Stratified stochastic variational inference for high-dimensional network factor model

Emanuele Aliverti[1] and Massimiliano Russo[2]

[1]*Department of Economics, University Ca' Foscari Venezia*

[2]*Harvard–MIT Center for Regulatory Science, Harvard Medical School and Department of Data Science*

Dana-Farber Cancer Institute

Abstract

There has been considerable recent interest in Bayesian modeling of high-dimensional networks via latent space approaches. When the number of nodes increases, estimation based on Markov Chain Monte Carlo can be extremely slow and show poor mixing, thereby motivating research on alternative algorithms that scale well in high-dimensional settings. In this article, we focus on the latent factor model, a widely used approach for latent space modeling of network data. We develop scalable algorithms to conduct approximate Bayesian inference via stochastic optimization. Leveraging sparse representations of network data, the proposed algorithms show massive computational and storage benefits, and allow to conduct inference in settings with thousands of nodes. An R package with an efficient C++ implementation of the proposed algorithms is provided.

1 INTRODUCTION

Network data are routinely collected and analyzed in different fields of science; for example, neuroscience (Bullmore and Sporns, 2009), genetics (Wu et al., 2009) and epidemiology (Keeling and Eames, 2005), among many others. Refer also to Newman (2018) for an introduction to network data and their analysis. One of the main goals in network data analysis is to characterize the geometry underlying node relationships, providing a parsimonious, yet flexible, representation of the connectivity patterns. This goal can be achieved modeling the connectivity architectures in terms of a low-dimensional latent structure, where the edges are represented as conditionally independent random variables given a set of latent coordinates (e.g., Hoff et al., 2002). Beside improving computation, these approaches provide concrete benefits in interpretation; for example, the latent structure can be related with observable covariates to improve the understanding of connectivity patterns, or divided into clusters to detect nodes that behave similarly in terms of their unobservable features (e.g., Aliverti and Durante, 2019). The increasing availability of network data has further motivated the development of novel latent structure models for networks, covering more complex settings such as dynamic networks (e.g., Durante and Dunson, 2014; Sewell and Chen, 2015), multilayer networks (e.g., Gollini and Murphy, 2016; D'Angelo et al., 2019) and populations of networks (Durante et al., 2017). Since the number of edges grows quadratically with the number of nodes, representing the latent structure of a network is computationally challenging, even for networks with few hundred nodes. This issue stimulates the development of novel methods and computational routines to accommodate large network structures efficiently, leveraging different network properties—such as sparsity or block structures—to facilitate computations. For example, large network data are often very sparse, with the number of observed edges being much smaller than the total number of possible connections. This feature allows to parsimoniously store network data via sparse matrix representations, or, equivalently, edge-lists formats (e.g., Csardi and Nepusz, 2006). However, estimation of latent structure models often requires edge-specific operations, lowering the benefits of these representations; see for example Ho et al. (2016) for related arguments. In addition, most inference procedures for latent space models are based on Markov Chain Monte Carlo (MCMC) algorithms (e.g., Hoff, 2021), and these methods scale poorly with the

number of nodes.

There have been some attempts to improve computational efficiency and mixing of MCMC for latent space model for networks. For example, Raftery et al. (2012) derive an unbiased estimator of the log-likelihood based on an informative subset of nodes, and successfully perform MCMC estimation in an application with roughly three thousands nodes (Raftery et al., 2012, Section 4.2). Although this approach effectively reduces the cost of each likelihood evaluation, inference is still computationally demanding; the algorithm requires a preliminary pilot MCMC run and a considerable amount of storage to perform Monte Carlo integration.

These computational issues motivate the development of scalable methods for approximate Bayesian inference, with Variational Bayes (VB) being a popular option. Specifically, VB algorithms approximate the posterior distribution via optimization, estimating the closest member (in Kullbak-Lielber divergence) within a pre-specified class of distributions. This class includes some restrictions to achieve computational tractability; for example, the resulting approximate posterior distribution factorizes in independent blocks of parameters (mean-field VB), or follows a specific parametric form (e.g., multivariate Gaussian). Refer to Blei et al. (2017) and references therein for more details. To date, the currently available VB routines for latent space models (Gollini and Murphy, 2016; Salter-Townshend and Murphy, 2013) are based on restrictive assumptions on the functional form of the variational distributions and rely on several approximations of the complete log-likelihood function (Gollini and Murphy, 2016). In addition, these algorithms encounter computational issues in high-dimensional settings, since they scale quadratically with the number of nodes (Salter-Townshend and Murphy, 2013).

In this paper, we focus on the Latent Factor Model (LFM, Hoff, 2005)—also referred to as random dot-model (Young and Scheinerman, 2007)—for binary undirected networks, a very general approach that can approximate several latent position models under mild conditions (Tang et al., 2013; Athreya et al., 2018). We propose scalable algorithms based on a stratified stochastic variational approximation for the LFM, which we refer to as SVILF in the following. From a computational perspective, SVILF explicitly relies on sparse network representations via edge-lists, which allows to scale up computations for networks consisting of thousands of nodes. Leveraging a

conditionally conjugate exponential family representation, we provide a unified framework for the LFM with logistic and probit link function and illustrate how SVILF can be directly implemented under both specifications.

