size n sampled from the MSBM with parameters (P, Q) . Following [14], we consider the number of clusters K as known and we say that MSBM parameters are identifiable in the parameter space $\mathcal{P} \times \mathcal{Q}$ if there exists $N \in \mathbb{N}$ such that

$$\forall P, P' \in \mathcal{P}, \forall Q, Q' \in \mathcal{Q}, \left[(\forall n \geq N, \mathbb{P}_{n,(P,Q)} = \mathbb{P}_{n,(P',Q')}) \implies (P, Q) \text{ and } (P', Q') \text{ are equal up to label switching}^1. \right]$$

We consider the following additional assumption.

Assumption A2. $D^2 = \min_{l \neq k} \|(Q_0)_{:,k} - (Q_0)_{:,l}\|_2^2 > 0$.

Provided that Assumptions A1 and A2 hold, Theorems 3.2 and 3.4 (cf. Sect. 3.2) prove that we are able to get consistent estimation (for the max norm) of the parameters P , π and Q of our model when

$$\alpha_n \log(n) \leq 1/L \quad \text{and} \quad \alpha_n \underset{n \rightarrow +\infty}{=} \omega\left(\frac{\log(n)}{n}\right). \quad (2.1)$$

In particular, denoting \mathcal{P} (resp. \mathcal{Q}) the set of $K \times K$ Markov kernels (resp. of connectivity matrices) satisfying Assumption A1 (resp. satisfying Assumption A2), the MSBM parameters are identifiable in $\mathcal{P} \times \mathcal{Q}$ provided that the conditions in equation (2.1) are satisfied. The first condition in equation (2.1) is a technical assumption related to the clustering algorithm considered in this paper (cf. [29]) while the second condition ensures that for n large enough, the clustering algorithm recovers exactly the hidden communities with high probability. Note that the paper [62] suggests that the condition $D^2 > 0$ may not be necessary for identifiability since in classical Hidden Markov Models, the additional temporal structure allows for identifiability even, say, when some states have exactly the same output distributions.

2.4. Error measure

Given two partitions $\hat{G} = (\hat{G}_1, \dots, \hat{G}_K)$ and $G = (G_1, \dots, G_K)$ of $[n]$ into K non-void groups, we define the proportion of non-matching points

$$\text{err}(\hat{G}, G) = \min_{\sigma \in \mathcal{S}_K} \frac{1}{2n} \sum_{k=1}^K \left| \hat{G}_k \Delta G_{\sigma(k)} \right|,$$

where $A \Delta B = (A \setminus B) \cup (B \setminus A)$ represents the symmetric difference between the two sets A and B , $|A|$ is the cardinality of the set A and \mathcal{S}_K represents the set of permutations on $\{1, \dots, K\}$. When \hat{G} is a partition estimating G , we refer to $\text{err}(\hat{G}, G)$ as the misclassification proportion (or error) of the clustering.

In this paper, $\hat{G} = (\hat{G}_1, \dots, \hat{G}_K)$ refers to the partition in K groups of the n nodes of the studied graph using the clustering algorithm presented in [29]. This clustering algorithm is presented in Appendix A and uses as input only the adjacency matrix of the graph and a predefined number of clusters K .

3. ESTIMATION PROCEDURES AND THEORETICAL RESULTS

3.1. Partial recovery bound for the MSBM

Using the clustering algorithm from [29] to infer the hidden communities, we provide a partial recovery bound in the Stochastic Block Model when the communities are assigned through a Markovian dynamic. In the following, $(\hat{C}_i)_{1 \leq i \leq n}$ and $(\hat{G}_k)_{k \in [K]}$ denote respectively the estimators of $(C_i)_{1 \leq i \leq n}$ and $(G_k)_{k \in [K]}$ provided by the Algorithm 1 from [29] which is described in Section A of the Appendix.

We define the signal-to-noise ratio as

$$S^2 := \frac{n\alpha_n\pi_m D^2}{L},$$

¹ (P, Q) and (P', Q') are equal up to label switching if there exists some permutation σ of $[K]$ such that $\forall i, j \in [K], P_{i,j} = P'_{\sigma(i),\sigma(j)}$ and $Q_{i,j} = Q'_{\sigma(i),\sigma(j)}$.

reminding that $\pi_m = \min_{c \in [K]} \pi(c)$, $\|Q_0\|_\infty \leq L$ and $D^2 = \min_{l \neq k} \|(Q_0)_{:,k} - (Q_0)_{:,l}\|_2^2$. The SNR should be understood as the ratio between *i*) the *signal* $\alpha_n^2 n \pi_m D^2$, which is an asymptotic lower bound on the minimal distance between two distinct centers Δ^2 defined as

$$\Delta^2 := \min_{k \neq j} \sum_l |G_l| (Q_{k,l} - Q_{j,l})^2 = \alpha_n^2 \sum_l |G_l| ((Q_0)_{k,l} - (Q_0)_{j,l})^2,$$

and *ii*) the *noise* $\alpha_n L$. We shed light on the fact that this quantity matches asymptotically the SNR from Theorem A.1 (cf. Section A in the Appendix) when π is the uniform distribution over $[K]$ and when the communities are assigned independently to each node according to the probability distribution π . Moreover, π_m can be related to standard quantities that measure how fast the chain converges to its stationary distribution π (see Section D-Proposition D.3 in the Appendix). The smaller π_m , the slower the convergence of the chain towards π and the smaller the SNR. Similarly to Theorem A.1, we prove with Theorem 3.1 that the misclassification error decays exponentially fast with respect to the SNR S^2 .

Theorem 3.1. *Let us recall that $\hat{G} = (\hat{G}_1, \dots, \hat{G}_K)$ is the partition in K groups of the n nodes of the graph obtained using the clustering algorithm from [29]. We assume that $\alpha_n \log(n) \leq 1/L$ and we consider that assumptions A1 and A2 of Section 2 are satisfied. Then there exist three constants $a, b, c > 0$ such that for any n satisfying*

$$n\alpha_n > a, \tag{3.1}$$

it holds with probability at least $1 - b/n^2$,

$$\text{err}(\hat{G}, G) \leq e^{-cS^2}.$$

In particular, it holds with probability at least $1 - b/n^2$,

$$-\log(\text{err}(\hat{G}, G)) = \Omega(n\alpha_n).$$

The constant a only depends on the parameters π and Q_0 , the constant b only depends on π , P and K while c is a universal constant.

The proof of Theorem 3.1 and the explicit expressions of the constants a and b are provided in Section C of the Appendix. Theorem 3.1 states that in the relatively sparse regime (*i.e.* when $\alpha_n \sim \log(n)/n$), we achieve a polynomial decay of the misclassification error with order $\pi_m D^2/L$. The greater the quantity $\pi_m D^2/L$ is, the faster the misclassification error decays. In particular, for $n\alpha_n/\log(n) \geq L/(c\pi_m D^2)$ it holds with high probability $\text{err}(\hat{G}, G) < 1/n$ which means that $\hat{G} = G$. The condition on the sparsity parameter α_n indicates that Theorem 3.1 can still be informative in the sparse regime (*i.e.* when $\alpha_n \sim 1/n$). Typically if $\lim_{n \rightarrow \infty} \alpha_n n > A$ for some $A > a$, then Theorem 3.1 ensures that for n large enough it holds with high probability, $\text{err}(\hat{G}, G) \leq e^{-cA\pi_m D^2/L}$.

3.2. Consistent parameter estimation

Note that Theorem 3.1 is a straightforward consequence of the work of [29]. Our methods from Sections 4 and 5 could be easily applied using your favorite clustering algorithm and we have decided to work with this recent SDP method for our simulations. In this section, we give estimates $\hat{\pi}$, \hat{P} and \hat{Q} of the parameters of our model, namely π , P and Q . We prove that they are consistent for the max norm when the average degree is of order $\log n$ or higher.

In Theorems 3.2, 3.3 and 3.4, we only specify on which parameters of the model the constants a, b, b' depend. In the Appendix, Lemmas C.3, C.4 and C.5 provide the explicit expressions of those constants with respect to the parameters of the model.

In Theorems 3.2, 3.3 and 3.4, the condition on the sparsity parameter α_n indicates that we get consistent estimation (for the max norm) respectively of the transition matrix, the stationary measure and the connectivity matrix in the relatively sparse regime (*i.e.* when $\alpha_n \sim \log(n)/n$) for n large enough provided that $\lim_{n \rightarrow \infty} n\alpha_n/\log(n) > a$.

3.2.1. The connectivity matrix

In the relatively sparse setting (*i.e.* when $\alpha_n \sim \log(n)/n$), Theorem 3.1 ensures that for n large enough it holds with high probability $\text{err}(\hat{G}, G) < 1/n$ which implies that the partition of the nodes is correctly recovered. In this case, a natural estimator for $Q_{k,l}$ (for $k, l \in [K]^2$) consists in computing the ratio between the number of edges between nodes with communities k and l and the maximum number of edges between nodes with communities k and l . For any $k, l \in [K]^2$,

$$\hat{Q}_{k,l} := \begin{cases} \frac{1}{|\hat{G}_k| \cdot |\hat{G}_l|} \sum_{i \in \hat{G}_k} \sum_{j \in \hat{G}_l} X_{i,j} & \text{if } k \neq l \\ \frac{1}{|\hat{G}_k| \cdot (|\hat{G}_k| - 1)} \sum_{i,j \in \hat{G}_k} X_{i,j} & \text{if } k = l \end{cases}.$$

At first glance, it would be tempting to state that $\hat{Q}_{k,l}$ is a sum of i.i.d. Bernoulli random variables. Actually, this is not true since the random variables $(\hat{G}_k)_{k \in [K]}$ depend on the random variables $(X_{i,j})_{i,j \in [n]}$. Taking this lack of independence carefully into account, one can show Theorem 3.2 which ensures the consistency of our estimate of the connectivity matrix for the norm $\|\cdot\|_\infty$.

Theorem 3.2. *Let us consider $\gamma > 0$. We assume that $\alpha_n \log(n) \leq 1/L$ and we consider that assumptions A1 and A2 of Section 2 are satisfied. Then there exist three constants $a, b, b' > 0$ such that for any n satisfying*

$$\frac{n\alpha_n}{\log(n)} \geq a \quad \text{and} \quad n > \left(\frac{\gamma+1}{\pi_m}\right)^2, \quad (3.2)$$

it holds with probability at least $1 - b(1/n^2 \vee \exp(-b'\gamma^2))$,

$$\|\hat{Q} - Q\|_\infty \leq \frac{\gamma}{\sqrt{n}}.$$

The constant a only depends on the parameters π and Q_0 , the constant b only depends on K while b' depends on π and P .

Theorem 3.2 is proved in details in Section C.2 and we only provide here the main arguments. To cope with the non-standard dependence structure of the entries of the matrix \hat{Q} previously highlighted, we consider in the proof of Theorem 3.2 the matrix \tilde{Q} defined by

$$\forall k, l \in [K], \quad \tilde{Q}_{k,l} := \begin{cases} \frac{1}{|G_k| \cdot |G_l|} \sum_{i \in G_k} \sum_{j \in G_l} X_{i,j} & \text{if } k \neq l \\ \frac{1}{|G_k| \cdot (|G_k| - 1)} \sum_{i,j \in G_k} X_{i,j} & \text{if } k = l \end{cases},$$

where the sum in $\tilde{Q}_{k,l}$ has binomial distribution as it is a sum of i.i.d. random variables with mean $Q_{k,l}$. The proof goes as follows.

- i) **Bounding** $\|\hat{Q} - \tilde{Q}\|_\infty$. First, we ensure that the clustering algorithm is recovering correctly the hidden communities with high probability in which case $\hat{Q} = \tilde{Q}$. This result is a direct consequence of Theorem 3.1 for α_n satisfying the first inequality in equation (3.2). The proof of Theorem 3.2 is then reduced to bound $\|\tilde{Q} - Q\|_\infty$.
- ii) **Bounding** $\|\tilde{Q} - Q\|_\infty$. Then, we aim at proving that $\tilde{Q}_{k,l}$ is close to $Q_{k,l} = \mathbb{E}[X_{i,j} \mid i \in C_k, j \in C_l]$ with high probability. For this, we need to choose n large enough so that the Markov chain $(C_i)_{i \in [n]}$ *mixed enough*, meaning that for any $k \in [K]$, $|G_k|$ is close to $\mathbb{E}[\sum_{i=1}^n \mathbb{1}_{C_i=k}] = \pi(k)n$. This is ensured for n satisfying the second inequality in equation (3.2).

3.2.2. The stationary distribution of the Markov chain

Thanks to the ergodic theorem, we know that the average number of visits in each state of the chain converges toward the stationary probability of the chain at this particular state. Stated otherwise, for all community $k \in [K]$, the average number of nodes with community k in the graph converges toward $\pi(k)$ as n tends to $+\infty$. Therefore we propose to estimate the stationary measure of the chain $(C_i)_{i \geq 1}$ with $\hat{\pi}$ defined by

$$\forall k \in [K], \quad \hat{\pi}_k := \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\hat{C}_i=k}.$$

Theorem 3.3 ensures the consistency of our estimate $\hat{\pi}$. Its proof can be found in Section C.3.

Theorem 3.3. *Let us consider $\gamma > 0$. We assume that $\alpha_n \log(n) \leq 1/L$ and we consider that assumptions A1 and A2 of Section 2 are satisfied. Then there exist three constants $a, b, b' > 0$ such that for any n satisfying*

$$\frac{n\alpha_n}{\log(n)} \geq a, \tag{3.3}$$

it holds with probability at least $1 - b(1/n^2 \vee \exp(-b'\gamma^2))$,

$$\|\hat{\pi} - \pi\|_\infty \leq \frac{\gamma}{\sqrt{n}}.$$

The constant a only depends on the parameters π and Q_0 , the constant b only depends on π , P and K whereas c is a universal constant.

Similarly to Theorem 3.2, the condition on α_n given by equation (3.3) ensures that the clustering algorithm recovers the correct hidden communities with high probability (*cf.* Thm. 3.1). Under this assumption, Theorem 3.3 is a direct consequence of the uniform ergodicity of the chain $(C_i)_{i \in [n]}$ using concentration inequality for Markov chains (*cf.* Appendix D or [34], Thm. 2).

3.2.3. The transition matrix of the Markov chain

We define $(Y_i)_{i \geq 1}$ a Markov Chain on $[K]^2$ by setting $Y_i = (C_i, C_{i+1})$. We define naturally the sequence $(\hat{Y}_i)_{i \geq 1}$ by $\hat{Y}_i = (\hat{C}_i, \hat{C}_{i+1})$. The transition kernel of the Markov Chain $(Y_i)_{i \geq 1}$ is $\mathcal{P}_{(k,l),(k',l')} = \mathbb{1}_{l=k'} P_{l,l'}$ and its stationary measure is given by μ such that $\forall k, l, \mu(k, l) = \pi(k)P_{k,l}$. We propose to estimate each entry of

the transition matrix P of the Markov chain $(C_i)_{i \geq 1}$ with

$$\forall k, l \in [K]^2, \quad \hat{P}_{k,l} := \frac{n}{n-1} \frac{\sum_{i=1}^{n-1} \mathbb{1}_{\hat{Y}_i=(k,l)}}{\sum_{i=1}^n \mathbb{1}_{\hat{C}_i=k}}.$$

Theorem 3.4. *Let us consider $\gamma > \frac{5K}{2\pi_m^2}$. We assume that $\alpha_n \log(n) \leq 1/L$ and we consider that assumptions A1 and A2 of Section 2 are satisfied. Then there exist three constants $a, b, b' > 0$ such that for any n satisfying*

$$\frac{n\alpha_n}{\log(n)} \geq a \quad \text{and} \quad n\alpha_n \geq \frac{a}{\gamma^2},$$

it holds with probability at least $1 - b \left[1/n^2 \vee \exp\left(-b'(\gamma - \frac{5K}{2\pi_m^2})^2\right) \right]$,

$$\|\hat{P} - P\|_\infty \leq \frac{\gamma}{\sqrt{n}}.$$

The constant a only depends on the parameters π and Q_0 , b depends only on K while the constant b' depends on π , P and K .

To prove Theorem 3.4, we consider the Markov chain $(Y_i)_{i \geq 1}$ built considering two consecutive states of the Markov chain $(C_i)_{i \geq 1}$. Stated otherwise, the state number i of the Markov chain used is formed by the couple of the communities of the nodes number i and number $i + 1$.

4. LINK PREDICTION

4.1. The plug-in approach

In this section, we use the underlying dynamic structure of the MSBM to solve link prediction problems. More precisely, considering a graph of size n with adjacency matrix X , we want to find an algorithm that predicts the absence or presence of an edge between nodes $n + 1$ and i for any $i \in [n]$. Such algorithm can be understood as a binary classifier and can be represented by a random vector valued in $\{0, 1\}^n$. As in binary classification, the risk of a classifier $\mathbf{g} = (g_i)_{i \in [n]}$ is given by

$$\begin{aligned} \mathcal{R}(\mathbf{g}) &:= \frac{1}{n} \sum_{i=1}^n \mathbb{P}(g_i \neq X_{i,n+1} \mid \mathbf{C}_{1:n}) \\ &= \frac{1}{n} \sum_{i=1}^n (1 - \eta_i(\mathbf{C}_{1:n})) \mathbb{E}\{\mathbb{1}_{g_i=1} \mid \mathbf{C}_{1:n}\} + \eta_i(\mathbf{C}_{1:n}) \mathbb{E}\{\mathbb{1}_{g_i=0} \mid \mathbf{C}_{1:n}\}, \end{aligned} \quad (4.1)$$

where $\mathbf{C}_{1:n} = (C_i)_{i \in [n]}$ and for all $i \in [n]$, η_i is the posterior probability function and is defined by

$$\eta_i(\mathbf{c}_{1:n}) = \mathbb{P}(X_{i,n+1} = 1 \mid \mathbf{C}_{1:n} = \mathbf{c}_{1:n}) = \sum_{k \in [K]} Q_{c_i, k} P_{c_n, k}. \quad (4.2)$$

Pushing further this analogy, we can define the classification error of some classifier \mathbf{g} by $L(\mathbf{g}) = \mathbb{E}[\mathcal{R}(\mathbf{g}, \mathbf{C}_{1:n})]$. Proposition 4.2 shows that the Bayes classifier – introduced in Definition 4.1 – is optimal for the risk defined in equation (4.1).

Definition 4.1. (Bayes classifier)

The Bayes classifier $\mathbf{g}^* = (g_i^*)_{i \in [n]}$ is defined by

$$\forall i \in [n], \quad g_i^* = \begin{cases} 1 & \text{if } \eta_i(\mathbf{C}_{1:n}) \geq \frac{1}{2} \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 4.2. (Optimality of the Bayes classifier for the risk \mathcal{R})

For any classifier \mathbf{g} which is $\sigma(\mathbf{C}_{1:n}, (X_{i,j})_{i,j \in [n]})$ -measurable, it holds for any $i \in [n]$,

$$\begin{aligned} & \mathbb{P}(g_i \neq X_{i,n+1} \mid \mathbf{C}_{1:n}) - \mathbb{P}(g_i^* \neq X_{i,n+1} \mid \mathbf{C}_{1:n}) \\ &= 2 \left| \eta_i(\mathbf{C}_{1:n}) - \frac{1}{2} \right| \mathbb{E} \left\{ \mathbb{1}_{g_i \neq g_i^*} \mid \mathbf{C}_{1:n} \right\}, \end{aligned}$$

which immediately implies that

$$\mathcal{R}(\mathbf{g}, \mathbf{C}_{1:n}) \geq \mathcal{R}(\mathbf{g}^*, \mathbf{C}_{1:n}) \text{ and therefore } L(\mathbf{g}) \geq L(\mathbf{g}^*).$$

The computation of the Bayes classifier is intractable since we only observe the adjacency matrix X while the communities $\mathbf{C}_{1:n}$ and the model parameters remain unknown. A reasonable approximation of the Bayes classifier is the MSBM classifier (cf. Def. 4.3).

Definition 4.3. (MSBM classifier) The MSBM classifier \mathbf{g}^{MSBM} is defined by

$$\forall i \in [n], \quad g_i^{MSBM} = \begin{cases} 1 & \text{if } \hat{\eta}_i(\hat{\mathbf{C}}_{1:n}) \geq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases},$$

where for all $\mathbf{c}_{1:n} \in [K]^n$ and for all $i \in [n]$,

$$\hat{\eta}_i(\mathbf{c}_{1:n}) = \sum_{k \in [K]} \hat{Q}_{c_i, k} \hat{P}_{c_n, k}. \quad (4.3)$$

In equation (4.3), \hat{Q} (resp. \hat{P}) corresponds to the estimate of the connectivity matrix (resp. of the transition matrix) presented in Section 3.2.

Proposition 4.4 shows that the MSBM classifier is consistent, meaning that given a training set, the probability of correct classification approaches – as the size of the training set increases – the best probability theoretically possible if the population distributions were fully known. Proposition 4.4 is proved in Section C.5 of the Appendix.

Proposition 4.4. (Consistency of the MSBM classifier) Let us consider $\gamma > \frac{5K}{2\pi_m^2}$. Assume that $\alpha_n \log(n) \leq 1/L$. Then there exist three constants $a, b, b' > 0$ such that for any n satisfying

$$\frac{n\alpha_n}{\log(n)} \geq a, \quad n\alpha_n \geq \frac{a}{\gamma^2} \quad \text{and} \quad n > \left(\frac{\gamma + 1}{\pi_m} \right)^2, \quad (4.4)$$

it holds with probability at least $1 - b \left[1/n \vee n \exp \left(-b' \left(\gamma - \frac{5K}{2\pi_m^2} \right)^2 \right) \right]$,

$$\forall i \in [n], \quad \left| \eta_i(\mathbf{C}_{1:n}) - \hat{\eta}_i(\hat{\mathbf{C}}_{1:n}) \right| \leq \frac{\gamma}{\sqrt{n}} (\alpha_n K L + 1). \quad (4.5)$$

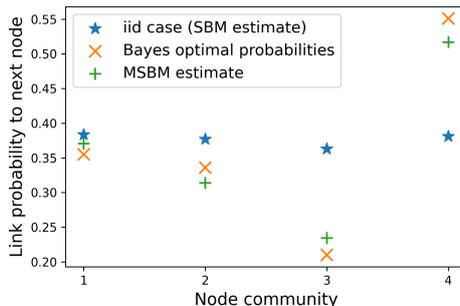


FIGURE 4. We plot the posterior probability functions considering that *i*) communities are assigned independently according to π (blue stars), or that *ii*) communities are assigned with a Markovian dynamic (orange crosses). The green plus represent the estimates of the orange crosses given by $\hat{\eta}_k(\hat{C}_n)$'s for $k \in [K]$.

Let us finally mention that the constant a only depends on the parameters π and Q_0 , b depends only on K while the constant b' depends on π , P and K .

Obtaining a non-trivial result from Proposition 4.4 may require to choose γ as function of n . Typically, choosing $\gamma = n^{1/4}$, we obtain that for n large enough, it holds with probability at least $1 - b/n$,

$$\forall i \in [n], \quad \left| \eta_i(\mathbf{C}_{1:n}) - \hat{\eta}_i(\hat{\mathbf{C}}_{1:n}) \right| \leq n^{-1/4} (\alpha_n K L + 1).$$

Since $\eta_i(\mathbf{c}_{1:n})$ only depends on c_i and c_n , we simply denote $\eta_i(\mathbf{c}_{1:n})$ by $\eta_{c_i}(c_n)$. We use an analogous slight abuse of notation for the $\hat{\eta}_i$'s and we denote $\hat{\eta}_i(\mathbf{c}_{1:n})$ by $\hat{\eta}_{c_i}(c_n)$ for any sequence $\mathbf{c}_{1:n} \in [K]^n$. Let us illustrate Proposition 4.4 with a specific numerical example. We consider $K = 4$ communities and we sample a random graph of size $n = 180$ from the MSBM using the matrices P and Q defined by equation (B.1) in the Appendix. We denote by π the stationary measure of the Markov kernel P . With Figure 4, we aim at comparing the posterior probabilities $(\eta_k(C_n))_{k \in [K]}$ (orange crosses) with *i*) the posterior probabilities obtained considering that communities have been assigned independently to each node using the probability measure π (blue stars), and *ii*) the estimates of the $\eta_k(C_n)$'s given by the $\hat{\eta}_k(\hat{C}_n)$'s (green plus). More precisely:

- ★ **i.i.d. case.** For each $k \in [K]$, the corresponding blue star represents $\eta_k^{\text{iid}} := \sum_{l \in [K]} \pi_l Q_{k,l}$.
- ✕ **Bayes optimal probabilities.** For each $k \in [K]$, the corresponding orange cross represents $\eta_k(C_n) = \sum_{l \in [K]} P_{C_n, l} Q_{k, l}$.
- + **MSBM estimate.** For each $k \in [K]$, the corresponding green plus represents $\hat{\eta}_k(\hat{C}_n)$.