2 METHODS

2.1 Latent Factor Model

Let $\mathbf{Y}$ denote an $n \times n$ binary symmetric adjacency matrix with elements $y_{ij} = y_{ji} = 1$ denoting the presence of an edge between node i and node j with $i = 2, \ldots, n$ and $j = 1, \ldots, i-1$, and $y_{ij} = y_{ji} = 0$ otherwise. The LFM for networks parametrizes the edges as conditional independent Bernoulli random variables given a set of latent positions $\mathbf{w}_i = (w_{i1}, \ldots, w_{iH}) \in \mathbb{R}^H$, $i = 1, \ldots, n$. This choice implies

$$(y_{ij} \mid \pi_{ij}) \sim \text{Ber}(\pi_{ij}), \quad g(\pi_{ij}) = \mathbf{w}_i^\mathsf{T} \mathbf{w}_j = \sum_{h=1}^{H} w_{ih} w_{jh}, \quad i = 2, \ldots, n, \, j = 1, \ldots, i-1, \quad (1)$$

where $g : [0, 1] \to \mathbb{R}$ is a monotone link function. Popular choices for g include the logit and probit link (e.g., Agresti, 2015). See also Young and Scheinerman (2007) for alternative specifications of the model outlined in Equation (1). In a Bayesian setting, we typically assign independent Gaussian priors to the latent factors, thereby letting $\mathbf{w}_i \sim \text{N}_H(\mathbf{a}_0, \mathbf{I}_H)$ for $i = 1, \ldots, n$. A prior mean $\mathbf{a}_0 = (a_{01}, \ldots, a_{0H})$ different from zero allows to center the factors around an expected network sparsity, accounting for the baseline probability of observing a connection. According to Equation (1), the probability of observing a connection between node i and node j depends on their latent positions $\mathbf{w}_i$ and $\mathbf{w}_j$. The more similar node i and node j are in the latent space, the more likely it is to observe an edge connecting them. The similarity between two nodes in the latent space is computed with the multiplicative effect $\mathbf{w}_i^\mathsf{T} \mathbf{w}_j$, and this measure is particularly helpful in characterizing transitivity and uncovering group structures (Hoff, 2021).

The LFM reduces the number of free parameters needed to characterize the network from $n(n-1)/2$ to nH, providing a low-dimensional representation of the connectivity patterns. Additionally, the LFM is a very flexible latent structure model for networks, since several latent position models

can be represented as Equation (1) when H grows (Tang et al., 2013). For these reasons, the LFM has been used as a building block for many generalizations involving, among others, time-varying and covariate-dependent networks (e.g., Durante and Dunson, 2014; Sewell and Chen, 2015). Therefore, efficient algorithms for the LFM are crucial to analyze large network data, which are routinely collected in different fields of application.

Recalling Equation (1), it is worth highlighting that the predictor $g(\pi_{ij})$ is linear in the latent factors. Focusing on a single factor $\mathbf{w}_i$, we can recast the model outlined in Equation (1) as a conditional binary regression, given the other factors $\mathbf{w}_j$, $j \neq i$. Denoting with $\mathbf{y}_i = (y_{i1}, \dots, y_{i\,i-1}, y_{i\,i+1}, \dots, y_{in})$ the i-th row of the adjacency matrix $\mathbf{Y}$, with $\boldsymbol{\pi}_i = (\pi_{i1}, \dots, \pi_{i\,i-1}, \pi_{i\,i+1}, \dots, \pi_{in})$ the $(n-1)$-variate vector of associated edge probabilities and with $\mathbf{W}_{-i}$ the $(n-1) \times H$ dimensional matrix obtained stacking the factors $\{\mathbf{w}_j\}_{j \neq i}$, it holds that

$$g(\boldsymbol{\pi}_i) = \mathbf{W}_{-i}\mathbf{w}_i, \tag{2}$$

where the link function $g(\cdot)$ in Equation (2) is applied element-wise.

Equation (2) allows one to rely on iterative algorithms for posterior inference under the LMF, considering n binary regressions where each factors acts, in turn, as a regression coefficient. Moreover, for some specific link functions g, data-augmentation schemes are available to further simplify posterior inference. Specifically, with the logit link function conditional conjugacy can be retrieved relying on the Pòlya-Gamma (PG) data augmentation strategy introduced by Polson et al. (2013), that leads to the following full conditionals distributions:

$$(z_{ij} \mid \mathbf{w}_i, \mathbf{w}_j, y_{ij}) \sim \text{PG}(1, \mathbf{w}_i^\intercal \mathbf{w}_j), \quad i = 2, \dots, n, \quad j = 1, \dots, i-1, \tag{3}$$

$$(\mathbf{w}_i \mid \mathbf{W}_{-i}, \mathbf{z}_i, \mathbf{y}_i) \sim \text{N}(\boldsymbol{\mu}_{\mathbf{w}_i}, \boldsymbol{\Sigma}_{\mathbf{w}_i}), \quad i = 1, \dots, n,$$

where

$$\boldsymbol{\Sigma}_{\mathbf{w}_i} = \left(\mathbf{W}_{-i}^\intercal \text{diag}(\mathbf{z}_i)\mathbf{W}_{-i} + \mathbf{I}_H\right)^{-1}, \quad \boldsymbol{\mu}_{\mathbf{w}_i} = \boldsymbol{\Sigma}_{\mathbf{w}_i}\left(\mathbf{W}_{-i}^\intercal(\mathbf{y}_i - 0.5 \cdot \mathbf{1}_{n-1}) + \mathbf{I}_H \mathbf{a}_0\right),$$

and $\mathbf{z}_i = (z_{i1}, \ldots, z_{i\,i-1}, z_{i\,i+1}, \ldots, z_{in})$. In Equation (3), PG$(c, d)$ denotes the density of a Pòlya-Gamma distribution with parameters c and d; refer to Polson et al. (2013) for more details.