Comparing the blue stars and the orange crosses on Figure 4, we see that considering an independent assignment of the communities lead to bad estimates of the Bayes optimal probabilities. The latter quantities are unknown in practice and Figure 4 shows that they can be correctly estimated by the $\hat{\eta}_k(\hat{C}_n)$'s, which confirms that the MSBM classifier from Definition 4.3 is consistent.

4.2. Reliable link prediction

Proposition 4.4 is an asymptotic result and its proof relies on the ability of the clustering algorithm to recover with high probability the correct partition of the nodes when n is large enough and satisfies the conditions from equation (4.4). In practice, we observe a graph with a finite number of nodes and the partition recovered by the algorithm is likely to contain some errors. In this case, using the plug-in approach of the previous section may give bad predictions. Indeed, the MSBM classifier only uses the two estimated communities \hat{C}_i and \hat{C}_n

TABLE 1. Classical inference problems in HMMs.

Inference problem	Filtering	Smoothing	Prediction	Likelihood
Quantity we aim at computing	$\mathbb{Q}(Z_t \mathbf{Y}_{1:t})$	$\mathbb{Q}(Z_t \mathbf{Y}_{1:s}), s > t$	$\mathbb{Q}(Z_t \mathbf{Y}_{1:s}), s < t$	$\mathbb{Q}(\mathbf{Y}_{1:T})$

to compute $\hat{\eta}_i(\hat{\mathbf{C}}_{1:n})$ without taking advantage of the complete sequence of recovered communities $\hat{\mathbf{C}}_{1:n}$. This makes the MSBM classifier heavily sensitive to local clustering errors.

To cope with this issue, we adopt a new perspective on our problem borrowed from the field of Hidden Markov Model (HMM). This section is divided into three parts: *i*) first we give a brief introduction to HMMs, then *ii*) we present the Baum-Welch algorithm which allows to conduct inference in HMMs, and finally *iii*) we explain how this algorithm can be used in our context to design a *reliable link prediction* method. By *reliable*, we mean that our method takes advantage of the complete sequence of estimated clusters $\hat{\mathbf{C}}_{1:n}$ and is less sensitive to local clustering errors compared to the plug-in approach from Section 4.1.

4.2.1. A brief introduction to HMMs for finite state spaces

Originally introduced by Baum and colleagues at the Institute for Defense Analyses in Princeton (*cf.* [8]), HMMs have found applications in different areas such as speech recognition [35], neurophysiology [28] or biology [19]. In the simpler version of HMMs, the system being modeled is assumed to be a homogeneous Markov chain $\mathbf{Z}_{1:T} = (Z_1, Z_1, \dots, Z_T)$ taking values in a finite space $[K]$ (for some $K \in \mathbb{N}$) with transition matrix P and initial probability distribution μ . The so-called hidden states (Z_1, \dots, Z_T) are not observed and we have only access to another sequence of random variables $\mathbf{Y}_{1:T} = (Y_1, \dots, Y_T)$ taking values in some finite set \mathcal{Y} called the observation space. The goal is to learn about $\mathbf{Z}_{1:T}$ by observing only $\mathbf{Y}_{1:T}$. To do so, we need to specify the way the hidden states influence the distribution of the observed states. We assume that the random variables (Y_1, \dots, Y_T) are independent conditional on the state sequence $\mathbf{Z}_{1:T}$, meaning that the probability mass function of $\mathbf{Y}_{1:T}$ conditional to $\mathbf{Z}_{1:T}$ is given by

$$\mathbb{Q}(\mathbf{Y}_{1:T} | \mathbf{Z}_{1:T}) = \prod_{t=1}^T \mathbb{Q}(Y_t | Z_t).$$

We consider that the HMM is homogeneous, which means that the conditional probability $\mathbb{Q}(Y_t | Z_t)$ does not depend on t . We write for $z \in [K]$ and $y \in \mathcal{Y}$, $O_{z,y} := \mathbb{Q}(Y_t = y | Z_t = z)$ where the $O_{z,y}$'s are called the emission probabilities. In Table 1, we give some classical inference problems in HMMs and we refer to [12] for further details. In this paper, we are particularly interested in the smoothing problem. In the next section, we present the general methodology to solve this inference task.

Notation: Considering μ a probability distribution on $[K]$ and P, O two Markov kernels on $[K]$, we will denote by \mathbb{Q}_θ with $\theta = (\mu, P, O)$ the probability under which the joint distribution of $(\mathbf{Y}_{1:T}, \mathbf{Z}_{1:T})$ factorizes according to the graph of a homogeneous HMM where

- $\mathbf{Z}_{1:T}$ (resp. $\mathbf{Y}_{1:T}$) are the hidden states (resp. the observed states),
- the emission probabilities are $O_{z,y} = \mathbb{Q}(Y_t = y | Z_t = z)$ for any $y \in \mathcal{Y}$, $z \in [K]$ and any $t \in [T]$,
- $\mathbf{Z}_{1:T}$ is a Markov chain with initial distribution μ and transition kernel P .

4.2.2. The Baum-Welch algorithm

We consider $\theta^* = (\mu^*, P^*, O^*)$ with μ^* a probability distribution on $[K]$ and P^*, O^* two Markov kernels on $[K]$. We consider that the joint distribution of $(\mathbf{Y}_{1:T}, \mathbf{Z}_{1:T})$ is \mathbb{Q}_{θ^*} that we simply denote by \mathbb{Q} . We focus on a

smoothing problem and we aim at computing

$$\mathbb{Q}(Z_t | \mathbf{Y}_{1:T}) = \frac{\mathbb{Q}(Z_t, \mathbf{Y}_{1:t})\mathbb{Q}(\mathbf{Y}_{t+1:T} | Z_t)}{\mathbb{Q}(\mathbf{Y}_{1:T})} \propto \underbrace{\mathbb{Q}(Z_t, \mathbf{Y}_{1:t})}_{=: \alpha_{Z_t}(t)} \underbrace{\mathbb{Q}(\mathbf{Y}_{t+1:T} | Z_t)}_{=: \beta_{Z_t}(t)}, \quad (4.6)$$

where the symbol \propto means that both sides are equal up to some factor that is a function of $\mathbf{Y}_{1:T}$. Hence, the probability distribution $\mathbb{Q}(\cdot | \mathbf{Y}_{1:T})$ can be obtained by normalizing $\mathbb{Q}(Z_t, \mathbf{Y}_{1:t})\mathbb{Q}(\mathbf{Y}_{t+1:T} | Z_t)$. The first term $\alpha_{Z_t}(t)$ can be obtained using a forward filtering while the second term $\beta_{Z_t}(t)$ is computed using a backward recursion. We detail both procedures in what follows.

• **Forward filtering.** In the filtering problem, we are interested in the conditional probability mass function $\mathbb{Q}(Z_t | \mathbf{Y}_{1:t})$ of the state Z_t conditional to the data observed up to time t . Note that, by Bayes rule, $\mathbb{Q}(Z_t | \mathbf{Y}_{1:t})$ can be obtained by normalizing $\mathbb{Q}(Z_t, \mathbf{Y}_{1:t}) = \alpha_{Z_t}(t)$. One can easily prove the following recursion formula

$$\alpha_{Z_t}(t) = \mathbb{Q}(Y_t | Z_t) \sum_{z_{t-1} \in [K]} \mathbb{Q}(Z_t | Z_{t-1} = z_{t-1}) \alpha_{z_{t-1}}(t-1) = O_{Z_t, Y_t} \sum_{z_{t-1} \in [K]} P_{z_{t-1}, Z_t} \alpha_{z_{t-1}}(t-1).$$

Hence, one can compute $\alpha_z(t)$ for all $z \in [K]$ and all $t \in [T]$ using the forward iterative procedure presented in Algorithm 1.

• **Backward procedure.** We are now interested in computing $\beta_{Z_t}(t) = \mathbb{Q}(\mathbf{Y}_{t+1:T} | Z_t)$. One can easily show that we have the following recursion

$$\beta_{Z_t}(t) = \sum_{z_{t+1} \in [K]} \mathbb{Q}(Y_{t+1} | Z_{t+1} = z_{t+1}) \mathbb{Q}(Z_{t+1} = z_{t+1} | Z_t) \beta_{z_{t+1}}(t+1).$$

We deduce that the backward recursion presented in Algorithm 2 allows to compute $\beta_z(t)$ for all $z \in [K]$ and all $t \in [T]$.

Algorithm 1 Forward filtering.

- 1: **for** $k = 1, \dots, K$ **set** $\alpha_k(1) = \mu_k O_{k, Y_1}$.
 - 2: **for** $t = 2, \dots, T$
 - 3: **for** $k = 1, \dots, K$
 - 4: $\alpha_k(t) = O_{k, Y_t} \sum_{l \in [K]} P_{l, k} \alpha_l(t-1)$.
-

Algorithm 2 Backward procedure.

- 1: **for** $k = 1, \dots, K$ **set** $\beta_k(T) = 1$.
 - 2: **for** $t = T, \dots, 2$
 - 3: **for** $k = 1, \dots, K$
 - 4: $\beta_k(t-1) = \sum_{l \in [K]} O_{l, Y_t} P_{k, l} \beta_l(t)$.
-

• **The Baum-Welch algorithm.** We have shown how to solve smoothing problems when the parameters of the HMMs given by $\theta := (\mu, P, O)$ are known. In practice, θ is unknown and a standard approach to bypass this difficulty consists in estimating θ using a maximum likelihood approach:

$$\theta^{\text{MLE}} = (\mu^{\text{MLE}}, P^{\text{MLE}}, O^{\text{MLE}}) \in \arg \max_{\theta} \{\log \mathbb{Q}_{\theta}(\mathbf{Y}_{1:T})\} = \arg \max_{\theta} \left\{ \log \sum_{\mathbf{z}_{1:T}} \mathbb{Q}_{\theta}(\mathbf{Y}_{1:T}, \mathbf{Z}_{1:T} = \mathbf{z}_{1:T}) \right\}.$$

However, this quantity is intractable since the distribution of $\mathbf{Z}_{1:T}$ is unknown. The Expectation-Maximization (EM) algorithm has been designed specifically for such situation where we aim at finding maximum likelihood estimates of parameters in statistical models, where the model depends on unobserved latent variables. Considering some parameters $\theta = (\mu, P, O)$ and $\theta' = (\mu', P', O')$ (with μ, μ' probability distributions on $[K]$ and

P, O, P', O' Markov kernels on $[K]$, we have

$$\begin{aligned} \log \sum_{\mathbf{z}_{1:T}} \mathbb{Q}_\theta(\mathbf{Y}_{1:T}, \mathbf{Z}_{1:T} = \mathbf{z}_{1:T}) &= \log \sum_{\mathbf{z}_{1:T}} \mathbb{Q}_\theta(\mathbf{Y}_{1:T}, \mathbf{Z}_{1:T} = \mathbf{z}_{1:T}) \frac{\mathbb{Q}_{\theta'}(\mathbf{Z}_{1:T} = \mathbf{z}_{1:T} \mid \mathbf{Y}_{1:T})}{\mathbb{Q}_{\theta'}(\mathbf{Z}_{1:T} = \mathbf{z}_{1:T} \mid \mathbf{Y}_{1:T})} \\ &\geq \sum_{\mathbf{z}_{1:T}} \log \left(\frac{\mathbb{Q}_\theta(\mathbf{Y}_{1:T}, \mathbf{Z}_{1:T} = \mathbf{z}_{1:T})}{\mathbb{Q}_{\theta'}(\mathbf{Z}_{1:T} = \mathbf{z}_{1:T} \mid \mathbf{Y}_{1:T})} \right) \mathbb{Q}_{\theta'}(\mathbf{Z}_{1:T} = \mathbf{z}_{1:T} \mid \mathbf{Y}_{1:T}) = Q(\theta \mid \theta') + H(\theta'), \end{aligned}$$

where we used Jensen's inequality and where

$$Q(\theta \mid \theta') := \mathbb{E}_{\mathbf{Z}_{1:T} \sim \mathbb{Q}_{\theta'}(\cdot \mid \mathbf{Y}_{1:T})} [\log \mathbb{Q}_\theta(\mathbf{Y}_{1:T}, \mathbf{Z}_{1:T})] \quad \text{and} \quad H(\theta') := -\mathbb{E}_{\mathbf{Z}_{1:T} \sim \mathbb{Q}_{\theta'}(\cdot \mid \mathbf{Y}_{1:T})} [\log \mathbb{Q}_{\theta'}(\mathbf{Z}_{1:T} \mid \mathbf{Y}_{1:T})].$$

From Jensen's inequality, we get that $H(\theta') \geq 0$ so that $Q(\theta \mid \theta')$ is a lower-bound of the log-likelihood $\log \mathbb{Q}_\theta(\mathbf{Y}_{1:T})$ for any θ' . Moreover, if we choose $\theta' := \theta^{\text{MLE}}$, one can easily check that $\theta^{\text{MLE}} \in \arg \max_\theta Q(\theta \mid \theta')$. Based on these observations, the EM algorithm seeks to find the MLE of the marginal likelihood by starting from an initial value of the parameters $\theta^{(0)} = (\mu^{(0)}, P^{(0)}, O^{(0)})$ and by iteratively applying these two steps:

- Expectation step (E step): Given the current value of the estimate of the parameters $\theta^{(m)} = (\mu^{(m)}, P^{(m)}, O^{(m)})$, we compute the posterior probability distribution considering that the parameters of the HMM are given by $\theta^{(m)}$, namely $\mathbb{Q}_{\theta^{(m)}}(\cdot \mid \mathbf{Y}_{1:T})$.
- Maximization step (M step): We find the parameter θ that maximizes $Q(\theta \mid \theta^{(m)})$ and we call it $\theta^{(m+1)}$.

When applied to HMMs, the EM algorithm is also called the Baum-Welch algorithm. In our framework, the M-step of the algorithm has a closed-form expression given by:

$$\forall k \in [K], \quad \mu_k^{(m+1)} = \gamma_k^{(m)}(1) \tag{4.7}$$

$$\forall k, l \in [K], \quad P_{k,l}^{(m+1)} = \frac{\sum_{t=1}^{T-1} \xi_{k,l}^{(m)}(t)}{\sum_{i=1}^{T-1} \gamma_k^{(m)}(t)} \quad \text{and} \quad O_{k,l}^{(m+1)} = \frac{\sum_{t=1}^T \mathbb{1}_{Y_t=l} \gamma_k^{(m)}(t)}{\sum_{t=1}^T \gamma_k^{(m)}(t)}, \tag{4.8}$$

where for all $k, l \in [K]$ and all $t \in [T]$,

$$\gamma_k^{(m)}(t) := \mathbb{Q}_{\theta^{(m)}}(Z_t = k \mid \mathbf{Y}_{1:T}) = \frac{\alpha_k^{(m)}(t) \beta_k^{(m)}(t)}{\sum_{b \in [K]} \alpha_b^{(m)}(t) \beta_b^{(m)}(t)}, \tag{4.9}$$

$$\begin{aligned} \xi_{k,l}^{(m)}(t) &:= \mathbb{Q}_{\theta^{(m)}}(Z_t = k, Z_{t+1} = l \mid \mathbf{Y}_{1:T}) = \frac{\mathbb{Q}_{\theta^{(m)}}(Z_t = k, Z_{t+1} = l, \mathbf{Y}_{1:T})}{\mathbb{Q}_{\theta^{(m)}}(\mathbf{Y}_{1:T})} \\ &= \frac{\alpha_k^{(m)}(t) P_{k,l}^{(m)} \beta_l^{(m)}(t+1) O_{l,Y_{t+1}}^{(m)}}{\sum_{b,c \in [K]} \alpha_c^{(m)}(t) P_{c,b}^{(m)} \beta_b^{(m)}(t+1) O_{b,Y_{t+1}}^{(m)}}. \end{aligned} \tag{4.10}$$

In equations (4.9) and (4.10), the quantities $\alpha_k^{(m)}(t)$ and $\beta_k^{(m)}(t)$ for $k \in [K]$ and $t \in [T]$ are the counterparts of the $\alpha_k(t)$'s and $\beta_k(t)$'s of the beginning of this section but computed using the probability distribution $\mathbb{Q}_{\theta^{(m)}}$. The expressions for the M-step given by equations (4.7) and (4.8) show that given $\theta^{(m)}$, the E-step of the Baum-Welch algorithm only requires to perform the forward and backward procedures (*cf.* Algorithms 1 and 2) to compute the $\alpha_k^{(m)}(t)$'s and the $\beta_k^{(m)}(t)$'s. We summarize the Baum-Welch algorithm with Algorithm 3.

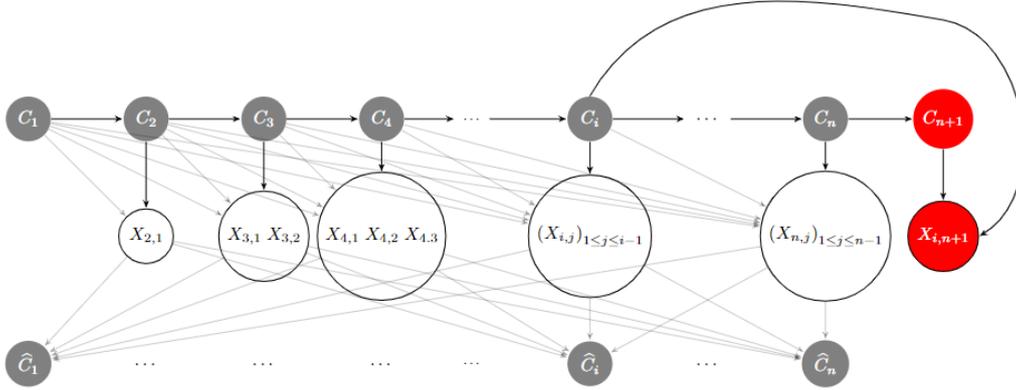


FIGURE 5. Graphical model presenting the conditional dependencies between random variables.

Algorithm 3 Baum-Welch algorithm.

-
- 1: **Initialization:** $\theta^{(0)} = (\mu^{(0)}, P^{(0)}, O^{(0)})$, $m_{\max} \geq 1$.
 - 2: $m \leftarrow 0$.
 - 3: **for** $m = 0 \dots m_{\max}$ **do**
 - 4: 1. Forward-Backward calculations:
 - 5: Compute $\alpha_k^{(m)}(t)$ for all $k \in [K]$ and all $t \in [T]$ using Algorithm 1 with $\theta^{(m)}$.
 - 6: Compute $\beta_k^{(m)}(t)$ for all $k \in [K]$ and all $t \in [T]$ using Algorithm 2 with $\theta^{(m)}$.
 - 7: 2. E-step:
 - 8: Compute $\gamma_k^{(m)}(t)$ for all $k \in [K]$ and all $t \in [T]$ using equation (4.9).
 - 9: Compute $\xi_{k,l}^{(m)}(t)$ for all $k, l \in [K]$ and all $t \in [T]$ using equation (4.10).
 - 10: 3. M-step:
 - 11: Compute $\mu_k^{(m+1)}$ for all $k \in [K]$ using equation (4.7).
 - 12: Compute $P_{k,l}^{(m+1)}$ and $O_{k,l}^{(m+1)}$ for all $k, l \in [K]$ using equation (4.8).
 - 13: $\theta^{(m+1)} \leftarrow (\mu^{(m+1)}, P^{(m+1)}, O^{(m+1)})$.
 - 14: **end for**
-

We showed how the Baum-Welch algorithm allows at the same time to provide estimates of the parameters of the model and of $\alpha_k(t)$ and $\beta_k(t)$ (for $k \in [K]$ and $t \in [T]$) that are particularly useful to solve smoothing problems (*cf.* Eq. (4.6)). Note that the convergence of the algorithm is ensured but it might stop at a local maximum of the log-likelihood. We refer to Chapters 3 and 10 of [12] for further details on the Baum-Welch and the EM algorithms.

4.2.3. From the Baum-Welch algorithm to reliable link prediction

As previously mentioned, the plug-in approach for link prediction described in Section 4.1 is highly sensitive to local clustering errors. We would like to use the complete sequence of estimated communities given by the clustering algorithm to estimate the probability that some node i in the observed graph of size n will be connected to the node $n+1$ (*i.e.* the upcoming node). Therefore, we aim at computing the following conditional probability $\mathbb{P}(X_{i,n+1} = 1 \mid \hat{\mathbf{C}}_{1:n})$. Figure 5 presents the graphical model expressing the conditional dependencies between the different random variables. Thanks to Figure 5, we get that:

- (*1) $X_{i,n+1}$ and $\hat{\mathbf{C}}_{1:n}$ are independent conditional to (C_i, C_{n+1}) ,
- (*2) C_{n+1} is independent of $(C_i, \hat{\mathbf{C}}_{1:n})$ conditional to C_n ,

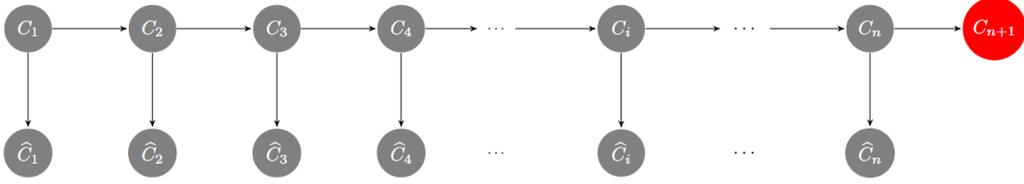


FIGURE 6. We assume that the joint distribution of $(\mathbf{C}_{1:n}, \hat{\mathbf{C}}_{1:n})$ factorizes according to the graph of a homogeneous HMM.

which implies that

$$\begin{aligned}
 & \mathbb{P}(X_{i,n+1} = 1 \mid \hat{\mathbf{C}}_{1:n}) \\
 = & \sum_{c_i, c_{n+1} \in [K]} \mathbb{P}(C_i = c_i, C_{n+1} = c_{n+1}, X_{i,n+1} = 1 \mid \hat{\mathbf{C}}_{1:n}) \\
 = & \sum_{c_i, c_{n+1} \in [K]} \frac{\mathbb{P}(C_i = c_i, C_{n+1} = c_{n+1})}{\mathbb{P}(\hat{\mathbf{C}}_{1:n})} \mathbb{P}(X_{i,n+1} = 1, \hat{\mathbf{C}}_{1:n} \mid C_i = c_i, C_{n+1} = c_{n+1}) \\
 \stackrel{\text{cf.} (*1)}{=} & \sum_{c_i, c_{n+1} \in [K]} \frac{\mathbb{P}(C_i = c_i, C_{n+1} = c_{n+1})}{\mathbb{P}(\hat{\mathbf{C}}_{1:n})} \underbrace{\mathbb{P}(X_{i,n+1} = 1 \mid C_i = c_i, C_{n+1} = c_{n+1})}_{=Q_{c_i, c_{n+1}}} \mathbb{P}(\hat{\mathbf{C}}_{1:n} \mid C_i = c_i, C_{n+1} = c_{n+1}) \\
 = & \sum_{c_i, c_{n+1} \in [K]} Q_{c_i, c_{n+1}} \mathbb{P}(C_i = c_i, C_{n+1} = c_{n+1} \mid \hat{\mathbf{C}}_{1:n}) \\
 = & \sum_{c_i, c_n, c_{n+1} \in [K]} Q_{c_i, c_{n+1}} \mathbb{P}(C_i = c_i, C_n = c_n, C_{n+1} = c_{n+1} \mid \hat{\mathbf{C}}_{1:n}) \\
 = & \sum_{c_i, c_n, c_{n+1} \in [K]} Q_{c_i, c_{n+1}} \mathbb{P}(C_i = c_i, C_n = c_n \mid \hat{\mathbf{C}}_{1:n}) \mathbb{P}(C_{n+1} = c_{n+1} \mid C_i = c_i, C_n = c_n, \hat{\mathbf{C}}_{1:n}) \\
 \stackrel{\text{cf.} (*2)}{=} & \sum_{c_i, c_n, c_{n+1} \in [K]} Q_{c_i, c_{n+1}} \mathbb{P}(C_i = c_i, C_n = c_n \mid \hat{\mathbf{C}}_{1:n}) P_{c_n, c_{n+1}}. \tag{4.11}
 \end{aligned}$$