Similarly, with the probit link function, conditional conjugacy is obtained adapting the data-augmentation proposed in Albert and Chib (1993), which introduces auxiliary observations distributed as truncated normal random variables. The full conditional distributions under this specification correspond to

$$(z_{ij} \mid \mathbf{w}_i, \mathbf{w}_j, y_{ij}) \sim \begin{cases} \text{TN}(\mathbf{w}_i^\intercal \mathbf{w}_j, 1, [0, +\infty]), & \text{if } y_{ij} = 1 \\ \text{TN}(\mathbf{w}_i^\intercal \mathbf{w}_j, 1, [-\infty, 0]), & \text{if } y_{ij} = 0 \end{cases} \quad i = 2, \ldots, n, \quad j = 1, \ldots, i-1, \quad (4)$$

$$(\mathbf{w}_i \mid \mathbf{W}_{-i}, \mathbf{z}_i, \mathbf{y}_i) \sim \text{N}_H(\boldsymbol{\mu}_{\mathbf{w}_i}, \boldsymbol{\Sigma}_{\mathbf{w}_i}), \qquad\qquad\qquad i = 1, \ldots, n,$$

where

$$\boldsymbol{\Sigma}_{\mathbf{w}_i} = \left(\mathbf{W}_{-i}^\intercal \mathbf{W}_{-i} + \mathbf{I}_H \right)^{-1}, \quad \boldsymbol{\mu}_{\mathbf{w}_i} = \boldsymbol{\Sigma}_{\mathbf{w}_i} \left(\mathbf{W}_{-i}^\intercal \mathbf{z}_i + \mathbf{I}_H \mathbf{a}_0 \right),$$

and with TN$(\mu, \sigma^2, [a, b])$ denoting a truncated normal distribution with parameters (μ, σ^2), restricted over the interval $[a, b]$. Posterior inference via MCMC relies on iterative sampling from Equation (3) or (4), constructing a Markov chain which has the joint posterior distribution $p(\mathbf{W}, \mathbf{z} \mid \mathbf{Y})$ as a limiting distribution (Gelfand and Smith, 1990). Some factors can be omitted in the expressions above; for example, $p(\mathbf{w}_i \mid \mathbf{W}_{-i}, \mathbf{z}_i, \mathbf{y}_i) = p(\mathbf{w}_i \mid \mathbf{W}_{-i}, \mathbf{z}_i)$ in Equation (4). However, this redundant notation helps to unify both algorithms under a general specification, as outlined in the following section.

2.2 Conditional conjugancy and Variational Inference

The conditional distributions for the augmented variables z_{ij} in Equation (3) and (4) can be expressed as

$$p(z_{ij} \mid \mathbf{w}_i, \mathbf{w}_j, y_{ij}) \propto \exp\left\{ \eta_{ij}(\mathbf{w}_i, \mathbf{w}_j) z_{ij} - \kappa(\eta_{ij}(\mathbf{w}_i, \mathbf{w}_j)) \right\} p(z_{ij} \mid y_{ij}),$$
$$i = 2, \ldots, n, j = 1, \ldots, i-1, \tag{5}$$

with $\kappa(\cdot)$ denoting the log-partition function. Equation (5) allows one to express the LFM with logit and probit link with the same exponential family structure. When g is the logit link, $\eta_{ij}(\mathbf{w}_i, \mathbf{w}_j) = -0.5(\mathbf{w}_i^\mathsf{T}\mathbf{w}_j)^2$ and $p(z_{ij} \mid y_{ij})$ corresponds to the density of a PG$(0, 1)$; when g is the probit link, $\eta_{ij}(\mathbf{w}_i, \mathbf{w}_j) = \mathbf{w}_i^\mathsf{T}\mathbf{w}_j$ and $p(z_{ij} \mid y_{ij})$ corresponds to the density of a TN$(0, 1, [-\infty, 0])$ if $y_{ij} = 0$ and TN$(0, 1, [0, \infty])$ if $y_{ij} = 1$.

Conditionally on the observed data, the augmented variables and the factors $\{\mathbf{w}_j\}_{j \neq i}$, the full-conditional distribution of each $\mathbf{w}_i$ corresponds to a multivariate Gaussian density,

$$p(\mathbf{w}_i \mid \mathbf{W}_{-i}, \mathbf{z}_i, \mathbf{y}_i) \propto \exp \left\{ \boldsymbol{\eta}_{i1} \left(\mathbf{W}_{-i}, \mathbf{y}_i, \mathbf{z}_i\right)^\mathsf{T} \mathbf{w}_i + \mathrm{vec}\left(\boldsymbol{\eta}_{i2}\left(\mathbf{W}_{-i}, \mathbf{z}_i\right)\right)^\mathsf{T} \cdot \mathrm{vec}\left(\mathbf{w}_i \mathbf{w}_i^\mathsf{T}\right) - \right.$$
$$\left. \kappa\left(\boldsymbol{\eta}_{i1}\left(\mathbf{W}_{-i}, \mathbf{y}_i, \mathbf{z}_i\right), \boldsymbol{\eta}_{i2}\left(\mathbf{W}_{-i}, \mathbf{z}_i\right)\right) \right\}, \quad i = 1, \dots, n, \tag{6}$$

where

$$\boldsymbol{\eta}_{i1}\left(\mathbf{W}_{-i}, \mathbf{y}_i, \mathbf{z}_i\right) = \mathbf{W}_{-i}^\mathsf{T}(\mathbf{y}_i - 0.5 \cdot \mathbf{1}_{n-1}) + \mathbf{I}_H \mathbf{a}_0, \quad \boldsymbol{\eta}_{i2}\left(\mathbf{W}_{-i}, \mathbf{z}_i\right) = -\frac{1}{2}\left(\mathbf{W}_{-i}^\mathsf{T} \mathrm{diag}(\mathbf{z}_i)\mathbf{W}_{-i} + \mathbf{I}_H\right),$$

with the logit link function, and

$$\boldsymbol{\eta}_{i1}\left(\mathbf{W}_{-i}, \mathbf{y}_i, \mathbf{z}_i\right) = \mathbf{W}_{-i}^\mathsf{T} \mathbf{z}_i + \mathbf{I}_H \mathbf{a}_0, \quad \boldsymbol{\eta}_{i2}\left(\mathbf{W}_{-i}, \mathbf{z}_i\right) = -\frac{1}{2}\left(\mathbf{W}_{-i}^\mathsf{T}\mathbf{W}_{-i} + \mathbf{I}_H\right),$$

with the probit link function.

This conditional representation of the LFM facilitates the development of algorithms that exploit maximization strategies for binary regression; for example, MAP optimization via EM algorithm or approximate inference via VB (Consonni and Marin, 2007; Durante and Rigon, 2019). As discussed in Section 1, the focus of the mean-field VB is on finding an approximation of the posterior distribution $p(\mathbf{W}, \mathbf{z} \mid \mathbf{Y})$ within a restricted class of densities $\mathcal{Q}$, specified as

$$\mathcal{Q} = \left\{ q(\mathbf{W}, \mathbf{z}) : q(\mathbf{W}, \mathbf{z}) = \prod_{i=1}^{n} q(\mathbf{w}_i; \boldsymbol{\lambda}_i) \prod_{i=2}^{n} \prod_{j=1}^{n-1} q(z_{ij}; \psi_{ij}) \right\}. \tag{7}$$

Note that each factor $\mathbf{w}_i$ is a function of its own variational parameters $\boldsymbol{\lambda}_i$; similarly, the augmented variables z_{ij} are function of the parameter ψ_{ij}. The optimal VB solution $q^\star(\mathbf{W}, \mathbf{z})$ corresponds to

the distribution within $\mathcal{Q}$ that minimizes the Kullback-Leibler (KL) divergence, defined as

$$\text{KL}\left[q(\mathbf{W},\mathbf{z})||p(\mathbf{W},\mathbf{z}\mid\mathbf{Y})\right] = \mathbb{E}_{q(\mathbf{W},\mathbf{z})}\left[\log q(\mathbf{W},\mathbf{z})\right] - \mathbb{E}_{q(\mathbf{W},\mathbf{z})}\left[\log p(\mathbf{W},\mathbf{z}\mid\mathbf{Y})\right], \quad q(\mathbf{W},\mathbf{z})\in\mathcal{Q}. \quad (8)$$

In practice, VB procedures maximize the related objective function

$$\text{ELBO}\left[q(\mathbf{W},\mathbf{z})\right] = \mathbb{E}_{q(\mathbf{W},\mathbf{z})}\left[\log p(\mathbf{W},\mathbf{z},\mathbf{Y})\right] - \mathbb{E}_{q(\mathbf{W},\mathbf{z})}\left[\log q(\mathbf{W},\mathbf{z})\right], \quad q(\mathbf{W},\mathbf{z}),\in\mathcal{Q}, \quad (9)$$

which corresponds to the negative KL up to an additive constant not depending on the parameters; see, for example, Blei et al. (2017) and Bishop (2006). The mean field assumption in Equation (7) and the conditionally conjugate representation facilitate a Coordinate Ascent Variational Inference (CAVI) routine to maximize Equation (9), where each distribution is iteratively optimized with respect to the others in an iterative fashion (e.g., Bishop, 2006). The variational distributions composing the optimal solution are in the same family of the full-conditional distributions outlined in Equations (3) and (4), and their parameters can be analytically expressed in terms of variational expectations (e.g., Blei et al., 2017). Under our specification,

$$\boldsymbol{\lambda}_{i1} = \mathbb{E}_{q(\mathbf{w}_{-i},\mathbf{z}_i)}\left[\boldsymbol{\eta}_{i1}\left(\mathbf{W}_{-i},\mathbf{y}_i,\mathbf{z}_i\right)\right], \quad \boldsymbol{\lambda}_{i2} = \mathbb{E}_{q(\mathbf{w}_{-i},\mathbf{z}_i)}\left[\boldsymbol{\eta}_{i2}\left(\mathbf{W}_{-i},\mathbf{z}_i\right)\right], \quad i=1,\ldots,n \quad (10)$$

and

$$\psi_{ij} = \mathbb{E}_{q(\mathbf{w}_i,\mathbf{w}_j)}\left[\eta_{ij}(\mathbf{w}_i,\mathbf{w}_j)\right], \quad i=2,\ldots,n,\ j=1,\ldots,i-1. \quad (11)$$

Algorithm 1 illustrates the CAVI algorithms for the LFM with logit link; refer to the Supplementary Materials for the probit link function. At each iteration t, variational expectations outlined in Equations (10) and (11) are taken with respect to the currently optimized distribution $q^{(t-1)}$, iterating until convergence. Note also that at each iteration the CAVI provides a monotone sequence

of the ELBO, and convergence to a local optimum is guaranteed (Blei et al., 2017; Bishop, 2006).