In Section 3.2, we proposed estimates \hat{Q} and \hat{P} of the model parameters Q and P with theoretical guarantees. Hence the challenge lies in the estimation of the probability mass function $\mathbb{P}(C_i, C_n \mid \hat{\mathbf{C}}_{1:n})$ to get an approximation of (4.11). The computation of $\mathbb{P}(C_i, C_n \mid \hat{\mathbf{C}}_{1:n})$ looks like a smoothing problem where the random variables $\hat{\mathbf{C}}_{1:n}$ are observed and the random variables $\mathbf{C}_{1:n}$ are hidden. The important difference with our discussion from Sections 4.2.1 and 4.2.2 is that the emission probabilities are not independent, *i.e.* \hat{C}_i is not independent from all other random variables (namely $(C_j, \hat{C}_j)_{j \neq i}$) conditional to C_i (*cf.* Fig. 5). Inspired by mean field approximation techniques (see for example [52]), we propose to simplify the complex dependence structure of our problem depicted in Figure 5 by assuming that the emission probabilities are independent and time invariant, meaning that the joint distribution of $(\mathbf{C}_{1:n}, \hat{\mathbf{C}}_{1:n})$ factorizes according to the graph of a homogeneous HMM (*cf.* Fig. 6). Denoting by \mathbb{Q} the probability distribution with model parameter (μ, P, Q) and where we assume the HMM structure presented in Figure 6, it holds

$$\mathbb{Q}(X_{i,n+1} = 1 \mid \hat{\mathbf{C}}_{1:n}) = \sum_{c_i, c_n, c_{n+1} \in [k]} Q_{c_i, c_{n+1}} P_{c_n, c_{n+1}} \mathbb{Q}(C_i = c_i, C_n = c_n \mid \hat{\mathbf{C}}_{1:n}), \tag{4.12}$$

where $\mathbb{Q}(C_i = c_i, C_j = c_j | \hat{\mathbf{C}}_{1:n}) \propto \mathbb{Q}(C_i = c_i, C_j = c_j, \hat{\mathbf{C}}_{1:n}) = \alpha_{c_i}(i) \chi_{c_i, c_j}^{(i,j)} \beta_{c_j}(j)$ with

$$\alpha_{c_i}(i) := \mathbb{Q}(C_i = c_i, \hat{\mathbf{C}}_{1:i}), \quad \chi_{c_i, c_j}^{(i,j)} := \mathbb{Q}(C_j = c_j, \hat{\mathbf{C}}_{i+1:j} | C_i = c_i) \quad \text{and} \quad \beta_{c_j}(j) := \mathbb{Q}(\hat{\mathbf{C}}_{j+1:n} | C_j = c_j),$$

and where the symbol \propto means that both side are equal up to some factor that is a function of $\hat{\mathbf{C}}_{1:n}$. Using the Baum-Welch algorithm, we obtain estimates $\hat{\alpha}_{c_i}(i)$'s of the $\alpha_{c_i}(i)$'s, estimates $\hat{\beta}_{c_j}(j)$'s of the $\beta_{c_j}(j)$'s and estimates $\hat{O}_{k,l}$'s of the $O_{k,l}$'s where for all $k, l \in [K]$, $O_{k,l} = \mathbb{Q}(\hat{C}_t = l | C_t = k)$ (where $t \in [n]$ is arbitrary since we assumed that the HMM is homogeneous). We can also estimate $\chi_{c_i, c_j}^{(i,j)}$ using a plug-in approach with

$$\hat{\chi}_{c_i, c_j}^{(i,j)} := \sum_{c_{i+1}, \dots, c_{j-1}} \hat{P}_{c_i, c_{i+1}} \hat{O}_{c_{i+1}, \hat{C}_{i+1}} \hat{P}_{c_{i+1}, c_{i+2}} \hat{O}_{c_{i+2}, \hat{C}_{i+2}} \dots \hat{P}_{c_{j-1}, c_j} \hat{O}_{c_j, \hat{C}_j}, \quad (4.13)$$

where we recall that \hat{P} is defined in Section 3.2.3². Recalling the definition of \hat{Q} given in Section 3.2.1 and using equation (4.12), we can then estimate $\eta_i(\mathbf{C}_{1:n})$ by

$$\hat{\eta}_i^R(\hat{\mathbf{C}}_{1:n}) := \sum_{c_i, c_n, c_{n+1} \in [K]} \hat{Q}_{c_i, c_{n+1}} \hat{P}_{c_n, c_{n+1}} \hat{\alpha}_{c_i}(i) \hat{\chi}_{c_i, c_n}^{(i,n)} \hat{\beta}_{c_n}(n). \quad (4.14)$$

Hence, the Reliable MSBM (RMSBM) classifier is defined by replacing $\hat{\eta}_i$ by $\hat{\eta}_i^R$ in the definition of the MSBM classifier (*cf.* Def. 4.3). Note that this approach is a heuristic because of the local optimum reached by the EM algorithm and also because of the dependence of the emission probabilities in our model (*cf.* Fig. 5). Algorithm 4 summarizes our method.

4.2.4. Numerical results

We perform our numerical experiment using the matrices P and Q provided in equation (B.1) of the Appendix. With Figure 7, we aim at showing the superiority of our reliable link prediction method compared to the plug-in one from Section 4.1 while with Figure 8, we aim at showing that the learned emission probabilities carry relevant information on the clustering algorithm used (and also on the studied network as we will see in Sect. 6.2).

In Figure 7, we consider a graph of size 70 and we plot the sorted list of the L_1 errors ($|\hat{\eta}_i(\hat{\mathbf{C}}_{1:n}) - \eta_i(\hat{\mathbf{C}}_{1:n})|$) _{$i \in [n]$} (*resp.* ($|\hat{\eta}_i^R(\hat{\mathbf{C}}_{1:n}) - \eta_i(\hat{\mathbf{C}}_{1:n})|$) _{$i \in [n]$}). Our reliable estimation of the posterior probabilities allows to get a significantly smaller variance compared to the plug-in approach. The RMSBM classifier gives smaller L_1 errors on the posterior probabilities. The difference is significant when the clustering algorithm fails to recover the complete partition of the nodes which leads to bad estimates for the plug-in approach.

In Figure 8, we work with a graph of size 120 and we plot the learned emission probabilities $\hat{O}_{k,l}$, $k, l \in [K]$. The ergodic theorem ensures that the first cluster is (asymptotically) the smaller. Indeed, the stationary measure of P from (B.1) is approximately [0.14, 0.22, 0.38, 0.26]. We observe that the errors made by the algorithm consist in assigning nodes from community 2, 3 or 4 to cluster 1. This means that the clustering algorithm from [29] tends to overestimate the size of small clusters. In the Appendix (*cf.* Sect. B.4), we conduct experiments in other settings that confirm this overestimation of the size of small clusters.

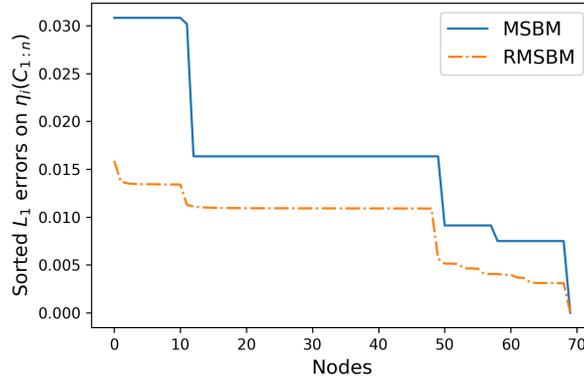
Notation: In the following, we will denote by $\hat{\mu}$ (*resp.* \hat{O}) the probability measure $\mu^{(m_{\max})}$ (*resp.* the matrix $O^{(m_{\max})}$) returned by the Baum-Welch algorithm (*cf.* Algorithm 3). Moreover, $\hat{\mathbb{Q}}$ will denote the probability measure under which

- the Markov chain $(C_i)_{i \geq 1}$ has initial distribution $\hat{\mu}$ and transition matrix \hat{P} (\hat{P} is defined in Sect. 3.2),

²Let us mention that we could have used the estimate of the transition matrix P provided by the Baum-Welch algorithm instead of the one given in Section 3.2.3. We prefer the latter since its computation does not rely on the approximation made by considering a homogeneous HMM.

Algorithm 4 Reliable link prediction.

- 1: **Data:** Adjacency matrix $X \in \{0, 1\}^{n \times n}$, $K \in \mathbb{N}$ with $K \geq 2$ and a number of iterations $m_{\max} \in \mathbb{N}$.
 - 2: Run the clustering algorithm on X with K clusters to get the estimated communities $\hat{\mathbf{C}}_{1:n}$.
 - 3: Compute the estimate \hat{Q} of the connectivity matrix from $\hat{\mathbf{C}}_{1:n}$ and X using Section 3.2.1.
 - 4: Compute the estimate \hat{P} of the Markov kernel from $\hat{\mathbf{C}}_{1:n}$ using Section 3.2.3.
 - 5: **Baum-Welch algorithm:**
 - 6: Initialization:
 - 7: $\mu^{(0)} := \mathbf{1}_K^\top$,
 - 8: $P^{(0)} := \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top$,
 - 9: $O^{(0)} := (1 - \epsilon) \text{Id}_K + \frac{\epsilon}{K-1} (\mathbf{1}_K \mathbf{1}_K^\top - \text{Id}_K)$, where $\epsilon \in (0, 1)$ (typically $\epsilon = 10^{-2}$).
 - 10: Working with the probability \mathbb{Q} (*i.e.* considering the assumption of homogeneous HMM), run the Baum-Welch algorithm (*cf.* Algorithm 3) with $\theta^{(0)} := (\mu^{(0)}, P^{(0)}, O^{(0)})$ and a number of iterations
 - 11: $m_i \hat{\alpha}_k(i) \leftarrow \alpha_k^{(m_{\max})}(i)$, $\forall k \in [K], \forall i \in [n]$.
 - 12: $\hat{\beta}_k(i) \leftarrow \beta_k^{(m_{\max})}(i)$, $\forall k \in [K], \forall i \in [n]$.
 - 13: $\hat{O}_{k,l} \leftarrow O_{k,l}^{(m_{\max})}$, $\forall k, l \in [K]$.
 - 14: Compute $\hat{\chi}_{k,l}^{(i,n)}$ for all $k, l \in [K]$ and all $i \in [n]$ using equation (4.13).
 - 15: Compute $\hat{\eta}_i^R(\hat{\mathbf{C}}_{1:n})$ for all $i \in [K]$ using equation (4.14).
 - 16: **Return:** $\{i \in [n] : \hat{\eta}_i^R(\hat{\mathbf{C}}_{1:n}) \geq \frac{1}{2}\}$ as the set of nodes from $[n]$ predicted to be connected to the node $n + 1$.
-


 FIGURE 7. Sorted list of L_1 errors between $\hat{\eta}_i(\hat{\mathbf{C}}_{1:n})$ (resp. $\hat{\eta}_i^R(\hat{\mathbf{C}}_{1:n})$) and $\eta_i(\hat{\mathbf{C}}_{1:n})$.

- $X_{i,j} \sim \text{Ber}(\hat{Q}_{C_i, C_j})$ for all $i, j \in [n]$, $i \neq j$ (\hat{Q} is defined in Sect. 3.2),
- we have the homogeneous HMM structure with emission probabilities the $\hat{O}_{k,l}$'s, *i.e.*

$$\hat{Q}(\hat{\mathbf{C}}_{1:n} | \mathbf{C}_{1:n}) = \prod_{i=1}^n \hat{Q}(\hat{C}_i | C_i) = \prod_{i=1}^n \hat{O}_{C_i, \hat{C}_i}.$$

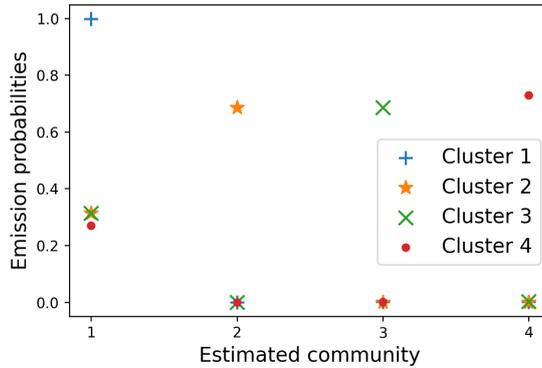


FIGURE 8. Plot of the learned emission probabilities $\hat{O}_{k,l}$, $k, l \in [K]$. Typically, on the first column we show $\hat{O}_{l,1}$ for $l \in [K]$.

5. COLLABORATIVE FILTERING

5.1. Reliable collaborative filtering

Let us now consider another prediction task for dynamic networks: collaborative filtering. Solving a collaborative filtering problem consists in inferring the community of one node of the graph when we have only partial information about how this node is connected to the rest of the graph. More precisely, we observe fully the graph at time m and for some $n > m$, we observe how the node n is connected (or not) to a subset of nodes $\mathcal{E} \subset [m]$, *i.e.* we have access to $(X_{i,n})_{i \in \mathcal{E}}$. Our goal is then to predict the community of node n : C_n . We propose to use the maximum a posteriori (MAP) estimator to tackle this problem. In the following, we present three different strategies to solve the collaborative filtering problem using the maximum a posteriori (MAP) estimator. The practical implementation of the different methods are presented in the next section.

- **The optimal MAP** has access to the hidden communities of the nodes in $[m]$ and to the model parameters. The community predicted for the node n by the optimal MAP is

$$\hat{C}_n^{OPT} \in \arg \max_{k \in [K]} \mathbb{P}(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \mathbf{C}_{1:m}).$$

- **The plug-in MAP** does not have access to the model parameters or to the hidden communities $\mathbf{C}_{1:n}$. It works with the probability $\mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})}$ where $\hat{\mu} := \mu^{(m_{\max})}$ is given by the Baum-Welch algorithm and where the Markov kernel \hat{P} and the connectivity matrix \hat{Q} are defined in Section 3.2. $\mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})}$ is the counterpart of \mathbb{P} where the Markov chain $\mathbf{C}_{1:m}$ has initial distribution $\hat{\mu}$ and transition matrix \hat{P} , and where the connectivity matrix Q is replaced by \hat{Q} . The plug-in MAP trusts the sequence of communities $\hat{\mathbf{C}}_{1:m}$ returned by the clustering algorithm and predicts the following community for the node n :

$$\hat{C}_n^{PI} \in \arg \max_{k \in [K]} \mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})}(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \mathbf{C}_{1:m} = \hat{\mathbf{C}}_{1:m}).$$

- **The Reliable MAP** does not know the model parameters or the hidden communities $\mathbf{C}_{1:n}$ either but it does not want to blindly rely on the estimated communities $\hat{\mathbf{C}}_{1:m}$ provided by the clustering algorithm. As a consequence, the Reliable MAP works with the probability $\hat{\mathbb{Q}}$ introduced in the previous section to be robust against possible clustering errors. The Reliable MAP predicts the following community for the

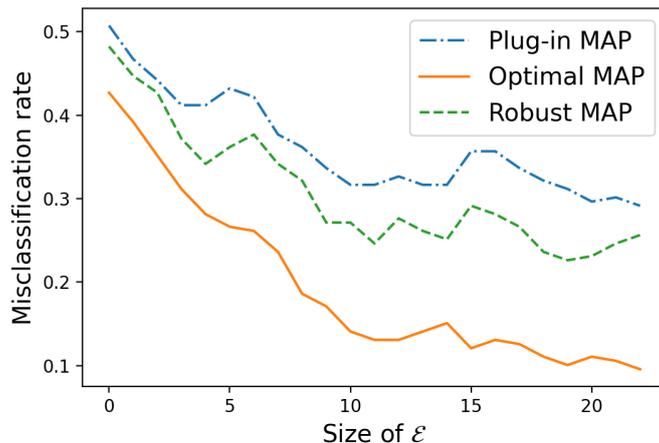


FIGURE 9. Misclassification error rate for the optimal MAP, the plug-in MAP and the Reliable MAP.

node n :

$$\hat{C}_n^R \in \arg \max_{k \in [K]} \hat{\mathbb{Q}} \left(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \hat{\mathbf{C}}_{1:m} \right).$$

Before describing in details how we compute \hat{C}_n^{OPT} , \hat{C}_n^{PI} and \hat{C}_n^R , we present some numerical results. We consider a random graph drawn from the MSBM using the matrices given in (B.1). We fully observe the graph until time $m = 100$ and we observe how the node $n = 120$ is connected to the nodes in \mathcal{E} where \mathcal{E} is equal to $\{m\}$, $\{m-1, m\}$, \dots or $\{m-25, \dots, m\}$. For those different choices of \mathcal{E} , we plot in Figure 9 the average error on the clustering of the node n using the optimal MAP, the plug-in MAP or the Reliable MAP. The plug-in MAP gives reasonable results but its misclassification rate is always lower bounded by the one of the Reliable MAP.

5.2. Computing the optimal MAP, the plug-in MAP and the Reliable MAP

In this section, we give the formula to compute the different estimates of the community of node n in the collaborative filtering problem tackled in the previous section. We denote $\mathcal{E} = \{i_1, \dots, i_S\}$ with $1 \leq i_1 < \dots < i_S \leq m$. In the following, the symbol \propto will be used in the sense that the considered quantities are equal up to some normalization factor.

- The optimal MAP selects $\hat{C}_n^{OPT} \in \arg \max_{k \in [K]} \mathbb{P}(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \mathbf{C}_{1:m})$ with

$$\begin{aligned} \mathbb{P}(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \mathbf{C}_{1:m}) &\propto \mathbb{P}(C_n = k, (X_{i,n})_{i \in \mathcal{E}} \mid \mathbf{C}_{1:m}) \\ &= \mathbb{P}((X_{i,n})_{i \in \mathcal{E}} \mid C_n = k, \mathbf{C}_{1:m}) \mathbb{P}(C_n = k \mid \mathbf{C}_{1:m}) \\ &= \prod_{i \in \mathcal{E}} Q_{C_i, k}^{X_{i,n}} (1 - Q_{C_i, k})^{1 - X_{i,n}} \cdot (P^{n-m})_{C_m, k}. \end{aligned}$$

- The plug-in MAP selects $\hat{C}_n^{PI} \in \arg \max_{k \in [K]} \mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})}(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \mathbf{C}_{1:m} = \hat{\mathbf{C}}_{1:m})$ with

$$\mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})}(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \mathbf{C}_{1:m} = \hat{\mathbf{C}}_{1:m})$$

$$\begin{aligned}
&\propto \mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})} \left(C_n = k, (X_{i,n})_{i \in \mathcal{E}} \mid \mathbf{C}_{1:m} = \hat{\mathbf{C}}_{1:m} \right) \\
&= \mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})} \left((X_{i,n})_{i \in \mathcal{E}} \mid C_n = k, \mathbf{C}_{1:m} = \hat{\mathbf{C}}_{1:m} \right) \mathbb{P}_{(\hat{\mu}, \hat{P}, \hat{Q})} \left(C_n = k \mid \mathbf{C}_{1:m} = \hat{\mathbf{C}}_{1:m} \right) \\
&= \prod_{i \in \mathcal{E}} \hat{Q}_{\hat{C}_i, k}^{X_{i,n}} (1 - \hat{Q}_{\hat{C}_i, k})^{X_{i,n}} \cdot \left(\hat{P}^{n-m} \right)_{\hat{C}_m, k}.
\end{aligned}$$

– The Reliable MAP selects $\hat{C}_n^R \in \arg \max_{k \in [K]} \hat{Q} \left(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \hat{\mathbf{C}}_{1:m} = \hat{\mathbf{C}}_{1:m} \right)$ with

$$\begin{aligned}
\hat{Q} \left(C_n = k \mid (X_{i,n})_{i \in \mathcal{E}}, \hat{\mathbf{C}}_{1:m} \right) &\propto \hat{Q} \left(C_n = k, (X_{i,n})_{i \in \mathcal{E}}, \hat{\mathbf{C}}_{1:m} \right) \\
&= \hat{Q} \left((X_{i,n})_{i \in \mathcal{E}}, \hat{\mathbf{C}}_{1:m} \mid C_n = k \right) \underbrace{\hat{Q} \left(C_n = k \right)}_{=:(*)}.
\end{aligned}$$

We have easily $\hat{Q} \left(C_n = k \right) = \sum_{l \in [K]} \hat{\mu}_l \left(\hat{P}^n \right)_{l,k}$. Moreover, denoting

$$\forall i \in [n], \forall c, k \in [K], \quad \hat{l}_{i,n}(c, k) = \hat{Q}_{c,k}^{X_{i,n}} (1 - \hat{Q}_{c,k})^{1-X_{i,n}},$$

we have

$$\begin{aligned}
(*) &= \hat{Q} \left((X_{i,n})_{i \in \mathcal{E}}, \hat{\mathbf{C}}_{1:m} \mid C_n = k \right) \\
&= \hat{Q} \left((X_{i_j,n})_{j \in [S]}, \hat{\mathbf{C}}_{1:m} \mid C_n = k \right) \\
&= \sum_{c_{i_1} \in [K]} \hat{Q} \left(C_{i_1} = c_{i_1}, \hat{\mathbf{C}}_{1:i_1} \right) \hat{Q}_{c_{i_1}, k}^{X_{i_1,n}} (1 - \hat{Q}_{c_{i_1}, k})^{X_{i_1,n}} \hat{Q} \left((X_{i_j,n})_{j \in \{2, \dots, S\}}, \hat{\mathbf{C}}_{i_1+1:m} \mid C_{i_1} = c_{i_1}, C_n = k \right).
\end{aligned}$$

Replacing $\hat{Q} \left(C_{i_1} = c_{i_1}, \hat{\mathbf{C}}_{1:i_1} \right)$ by its estimate $\hat{\alpha}_{c_{i_1}}(i_1)$ provided by the Baum-Welch algorithm, we have

$$\begin{aligned}
(*) &\approx \sum_{c_{i_1} \in [K]} \hat{\alpha}_{c_{i_1}}(i_1) \hat{Q}_{c_{i_1}, k}^{X_{i_1,n}} (1 - \hat{Q}_{c_{i_1}, k})^{X_{i_1,n}} \\
&\quad \cdot \hat{Q} \left((X_{i_j,n})_{j \in \{2, \dots, S\}}, \hat{\mathbf{C}}_{i_1+1:m} \mid C_{i_1} = c_{i_1}, C_n = k \right) \\
&= \sum_{c_{i_1}, c_{i_2} \in [K]} \hat{\alpha}_{c_{i_1}}(i_1) \hat{l}_{i_1,n}(c_{i_1}, k) \hat{Q} \left(\hat{\mathbf{C}}_{i_1+1:i_2}, C_{i_2} = c_{i_2} \mid C_{i_1} = c_{i_1} \right) \hat{l}_{i_2,n}(c_{i_2}, k) \\
&\quad \cdot \hat{Q} \left((X_{i_j,n})_{j \in \{3, \dots, S\}}, \hat{\mathbf{C}}_{i_2+1:m} \mid C_{i_2} = c_{i_2}, C_n = k \right).
\end{aligned}$$

Once again, we can replace $\hat{Q} \left(\hat{\mathbf{C}}_{i_1+1:i_2}, C_{i_2} = c_{i_2} \mid C_{i_1} = c_{i_1} \right)$ by its estimate $\hat{\chi}_{c_{i_1}, c_{i_2}}^{(i_1, i_2)}$ given by the Baum-Welch algorithm (see Eq. (4.13)). Iterating this procedure, one can derive the following approximation of (*):

$$\begin{aligned}
(*) &\approx \sum_{c_{i_1}, c_{i_2} \in [K]} \hat{\alpha}_{c_{i_1}}(i_1) \hat{l}_{i_1,n}(c_{i_1}, k) \hat{\chi}_{c_{i_1}, c_{i_2}}^{(i_1, i_2)} \hat{l}_{i_2,n}(c_{i_2}, k) \hat{Q} \left((X_{i_j,n})_{j \in \{3, \dots, S\}}, \hat{\mathbf{C}}_{i_2+1:m} \mid C_{i_2} = c_{i_2}, C_n = k \right) \\
&= \dots
\end{aligned}$$

$$\approx \sum_{c_{i_1}, \dots, c_{i_S} \in [K]} \hat{\alpha}_{c_{i_1}}(i_1) \prod_{j=1}^{S-1} \left(\hat{l}_{i_j, n}(c_{i_j}, k) \hat{\chi}_{c_{i_j}, c_{i_{j+1}}}^{(i_j, i_{j+1})} \right) \hat{\beta}_{c_{i_S}}(i_S) \hat{l}_{i_S, n}(c_{i_S}, k).$$

Thanks to the previous computations, we deduce that the estimated community for node n provided by the Reliable MAP is

$$\arg \max_{k \in [K]} \sum_{l \in [K]} \hat{\mu}_l \left(\hat{P}^n \right)_{l, k} \cdot \sum_{c_{i_1}, \dots, c_{i_S} \in [K]} \hat{\alpha}_{c_{i_1}}(i_1) \prod_{j=1}^{S-1} \left(\hat{l}_{i_j, n}(c_{i_j}, k) \hat{\chi}_{c_{i_j}, c_{i_{j+1}}}^{(i_j, i_{j+1})} \right) \hat{\beta}_{c_{i_S}}(i_S) \hat{l}_{i_S, n}(c_{i_S}, k).$$

6. IMPLEMENTATION AND EXPERIMENTS

6.1. Complexity and implementation

The clustering algorithm used is a SDP method and, as a consequence, its time complexity scales with n^3 , while the complexity of the Baum-Welch algorithm is of order K^2n . From here, computing $\hat{\eta}_i(\hat{\mathbf{c}}_{1:n})$ for all $i \in [n]$ requires K^3n^2 operations (see Sect. 3.b of the notebook *experiments.ipynb*). Regarding the collaborative filtering task, the Reliable MAP estimator from Section 5.1 has a time complexity of order K^4n^2 (see method *collaborative_filtering_robustMAP* in the file *markovianSBM/BaumWelch.py*).