Algorithm 1: CAVI for LFM with logit link.

Initialize $\left\{\mathbf{S}_1^{(1)}, \ldots, \mathbf{S}_n^{(1)}\right\}$ and $\left\{\boldsymbol{\mu}_1^{(1)}, \ldots, \boldsymbol{\mu}_n^{(1)}\right\}$.

for $t = 2$ until convergence **do**

[1] $q^{(t)}(\mathbf{w}_i)$ is the density of a $\mathrm{N}_H(\boldsymbol{\mu}_i^{(t)}, \boldsymbol{\Sigma}_i^{(t)})$ with $\boldsymbol{\mu}_i^{(t)} = [-2\boldsymbol{\lambda}_{i2}^{(t)}]^{-1}\boldsymbol{\lambda}_{i1}^{(t)}$, $\boldsymbol{\Sigma}_i^{(t)} = [-2\boldsymbol{\lambda}_{i2}^{(t)}]^{-1}$

and natural parameters

$$\boldsymbol{\lambda}_{i1}^{(t)} = \sum_{j \neq i} \boldsymbol{\mu}_j^{(t-1)} [y_{ij} - 0.5] + \mathbf{I}_H \mathbf{a}_0, \quad \boldsymbol{\lambda}_{i2}^{(t)} = -0.5 \left(\sum_{j \neq i} \mathbf{S}_j^{(t-1)} \cdot \bar{z}_{ij}^{(t-1)} + \mathbf{I}_H \right),$$

for $i = 1, \ldots, n$. In the above expression, $\bar{z}_{ij}^{(t-1)} = 0.5[\xi_{ij}^{(t-1)}]\tanh(0.5\xi_{ij}^{(t-1)})$, where

$$\mathbf{S}_j^{(t-1)} = \left[\boldsymbol{\Sigma}_j^{(t-1)} + \left[\boldsymbol{\mu}_j^{(t-1)}\right] \left[\boldsymbol{\mu}_j^{(t-1)}\right]^\mathsf{T} \right], \quad \xi_{ij}^{(t-1)} = \left[\mathrm{vec}\left(\mathbf{S}_i^{(t-1)}\right)^\mathsf{T} \mathrm{vec}\left(\mathbf{S}_j^{(t-1)}\right) \right]^{\frac{1}{2}}$$

[2] $q^{(t)}(z_{ij})$ is the density of a $\mathrm{PG}(1, \xi_{ij}^{(t)})$, for $i = 2, \ldots, n$ and $j = 1, \ldots, i-1$.

Output $q^\star(\mathbf{W}, \mathbf{z}) = \prod_{i=1}^n q^\star(\mathbf{w}_i) \prod_{i=2}^n \prod_{j=1}^{n-1} q^\star(z_{ij})$.

---

3 STRATIFIED STOCHASTIC VARIATIONAL INFERENCE

The CAVI algorithms introduced in Section 2.2 provide efficient routines to perform approximate Bayesian inference under the LFM. However, when the number of nodes n is extremely large, computational issues might drastically limit the analysis. For example, in Step [1] of Algorithm 1, the updates of each natural parameter $\boldsymbol{\lambda}_i$ involve summing over $(n-1)$ terms, with $i = 1, \ldots, n$. Also, conditional conjugacy simplifies the derivation of analytical results, but it requires to update $n(n-1)/2$ augmented observations even when the dimension H of the latent space is small, therefore exacerbating computational and storage issues. Although variational routines generally require significantly less iterations than MCMC to reach convergence, the overall complexity of the CAVI algorithms is still $\mathcal{O}(n^2)$; therefore, CAVI provides a viable solution only in settings where n is moderately large (Salter-Townshend and Murphy, 2013). We introduce a novel approach that allows to approximate the posterior distribution of the LFM for high-dimensional network-data.

A scalable generalization of classical CAVI is provided by Stochastic Variational Inference (SVI, Hoffman et al., 2013), where stochastic optimization (Robbins and Monro, 1951) is used to reduce

the computation cost of VB routines. We follow a similar perspective and develop a stochastic VB algorithm specifically tailored for sparse network data. Following Hoffman et al. (2013), it is useful to rewrite the CAVI routines outlined in Algorithms 1 as the solutions of a system of estimating equations, obtained computing the derivatives of the ELBO with respect to the parameters of the variational distributions.

For each fixed factor $q(\mathbf{w}_i; \boldsymbol{\lambda}_i)$, computing the gradient of the ELBO outlined in Equation (9) with respect to $\boldsymbol{\lambda}_i$ and equating to 0 leads to the following estimating equations.

$$\mathbb{E}_{q(\mathbf{w}_{-i}, \mathbf{z}_i)} \left[\boldsymbol{\eta}_{i1} \left(\mathbf{w}_{-i}, \mathbf{y}_i, \mathbf{z}_i \right) \right] - \boldsymbol{\lambda}_{i1} = 0, \quad \mathbb{E}_{q(\mathbf{w}_{-i}, \mathbf{z}_1)} \left[\boldsymbol{\eta}_{i2} \left(\mathbf{w}_{-i}, \mathbf{z}_i \right) \right] - \boldsymbol{\lambda}_{i2} = 0. \tag{12}$$

See Hoffman et al. (2013, Section 2.3) for a formal proof. Step [1] of Algorithm 1 is obtained replacing the natural parameters with the quantities outlined in Equation (6), expanding expectations with the resulting Gaussian or Pòlya-Gamma or moments, and solving for $\boldsymbol{\lambda}_{i1}$ and $\boldsymbol{\lambda}_{i2}$.

Faster algorithms can be obtained replacing the gradient in Equation (12) with a computationally cheaper estimate (Robbins and Monro, 1951). In particular, we develop an efficient version of this algorithm based on an informative subset of nodes. We focus on the update of a generic factor $\mathbf{w}_i$, and denote with $\mathcal{J}_{i0} = \{j : y_{ij} = 0\}$ the set of indices associated with nodes not connected with node i. Similarly, we define $\mathcal{J}_{i1} = \{j : y_{ij} = 1\}$, with $|\mathcal{J}_{i0}| = n_{i0}$ and $|\mathcal{J}_{i1}| = n_{i1} = n - n_{i0}$; note that n_{i1} corresponds to the degree of node i. Adapting Durante and Rigon (2019), Equation (12) can be easily decomposed into the contribution of the n_{i1} nodes connected with i and the remaining as follows.

$$\sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \left(y_{ij} - 0.5 \right) + \sum_{j \in \mathcal{J}_{i0}} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \left(y_{ij} - 0.5 \right) + \mathbf{I}_H \mathbf{a}_0 - \boldsymbol{\lambda}_{i1},$$

$$-0.5 \left(\sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + \sum_{j \in \mathcal{J}_{i0}} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + \mathbf{I}_H \right) - \boldsymbol{\lambda}_{i2}. \tag{13}$$

Similarly, with the probit link

$$\sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + \sum_{j \in \mathcal{J}_{i0}} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + \mathbf{I}_H \mathbf{a}_0 - \boldsymbol{\lambda}_{i1},$$

$$-0.5 \left(\sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] + \sum_{j \in \mathcal{J}_{i0}} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] + \mathbf{I}_H \right) - \boldsymbol{\lambda}_{i2}. \tag{14}$$

Our strategy relies on noisy estimates of Equations (13) and (14), constructed using an informative subset of nodes $\mathcal{J}_i^\star = \mathcal{J}_{i1} \cup \mathcal{J}_{i0}^\star$, where $\mathcal{J}_{i0}^\star \subset \mathcal{J}_{i0}$ denotes a sample from $\mathcal{J}_{i0}$, with $|\mathcal{J}_0| = n_{i0}^\star \ll n_{i0}$. Therefore, the update for each factor $\mathbf{w}_i$ relies on the random subset $\mathcal{J}_i^\star$, which consists of all the nodes connected with node i and a smaller subset of not connected nodes. This approach implicitly assumes that nodes $j \in \mathcal{J}_{i0}$ not connected with i provide little information about its position in the latent space, and therefore we can estimate the contribution of all these disconnected nodes relying only on few units.

A simple strategy leading to a computationally cheap estimate of Equations (13) and (14) is to construct the set $\mathcal{J}_{i0}^\star$ relying on random sampling, where each unit $j \in \mathcal{J}_{i0}$ is included in the sample with probability n_{i0}^{-1}. This choice leads to the following unbiased estimator of Equation (13), denoted as $[B(\boldsymbol{\lambda}_{i1}), B(\boldsymbol{\lambda}_{i2})]$ and corresponding to the discrete random variable taking values

$$B_{\mathcal{J}_i^\star}(\boldsymbol{\lambda}_{i1}) = \sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \left(y_{ij} - 0.5 \right) + r_{i0} \sum_{j \in \mathcal{J}_{i0}^\star} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \left(y_{ij} - 0.5 \right) + \mathbf{I}_H \mathbf{a}_0 - \boldsymbol{\lambda}_{i1},$$

$$B_{\mathcal{J}_i^\star}(\boldsymbol{\lambda}_{i2}) = -0.5 \Big(\sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + \tag{15}$$

$$r_{i0} \sum_{j \in \mathcal{J}_{i0}^\star} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + \mathbf{I}_H \Big) - \boldsymbol{\lambda}_{i2},$$

for each possible $\mathcal{J}_{i0}^\star \subset \mathcal{J}_{i0}$, with $r_{i0} = n_{i0}/n_{i0}^\star$. Similarly, for Equation (14)

$$B_{\mathcal{J}_i^\star}(\boldsymbol{\lambda}_{i1}) = \sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + r_{i0} \sum_{j \in \mathcal{J}_{i0}^\star} \mathbb{E}_{q(\mathbf{w}_j)}[\mathbf{w}_j] \cdot \mathbb{E}_{q(z_{ij})}[z_{ij}] + \mathbf{I}_H \mathbf{a}_0 - \boldsymbol{\lambda}_{i1},$$

$$B_{\mathcal{J}_i^\star}(\boldsymbol{\lambda}_{i2}) = -0.5 \left(\sum_{j \in \mathcal{J}_{i1}} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] + r_{i0} \sum_{j \in \mathcal{J}_{i0}^\star} \mathbb{E}_{q(\mathbf{w}_j)} \left[\mathbf{w}_j \mathbf{w}_j^\mathsf{T} \right] + \mathbf{I}_H \right) - \boldsymbol{\lambda}_{i2}. \tag{16}$$