6.2. Inferring the number of communities

In this section, we propose a heuristic based on the learned emission probabilities from the Baum-Welch algorithm to estimate the number of communities K of our model. The proposed approach consists in running the Baum-Welch algorithm for a finite list of possible number of clusters $\{K_{min}, \dots, K_{max}\} = \mathcal{K} \subset \mathbb{N}^*$. For each $K \in \mathcal{K}$, we denote $\hat{O}^{(K)}$ the matrix of emission probabilities learned by Algorithm 4 when we consider that the number of communities is K . For any $K \in \mathcal{K}$, we define

$$M^{(K)} := \max_{k, l \in [K], k \neq l} \left\{ \hat{O}_{l, k}^{(K)} + \hat{O}_{k, l}^{(K)} \right\}.$$

For any $K \in \mathcal{K}$ and any $k, l \in [K]$, $k \neq l$, $\hat{O}_{l, k}^{(K)} + \hat{O}_{k, l}^{(K)}$ represents the probability that the clustering algorithm predicts community k or l if the true cluster is the other one.

- When K is less than or equal to the true number of clusters, $M^{(K)}$ stays small as soon as the graph is large enough and as the clustering algorithm used is efficient.

Typically if the observed graph has 4 hidden communities, we expect the clustering algorithm run with $K = 3$ to return a partition of the nodes that is - in the ideal case - of the form $(G_1, G_2, G_3 \cup G_4)$. In this ideal case, it holds $M^{(3)} = 0$. Stated otherwise, we believe that the clustering algorithm will merge true clusters to output a partition of nodes with only K groups while the true underlying structure of the graph contains more than K clusters.

- When K becomes greater than the true number of clusters, $M^{(K)}$ is larger compared to the previous case because at least one true cluster will be arbitrarily split in two different groups by the clustering algorithm.

For example, let us consider that the observed graph has 4 hidden communities and that we run the clustering algorithm with $K = 5$. In this case, a possible way to interpret the latent structure of the graph into 5 groups is to consider the partition of the nodes given by $(G_1, G_2, G_3, G_4^{(1)}, G_4^{(2)})$ where $G_4^{(1)} \cup G_4^{(2)} = G_4$. Stated otherwise, we have arbitrarily split one true cluster in two groups. Considering that the true labels of the nodes is given by $(G_1, G_2, G_3, G_4^{(1)}, G_4^{(2)})$, we understand that the clustering algorithm will have difficulties to obtain the label 4 for nodes in $G_4^{(1)}$ and the label 5 for nodes in $G_4^{(2)}$.

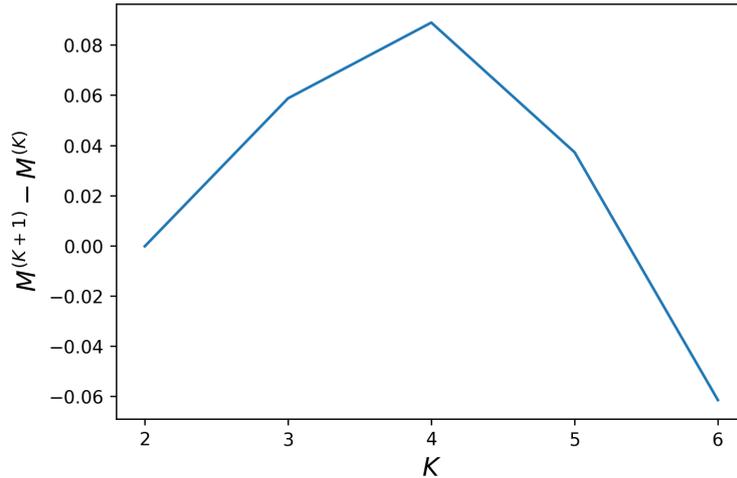


FIGURE 10. $K \mapsto M^{(K+1)} - M^{(K)}$ working with the connectivity matrix Q from (B.1).

because nodes belonging to $G_4 = G_4^{(1)} \cup G_4^{(2)}$ have the same connection probabilities. Hence, we expect clustering errors leading to $\hat{O}_{4,5} + \hat{O}_{5,4} > 0$ and thus $M^{(K)} > 0$.

Based on this remark, we propose to estimate the number of communities by choosing the value $K \in \mathcal{K}$ leading to the larger positive jump of the function $K \mapsto M^{(K)}$ namely

$$\hat{K} \in \arg \max_{K \in \{K_{min}, \dots, K_{max}-1\}} \left\{ M^{(K+1)} - M^{(K)} \right\}.$$

First we test our method with a graph of size $n = 110$ and with $K = 4$ communities using the transition kernel P and the connectivity matrix Q defined by (B.1). Figure 10 shows that our approach allows to estimate the correct number of communities $K = 4$.

Example of application. We also test our procedure on a real network corresponding to American football games between Division IA colleges during regular season Fall 2000. Two teams are connected if they played against each other. The nodes have values that indicate which conferences the corresponding team belongs to. We worked with 6 different conferences ³. Figure 11 shows that our procedure infers the correct number of communities.

Related literature. Standard model selection methods such as the Bayesian Information Criterion (BIC) or the Akaike Information Criterion (AIC) (*cf.* [11]) are based on the likelihood of the observed data, which is intractable in the MSBM (and in particular in the SBM). To tackle this issue, [48] and [23] used the Integrated Completed Likelihood (ICL) criterion in the SBM. The ICL method is based on an asymptotic approximation of the integrated complete-data likelihood. The ICL was the first model-based criterion developed for SBM. Let us mention that other approaches have been explored such as the Integrated Likelihood Variational Bayes (*cf.* [42]) where the authors aim at proposing a model selection method that relies on a non-asymptotic approximation of the marginal likelihood.

Comparison with the ICL criterion. We consider $K = 4$ communities and we sample graphs from the MSBM model using the transition kernel P and the connectivity matrix Q defined by (B.1). For each

³namely Atlantic Coast, Big East, Big Ten, Big Twelve, Conference USA, Mid-American.

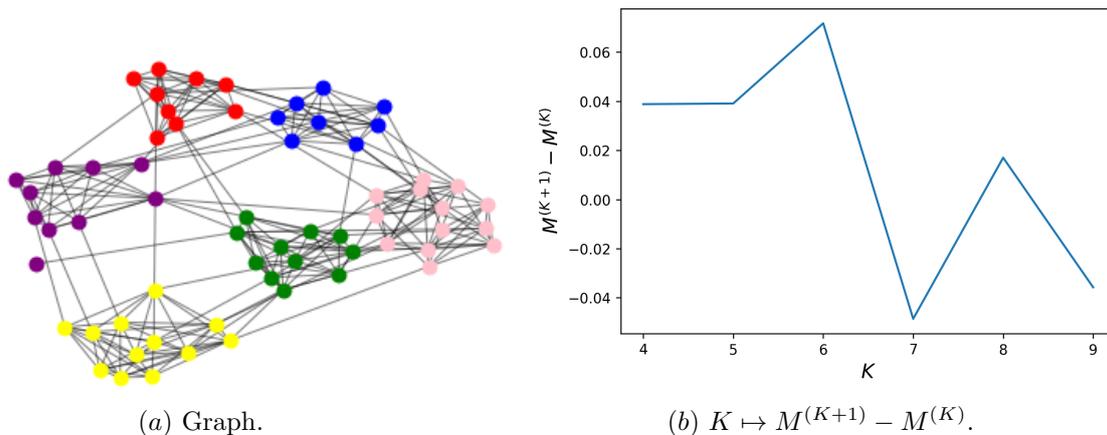


FIGURE 11. We test our model selection method on the *football network* from the Networkx python package.

$n \in \{30, 40, 60, 80\}$, we sample randomly 100 graphs of size n and we use the estimated number of hidden clusters using our model selection method and using the ICL criterion. For each possible cluster size $K \in \mathbb{N}$, we compute the proportion of graphs for which we estimate a number of hidden communities equal to K . Figure 12 shows the result of these experiments. [48] noticed that the ICL criterion tends to underestimate the number of classes when dealing with small networks and we recover similar results with our experiments as shown in Figure 12.(a) and (b). These experiments show that the model selection proposed in this paper should be preferred for small graphs. When the size of the graph is getting larger, the ICL criterion should be preferred to infer the number of hidden communities. We used the Python package *sparsebm* to obtain the results when applying the ICL criterion.

6.3. Application on real data

Migratory animals are essential components of the ecosystems that support all life on Earth. By acting as pollinators and seed distributors they contribute to ecosystem structure and function. They provide food for other animals and regulate the number of species in ecosystems. Migratory animals are potentially very effective indicators of environmental changes that affect us all.

In [38], the authors proposed a periodic Markov model on a spatial migration network to formally describe the process of animal migration on the population level. They built their dataset using the Movebank data repository (see [39]) that provides historic of animal movements. We propose to test our approach on this dataset. The data is publicly available [here](#)⁴ and our experiments can be reproduced with the notebook *experiments.ipynb*.

Description of the dataset. The dataset presents the locations of several white-fronted geese with timestamps. The animals have been tracked from 2006 to 2010. Each location can be associated with a class using classes defined from Argos User’s Manual 2011. We refer to [38] for details.

We focus on one specific white-fronted goose and we keep the list of its chronological locations between 2006 and 2010 for four location classes. Nodes correspond to the entries of the previous sequence of locations of the animal while communities are the classes associated to each location. In our network, we connect two nodes if the distance between the corresponding precise locations (given with latitude and longitude coordinates) is smaller than some specified threshold.

Results. With Figure 13(a), we show that the model selection method of Section 6.2 allows to retrieve the correct number of clusters on our dataset. In order to evaluate the performance of our reliable link prediction method, we compute the transition matrix P and the connection matrix Q associated with our network. More

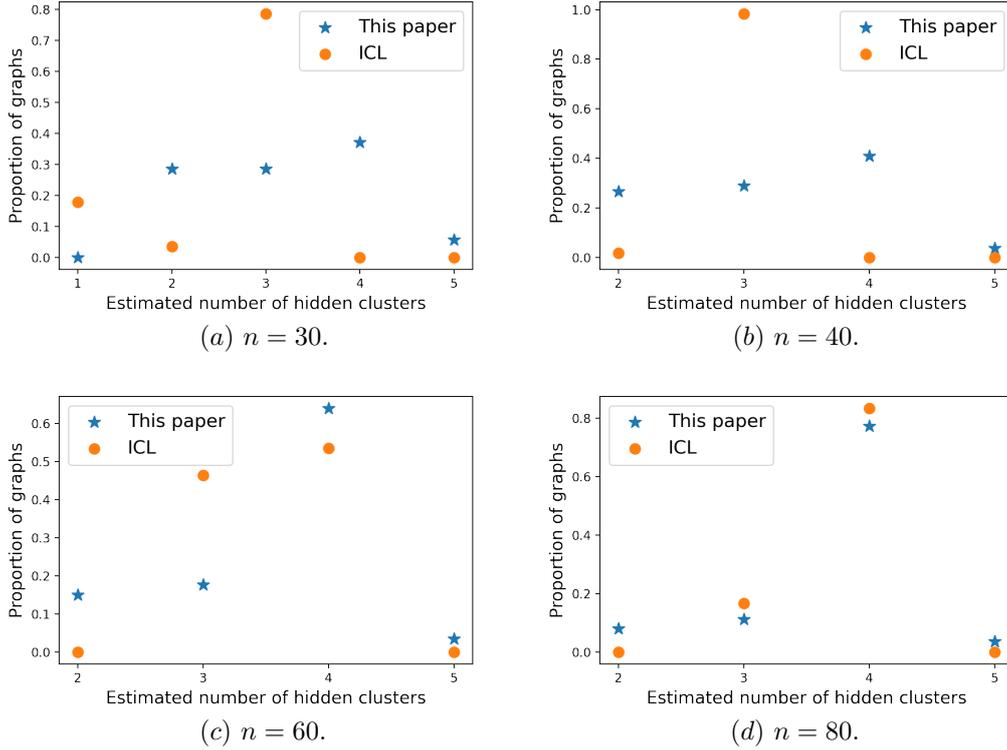


FIGURE 12. Comparison of the model selection method proposed in this paper and the ICL criterion. We show the proportion of graphs leading to a specific estimate for the number of hidden communities.

precisely, we define $\forall k, l \in [K]$,

$$Q_{k,l} := \begin{cases} \frac{1}{|G_k| \cdot |G_l|} \sum_{i \in G_k} \sum_{j \in G_l} X_{i,j} & \text{if } k \neq l \\ \frac{1}{|G_k| \cdot (|G_k| - 1)} \sum_{i,j \in G_k} X_{i,j} & \text{if } k = l \end{cases} \quad \text{and} \quad P_{k,l} := \frac{n}{n-1} \frac{\sum_{i=1}^{n-1} \mathbb{1}_{(C_i, C_{i+1})=(k,l)}}{\sum_{i=1}^n \mathbb{1}_{C_i=k}},$$

where X is the adjacency matrix, C_i is the community of node $i \in [n]$ and G_k is the set of nodes with label $k \in [K]$. We use these matrices to compute the posterior probabilities $\left(\eta_i(\hat{\mathbf{C}}_{1:n})\right)_{i \in [n]}$ (see Eq. (4.2)) and we can compare them with the estimations given by the plug-in approach and the reliable approach from Section 4. Figure 13.(b) shows that the reliable approach allows to significantly improve the estimate of the posterior probabilities.

Comments. On simulated data with a small number of clusters, when n gets larger, the clustering algorithm will recover (almost) perfectly the true partition. In that case, it is clear that the reliable version cannot improve drastically the plug in method since this latter (almost) coincides with the Bayes classifier. However, real datasets never fit a particular model and recovering the true partition is really unlikely even for very large graphs. In such cases, our method is of great interest to provide reliable estimations for link prediction despite clustering errors.

⁴<https://www.datarepository.movebank.org/handle/10255/move.747>

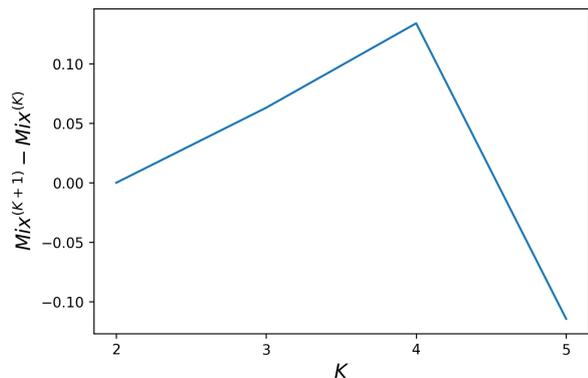
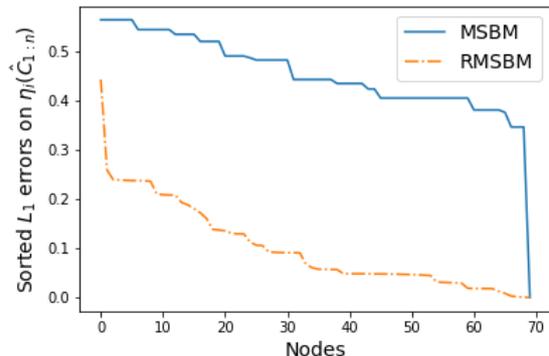
(a) $K \mapsto M^{(K+1)} - M^{(K)}$.(b) L^1 errors between $\hat{\eta}_i(\hat{C}_{1:n})$ (resp. $\hat{\eta}_i^R(\hat{C}_{1:n})$) and $\eta_i(\hat{C}_{1:n})$.

FIGURE 13. We test our model on the bird migrations dataset from [38].

7. CONCLUSION

We introduce a growth model for community-based random graphs and we addressed fundamental questions for practical applications in this setting. We propose a general approach to get reliable estimation of the probability of connection between a future node and the nodes already present in the graph. We show how the emission probabilities learned by our method can be used to derive interesting properties on the clustering algorithm and to estimate the number of communities. Our approach is shown to improve the results of the plug-in method for link prediction and collaborative filtering problems. Let us mention that we chose a recent SDP method to derive theoretical guarantees but all our work is directly transferable to other clustering algorithms.

Our approach could also be used to study the performance of the algorithm in the presence of spurious links in the graph by looking at the estimated matrix of emission probabilities \hat{O} . This could be an interesting direction of research to choose among the wide variety of clustering algorithms for practical cases where the graph may contain some fake links.

Acknowledgements. This work was supported by a grant from Région Ile-de-France.

APPENDIX

Guidelines for the Appendix

Section A: Algorithms

In Section A, we present the SDP algorithm used to recover the communities. We give theoretical elements allowing the reader to understand this approach.

Section B: Experiments

We provide additional experiments. In particular, we present our Markovian hypothesis testing procedure and an application of our methods on real data. We also give details about a possible application on real data for items recommendation or for the analysis of tumor growth.

Sections C and D: Proofs and Complements

Section C contains the proofs of the main statements of our paper while Section D is a complement recalling basic definitions and properties on Markov chains.

Section E: The Baum-Welch algorithm with information gap

We provide further details on our reliable prediction methods when information is missing in the observed network.

Section F: Notations

APPENDIX A. CLUSTERING ALGORITHM

A.1 Partial recovery bound in SBMs with fixed assignment of the communities

In [29], the authors introduce a relaxed version of the K -means algorithms on the columns of the adjacency matrix. One specificity of their algorithm is the fact that they are working with the square of the adjacency matrix. This choice allows them to tackle problems outside of the assortative setting and with a wide set of possible connectivity matrices Q contrary to previous works.

Theorem A.1 presents the result of Verzelen and Giraud in the SBM framework with a connectivity matrix $Q = \alpha_n Q_0$.

Theorem A.1. [29, Theorem 2]

Assume that $\|Q_0\|_\infty \leq L$. The size of the community $k \in [K]$ will be denoted m_k . The size of the smallest community will be denoted m . We define the signal-to-noise ratio $s^2 = \Delta^2/(\alpha_n L)$, where $\Delta^2 = \min_{k \neq j} \Delta_{k,j}^2$ with

$$\Delta_{k,j}^2 = \sum_l m_l (Q_{k,l} - Q_{j,l})^2 = \alpha_n^2 \sum_l m_l ((Q_0)_{k,l} - (Q_0)_{j,l})^2.$$

Then, there exist three positive constants c, c', c'' , such that for any $1/m \leq \alpha_n L \leq 1/\log(n)$,

$$\frac{1}{m} \leq \beta \leq \beta(\alpha_n L) := \frac{K^3}{n} e^{4n\alpha_n L}$$

and

$$s^2 \geq c'' n/m,$$

with probability at least $1 - c/n^2$,

$$\text{err}(\hat{G}, G) \leq e^{-c' s^2}.$$

In particular, since

$$s^2 = \frac{\alpha_n \min_{k \neq j} \sum_{l \in [K]} m_l ((Q_0)_{k,l} - (Q_0)_{j,l})^2}{L} \geq \frac{\alpha_n m D^2}{L},$$

we get that with probability at least $1 - c/n^2$,

$$-\log(\text{err}(\hat{G}, G)) = \Omega(m\alpha_n).$$

A.2 Presentation of the SDP-based clustering algorithm

In this Section, we present how we estimate the partition of the nodes \hat{G} when communities are assigned using a Markovian dynamic. Our main result Theorem 3.1 shows that we are able to achieve

$$-\log \text{err}(\hat{G}, G) = \Omega(n\alpha_n).$$

Stated otherwise, we get a misclassification error that decays exponentially fast with respect to $n\alpha_n$. We recover the convergence rate recently proved in [29] in the standard SBM⁵ when the size of the smallest cluster scales linearly with n like in our case. To reach this result, we use the SDP algorithm proposed by Giraud and Verzelen in [29]. In the following, we expose how the method works.

Suppose the community of each node in the graph has been assigned. In all this subsection, all the communities are considered fixed. We denote X the adjacency matrix of the graph and we refer to Theorem A.1 for the definition of $(m_k)_k$ and m . [29] are interested in solving optimization problem similar to the following

$$\begin{aligned} & \max_{B \in \mathcal{C}'} \langle X, B \rangle \quad \text{with} & (A.1) \\ \mathcal{C}' := & \{B : \text{PSD}, B_{k,l} \geq 0, |B|_1 = \sum_k m_k^2\}, \end{aligned}$$

where PSD means that B is positive semidefinite and where $|\cdot|_1$ is the element-wise l_1 norm, namely the sum of the absolute values of all entries of a given matrix.

We remind that, dealing with two communities, when the values of the probability matrix Q are a constant p on the diagonal and another constant q off the diagonal with $p > q$, we are in the assortative case. In the assortative setting, optimization problems like (A.1) have been widely used to recover communities, see [16, 26, 30, 32, 55]. Those SDP programs are trying to maximize the probability of connection between nodes belonging to the same community. Therefore, they cannot be used directly to solve community detection outside of the assortative framework.