11

From a storage perspective, this representation allows to explicitly rely on the edge-list format of the network, where only the distinct pairs (i, j) associated with an edge are stored in memory. In its original formulation, SVI relies on sampling one single observation per iteration, and updating the gradient with the resulting estimate. Instead, for SVILF we recommend letting $n_{i0}^\star = \min(n_{i0}, \lfloor \gamma n_{i1} \rfloor)$, with $\gamma > 1$ and where $\lfloor \gamma n_{i1} \rfloor$ denotes the greatest integer less than or equal to γn_{i1}. In sparse network settings, $n_{i0}^\star \gg n_{i1}^\star$ and the random subset $\mathcal{J}_{i0}^\star$ will generally consist of $n_{i0}^\star = \lfloor \gamma n_{i1} \rfloor$ elements. Larger values of γ leads to estimators $[B(\boldsymbol{\lambda}_{i1}), B(\boldsymbol{\lambda}_{i2})]$ with smaller variance, while small values of γ lead to more efficient computation. In our experience, letting $\gamma \in [1, 5]$ provides good performance in a large number of applications; see sections 4 and 5.

The simple form of the proposed estimator allows the direct applications of the stochastic approximation method proposed in Robbins and Monro (1951) to solve Equation (12) via iterative updates as

$$
\begin{aligned}
\boldsymbol{\lambda}_{i1}^{(t)} &= (1 - \rho_t)\boldsymbol{\lambda}_{i1}^{(t-1)} + \rho_t B_{\mathcal{J}_t^\star}(\boldsymbol{\lambda}_{i1}^{(t-1)}), \\
\boldsymbol{\lambda}_{i2}^{(t)} &= (1 - \rho_t)\boldsymbol{\lambda}_{i2}^{(t-1)} + \rho_t B_{\mathcal{J}_t^\star}(\boldsymbol{\lambda}_{i2}^{(t-1)}),
\end{aligned}
\tag{17}
$$

where $[B_{\mathcal{J}_t^\star}(\boldsymbol{\lambda}_{i1}^{(t-1)}), B_{\mathcal{J}_t^\star}(\boldsymbol{\lambda}_{i2}^{(t-1)})]$ denotes a draw from the estimators outlined in Equation (15) and (16) evaluated at $(\boldsymbol{\lambda}_{i1}^{(t-1)}, \boldsymbol{\lambda}_{i2}^{(t-1)})$, and ρ_t denotes a sequence of step size such that $\sum_t \rho_t = +\infty$ and $\sum_t \rho_t^2 < +\infty$ (Robbins and Monro, 1951). A standard choice is to let $\rho_t = (t + \alpha)^{-\beta}$, with $\alpha > 0$ and $0.5 < \beta < 1$ (Hoffman et al., 2013; Durante and Rigon, 2019). Recalling Appendix A of Hoffman et al. (2013), in Equation (17) parameters $\boldsymbol{\lambda}_i$ are simultaneously updated to guarantee converge to the solution of (12). These optima are analytically obtained for each $q(z_{ij})$, and iteratively for $q(\mathbf{w}_i)$, conditioning on the values $\{\mathbf{w}_j\}_{j \neq i}$ at the previous iteration, with $i = 1, \ldots, n$.

As discussed in Raftery et al. (2012) within a different context, uniform subsampling of disconnected nodes might lead to a poor representation of the network structure, due to the heterogeneity in this sub-population. In our approach, uniform subsampling often leads to good performance; see Sections 4 and 5. As an alternative, we also propose an adaptive sampling mechanism, which relies on the currently estimated network structure to draw an informative sample of nodes. This

adaptive strategy samples, at iteration t, each node $j \in \mathcal{J}_{i0}^{\star}$ with a probability proportional to

$$g^{-1}\left(\left[\boldsymbol{\mu}_i^{(t-1)}\right]^{\mathsf{T}}\left[\boldsymbol{\mu}_j^{(t-1)}\right]\right), \quad \boldsymbol{\mu}_j^{(t-1)} := \mathbb{E}_{q^{(t-1)}(\mathbf{w}_j)}[\mathbf{w}_j], \quad j \in \mathcal{J}_{i0}^{\star}, \tag{18}$$

which corresponds to the current prediction for the probability to observe and edge between node i and node j, according to the underlying LFM. Therefore, we expect that nodes $j \in \mathcal{J}_{i0}^{\star}$ more similar to node i—according to the latent structure specification—are more likely to be sampled and contribute to the update of $q(\mathbf{w}_i; \boldsymbol{\lambda}_i)$ at iteration t. From a computational perspective, this strategy requires an additional loop over $j \in \mathcal{J}_{i0}^{\star}$ to compute all the products outlined in Equation (18). This operation might increase the computational time in high-dimensional settings, without significantly affecting the storage; see section 4 for an empirical evaluation. With the logit link, the proposed adaptive sampling can be used modifying the factor r_{i0} in Equation (13) into $m_{i0}/m_{i0}^{\star}$, with

$$m_{i0} := \sum_{j \in \mathcal{J}_{i0}} \left[1 + \exp\left(-\left[\boldsymbol{\mu}_i^{(t-1)}\right]^{\mathsf{T}}\left[\boldsymbol{\mu}_j^{(t-1)}\right]\right)\right]^{-1}, \quad m_{i0}^{\star} := \sum_{j \in \mathcal{J}_{i0}^{\star}} \left[1 + \exp\left(-\left[\boldsymbol{\mu}_i^{(t-1)}\right]^{\mathsf{T}}\left[\boldsymbol{\mu}_j^{(t-1)}\right]\right)\right]^{-1}. \tag{19}$$