[53] showed that any partition G of $[n]$ can be uniquely represented by a $n \times n$ matrix $B^* \in \mathbb{R}^{n \times n}$ defined by $\forall i, j \in [n]$,

$$B_{i,j}^* = \begin{cases} \frac{1}{m_k} & \text{if } i \text{ and } j \text{ belong to community } k \\ 0 & \text{otherwise.} \end{cases}$$

The set of such matrices B^* that can be built from a particular partition of $[n]$ in K groups is defined by

$$\mathcal{S} = \{B \in \mathbb{R}^{n \times n} : B^\top = B, B^2 = B, \text{Tr}(B) = K, \\ B\mathbf{1} = \mathbf{1}, B \geq 0\},$$

where $\mathbf{1} \in \mathbb{R}^n$ is the n -dimensional vector with all entries equal to one and where $B \geq 0$ means that all entries of B are nonnegative. [53] proved that solving the K -means problem

$$\text{Crit}(G) = \sum_{k=1}^K \sum_{i \in G_k} \left\| X_{:,i} - \frac{1}{|G_k|} \sum_{j \in G_k} X_{:,j} \right\|_2^2,$$

is equivalent to

$$\max_{B \in \mathcal{S}} \langle XX^\top, B \rangle. \quad (A.2)$$

Writing B^* an optimal solution of (A.2), an optimal solution for the K -means problem is obtained by gathering indices $i, j \in [n]$ such that $B_{i,j}^* \neq 0$. The set \mathcal{S} is not convex and the authors of [29] propose the following

⁵See Theorem A.1.

relaxation of problem (A.2)

$$\begin{aligned} \hat{B} \in \arg \max_{B \in \mathcal{C}_\beta} \langle XX^\top, B \rangle \text{ with} & \quad (\text{A.3}) \\ \mathcal{C}_\beta := \{B \in \mathbb{R}^{n \times n} : \text{symmetric, } \text{Tr}(B) = K, \\ & B\mathbf{1} = \mathbf{1}, 0 \leq B \leq \beta\}, \end{aligned}$$

where $K/n \leq \beta \leq 1$. The constraint $B \leq \beta$ allows to deal with sparse graphs. Indeed, when $\alpha_n = o(\log(n)/n)$, solving (A.3) without this constraint will produce unbalanced partition.

At this step, we cannot ensure that \hat{B} belongs to \mathcal{S} and a final refinement is necessary to end up with a clustering of the nodes of the graph. This final rounding step is achieved by running a K -medoid algorithm on the rows of \hat{B} . Given a partition $\{G_1, \dots, G_k\}$ of the n nodes of the graph into K communities, we define the related membership matrix $A \in \mathbb{R}^{n \times K}$ where $A_{i,k} = \mathbf{1}_{i \in G_k}$. Working on the rows of \hat{B} , a K -medoid algorithm tries to find efficiently a pair (\hat{A}, \hat{M}) with $\hat{A} \in \mathcal{A}_K$, $\hat{M} \in \mathbb{R}^{K \times n}$, $\text{Rows}(\hat{M}) \subset \text{Rows}(\hat{B})$ satisfying for some $\rho > 0$

$$|\hat{A}\hat{M} - \hat{B}|_1 \leq \rho \min_{A \in \mathcal{A}_K, \text{Rows}(M) \subset \text{Rows}(\hat{B})} |AM - \hat{B}|_1, \quad (\text{A.4})$$

where \mathcal{A}_K is the set of all possible membership matrices and $\text{Rows}(\hat{B})$ the set of all rows of \hat{B} . The K -medoids algorithm proposed in [15] gives in polynomial time a pair (\hat{A}, \hat{M}) satisfying the inequality (A.4) with $\rho = 7$. From \hat{A} we are able to define the final partition of the nodes of the graph by setting

$$\forall k \in [K], \quad \hat{G}_k = \{i \in [n] : \hat{A}_{i,k} = 1\}.$$

Remark.

As highlighted in [29], the parameter β can not be computed since L is unknown. Verzelen and Giraud propose to set β to value $\hat{\beta} = \frac{K^3}{n} e^{2nd_X} \wedge 1$, where d_X denotes the density of the graph. We end up with the Algorithm 5 to estimate the communities in the SBM.

Algorithm 5 Algorithm to estimate the partition of the nodes of the graph.

Data: Adjacency matrix X of a graph $G = (V, E)$, Number of communities K .

- 1: Compute the density of the graph $d_X = \frac{2|E|}{n(n-1)}$ and set $\hat{\beta} = \frac{K^3}{n} e^{2nd_X} \wedge 1$.
 - 2: Find $\hat{B} \in \arg \max_{B \in \mathcal{C}_{\hat{\beta}}} \langle XX^\top, B \rangle$ (using for example the interior-point method).
 - 3: Run the K -medoids algorithm from [15] on the rows of \hat{B} . Note $\hat{A} \in \{0, 1\}^{n \times K}$ the membership matrix obtained.
 - 4: Define $\forall k \in [K], \quad \hat{G}_k = \{i \in [n] : \hat{A}_{i,k} = 1\}$ and $\forall i \in [n], \quad \hat{C}_i = k$ where $k \in [K]$ is such that $\hat{A}_{i,k} = 1$.
-

APPENDIX B. ADDITIONAL EXPERIMENTS

B.1 Markovian dynamic testing

We illustrate our model on a toy example with $K = 4$ communities, with the transition matrix P and the connectivity matrix Q defined by

$$P = \begin{bmatrix} 0.1 & 0.3 & 0.5 & 0.1 \\ 0.45 & 0.15 & 0.2 & 0.2 \\ 0.15 & 0.3 & 0.1 & 0.45 \\ 0.25 & 0.3 & 0.1 & 0.35 \end{bmatrix} \quad \text{and} \quad Q = \begin{bmatrix} 0.22 & 0.48 & 0.29 & 0.44 \\ 0.48 & 0.61 & 0.18 & 0.15 \\ 0.29 & 0.18 & 0.08 & 0.87 \\ 0.44 & 0.15 & 0.87 & 0.27 \end{bmatrix}. \quad (\text{B.1})$$

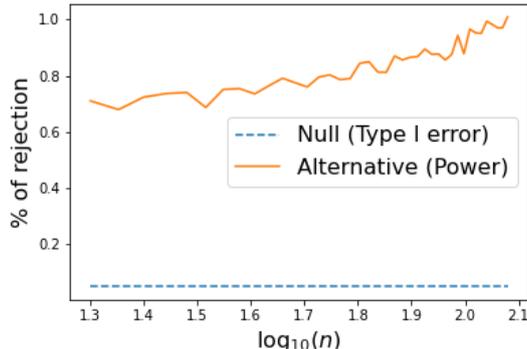


FIGURE B.1. Power of our hypothesis test with level 5%.

We propose a hypothesis test to statistically distinguish between an independent assignment of the communities with the distribution π and a Markovian assignment with a non-trivial dependence structure. More precisely, we consider the null \mathbb{H}_0 : *communities are independently assigned with distribution π* where π denotes the stationary distribution of the transition matrix P from (B.1). Our test is based on estimate \hat{P} of the

transition matrix. The null can be rephrased as \mathbb{H}_0 : $P = P^0$ where $P^0 := \begin{bmatrix} \pi \\ \vdots \\ \pi \end{bmatrix}$. One can use any *black-box*

goodness-of-fit test comparing \hat{P} to P^0 . Figure B.1 shows the power of this hypothesis test with level 5% (Type I error) and using the χ^2 -test described by ([10], Sect. 2.4). We choose alternative given by the matrices defined in (B.1). Rejection region is calibrated (*i.e.*, threshold of the χ^2 -test) by *Monte Carlo simulations under the null*. It allows us to control Type I error as depicted by dotted blue line. We run our algorithm to estimate the transition matrix from which we compute the χ^2 -test statistic namely

$$S_n := \sum_{1 \leq k, l \leq K} \hat{G}_k \frac{(\hat{P}_{k,l} - \pi_l)^2}{\pi_l} \text{ with } \hat{G}_k = \sum_{i=1}^n \mathbb{1}_{\hat{C}_i=k}.$$

S_n is known to be asymptotically distributed as a χ^2 random variable with $K(K-1)$ degrees of freedom. Figure B.1 shows that for graphs of size larger than 100, the rejection rate is almost 1 under the alternative (Type II error is almost zero), the test is very powerful.

B.2 Experiments with 2 communities

We test our algorithm on a toy example with $K = 2$ communities, $\alpha_n = 1$ and with the following matrices:

$$P = \begin{bmatrix} 0.2 & 0.8 \\ 0.6 & 0.4 \end{bmatrix} \text{ and } Q_0 = \begin{bmatrix} 0.8 & 0.2 \\ 0.1 & 0.3 \end{bmatrix}. \quad (\text{B.2})$$

The Figure B.2 shows the evolution of the max norm of the difference between the true transition matrix P and our estimate \hat{P} when the size of the graph is increasing. For each point, the bar represents the standard deviation of the max norm error computed over thirty randomly generated graphs with the same number of nodes and using the matrices P and Q defined by (B.2). Those numerical results are consistent with Theorem 3.4: we recover the parametric convergence rate with our estimator of the transition matrix.

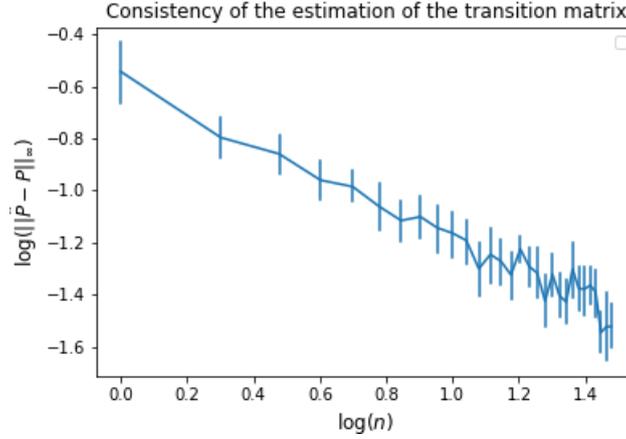


FIGURE B.2. We plot the log of the max norm of the difference between the true transition matrix P and our estimate \hat{P} according to the log of the number of nodes in the graph.

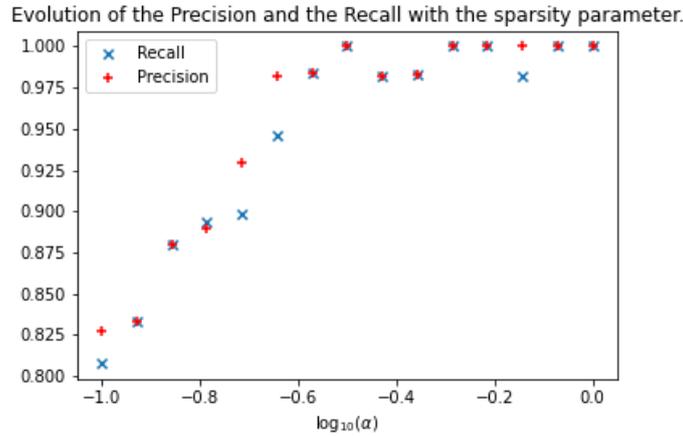


FIGURE B.3. Precision and recall of the studied binary classification problem when $K = 2$ and $Q = \alpha Q_0$.

With Figure B.3, we shed light on the influence of the average degree of the nodes on the performance of our algorithm. We plot the recall and the precision of the output of our algorithm with a graph sampled from SBM with a Markovian assignment of the communities using $n = 100$ nodes, a transition matrix P defined in (B.2) and a connectivity matrix $Q = \alpha Q_0$ where Q_0 is defined in (B.2) and α varies on a log scale between 0.1 and 1. We show the recall and the precision with respect to the \log_{10} of the sparsity parameter α . We recall that in a binary classification problem, the precision is the ratio between the number of examples labeled 1 that belong to class 1 and the number of examples labeled 1. The recall is the ratio between the number of examples labeled 1 that belong to class 1 and the number of examples that belong to class 1. In our context, those definitions read as

$$\text{precision} = \frac{\sum_{i=1}^n \mathbb{1}\{\hat{C}_i = 1, C_i = 1\}}{\sum_{i=1}^n \mathbb{1}\{\hat{C}_i = 1\}} \quad \text{and} \quad \text{recall} = \frac{\sum_{i=1}^n \mathbb{1}\{\hat{C}_i = 1, C_i = 1\}}{\sum_{i=1}^n \mathbb{1}\{C_i = 1\}}.$$

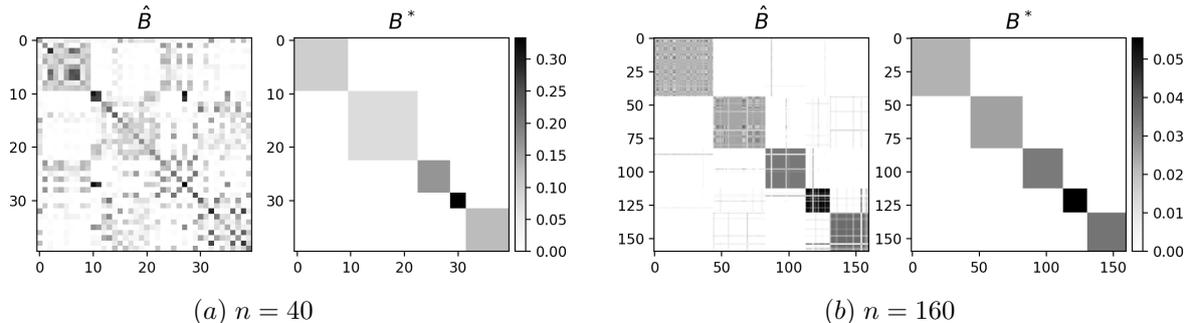


FIGURE B.4. Visualization of the matrix B^* solution of equation (A.1) and its approximation \hat{B} obtained by solving the SDP of equation (A.3) for graphs with 40 nodes (Fig. (a)) or 160 nodes (Fig. (b)).

B.3 Experiments with 5 communities

We test our algorithm on a toy example with $K = 5$ communities, with the transition matrix P and the connectivity matrix Q defined by

$$P = \begin{bmatrix} 0.1 & 0.3 & 0.5 & 0.01 & 0.09 \\ 0.55 & 0.15 & 0.1 & 0.05 & 0.15 \\ 0.15 & 0.3 & 0.1 & 0.2 & 0.25 \\ 0.15 & 0.05 & 0.1 & 0.5 & 0.2 \\ 0.2 & 0.3 & 0.1 & 0.05 & 0.35 \end{bmatrix} \quad \text{and} \quad Q = \begin{bmatrix} 0.6 & 0.1 & 0.15 & 0.1 & 0.2 \\ 0.2 & 0.5 & 0.35 & 0.1 & 0.4 \\ 0.4 & 0.15 & 0.6 & 0.25 & 0.05 \\ 0.4 & 0.1 & 0.1 & 0.2 & 0.55 \\ 0.3 & 0.35 & 0.2 & 0.1 & 0.7 \end{bmatrix}. \quad (\text{B.3})$$

We order the nodes of the graph such that the true partition of the nodes is given by $G_1 = \{1, \dots, m_1\}$, $G_2 = \{m_1 + 1, \dots, m_1 + m_2\}$, \dots , $G_5 = \{\sum_{j=1}^4 m_j + 1, \dots, n\}$. Figure B.4 shows the matrix B^* solution of equation (A.1) and its approximation \hat{B} obtained by solving the SDP of equation (A.3). Thanks to the node ordering, the matrix B^* has a block diagonal structure where each entry of one block is equal to the inverse of the size of the associated cluster. Figure B.4.(a) allows us to compare the matrices B^* and \hat{B} when the number of nodes in the graph is equal to 40 while Figure B.4.(b) deals with a graph of size 160.

For a graph sampled with a size equal to 40, Figure B.4.(a) shows us that the SDP algorithm defined in Algorithm 1 is able to capture relevant information about the clustering of the nodes in communities 1, 2 and 5. However, we see that using a number of nodes equal to 40 is not enough to distinguish nodes belonging to community 3 or 4. Figure B.4.(b) proves that increasing the size of the graph (*i.e.* for $n = 160$) allows to solve this issue. One can easily guess that running a K -medoid algorithm on the rows of the matrix \hat{B} plotted in Figure B.4.(b) will lead to an accurate clustering of the nodes of the graph. Figure B.5 shows that the log of the misclassification error decreases linearly with the size of the graph.

B.4 Overestimation of the size of small clusters using the algorithm of Section A

In Figure 8 of the main paper (*cf.* Sect. 4.2.4), we showed the learned emission probabilities $\hat{O}_{k,l}$, $k, l \in [K]$ and we concluded that the size of the smaller cluster was overestimated by the clustering algorithm from [29]. In this section, we conduct additional experiments to check if we reach the same conclusion. We consider two communities and a probability vector $\pi = [\gamma, 1 - \gamma]$ with $\gamma \in (0, 0.5)$. We consider a SBM with 100 nodes with an independent assignment of the communities according to the vector π . The connectivity matrix is defined by $Q = \begin{bmatrix} 0.8 & 0.05 \\ 0.05 & 0.8 \end{bmatrix}$. Varying the value of γ between 0.01 and 0.4, we compute the clustering of the nodes (denoted (\hat{G}_1, \hat{G}_2)) obtained with the SDP algorithm from [29]. We compute the ratio between the size of the estimated first cluster (namely $|\hat{G}_1|$) and the size of the true first cluster (namely $|G_1|$). Figure B.6 gives the result we

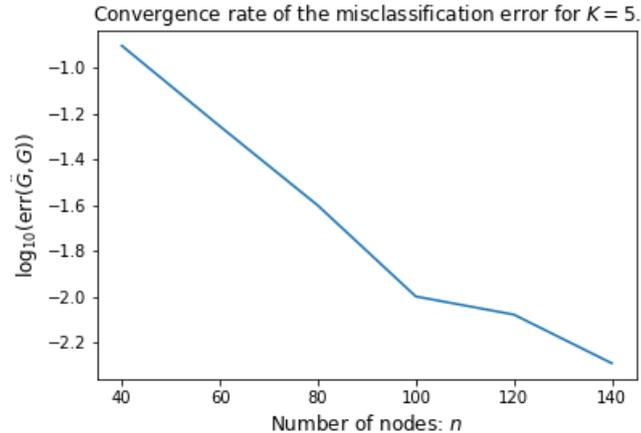


FIGURE B.5. Log of the misclassification error as a function of the size of the graphs sampled.

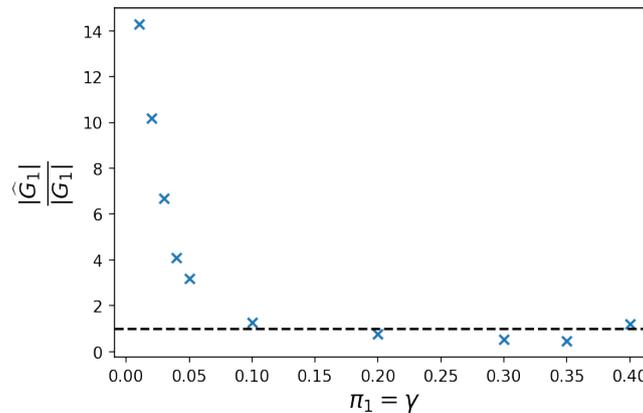


FIGURE B.6. Ratio between the size of the estimated first cluster and the true size of the first cluster. We vary the parameter γ which is such that $n\gamma$ is the expected value of the size of the first cluster.

obtain. The horizontal dotted black line corresponds to the value 1. Hence, crosses below that line show cases where the size of the first cluster is underestimated while crosses above that line show cases where the size of the first cluster is overestimated. These experiments go in the same direction as the one from Section 4.2.4: the size of small clusters tend to be overestimated by the SDP algorithm from [29].

B.5 Other potential application on real data: The example of recommendation system

In this section, we give more details on another possible application of our model for recommendation system as mentioned in the introduction. Let us remind the framework of our example. We suppose that we have access to the online purchases of different customers. For each of them, we know the dates and the product IDs of each of their purchases. Our goal is threefold: *i*) learn the category of product sold by each url *ii*) learn the purchasing behavior of each customer *iii*) use this information to suggest relevant new products to each customer. For each customer U , we have a network where nodes are product IDs (ordered by timestamp of purchase). We connect two products i and j if the ratio $W_{i,j}/\sqrt{w_i \cdot w_j}$ is larger than some threshold $\tau \in (0, 1)$, where $W_{i,j}$ is the number of clients in the dataset who bought both the products i and j , and w_i is the number of clients who bought the product i . We can proceed as follows.

1. Running our algorithm, we can infer the number of different categories of products bought by U using our heuristic from Section 6.2.
2. Then, we can learn both the category of each product and the transition matrix \hat{P} which gives the purchasing behavior of U .
3. To recommend a new product to the client U , one can use the purchasing behaviour of the client V who shares the largest number of common purchased products with U .

Let us finally mention that the connection probabilities learned from client V for categories unseen so far by U could be used as an initialization of a stochastic bandit algorithm for recommendation of products for U .

B.6 Other potential application on real data: The example of the tumor growth

To better characterize tumour heterogeneity and predict the potential response of the constituent parts of a tumour to drugs, biologists often rely on single-cell RNA sequencing data (*cf.* [21]). RNA sequencing allows the measurement of gene expression in thousands of cells in a single experiment. To identify the cell types and cell states present in a sample, unsupervised clustering are often used (*cf.* [45]).

During the cell cycle, a cell increases in size, replicates its DNA and splits into daughter cells. Hence, a given sample of tumor is structured according to an underlying evolutionary history of cancer cells. It has been reported that incorporating the pathway information (*i.e.* the parent-child hierarchical structure) is key to reach better drug prediction (*cf.* [65]).

In this context, an interesting application of the MSBM model would be to follow the ex-vivo evolution of a cell. Namely, starting from a specific cell, a biologist could choose a daughter of this cell, and a daughter of this one, etc. Stated otherwise, the biologist will follow a specific path in the genealogical tree where the root is the starting cell. Performing the RNA sequencing of each cell along the path, he could build a graph as follows:

- each node is a cell in the path selected in the genealogical tree,
- the order of the nodes is simply given by the timestep of appearance of the cells,
- the RNA sequencing gives for all cell i a vector $x^{(i)} \in \mathbb{R}^d$ where each entry $x_s^{(i)}$ corresponds to the expression level of gene s in cell i . One can then build a graph by simply putting an edge between cell i and cell j with probability equal to $\frac{1}{2}(\text{cor}(x^{(i)}, x^{(j)}) + 1)$ where $\text{cor}(x^{(i)}, x^{(j)})$ is the correlation between vector $x^{(i)}$ and $x^{(j)}$.

Related literature.

Let us mention that this approach to tackle the problem of cell clustering has already been proposed in the literature such as in [33] where the authors build a cell-cell similarity metric (based on the optimal transport cost between the probability distributions given by $\frac{x^{(i)}}{\|x^{(i)}\|_2}$ for $i \in [n]$). They use a spectral method to infer clusters considering a weighted graph whose adjacency matrix is given by $A = 1 - D$ where D is a normalized version of the matrix gathering the optimal transport costs between the pairs of cells previously computed (the normalization simply ensures that the maximum entry of D is 1, so that weights in the graph are all non negative).

Use of the MSBM in this context.

The biologist could:

- Use our model selection method from Section 6.2 to infer the number of hidden communities K in the graph.
- Run a clustering algorithm on the graph to obtain K groups of cells.
- Estimate the Markov kernel P to understand the evolution process of cancer cells and answering questions such that: "What is the probability that the daughter of a cancer cell of type A is of type A, B, C, ...?"
- By estimating the probability of connection between any node $i \in [n]$ and the upcoming node $n + 1$ as explained in Section 4.2, the biologist would get the correlation between the gene expression profile of cell i and the future cell $n + 1$.

APPENDIX C. PROOFS

C.1 Proof of Theorem 3.1

Lemma C.1 provides a more complete version of Theorem 3.1 by giving explicitly the constants.

Lemma C.1. *Let us consider the three positive constants c , c' and c'' involved in Theorem A.1.*

Assume that $\alpha_n \log(n) \leq 1/L$ and that $n\alpha_n > \max\left(\frac{4Lc''}{\pi_m^2 D^2}, \frac{2}{L\pi_m}\right)$. Then it holds

$$\mathbb{P}\left(\text{err}(\hat{G}, G) > \exp\left(-\frac{c'S^2}{2}\right)\right) \leq \frac{c}{n^2} + 2K \exp\left(-\frac{n\pi_m^2}{2A_1 + 4A_2\pi_m}\right),$$

where $S^2 = \frac{n\alpha_n\pi_mD^2}{L}$ and where A_1 and A_2 are constants that only depend on the Markov chain $(C_i)_{i \geq 1}$ with $A_1 := \frac{1 + (\lambda_+ \vee 0)}{1 - (\lambda_+ \vee 0)}$ and $A_2 := \frac{1}{3}\mathbb{1}_{\lambda_+ \leq 0} + \frac{5}{1 - \lambda_+}\mathbb{1}_{\lambda_+ > 0}$. Here $1 - \lambda_+$ is the right L^2 spectral gap of the Markov chain $(C_i)_{i \geq 1}$ (see Def. D.1 in Sect. D).

Remarks.

- The fact that $\pi_m > 0$ is a direct consequence of the positive recurrent property of the Markov chain.
 - The second term in the right hand side of the inequality from Lemma C.1 comes from the concentration of the average number of visits of the Markov chain towards the stationary distribution of the chain. The first term in this inequality corresponds to the bound from Theorem A.1 when communities have been assigned.
- Recalling that $\|Q\|_\infty$ is upper bounded by $\alpha_n L$, the condition $\alpha_n \log(n) \leq 1/L$ enforces the signal to noise ratio defined by Giraud and Verzelen $s^2 := \Delta^2/(\alpha_n L)$ (see Thm. A.1) to be larger than $\Delta^2 \cdot \log(n)$. Another way to interpret this condition is to say that it enforces the expected degree of all nodes of the graph to be smaller than $n/\log(n)$.
- In order to get some intuition on the conditions on n in the previous theorem, keep in mind that asymptotically, the size of the smallest community in the graph will be $n\pi_m$.
 - The condition $n > \frac{4Lc''}{\alpha_n\pi_m^2 D^2}$ can be read as $(n\pi_m)\alpha_n D^2/L > \frac{4c''}{\pi_m} = 4c'' \cdot \frac{n}{n\pi_m}$. Asymptotically, $(n\pi_m)\alpha_n D^2/L$ provides a lower bound on the signal-to-noise ratio defined in Theorem A.1. This shows that the condition $n > \frac{4Lc''}{\alpha_n\pi_m^2 D^2}$ is related to the constraint $s^2 \gtrsim n/m$ of Theorem A.1.
 - The condition $n > \frac{2}{\alpha_n L \pi_m}$ can be read as $\frac{1}{n\pi_m} < \alpha_n L/2$. This shows that the condition $n > \frac{2}{\alpha_n L \pi_m}$ is related to the constraint $1/m < \alpha_n L$ from Theorem A.1.

The proof of Lemma C.1 is based on the following Lemma which is proved at the end of this subsection.

Lemma C.2. *We consider c, c' and c'' the three numerical constants involved in Theorem A.1.*

Let us consider $0 < t < \pi_m$. Assume that $\alpha_n L \leq 1/\log(n)$. Then for any $\epsilon > 0$ and n large enough such that:

$$n(\pi_m - t) \geq \begin{cases} \frac{L \log(1/\epsilon)}{c' \alpha_n D^2} & (i) \\ \left(\frac{c'' n L}{\alpha_n D^2}\right)^{1/2} & (ii) \\ 1/(\alpha_n L) & (iii) \end{cases}$$

it holds

$$\mathbb{P}\left(\text{err}(\hat{\mathbf{G}}, \mathbf{G}) > \epsilon\right) \leq \frac{c}{n^2} + 2K \exp\left(-\frac{nt^2}{2(A_1/4 + A_2t)}\right),$$

where A_1 and A_2 are constants defined in Theorem 3.1.

Note that the only constraint on ϵ is given by the condition (i) which is equivalent to

$$\epsilon \geq \exp\left(-\frac{c'D^2n\alpha_n(\pi_m - t)}{L}\right).$$

In order to get the tighter result possible, we want to choose $\epsilon = \exp\left(-\frac{c'D^2n\alpha_n(\pi_m - t)}{L}\right)$ which leads to

$$t = \pi_m - \frac{L \log(1/\epsilon)}{c'D^2n\alpha_n}.$$

The condition $t > 0$ is then equivalent to

$$\pi_m > \frac{L \log(1/\epsilon)}{c'D^2n\alpha_n} \Leftrightarrow \exp(-\pi_m n \alpha_n c' D^2 / L) < \epsilon.$$

The condition (ii) is equivalent to

$$n(\pi_m - t) = \frac{L \log(1/\epsilon)}{c'\alpha_n D^2} \geq \left(\frac{c''nL}{\alpha_n D^2}\right)^{1/2} \Leftrightarrow \exp\left(-c'\sqrt{\frac{D^2 c''n\alpha_n}{L}}\right) \geq \epsilon.$$

The condition (iii) is equivalent to

$$n(\pi_m - t) = \frac{L \log(1/\epsilon)}{c'\alpha_n D^2} \geq (1/\alpha_n L) \Leftrightarrow \exp\left(-\frac{c'D^2}{L^2}\right) \geq \epsilon.$$

One can easily prove that for $n\alpha_n > \max\left(\frac{4Lc''}{\pi_n^2 D^2}, \frac{2}{L\pi_m}\right)$, $\epsilon := \exp\left(-\frac{\pi_m n \alpha_n c' D^2}{2L}\right)$ satisfies the three conditions above. This gives Lemma C.1 from Lemma C.2.

Proof of Lemma C.2. Using Theorem D.2 (cf. Appendix D or Theorem 2 of [34]), we get that

$$\forall c \in [K], \forall t > 0, \mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=c} - \pi(c)\right| \geq t\right) \leq 2 \exp\left(-\frac{nt^2}{2(A_1\sigma_c^2 + A_2t)}\right) \quad (\text{C.1})$$

where $A_1 = \frac{1 + (\lambda_+ \vee 0)}{1 - (\lambda_+ \vee 0)}$, $A_2 = \frac{1}{3} \mathbb{1}_{\lambda_+ \leq 0} + \frac{5}{1 - \lambda_+} \mathbb{1}_{\lambda_+ > 0}$ and $\sigma_c^2 = \pi(c)(1 - \pi(c))$.

We deduce that for all $t > 0$,

$$\mathbb{P}\left(\bigcup_c \left\{\left|\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=c} - \pi(c)\right| \geq t\right\}\right) \leq 2K \exp\left(-\frac{nt^2}{2(A_1\sigma^2 + A_2t)}\right),$$

where $\sigma^2 := \max_c \sigma_c^2$ ($\leq 1/4$). We define $\Omega^c := \bigcup_c \left\{ \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=c} - \pi(c) \right| \geq t \right\}$ and we recall $\pi_m = \min_c \pi(c)$ and $D^2 = \min_{j \neq k} \sum_l ((Q_0)_{k,l} - (Q_0)_{j,l})^2$.

Suppose that $0 < t < \pi_m$ and that n is large enough to satisfy (i), (ii) and (iii). Then it holds

$$\begin{aligned}
& \mathbb{P} \left(\text{err}(\hat{G}, G) > \epsilon \right) \\
&= \mathbb{P} \left(\{ \text{err}(\hat{G}, G) > \epsilon \} \cap \Omega \right) + \mathbb{P} \left(\{ \text{err}(\hat{G}, G) > \epsilon \} \cap \Omega^c \right) \\
&\leq \mathbb{P} \left(\{ \text{err}(\hat{G}, G) > \epsilon \} \cap \Omega \right) + 2K \exp \left(-\frac{nt^2}{2(A_1\sigma^2 + A_2t)} \right) \\
&= \mathbb{P} \left(\text{err}(\hat{G}, G) > \epsilon \mid \Omega \right) \times \mathbb{P}(\Omega) + 2K \exp \left(-\frac{nt^2}{2(A_1\sigma^2 + A_2t)} \right). \tag{C.2}
\end{aligned}$$

We denote by M the random variable that gives the size of the smallest cluster: $M := \min_{k \in [K]} m_k$. Condition (i) is equivalent to

$$\epsilon \geq \exp \left(-c' \frac{n\alpha_n(\pi_m - t)D^2}{L} \right).$$

Since on the event Ω we have $n(\pi_m - t) \leq M$, we get that on Ω it holds

$$\epsilon \geq \exp \left(-c' \frac{M\alpha_n D^2}{L} \right) \geq \exp(-c's^2), \tag{C.3}$$

where $s^2 = \Delta^2/(\alpha_n L)$ with $\Delta^2 = \min_{k \neq j} \Delta_{k,j}^2$ and $\Delta_{k,j}^2 = \sum_l m_l (Q_{k,l} - Q_{j,l})^2$. The last inequality comes from the fact that $\Delta^2 \geq M\alpha_n^2 D^2$. Using (C.2) we get that

$$\begin{aligned}
\mathbb{P} \left(\text{err}(\hat{G}, G) > \epsilon \right) &\leq \mathbb{P} \left(\text{err}(\hat{G}, G) > \epsilon \mid \Omega \right) + 2K \exp \left(-\frac{nt^2}{2(A_1\sigma^2 + A_2t)} \right) \\
&\leq \mathbb{P} \left(\text{err}(\hat{G}, G) > e^{-c's^2} \mid \Omega \right) + 2K \exp \left(-\frac{nt^2}{2(A_1\sigma^2 + A_2t)} \right).
\end{aligned}$$

We note that on Ω :

– Condition (ii) gives

$$M^2 \geq \frac{c''nL}{\alpha_n D^2} \Leftrightarrow \frac{M\alpha_n D^2}{L} \geq c''n/M,$$

which implies that $s^2 = \frac{\Delta^2}{\alpha_n L} \geq c''n/M$ since $\Delta^2 \geq M\alpha_n^2 D^2$.

– Condition (iii) gives

$$\frac{1}{M} \leq \alpha_n L.$$

Applying Theorem A.1 from Verzelen and Giraud, we get that

$$\mathbb{P}\left(\text{err}(\hat{G}, G) > e^{-c's^2}|\Omega\right) \leq \frac{c}{n^2}.$$

Finally we obtain using equation (C.3) that

$$\mathbb{P}\left(\text{err}(\hat{G}, G) > \epsilon\right) \leq \frac{c}{n^2} + 2K \exp\left(-\frac{nt^2}{2(A_1\sigma^2 + A_2t)}\right).$$

□

C.2 Proof of Theorem 3.2

We start by proving Lemma C.3 which enriches the statement of Theorem 3.2 by giving explicitly the constants.

Lemma C.3. *We consider c, c' and c'' the three numerical constants involved in Theorem A.1.*

Assume that $\alpha_n \log(n) \leq \frac{1}{L}$ and that $n\alpha_n > \max\left(\frac{4Lc''}{\pi_m^2 D^2}, \frac{2}{L\pi_m}, \frac{2L \log(n)}{\pi_m c' D^2}\right)$. Then for all $0 < t < \pi_m - \frac{1}{n}$, it holds

$$\begin{aligned} & \mathbb{P}\left(\|\hat{Q} - Q\|_\infty > t\right) \\ & \leq K(K+1) \exp\left(-\frac{(n\pi_m - nt - 1)^2 t^2}{\frac{1}{2} + \frac{2}{3}t}\right) + \frac{c}{n^2} + 2K \exp\left(-\frac{nt^2}{2(A_1/4 + A_2t)}\right). \end{aligned}$$

Proof of Lemma C.3.

– Preliminary 1

Using the standard Bernstein's inequality for independent random variables, we get that for all $k, l \in [K]^2$ with $k \neq l$, and for all $t > 0$, it holds

$$\mathbb{P}\left(\left|\frac{1}{|G_k| \cdot |G_l|} \sum_{i \in G_k} \sum_{j \in G_l} X_{i,j} - Q_{k,l}\right| \geq t\right) \leq 2 \exp\left(-\frac{|G_k| \cdot |G_l| t^2}{2(Q_{k,l}(1 - Q_{k,l}) + t/3)}\right)$$

and for all $k \in [K], t > 0$, it holds

$$\mathbb{P}\left(\left|\frac{1}{|G_k| \cdot (|G_k| - 1)} \sum_{i,j \in G_k, i \neq j} X_{i,j} - Q_{k,k}\right| \geq t\right) \leq 2 \exp\left(-\frac{|G_k| \cdot (|G_k| - 1) t^2}{2(Q_{k,k}(1 - Q_{k,k}) + t/3)}\right).$$

– Preliminary 2

We define the event $N := \left\{\text{err}(\hat{G}, G) < \exp\left(-\frac{\pi_m n \alpha_n c' D^2}{2L}\right)\right\}$. Note that on N , the partition of the clusters is correctly recovered thanks to the condition $n\alpha_n > \frac{2L \log(n)}{\pi_m c' D^2}$.

– Preliminary 3

Using Theorem D.2 (cf. Appendix D or Theorem 2 of [34]), we get that

$$\forall c \in [K], \forall t > 0, \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=c} - \pi(c) \right| \geq t \right) \leq 2 \exp \left(-\frac{nt^2}{2(A_1/4 + A_2t)} \right).$$

We deduce that for all $t > 0$,

$$\mathbb{P} \left(\bigcup_c \left\{ \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=c} - \pi(c) \right| \geq t \right\} \right) \leq 2K \exp \left(-\frac{nt^2}{2(A_1/4 + A_2t)} \right).$$

We define $\Omega^c := \bigcup_{c \in [K]} \left\{ \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=c} - \pi(c) \right| \geq t \right\}$.

Let us define for any $k, l \in [K]^2$,

$$\tilde{Q}_{k,l} = \frac{1}{|G_k| \cdot |G_l|} \sum_{i \in G_k} \sum_{j \in G_l} X_{i,j}.$$

Then, considering $0 < t < \pi_m - \frac{1}{n}$, we have

$$\begin{aligned} & \mathbb{P} \left(\|\hat{Q} - Q\|_\infty > t \right) \\ \leq & \mathbb{P} \left(\bigcup_{k,l \in [K]^2, k \leq l} \{|\hat{Q}_{k,l} - Q_{k,l}| > t\} \mid \Omega \right) + \mathbb{P}(\Omega^c) \\ = & \mathbb{P} \left(\left[\bigcup_{k,l \in [K]^2, k \leq l} \{|\hat{Q}_{k,l} - Q_{k,l}| > t\} \right] \cap N \mid \Omega \right) + \mathbb{P} \left(\left[\bigcup_{k,l \in [K]^2, k \leq l} \{|\hat{Q}_{k,l} - Q_{k,l}| > t\} \right] \cap N^c \mid \Omega \right) + \mathbb{P}(\Omega^c) \\ \leq & \mathbb{P} \left(\bigcup_{k,l \in [K]^2, k \leq l} \left(\{|\hat{Q}_{k,l} - Q_{k,l}| > t\} \cap N \right) \mid \Omega \right) + \mathbb{P}(N^c \mid \Omega) + \mathbb{P}(\Omega^c) \\ \leq & \sum_{k,l \in [K]^2, k \leq l} \mathbb{P} \left(\{|\hat{Q}_{k,l} - Q_{k,l}| > t\} \cap N \mid \Omega \right) + \mathbb{P}(N^c \mid \Omega) + \mathbb{P}(\Omega^c) \\ = & \sum_{k,l \in [K]^2, k \leq l} \mathbb{E} \left(\mathbb{1}_{\{|\hat{Q}_{k,l} - Q_{k,l}| > t\}} \mathbb{1}_N \mid \Omega \right) + \mathbb{P}(N^c \mid \Omega) + \mathbb{P}(\Omega^c). \end{aligned}$$

At this point, one can notice that it holds for any $k, l \in [K]^2$,

$$\mathbb{1}_{\{|\hat{Q}_{k,l} - Q_{k,l}| > t\}} \mathbb{1}_N = \mathbb{1}_{\{|\tilde{Q}_{k,l} - Q_{k,l}| > t\}} \mathbb{1}_N. \quad (\text{C.4})$$

Indeed,

- either it holds $\mathbb{1}_N = 0$ in which case the equality from equation (C.4) is trivially true,
- or $\mathbb{1}_N = 1$ (meaning that the event N holds) which implies that $\hat{G}_s = G_s$ for all $s \in [K]$ (see Preliminary 2), leading to $\hat{Q}_{k,l} = \tilde{Q}_{k,l}$.

Hence for any $0 < t < \pi_m - \frac{1}{n}$ we have,

$$\begin{aligned}
& \mathbb{P} \left(\|\hat{Q} - Q\|_\infty > t \right) \\
& \leq \sum_{k,l \in [K]^2, k \leq l} \mathbb{E} \left(\mathbb{1}_{\{|\tilde{Q}_{k,l} - Q_{k,l}| > t\}} \mathbb{1}_N \mid \Omega \right) + \mathbb{P}(N^c \mid \Omega) + \mathbb{P}(\Omega^c) \\
& \text{and using preliminary 3,} \\
& \leq \sum_{k,l \in [K]^2, k \leq l} \mathbb{P} \left(\{|\tilde{Q}_{k,l} - Q_{k,l}| > t\} \mid \Omega \right) + \mathbb{P}(N^c \mid \Omega) \\
& \quad + 2K \exp \left(-\frac{nt^2}{2(A_1/4 + A_2t)} \right) \\
& \leq 2 \sum_{1 \leq k \leq l \leq K} \exp \left(-\frac{n(\pi(k) - t) \cdot (n\pi(l) - nt - 1)t^2}{2(Q_{k,l}(1 - Q_{k,l}) + t/3)} \right) + \frac{c}{n^2} + \\
& \quad 2K \exp \left(-\frac{nt^2}{2(A_1/4 + A_2t)} \right),
\end{aligned}$$

where we used that $\mathbb{P}(N^c \mid \Omega) \leq \frac{c}{n^2}$ (shown in the proof of Thm. 3.1).

□

Proof of Theorem 3.2. Let us consider $\gamma > 0$ and let us define $t = \frac{\gamma}{\sqrt{n}}$. Considering that

$$\frac{n\alpha_n}{\log(n)} \geq a \quad \text{with} \quad a := \frac{4Lc''}{c'\pi_m^2 D^2} \vee \frac{2L}{c'\pi_m D^2} \vee \frac{2}{L\pi_m},$$

we ensure that $n\alpha_n$ satisfies the conditions of Lemma C.3.

Now let us look into the condition $t = \frac{\gamma}{\sqrt{n}} < \pi_m - \frac{1}{n}$ of Lemma C.3. We will ask t to satisfy the stronger condition

$$t = \frac{\gamma}{\sqrt{n}} < \frac{\pi_m}{2} - \frac{1}{n} \quad \Leftrightarrow \quad 0 < \frac{\pi_m}{2}n - \gamma\sqrt{n} - 1. \quad (\text{C.5})$$

Studying the polynomial function $f : x \mapsto \frac{\pi_m}{2}x^2 - \gamma x - 1$, one can find that the zeros of f are

$$x_1 := \frac{\gamma - \sqrt{\gamma^2 + 2\pi_m}}{\pi_m} \quad \text{and} \quad x_2 := \frac{\gamma + \sqrt{\gamma^2 + 2\pi_m}}{\pi_m} \leq \frac{2\gamma + \sqrt{2\pi_m}}{\pi_m}.$$

We deduce that considering that

$$n > 4 \left(\frac{\gamma + 1}{\pi_m} \right)^2, \quad (\text{C.6})$$

which implies that $\sqrt{n} > \frac{2\gamma + \sqrt{2\pi_m}}{\pi_m}$, we guarantee that $\gamma/\sqrt{n} < \pi_m - 1/n$. Applying Lemma C.3, we get that with probability at least

$$1 - \left[(K^2 + K) \exp \left(\frac{-(n\pi_m - \gamma\sqrt{n} - 1)^2 \frac{\gamma^2}{n}}{\frac{1}{2} + \frac{2}{3} \frac{\gamma}{\sqrt{n}}} \right) + \frac{c}{n^2} + 2K \exp \left(\frac{-\gamma^2}{2(A_1/4 + A_2 \frac{\gamma}{\sqrt{n}})} \right) \right],$$

it holds $\|\hat{Q} - Q\|_\infty \leq \gamma/\sqrt{n}$.

Thanks to equations (C.6) and (C.5), we have $(n\pi_m - \gamma\sqrt{n} - 1)^2 = n^2(\pi_m - \gamma/\sqrt{n} - 1/n)^2 \geq n^2\pi_m^2/4$ and $\gamma/\sqrt{n} \leq \pi_m/2$. We deduce that defining

$$b := c + (2K(K+1)) \quad \text{and} \quad b' := \frac{1}{2(A_1/4 + A_2\pi_m)} \wedge \frac{\pi_m^2}{2 + \frac{4}{3}\pi_m},$$

it holds with probability at least $1 - b(1/n^2 \vee \exp(-b'\gamma^2))$

$$\|\hat{Q} - Q\|_\infty \leq \frac{\gamma}{\sqrt{n}}.$$

□

C.3 Proof of Theorem 3.3

Lemma C.4 provides a more complete version of Theorem 3.3 by giving explicitly the constants.

Lemma C.4. *We consider c, c' and c'' the three numerical constants involved in Theorem A.1.*

Assume that $\alpha_n \log(n) \leq 1/L$ and that $n\alpha_n > \max\left(\frac{4Lc''}{\pi_m^2 D^2}, \frac{2}{L\pi_m}, \frac{2L \log(n)}{\pi_m c' D^2}\right)$. Then for all $t > 0$, it holds

$$\mathbb{P}(\|\hat{\pi} - \pi\|_\infty > t) \leq 2K \exp\left(-\frac{nt^2}{2(A_1/4 + A_2t)}\right) + \frac{c}{n^2} + 2K \exp\left(-\frac{n\pi_m^2}{2A_1 + 4A_2\pi_m}\right).$$

Proof of Lemma C.4. Using Theorem D.2 (cf. Appendix D or Theorem 2 of [34]), we get that

$$\forall c \in [K], \forall t > 0, \quad \mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=c} - \pi(c)\right| \geq t\right) \leq 2 \exp\left(-\frac{nt^2}{2(A_1\sigma_c^2 + A_2t)}\right)$$

where $A_1 = \frac{1 + (\lambda_+ \vee 0)}{1 - (\lambda_+ \vee 0)}$, $A_2 = \frac{1}{3}\mathbb{1}_{\lambda_+ \leq 0} + \frac{5}{1 - \lambda_+}\mathbb{1}_{\lambda_+ > 0}$ and $\sigma_c^2 = \pi(c)(1 - \pi(c)) \leq 1/4$.

We define the event $N := \left\{\text{err}(\hat{G}, G) < \exp\left(-\frac{\pi_m n \alpha_n c' D^2}{2L}\right)\right\}$. Note that on N , the partition of the clusters is correctly recovered thanks to the condition $n\alpha_n > \frac{2L \log(n)}{\pi_m c' D^2}$. Then,

$$\begin{aligned} & \mathbb{P}\left(\bigcup_{k \in [K]} \{|\hat{\pi}(k) - \pi(k)| > t\}\right) \\ & \leq \mathbb{P}\left(\bigcup_{k \in [K]} \{|\hat{\pi}(k) - \pi(k)| > t\} \mid N\right) + \mathbb{P}(N^c) \\ & = \mathbb{P}\left(\bigcup_{k \in [K]} \left\{\left|\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=k} - \pi(k)\right| > t\right\} \mid N\right) + \mathbb{P}(N^c) \end{aligned}$$

$$\leq 2K \exp\left(-\frac{nt^2}{2(A_1/4 + A_2t)}\right) + \frac{c}{n^2} + 2K \exp\left(-\frac{n\pi_m^2}{2A_1 + 4A_2\pi_m}\right),$$

where we apply Lemma C.1 in the last inequality. \square

C.4 Proof of Theorem 3.4

We will prove a more accurate result with Lemma C.5.

Lemma C.5. *Let us consider $\gamma > \frac{5K}{2\pi_m^2}$.*

Assume that $\alpha_n \log(n) \leq 1/L$, that $n\alpha_n > \max\left(\frac{4Lc''}{\pi_m^2 D^2}, \frac{4}{L\pi_m}, \frac{2L \log(n)}{\pi_m c' D^2}\right)$ and that $\sqrt{n} > \frac{2}{\pi_m}(1 + \pi_m^2 \gamma/5)$.

Then it holds

$$\begin{aligned} & \mathbb{P}\left(\|\hat{P} - P\|_\infty \geq \frac{\gamma}{\sqrt{n}}\right) \\ & \leq 2K^2 \exp\left(-\frac{\left(\frac{\pi_m^2 \gamma}{5K} - \frac{1}{2}\right)^2}{2(B_1/4 + B_2 \frac{\pi_m^2 \gamma}{5K\sqrt{n}})}\right) + \frac{c}{n^2} + 2K \exp\left(-\frac{n\pi_m^2}{8A_1\sigma^2 + 4A_2\pi_m}\right), \end{aligned}$$

where B_1 and B_2 depend only on the Markov chain and are defined by $B_1 := \frac{1 + (\xi_+ \vee 0)}{1 - (\xi_+ \vee 0)}$ and $B_2 := \frac{1}{3} \mathbb{1}_{\xi_+ \leq 0} + \frac{5}{1 - \xi_+} \mathbb{1}_{\xi_+ > 0}$. Here $1 - \xi_+$ is the right L^2 spectral gap of the Markov chain $(Y_i)_{i \geq 1}$ (see Definition D.1 in Section D).

Remarks.

- The first term in the right hand side of the inequality in Lemma C.5 is due to the concentration of the average number of visits of the chain $(Y_i)_{i \geq 1}$ (defined in Section 3.2.3 of this chapter) towards its stationary distribution. The two last terms of the inequality correspond to the bound guaranteeing the recovery of the true partitions with a direct application of Theorem 3.1.
- The condition $n\alpha_n > \frac{2L \log(n)}{\pi_m c' D^2}$ ensures that $\exp\left(-\frac{\pi_m n \alpha_n c' D^2}{2L}\right) < \frac{1}{n}$. Theorem 3.1 will then guarantee that we recover perfectly the partition of the communities.
- Expecting the accuracy γ/\sqrt{n} , the condition $\sqrt{n} > \frac{2}{\pi_m}(1 + \pi_m^2 \gamma/5)$ ensures that the Markov chain $(C_i)_{i \geq 1}$ has visited enough each state $k \in [K]$ to guarantee the convergence of the average number of visits toward the stationary distribution.

Proof of Lemma C.5.

I. Concentration of the average number of visits for $(Y_i)_{i \geq 1}$.

We recall that $(Y_i)_{i \geq 1}$ is a Markov Chain on $[K]^2$ defined by $Y_i = (C_i, C_{i+1})$.

Then using again Theorem D.2, we get that $\forall t > 0, \forall k, l \in [K]^2$,

$$\mathbb{P}\left(\left|\frac{1}{n-1} \sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)} - \pi(k)P_{k,l}\right| \geq t\right) \leq 2 \exp\left(-\frac{nt^2}{2(B_1/4 + B_2t)}\right),$$

II. First step toward the Theorem.

We define the event $N := \left\{ \text{err}(\hat{G}, G) < \exp\left(-\frac{\pi_m n \alpha_n c' D^2}{2L}\right) \right\}$. Note that on N , the partition of the clusters is correctly recovered thanks to the condition $n \alpha_n > \frac{2L \log(n)}{\pi_m c' D^2}$. Let $\gamma > \frac{5K}{2\pi_m^2}$ and let us define

$$r = \frac{\zeta}{\sqrt{n}} \quad \text{with } \zeta = \frac{\pi_m^2 \gamma}{5K} - \frac{1}{2} > 0,$$

$$\text{and } \Gamma = \bigcap_{k,l} \left\{ \left| \frac{1}{n-1} \sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)} - \pi(k)P_{k,l} \right| < r \right\}.$$

Then,

$$\begin{aligned} & \mathbb{P} \left(\bigcup_{k,l} \left\{ \left| \hat{P}_{k,l} - P_{k,l} \right| \geq \frac{\gamma}{\sqrt{n}} \right\} \right) \\ & \leq \underbrace{\mathbb{P} \left(\bigcup_{k,l} \left\{ \left| \hat{P}_{k,l} - P_{k,l} \right| \geq \frac{\gamma}{\sqrt{n}} \right\} \mid N, \Gamma \right)}_{(*)} \mathbb{P}(N) \mathbb{P}(\Gamma \mid N) + \mathbb{P}(\Gamma^c) + \mathbb{P}(N^c). \end{aligned}$$

Note that the condition $\sqrt{n} > \frac{2}{\pi_m} (1 + \pi_m^2 \gamma / 5)$ of Lemma C.5 implies that

$$\sqrt{n} > \frac{2}{\pi_m} (1 + K\zeta). \quad (\text{C.7})$$

III. We prove that $(*)$ is zero.

In this third step of the proof, we are going to show that conditionally on the event $N \cap \Gamma$, the infinite norm between our estimate of the transition matrix \hat{P} and P is smaller than γ/\sqrt{n} .

1 We split $(*)$ in two terms.

$$\begin{aligned} & \mathbb{P} \left(\bigcup_{k,l} \left\{ \left| \hat{P}_{k,l} - P_{k,l} \right| \geq \frac{\gamma}{\sqrt{n}} \right\} \mid N, \Gamma \right) \\ & = \mathbb{P} \left(\bigcup_{k,l} \left\{ \left| \hat{P}_{k,l} - \frac{\sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)}}{(n-1)\pi(k)} + \frac{\sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)}}{(n-1)\pi(k)} - P_{k,l} \right| \geq \frac{\gamma}{\sqrt{n}} \right\} \mid N, \Gamma \right) \end{aligned}$$

2 We show that on Γ : $\left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=k} - \pi(k) \right| \leq \frac{1}{n} + Kr$. Here we show that a concentration of the average number of visits for $(Y_i)_{i \geq 1}$ gives for free a concentration result of the average number of visits for $(C_i)_{i \geq 1}$.

Note that on the event Γ :

$$\begin{aligned}
\bullet \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=k} &= \frac{1}{n} \sum_{l=1}^K \sum_{i=1}^{n-1} \mathbb{1}_{C_i=k, C_{i+1}=l} \\
&= \frac{n-1}{n} \sum_{l=1}^K \frac{1}{n-1} \sum_{i=1}^{n-1} \mathbb{1}_{C_i=k, C_{i+1}=l} \\
&\geq \frac{n-1}{n} \sum_{l=1}^K (\pi(k)P_{k,l} - r) \\
&= \frac{n-1}{n} (\pi(k) - Kr).
\end{aligned}$$

Hence $\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=k} - \pi(k) \geq -\frac{\pi(k)}{n} - \frac{n-1}{n} Kr \geq -\left(\frac{1}{n} + Kr\right)$.

$$\begin{aligned}
\bullet \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=k} &\leq \frac{1}{n} \sum_{l=1}^K \sum_{i=1}^{n-1} \mathbb{1}_{C_i=k, C_{i+1}=l} + \frac{1}{n} \\
&= \frac{n-1}{n} \sum_{l=1}^K \frac{1}{n-1} \sum_{i=1}^{n-1} \mathbb{1}_{C_i=k, C_{i+1}=l} + \frac{1}{n} \\
&\leq \frac{n-1}{n} \sum_{l=1}^K (\pi(k)P_{k,l} + r) + \frac{1}{n} \\
&\leq \pi(k) + Kr + \frac{1}{n}.
\end{aligned}$$

Hence $\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=k} - \pi(k) \leq \frac{1}{n} + Kr$.

We deduce then that on Γ , $\left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{C_i=k} - \pi(k) \right| \leq \frac{1}{n} + Kr$.

3 We show that the first term from **1** is zero. In the following, we show that the definition of ζ with the condition $\sqrt{n} > \frac{2}{\pi_m} (1 + \pi_m^2 \gamma / 5)$ implies that the first term in **1** is zero.

$$\begin{aligned}
&\mathbb{P} \left(\left| \hat{P}_{k,l} - \frac{1}{n-1} \frac{\sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)}}{\pi(k)} \right| \geq \frac{\gamma}{2\sqrt{n}} \mid N, \Gamma \right) \\
&= \mathbb{P} \left(\frac{1}{n-1} \sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)} \left| \frac{n}{\sum_{i=1}^n \mathbb{1}_{C_i=k}} - \frac{1}{\pi(k)} \right| \geq \frac{\gamma}{2\sqrt{n}} \mid N, \Gamma \right) \\
&\leq \mathbb{P} \left((r + \pi(k)P_{k,l}) \left| \frac{n}{\sum_{i=1}^n \mathbb{1}_{C_i=k}} - \frac{1}{\pi(k)} \right| \geq \frac{\gamma}{2\sqrt{n}} \mid N, \Gamma \right) \quad (\text{by definition of } \Gamma) \\
&= \mathbb{P} \left(\left| \frac{n\pi(k) - \sum_{i=1}^n \mathbb{1}_{C_i=k}}{\pi(k) \sum_{i=1}^n \mathbb{1}_{C_i=k}} \right| \geq \frac{1}{2r + 2\pi(k)P_{k,l}} \cdot \frac{\gamma}{\sqrt{n}} \mid N, \Gamma \right) \text{ and using } \mathbf{2},
\end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{P} \left(\frac{\frac{1}{n} + Kr}{\pi(k)(\pi(k) - \frac{1}{n} - Kr)} \geq \frac{1}{2r + 2\pi(k)P_{k,l}} \cdot \frac{\gamma}{\sqrt{n}} \mid N, \Gamma \right) \\
&\leq \mathbb{P} \left(\frac{\frac{1}{n} + Kr}{\pi_m(\pi_m - \frac{1}{n} - Kr)} \geq \frac{1}{2r + 2} \cdot \frac{\gamma}{\sqrt{n}} \mid N, \Gamma \right) \text{ and since } r = \frac{\zeta}{\sqrt{n}}, \\
&= \mathbb{P} \left(\frac{\frac{1}{n} + K\frac{\zeta}{\sqrt{n}}}{\pi_m(\pi_m - \frac{1}{n} - K\frac{\zeta}{\sqrt{n}})} \geq \frac{1}{2\frac{\zeta}{\sqrt{n}} + 2} \cdot \frac{\gamma}{\sqrt{n}} \mid N, \Gamma \right) \\
&= \mathbb{P} \left(\frac{(\frac{1}{n} + K\frac{\zeta}{\sqrt{n}})(2\zeta + 2\sqrt{n})}{\pi_m(\pi_m - \frac{1}{n} - K\frac{\zeta}{\sqrt{n}})} \geq \gamma \mid N, \Gamma \right) \\
&\leq \mathbb{P} \left(\frac{(\frac{1}{n} + K\frac{\zeta}{\sqrt{n}})(2\zeta + 2\sqrt{n})}{\pi_m(\pi_m - \frac{1}{\sqrt{n}}(1 + K\zeta))} \geq \gamma \mid N, \Gamma \right). \tag{C.8}
\end{aligned}$$

Since from (C.7), $\sqrt{n} \geq \frac{2}{\pi_m}(1 + K\zeta)$, we have

$$\frac{\pi_m}{2} \leq \pi_m - \frac{1}{\sqrt{n}}(1 + K\zeta),$$

which leads to

$$\mathbb{P} \left(\frac{(\frac{1}{n} + K\frac{\zeta}{\sqrt{n}})(2\zeta + 2\sqrt{n})}{\pi_m(\pi_m - \frac{1}{\sqrt{n}}(1 + K\zeta))} \geq \gamma \mid N, \Gamma \right) \leq \mathbb{P} \left(\frac{2(\frac{1}{\sqrt{n}} + K\zeta)(\frac{\zeta}{\sqrt{n}} + 1)}{\pi_m^2/2} \geq \gamma \mid N, \Gamma \right).$$

Moreover, since from (C.7) and the fact that $\pi_m \in (0, 1)$, $\sqrt{n} \geq \frac{2}{\pi_m}(1 + K\zeta) > 2K\zeta$, it holds

$$\frac{\zeta}{\sqrt{n}} < \frac{1}{2K} < \frac{1}{4}.$$

Coming back to (C.8), we finally get

$$\begin{aligned}
&\mathbb{P} \left(\left| \hat{P}_{k,l} - \frac{1}{n-1} \frac{\sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)}}{\pi(k)} \right| \geq \frac{\gamma}{2\sqrt{n}} \mid N, \Gamma \right) \\
&\leq \mathbb{P} \left(\frac{2(\frac{1}{\sqrt{n}} + K\zeta)(\frac{\zeta}{\sqrt{n}} + 1)}{\pi_m^2/2} \geq \gamma \mid N, \Gamma \right) \\
&\leq \mathbb{P} \left(\frac{5(\frac{1}{\sqrt{n}} + K\zeta)}{\pi_m^2} \geq \gamma \mid N, \Gamma \right) \\
&= 0.
\end{aligned}$$

The last equality is due to the definition of ζ . Indeed,

$$\zeta = \frac{\gamma\pi_m^2}{5K} - \frac{1}{2} < \frac{\gamma\pi_m^2}{5K} - \frac{1}{K\sqrt{n}} \text{ leading to } \frac{5(\frac{1}{\sqrt{n}} + K\zeta)}{\pi_m^2} < \gamma.$$

4 We show that the second term from **1** is zero.

$$\begin{aligned}
& \mathbb{P} \left(\left| \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{\mathbb{1}_{Y_i=(k,l)}}{\pi(k)} - P_{k,l} \right| \geq \frac{\gamma}{2\sqrt{n}} \mid N, \Gamma \right) \\
&= \mathbb{P} \left(\left| \frac{1}{n-1} \sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)} - \pi(k)P_{k,l} \right| \geq \pi(k) \frac{\gamma}{2\sqrt{n}} \mid N, \Gamma \right) \\
&\leq \mathbb{P} \left(\left| \frac{1}{n-1} \sum_{i=1}^{n-1} \mathbb{1}_{Y_i=(k,l)} - \pi(k)P_{k,l} \right| \geq \pi_m \frac{\gamma}{2\sqrt{n}} \mid N, \Gamma \right) \\
&= 0,
\end{aligned}$$

where the last equality comes from the definition of $r = \zeta/\sqrt{n}$ and the definition of Γ because

$$\zeta = \frac{\pi_m^2 \gamma}{5K} - \frac{1}{2} < \frac{\pi_m^2 \gamma}{5K} \leq \frac{\pi_m \gamma}{2}.$$

IV. Conclusion.

$$\begin{aligned}
& \mathbb{P} \left(\bigcup_{k,l} \left\{ \left| \hat{P}_{k,l} - P_{k,l} \right| \geq \frac{\gamma}{\sqrt{n}} \right\} \right) \\
&\leq \mathbb{P} \left(\bigcup_{k,l} \left\{ \left| \hat{P}_{k,l} - P_{k,l} \right| \geq \frac{\gamma}{\sqrt{n}} \right\} \mid N, \Gamma \right) \mathbb{P}(N) \mathbb{P}(\Gamma \mid N) + \mathbb{P}(\Gamma^c) + \mathbb{P}(N^c) \\
&= \mathbb{P}(\Gamma^c) + \mathbb{P}(N^c) \\
&\leq 2K^2 \exp \left(-\frac{nr^2}{2(B_1/4 + B_2r)} \right) + \frac{c}{n^2} + 2K \exp \left(-\frac{n\pi_m^2}{8A_1\sigma^2 + 4A_2\pi_m} \right) \\
&\leq 2K^2 \exp \left(-\frac{\left(\frac{\pi_m^2 \gamma}{5K} - \frac{1}{2} \right)^2}{2(B_1/4 + B_2 \frac{\pi_m^2 \gamma}{5K \sqrt{n}})} \right) + \frac{c}{n^2} + 2K \exp \left(-\frac{n\pi_m^2}{8A_1\sigma^2 + 4A_2\pi_m} \right),
\end{aligned}$$

where we apply Lemma C.1 in the last inequality. \square

C.5 Proof of Proposition 4.4

Let us consider $\gamma > \frac{5K}{2\pi_m^2}$. We assume that the conditions (4.4) of Proposition 4.4 are satisfied and we deduce from Theorems 3.4 and 3.2 that there exists three constants $a, b, b' > 0$ and some event \mathcal{E}_a satisfying

$$\mathbb{P}(\mathcal{E}_a) \geq 1 - b \left[1/n^2 \vee \exp \left(-b' \left(\gamma - \frac{5K}{2\pi_m^2} \right)^2 \right) \right],$$

such that it holds on \mathcal{E}_a ,

$$\|\hat{P} - P\|_\infty \vee \|\hat{Q} - Q\|_\infty \leq \frac{\gamma}{\sqrt{n}}.$$

Taking a close look at the proof of Theorems 3.2 and 3.4, one can see that the constant a is chosen so that on \mathcal{E}_a it holds $\text{err}(\hat{G}, G) < 1/n$ meaning that the clustering algorithm has recovered correctly the partition of the nodes, *i.e.* $\mathbf{C}_{1:n} = \hat{\mathbf{C}}_{1:n}$. Hence we can focus only on $|\eta_i(\mathbf{C}_{1:n}) - \hat{\eta}_i(\mathbf{C}_{1:n})|$ to prove Proposition 4.4.

For any $i \in [n]$ we have

$$\begin{aligned}
& |\eta_i(\mathbf{C}_{1:n}) - \hat{\eta}_i(\mathbf{C}_{1:n})| \\
&= \left| \sum_{k \in [K]} P_{C_i,k} Q_{C_i,k} - \sum_{k \in [K]} \hat{P}_{C_i,k} \hat{Q}_{C_i,k} \right| \\
&\leq \sum_{k \in [K]} \left| P_{C_i,k} Q_{C_i,k} - \hat{P}_{C_i,k} Q_{C_i,k} + \hat{P}_{C_i,k} Q_{C_i,k} - \hat{P}_{C_i,k} \hat{Q}_{C_i,k} \right| \\
&\leq \sum_{k \in [K]} \left| P_{C_i,k} - \hat{P}_{C_i,k} \right| \times |Q_{C_i,k}| + \sum_{k \in [K]} \left| Q_{C_i,k} - \hat{Q}_{C_i,k} \right| \times \left| \hat{P}_{C_i,k} \right| \\
&\leq \|\hat{P} - P\|_\infty \sum_{k \in [K]} Q_{C_i,k} + \|Q - \hat{Q}\|_\infty \sum_{k \in [K]} \hat{P}_{C_i,k} \\
&\leq \|\hat{P} - P\|_\infty K \alpha_n L + \|Q - \hat{Q}\|_\infty,
\end{aligned}$$

where we used that $\|Q\|_\infty = \alpha_n \|Q_0\|_\infty \leq \alpha_n L$ and the fact that \hat{P} is a stochastic matrix.

We deduce that for any $i \in [n]$, it holds with probability at least $1 - b \left[1/n^2 \vee \exp\left(-b'(\gamma - \frac{5K}{2\pi_n^2})^2\right) \right]$,

$$|\eta_i(\mathbf{C}_{1:n}) - \hat{\eta}_i(\mathbf{C}_{1:n})| \leq \frac{\gamma}{\sqrt{n}} (\alpha_n K L + 1).$$

Using a union bound concludes the proof.

APPENDIX D. SPECTRAL GAPS AND MIXING TIMES FOR MARKOV CHAINS

This section is largely inspired from [34] of Section 2.1.

D.1 Spectral gap

We consider a state space E and a sigma-algebra Σ on E which is a standard Borel space. We denote by $(X_i)_{i \geq 1}$ a Markov chain on the state space (E, Σ) with stationary distribution π .

For any real-valued, Σ -measurable function $h : E \rightarrow \mathbb{R}$, we define $\pi(h) := \int h(x) \pi(dx)$. The set

$$\mathcal{L}^2(E, \Sigma, \pi) := \{h : \pi(h^2) < \infty\}$$

is a Hilbert space endowed with the inner product

$$\langle h_1, h_2 \rangle_\pi = \int h_1(x) h_2(x) \pi(dx), \quad \forall h_1, h_2 \in \mathcal{L}^2(E, \Sigma, \pi).$$

The map

$$\|\cdot\|_\pi : h \in \mathcal{L}^2(E, \Sigma, \pi) \mapsto \|h\|_\pi = \sqrt{\langle h, h \rangle_\pi},$$

is a norm on $\mathcal{L}^2(E, \Sigma, \pi)$. $\|\cdot\|_\pi$ naturally allows to define the norm of a linear operator T on $\mathcal{L}^2(E, \Sigma, \pi)$ as

$$N_\pi(T) = \sup\{\|Th\|_\pi : \|h\|_\pi = 1\}.$$

To each transition probability kernel $P(x, B)$ with $x \in E$ and $B \in \Sigma$ invariant with respect to π , we can associate a bounded linear operator $h \mapsto \int h(y)P(\cdot, dy)$ on $\mathcal{L}^2(E, \Sigma, \pi)$. Denoting this operator P , we get

$$Ph(x) = \int h(y)P(x, dy), \quad \forall x \in E, \quad \forall h \in \mathcal{L}^2(E, \Sigma, \pi).$$

Denoting by P^* the adjoint or time-reversal operator of the Markov operator P , we can define the self-adjoint operator $R = (P + P^*)/2$. Let $\mathcal{L}_0^2(\pi) := \{h \in \mathcal{L}^2(E, \Sigma, \pi) : \pi(h) = 0\}$. The spectrum of a self-adjoint Markov operator like R acting on $\mathcal{L}_0^2(\pi)$ is contained in $[-1, +1]$. The gap between 1 and the maximum of the spectrum of R is called the right \mathcal{L}^2 -spectral gap of P .

Definition D.1. (Right \mathcal{L}^2 -spectral gap) A Markov operator P has right \mathcal{L}^2 -spectral gap $1 - \lambda_+(R)$ if the operator $R = (P + P^*)/2$ satisfies

$$\lambda_+(R) := \sup\{s : s \in \text{spectrum of } R \text{ acting on } \mathcal{L}_0^2(\pi)\} < 1.$$

Theorem D.2 shows that the existence of a non-zero right \mathcal{L}^2 -spectral gap is of particular interest to prove concentration result for empirical processes. Theorem D.2 is a key tool for the proofs of the theoretical results of this paper.

Theorem D.2. ([34], Thm. 2)

Suppose that the sequence $(X_i)_{i \geq 1}$ is a Markov chain with stationary distribution π and non-zero right \mathcal{L}^2 -spectral gap $1 - \lambda_+ > 0$ (see Def. D.1). We assume further that X_1 is distributed according to π . Let us consider some $n \in \mathbb{N} \setminus \{0\}$ and a real valued function $f : E \rightarrow \mathbb{R}$ such that $\int f(x)d\pi(x) = 0$ and $\sup_{x \in E} |f(x)| < c$ for some constant $c > 0$ independent of n . Let $\sigma^2 = \int f^2(x)d\pi(x)$. Then for any $\epsilon \geq 0$ it holds

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n f(X_i) \geq \epsilon\right) \leq \exp\left(-\frac{n\epsilon^2}{A_2\sigma^2 + A_1c\epsilon}\right),$$

where $A_2 := \frac{1+\lambda_+\sqrt{0}}{1-\lambda_+\sqrt{0}}$ and $A_1 := \frac{1}{3}\mathbb{1}_{\lambda_+ \leq 0} + \frac{5}{1-\lambda_+}\mathbb{1}_{\lambda_+ > 0}$.

D.2 Signal to noise ratio and Markov mixing times

In this subsection, we want to briefly highlights that the SNR defined by $S^2 := \frac{n\alpha_n\pi_m D^2}{L}$ carries information on the ergodicity of the chain through π_m . The mixing time and the spectral gap are two widely used quantities to measure how fast an ergodic Markov chain will converge to its stationary distribution π . Proposition D.3 states a direct connection between π_m and the latter mentioned coefficients.

Proposition D.3. ([44], Thm. 12.3)

In the following, we denote $\|\cdot\|_{TV}$ the total variation norm. Let P be the transition matrix of a reversible, irreducible Markov chain with state space and for $0 < \epsilon < 1$, let

$$t_{\text{mix}}(\epsilon) := \min\{t > 0 : \sup_x \|P^t(x, \cdot) - \pi\|_{TV} \leq \epsilon\},$$

be the mixing time of the chain. Then it holds

$$t_{\text{mix}}(\epsilon) \leq \log((\epsilon\pi_m)^{-1})/(1 - \lambda_+),$$

where $1 - \lambda_+$ is the right \mathcal{L}_2 -spectral gap of the chain from Definition D.1.

APPENDIX E. THE BAUM-WELCH ALGORITHM WITH INFORMATION GAP

In Sections 4.2.3 and 5, we have presented a reliable approach to solve link prediction or a collaborative filtering problem when we fully observe the graph at time n and when we want to perform some prediction involving future nodes. We propose to consider a more general framework considering that we fully observe the graph at time $n + \delta$ ($\delta \in \mathbb{N}^*$) but we consider that edges involving nodes between time T (with $T < n$) and time n are not reliable. Note that the simpler framework addressed in the paper is simply recovered by taking $n = T + 1$. Hence, we want only to take into account the edges involving pairs of nodes in $\{1, \dots, T, n, \dots, n + \delta\}$. We denote $E_{T,n,\delta}$ this set of edges. We describe the Baum-Welch algorithm in this framework. Running the clustering algorithm on the graph $G = (\{1, \dots, T, n, \dots, n + \delta\}, E_{T,n,\delta})$, we find sequences of estimates for the communities $\hat{\mathbf{C}}_{1:T}, \hat{\mathbf{C}}_{n:n+\delta}$. In the following, we will consider by abuse of notations that for any $j \geq T + 1$, the sequence $\hat{\mathbf{C}}_{1:j}$ represents the sequence $(\hat{\mathbf{C}}_i, i \in [j] \setminus \{T + 1, \dots, n - 1\})$.

The Baum-Welch algorithm consists in a forward and a backward procedure followed by an update step that we describe below. Denoting $\mathbf{1}_K = (1, 1, \dots, 1)^\top \in \mathbb{R}^K$, $\theta^{(0)} = (\mu^{(0)}, P^{(0)}, O^{(0)})$ is initialized as follows

$$\begin{aligned} P^{(0)} &= \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top, \\ \mu^{(0)} &= \mathbf{1}_K^\top, \\ O^{(0)} &= (1 - \epsilon) \text{Id}_K + \frac{\epsilon}{K - 1} (\mathbf{1}_K \mathbf{1}_K^\top - \text{Id}_K), \end{aligned}$$

where $\epsilon \in (0, 1)$ (typically $\epsilon = 10^{-2}$).

– Forward procedure

Let us recall that we have denoted $\alpha_k^{(m)}(i) = \mathbb{Q}_{\theta^{(m)}}(\hat{\mathbf{C}}_{1:i}, C_i = k \mid \theta)$ the probability of seeing the observations $\hat{C}_1, \dots, \hat{C}_i$ and being in state k at time i . This is found recursively with

$$\begin{aligned} \forall k \in [K], \quad \alpha_k^{(m)}(1) &= \mu_k^{(m)} O_{k, \hat{C}_1}^{(m)} \\ \forall k \in [K], \forall i \in [n], \quad \alpha_k^{(m)}(i) &= \begin{cases} \sum_{l \in [K]} \alpha_l^{(m)}(T) ((P^{(m)})^{i-T})_{l,k} & \text{if } T < i \leq n \\ O_{k, \hat{C}_i}^{(m)} \sum_{l \in [K]} \alpha_l^{(m)}(i-1) P_{l,k}^{(m)} & \text{otherwise.} \end{cases} \end{aligned}$$

– Backward procedure

Let us recall that we have denoted $\beta_k^{(m)}(i) = \mathbb{Q}_{\theta^{(m)}}(\hat{\mathbf{C}}_{i+1:n+\delta} \mid C_i = k, \theta)$ the probability of the ending partial sequence $\hat{\mathbf{C}}_{i+1:n+\delta}$ given starting in state k at time i . This is found recursively with

$$\begin{aligned} \forall k \in [K], \quad \beta_k^{(m)}(n) &= 1 \\ \forall k \in [K], \forall i \in [n], \quad \beta_k^{(m)}(i) &= \begin{cases} \sum_{l \in [K]} \beta_l^{(m)}(n-1) ((P^{(m)})^{n-1-i})_{k,l} & \text{if } T \leq i \leq n-2 \\ \sum_{l \in [K]} \beta_l^{(m)}(i+1) P_{k,l}^{(m)} O_{l, \hat{C}_{i+1}}^{(m)} & \text{otherwise.} \end{cases} \end{aligned}$$

– Update step

We can first update the temporary variables $\gamma^{(m)}$ and $\xi^{(m)}$ defined below. The probability of being in state k at time i given the observed sequence $\hat{\mathbf{C}}_{1:n+\delta}$ and the parameters $\theta^{(m)}$ is denoted $\gamma_k^{(m)}(i)$ with

$$\forall k \in [K], \forall i \in [n], \quad \gamma_k^{(m)}(i) = \mathbb{Q}_{\theta^{(m)}}(C_i = k | \hat{\mathbf{C}}_{1:n+\delta}) = \frac{\alpha_k^{(m)}(i)\beta_k^{(m)}(i)}{\sum_{l \in [K]} \alpha_l^{(m)}(i)\beta_l^{(m)}(i)}.$$

The probability of being in state k and l at times i and $i+1$ respectively given the observed sequence $\hat{\mathbf{C}}_{1:n+\delta}$ and parameters $\theta^{(m)}$ is denoted $\xi_{k,l}^{(m)}(i)$ with for all $k, l \in [K]$ and for all $i \in [n]$,

$$\xi_{k,l}^{(m)}(i) = \mathbb{Q}_{\theta^{(m)}}(C_i = k, C_{i+1} = l | \hat{\mathbf{C}}_{1:n+\delta}) = \frac{\mathbb{Q}_{\theta^{(m)}}(C_i = k, C_{i+1} = l, \hat{\mathbf{C}}_{1:n+\delta})}{\mathbb{Q}_{\theta^{(m)}}(\hat{\mathbf{C}}_{1:n+\delta})}.$$

$$\text{Hence, } \xi_{k,l}^{(m)}(i) = \begin{cases} \frac{\alpha_k^{(m)}(i)P_{k,l}^{(m)}\beta_l^{(m)}(i+1)}{\sum_{c,b \in [K]} \alpha_c^{(m)}(i)P_{c,b}^{(m)}\beta_b^{(m)}(i+1)} & \text{if } T \leq i \leq n-2 \\ \frac{\alpha_k^{(m)}(i)P_{k,l}^{(m)}\beta_l^{(m)}(i+1)O_{l,\hat{C}_{i+1}}^{(m)}}{\sum_{c,b \in [K]} \alpha_c^{(m)}(i)P_{c,b}^{(m)}\beta_b^{(m)}(i+1)O_{b,\hat{C}_{i+1}}^{(m)}} & \text{otherwise.} \end{cases}$$

The parameters of the hidden Markov model can now be updated.

$$\begin{aligned} \forall k \in [K], \quad \mu_k^{(m+1)} &= \gamma_k^{(m)}(1), \\ \forall k, l \in [K], \quad P_{k,l}^{(m+1)} &= \frac{\sum_{i=1}^{n-1} \xi_{k,l}^{(m)}(i)}{\sum_{i=1}^{n-1} \gamma_k^{(m)}(i)}, \\ \forall k, l \in [K], \quad O_{k,l}^{(m+1)} &= \frac{\sum_{i=1}^n \mathbb{1}_{\hat{C}_i=l} \gamma_k^{(m)}(i)}{\sum_{i=1}^n \gamma_k^{(m)}(i)}. \end{aligned}$$

APPENDIX F. NOTATIONS

Standard MSBM

$Q \in [0, 1]^{K \times K}$	Connectivity matrix.
$P \in [0, 1]^{K \times K}$	Markov kernel.
$X \in \{0, 1\}^{n \times n}$	Adjacency matrix of the observed graph of size n .
$\mathbf{C}_{1:n}$	Sequence of hidden communities of the nodes 1 to n .
$\hat{\mathbf{C}}_{1:n}$	Sequence of estimates of the communities of the nodes 1 to n given by the clustering algorithm from a given number of communities K and from the adjacency matrix X .
$\eta_i(\mathbf{C}_{1:n})$	Posterior probability of having a connection between node i and the upcoming node (the node $n+1$), namely $\eta_i(\mathbf{C}_{1:n}) = \mathbb{P}(X_{i,n+1} = 1 \mathbf{C}_{1:n})$.

\mathbb{P}	Probability distribution (or probability mass function depending on the context) for which $\mathbf{C}_{1:n}$ is a Markov chain with initial distribution π and Markov kernel P , and for all $i, j \in [n]$, $X_{i,j} \sim \text{Bern}(Q_{C_i, C_j})$. The dependence between the random variables $\mathbf{C}_{1:n}$, $(X_{i,j})_{1 \leq i, j \leq n}$ and $\hat{\mathbf{C}}_{1:n}$ is described by Figure 5.
\hat{P}	Estimate of the Markov kernel P given in Section 3.2.
\hat{Q}	Estimate of the connectivity matrix Q given in Section 3.2.
$\mathbb{P}(\mu', P', Q')$	For any probability distribution μ' on $[K]$, any Markov kernel $P' \in [0, 1]^{K \times K}$ and any connectivity matrix $Q' \in [0, 1]^{K \times K}$, $\mathbb{P}(\mu', P', Q')$ is the probability distribution (or probability mass function depending on the context) for which $\mathbf{C}_{1:n}$ is a Markov chain with initial distribution μ' and Markov kernel P' , and for all $i, j \in [n]$, $X_{i,j} \sim \text{Bern}(Q'_{C_i, C_j})$. The dependence between the random variables $\mathbf{C}_{1:n}$, $(X_{i,j})_{1 \leq i, j \leq n}$ and $\hat{\mathbf{C}}_{1:n}$ is described by Figure 5.

HMM approximation

\mathbb{Q} $O_{k,l}$	\mathbb{Q} is the probability distribution (or probability mass function depending on the context) for which $\mathbf{C}_{1:n}$ is a Markov chain with initial distribution π and Markov kernel P , and for all $i, j \in [n]$, $X_{i,j} \sim \text{Bern}(Q_{C_i, C_j})$. Moreover, the joint distribution of $(\mathbf{C}_{1:n}, \hat{\mathbf{C}}_{1:n})$ factorizes according to the graph of a homogeneous HMM with emission probabilities $O_{k,l} = \mathbb{P}(\hat{C}_1 = l \mid C_1 = k)$ for $k, l \in [K]$.
$\alpha_k(i)$	For any $k \in [K]$ and $i \in [n]$, $\alpha_k(i) = \mathbb{Q}(\hat{\mathbf{C}}_{1:i}, C_i = k)$.
$\beta_k(j)$	For any $k \in [K]$ and $j \in [n]$, $\beta_k(j) = \mathbb{Q}(\hat{\mathbf{C}}_{j+1:n} \mid C_j = k)$.
$\chi_{k,l}^{(i,j)}$	For any $k, l \in [K]$ and $i, j \in [n]$ with $i < j$, $\chi_{k,l}^{(i,j)} = \mathbb{Q}(C_j = l, \hat{\mathbf{C}}_{i+1:j} \mid C_i = k)$.

HMM approximation & Estimates from the Baum-Welch algorithm

$\hat{\alpha}_k(i)$	For any $k \in [K]$ and $i \in [n]$, $\hat{\alpha}_k(i)$ is the estimate of $\alpha_k(i)$ given by the Baum-Welch algorithm.
$\hat{\beta}_k(i)$	For any $k \in [K]$ and $i \in [n]$, $\hat{\beta}_k(i)$ is the estimate of $\beta_k(i)$ given by the Baum-Welch algorithm.
$\hat{\chi}_{k,l}^{(i,j)}$	For any $k, l \in [K]$ and $i, j \in [n]$ with $i < j$, $\hat{\chi}_{k,l}^{(i,j)}$ is the estimate of $\chi_{k,l}^{(i,j)}$ computed using equation (4.13).
$\hat{\mu}$	Estimate of the initial distribution of the Markov chain $\mathbf{C}_{1:n}$ obtained using the Baum-Welch algorithm.
$\hat{O}_{k,l}$	Estimate of the emission probability $O_{k,l}$ given by the Baum-Welch algorithm.

$\hat{\mathbb{Q}}$	Probability distribution (or probability mass function depending on the context) for which $\mathbf{C}_{1:n}$ is a Markov chain with initial distribution $\hat{\mu}$ and Markov kernel \hat{P} , and for all $i, j \in [n]$, $X_{i,j} \sim \text{Bern}(\hat{Q}_{C_i, C_j})$. Moreover, the joint distribution of $(\mathbf{C}_{1:n}, \hat{\mathbf{C}}_{1:n})$ factorizes according to the graph of a homogeneous HMM with emission probabilities $\hat{O}_{k,l}$.
$\hat{\eta}_i^R(\hat{\mathbf{C}}_{1:n})$	Estimate of the posterior probability $\eta_i(\mathbf{C}_{1:n})$. $\hat{\eta}_i^R(\hat{\mathbf{C}}_{1:n})$ is defined in equation (4.14).

REFERENCES

- [1] E. Abbe, Community detection and stochastic block models: recent developments. *J. Mach. Learning Res.* **18** (2017) 1–86.
- [2] E. Abbe and C. Sandon, Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery, in 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, IEEE (2015).
- [3] E. Abbe and C. Sandon, Achieving the KS threshold in the general stochastic block model with linearized acyclic belief propagation (2016).
- [4] L.A. Adamic and E. Adar, Friends and neighbors on the Web. *Social Netw.* **25** (2003) 211–230.
- [5] R. Ahmad and K.S. Xu, Effects of contact network models on stochastic epidemic simulations, in International Conference on Social Informatics. Springer (2017).
- [6] E. Armengol *et al.*, Evaluating link prediction on large graphs, in Artificial intelligence research and development: proceedings of the 18th international conference of the Catalan association for artificial intelligence, vol. 277 (2015).
- [7] A.-L. Barabási, Scale-free networks: a decade and beyond. *Science* **325** (2009) 412–413.
- [8] L.E. Baum and T. Petrie, Statistical inference for probabilistic functions of finite state Markov chains, *Ann. Math. Stat.* **37** (1966) 1554–1563.
- [9] Q. Berthet and N. Baldin, Statistical and Computational Rates in Graph Logistic Regression, in S. Chiappa and R. Calandra (editors), Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics. *Proceedings of Machine Learning Research*, vol. 108, PMLR (2020) 2719–2730.
- [10] F. Bickenbach, E. Bode *et al.*, Markov or not Markov—This should be a question, *Tech. rep.*, Kiel working paper (2001).
- [11] K.P. Burnham and D.R. Anderson, Multimodel inference: understanding AIC and BIC in model selection. *Sociolog. Methods Res.* **33** (2004) 261–304.
- [12] O. Cappé, E. Moulines and T. Rydén, Inference in Hidden Markov Models, in Proceedings of EUSFLAT conference (2009) 14–16.
- [13] A. Celisse, J.-J. Daudin and L. Pierre, Consistency of maximum-likelihood and variational estimators in the stochastic block model. *Electr. J. Stat.* **6** (2012) 1847–1899.
- [14] A. Celisse, J.-J. Daudin and L. Pierre, Consistency of maximum-likelihood and variational estimators in the Stochastic Block Model. *Electr. J. Stat.* **6** (2012) 1847–1899.
- [15] M. Charikar, S. Guha, E. Tardos and D. Shmoys, A constant-factor approximation algorithm for the k-median problem. *J. Comput. Syst. Sci.* **65** (2002) 129–149.
- [16] Y. Chen and J. Xu, Statistical-Computational Tradeoffs in Planted Problems and Submatrix Localization with a Growing Number of Clusters and Submatrices (2014).
- [17] Z. Chen and S. Wang, A review on matrix completion for recommender systems. *Knowl. Inf. Syst.* **64** (2022) 1–34.
- [18] S. Chin, A. Rao and V. Vu, Stochastic Block Model and Community Detection in the Sparse Graphs: A spectral algorithm with optimal rate of recovery (2015).
- [19] G.A. Churchill, Stochastic models for heterogeneous DNA sequences. *Bull. Math. Biol.* **51** (1989) 79–94.
- [20] A. Clauset, C. Moore and M.E.J. Newman, Hierarchical structure and the prediction of missing links in networks. *Nature* **453** (2008) 98–101.
- [21] J.C. Costello, L.M. Heiser, E. Georgii, M. Gönen, M.P. Menden, N.J. Wang, M. Bansal, M. Ammad-Ud-Din, P. Hintsanen, S.A. Khan *et al.*, A community effort to assess and improve drug sensitivity prediction algorithms, *Nat. Biotechnol.* **32** (2014) 1202–1212.
- [22] S. Das and S.K. Das, A probabilistic link prediction model in time-varying social networks, in *2017 IEEE International Conference on Communications (ICC)* (2017) 1–6.
- [23] J.-J. Daudin, F. Picard and S. Robin, A mixture model for random graphs. *Stat. Comput.* **18** (2008) 173–183.
- [24] A. Divakaran and A. Mohan, Temporal link prediction: a survey. *New Generat. Comput.* **38** (2020) 213–258.
- [25] Q. Duchemin and Y. De Castro, Markov random geometric graph, MRGG: a growth model for temporal dynamic networks. *Electr. J. Stat.* **16** (2022) 671–699.
- [26] Y. Fei and Y. Chen, Exponential error rates of SDP for block models: beyond Grothendieck’s inequality. *IEEE Trans. Inf. Theory* **PP** (2017).
- [27] X. Feng, J. Zhao and K. Xu, Link prediction in complex networks: a clustering perspective. *Eur. Phys. J. B* **85** (2012) 1–9.
- [28] D.R. Fredkin and J.A. Rice, Correlation functions of a function of a finite-state Markov process with application to channel kinetics. *Math. Biosci.* **87** (1987) 161–172.

- [29] C. Giraud and N. Verzelen, Partial recovery bounds for clustering with the relaxed K -means. *Math. Stat. Learn.* **1** (2019) 317–374.
- [30] O. Guédon and R. Vershynin, Community detection in sparse networks via Grothendieck’s inequality. *Probab. Theory Related Fields* **165** (2014).
- [31] R. Guimerà and M. Sales-Pardo, Missing and spurious interactions and the reconstruction of complex networks. *Proc. Natl. Acad. Sci.* **106** (2009) 22073–22078.
- [32] B. Hajek, Y. Wu and J. Xu, Semidefinite Programs for Exact Recovery of a Hidden Community (2016).
- [33] G.-J. Huizing, G. Peyré and L. Cantini, Optimal Transport improves cell–cell similarity inference in single-cell omics data. *bioRxiv* (2021).
- [34] B. Jiang, Q. Sun and J. Fan, Bernstein’s inequality for general Markov chains. arXiv preprint [arXiv:1805.10721](https://arxiv.org/abs/1805.10721) (2018).
- [35] B.H. Juang and L.R. Rabiner, Hidden Markov models for speech recognition. *Technometrics* **33** (1991) 251–272.
- [36] B. Karrer and M. Newman, Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83** (2011).
- [37] N. Keriven and S. Vaiteer, Sparse and Smooth: improved guarantees for Spectral Clustering in the Dynamic Stochastic Block Model (2020).
- [38] A. Kölzsch, E. Kleyheeg, H. Kruckenberg, M. Kaatz and B. Blasius, A periodic Markov model to formalize animal migration on a network. *Royal Soc. Open Sci.* **5** (2018) 180438.
- [39] H. Kruckenberg, G. Müskens and B. Ebbinge, Data from: a periodic Markov model to formalise animal migration on a network [white-fronted goose data] (2018).
- [40] A. Kumar, S.S. Singh, K. Singh and B. Biswas, Link prediction techniques, applications, and performance: a survey. *Physica A* **553** (2020) 124289.
- [41] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tompkins and E. Upfal, The Web as a Graph, in Proceedings of the Nineteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS ’00, Association for Computing Machinery, New York, NY, USA (2000) 1–10.
- [42] P. Latouche, E. Birmele and C. Ambroise, Variational Bayesian inference and complexity control for stochastic block models. *Stat. Modell.* **12** (2012) 93–115.
- [43] J. Lei and A. Rinaldo, Consistency of spectral clustering in stochastic block models. *Ann. Stat.* **43** (2015).
- [44] Levin, Markov chains and mixing times, American Mathematical Soc. (2017).
- [45] J.H. Levine, E.F. Simonds, S.C. Bendall, K.L. Davis, D.A. El-ad, M.D. Tadmor, O. Litvin, H.G. Fienberg, A. Jager, E.R. Zunder *et al.*, Data-driven phenotypic dissection of AML reveals progenitor-like cells that correlate with prognosis. *Cell* **162** (2015) 184–197.
- [46] X. Li, N. Du, H. Li, K. Li, J. Gao and A. Zhang, A deep learning approach to link prediction in dynamic networks, in Proceedings of the 2014 SIAM International conference on data mining. SIAM (2014) 289–297.
- [47] X. Ma, P. Sun and Y. Wang, Graph regularized nonnegative matrix factorization for temporal link prediction in dynamic networks. *Physica A* **496** (2018) 121–136.
- [48] M. Mariadassou, S. Robin and C. Vacher, Uncovering latent structure in valued graphs: a variational approach. *Ann. Appl. Stat.* **4** (2010) 715–742.
- [49] C. Matias and V. Miele, Statistical clustering of temporal networks through a dynamic stochastic block model. *J. Roy. Stat. Soc.: Ser. B (Statistical Methodology)* **79** (2015) 1119–1141.
- [50] S.A. Morris and M.L. Goldstein, Manifestation of research teams in journal literature: a growth model of papers, authors, collaboration, coauthorship, weak ties, and Lotka’s law. *J. Assoc. Inf. Sci. Technol.* **58** (2007) 1764–1782.
- [51] M.E. Newman, Clustering and preferential attachment in growing networks. *Phys. Rev. E* **64** (2001) 025102.
- [52] M. Opper and D. Saad, Advanced Mean Field Methods: Theory and Practice. The MIT Press (2001).
- [53] J. Peng and Y. Wei, Approximating K -means-type clustering via semidefinite programming. *SIAM J. Optim.* **18** (2007) 186–205.
- [54] M. Pensky and T. Zhang, Spectral clustering in the dynamic stochastic block model. *Electr. J. Stat.* **13** (2017) 678–709.
- [55] W. Perry and A.S. Wein, A semidefinite program for unbalanced multisection in the stochastic block model (2015).
- [56] P. Sarkar, D. Chakrabarti and M. Jordan, Nonparametric link prediction in dynamic networks. arXiv preprint [arXiv:1206.6394](https://arxiv.org/abs/1206.6394) (2012).
- [57] R.R. Sarukkai, Link prediction and path analysis using Markov chains. *Comput. Netw.* **33** (2000) 377–386.
- [58] O. Shchur and S. Günnemann, Overlapping Community Detection with Graph Neural Networks (2019).
- [59] A.L. Smith, D.M. Asta and C.A. Calder, The geometry of continuous latent space models for network data. *Statist. Sci.* **34** (2019) 428–453.
- [60] P. Wang, B. Xu, Y. Wu and X. Zhou, Link prediction in social networks: the state-of-the-art. *Sci. China Inf. Sci.* **58** (2015) 1–38.
- [61] D.J. Watts and S.H. Strogatz, Collective dynamics of ‘small-world’ networks. *Nature* **393** (1998) 440–442.
- [62] R. Weiss and B. Nadler, Learning parametric-output HMMs with two aliased states, in F. Bach and D. Blei (editors), Proceedings of the 32nd International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 37, PMLR, Lille, France (2015) 635–644.
- [63] Z. Wu and Y. Chen, Link prediction using matrix factorization with bagging, in 2016 IEEE/ACIS 15th International Conference on Computer and Information Science (ICIS), IEEE (2016) 1–6.
- [64] K.S. Xu, Stochastic block transition models for dynamic networks. *CoRR* (2014).

- [65] M. Yang, J. Simm, C.C. Lam, P. Zakeri, G.J. van Westen, Y. Moreau and J. Saez-Rodriguez, Linking drug target and pathway activation for effective therapy using multi-task learning. *Sci. Rep.* **8** (2018) 1–10.
- [66] T. Yang, Y. Chi, S. Zhu, Y. Gong and R. Jin, Detecting communities and their evolutions in dynamic social networks – a Bayesian approach. *Mach. Learn.* **82** (2011) 157–189.
- [67] W. Yuan, K. He, D. Guan, L. Zhou and C. Li, Graph kernel based link prediction for signed social networks. *Inf. Fusion* **46** (2019) 1–10.
- [68] X. Zhang, X. Wang, C. Zhao, D. Yi and Z. Xie, Degree-corrected stochastic block models and reliability in networks. *Physica A* **393** (2014) 553–559.
- [69] T. Zhou, L. Lü and Y.-C. Zhang, Predicting missing links via local Information. *Eur. Phys. J. B* **71** (2009) 623–630.



Please help to maintain this journal in open access!

This journal is currently published in open access under the Subscribe to Open model (S2O). We are thankful to our subscribers and supporters for making it possible to publish this journal in open access in the current year, free of charge for authors and readers.

Check with your library that it subscribes to the journal, or consider making a personal donation to the S2O programme by contacting subscribers@edpsciences.org.

More information, including a list of supporters and financial transparency reports, is available at <https://edpsciences.org/en/subscribe-to-open-s2o>.