Similarly, for the probit link function we set

$$m_{i0} := \sum_{j \in \mathcal{J}_{i0}} \Phi\left(\left[\boldsymbol{\mu}_i^{(t-1)}\right]^{\mathsf{T}}\left[\boldsymbol{\mu}_j^{(t-1)}\right]\right), \quad m_{i0}^{\star} := \sum_{j \in \mathcal{J}_{i0}^{\star}} \Phi\left(\left[\boldsymbol{\mu}_i^{(t-1)}\right]^{\mathsf{T}}\left[\boldsymbol{\mu}_j^{(t-1)}\right]\right), \tag{20}$$

where $\Phi(x)$ denotes the cumulative distribution function of a standard Gaussian evaluated at x. Pseudo code illustrating SVILF is reported in Algorithms 2 and 3, and an R package implementing the methods is available at `github.com/emanuelealiverti/svilf` and in the Supplementary Materials.

Algorithm 2: SVILF for LFM with logit link.

Initialize $\left\{ \boldsymbol{\lambda}_1^{(1)}, \ldots, \boldsymbol{\lambda}_n^{(1)} \right\}$ and set step size sequence ρ_t.

for $t = 2$ until convergence **do**

    Sample a permutation σ of $\{1, \ldots, n\}$ uniformly

    for $i = \sigma(1), \ldots, \sigma(n)$ **do**

        [1] Sampling Construct the random set $\mathcal{J}_i^\star = \mathcal{J}_{i1} \cup \mathcal{J}_{i0}^\star$ using uniform or adaptive subsampling

        [2] Local optimization Compute the locally optimized densities for $z_{ij}, j \in \mathcal{J}_i^\star$ leading to Pòlya-Gamma densities with natural parameters

$$\psi_{ij}\left(\boldsymbol{\lambda}_i^{(t-1)}, \boldsymbol{\lambda}_j^{(t-1)} \right) = -0.5 \left[\mathrm{vec}\left(\mathbf{S}_i^{(t-1)} \right)^\mathsf{T} \mathrm{vec}\left(\mathbf{S}_j^{(t-1)} \right) \right],$$

        with

$$\mathbf{S}_j^{(t-1)} = \left(-2\boldsymbol{\lambda}_{j2}^{(t-1)} \right)^{-1} + \left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)} \right)^{-1} \boldsymbol{\lambda}_{j1}^{(t-1)} \right]^\mathsf{T} \left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)} \right)^{-1} \boldsymbol{\lambda}_{j1}^{(t-1)} \right].$$

        [2] Global optimization. Update the global parameters leveraging stochastic optimization.

$$\boldsymbol{\lambda}_{i1}^{(t)} = (1 - \rho_t)\boldsymbol{\lambda}_{i1}^{(t-1)} + \rho_t \Bigg\{ \sum_{j \in \mathcal{J}_{i1}} \left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)} \right)^{-1} \boldsymbol{\lambda}_{j1}^{(t-1)} \right] (y_{ij} - 0.5) +$$

$$r_{i0} \sum_{j \in \mathcal{J}_{i0}^\star} \left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)} \right)^{-1} \boldsymbol{\lambda}_{j1}^{(t-1)} \right] (y_{ij} - 0.5) + \mathbf{I}_H \mathbf{a}_0 \Bigg\}$$

$$\boldsymbol{\lambda}_{i2}^{(t)} = (1 - \rho_t)\boldsymbol{\lambda}_{i2}^{(t-1)} - \rho_t 0.5 \left(\sum_{j \in \mathcal{J}_{i1}} \mathbf{S}_j^{(t-1)} \cdot \bar{z}_{ij}^{(t-1)} + r_{i0} \sum_{j \in \mathcal{J}_{i0}^\star} \mathbf{S}_j^{(t-1)} \cdot \bar{z}_{ij}^{(t-1)} + \mathbf{I}_H \right),$$

        with

$$\bar{z}_{ij}^{(t-1)} = \left[\psi_{ij}\left(\boldsymbol{\lambda}_i^{(t-1)}, \boldsymbol{\lambda}_j^{(t-1)} \right) \right]^{-1} \tanh\left[\psi_{ij}\left(\boldsymbol{\lambda}_i^{(t-1)}, \boldsymbol{\lambda}_j^{(t-1)} \right) \right]$$

        and with $r_{i0} = n_{i0}/n_{i0}^\star$ in case of uniform sampling and $r_i = m_{i0}/m_{i0}^\star$ for the adaptive version, as outlined in Equation (19). Therefore, the approximating density for $\mathbf{w}_i$ is Gaussian with mean $\boldsymbol{\mu}_i^{(t)}$ and covariance $\boldsymbol{\Sigma}_i^{(t)}$, where

$$\boldsymbol{\mu}_i^{(t)} = \left(-2\boldsymbol{\lambda}_{i2}^{(t)} \right)^{-1} \boldsymbol{\lambda}_{i1}^{(t)}, \quad \boldsymbol{\Sigma}_i^{(t)} = \left(-2\boldsymbol{\lambda}_{i2}^{(t)} \right)^{-1}.$$

Output $q^\star(\mathbf{W}) = \prod_{i=1}^n q^\star(\mathbf{w}_i)$.

Algorithm 3: SVILF for LFM with probit link.

Initialize $\left\{ \boldsymbol{\lambda}_1^{(1)}, \ldots, \boldsymbol{\lambda}_n^{(1)} \right\}$ and set step size sequence ρ_t.

for $t = 2$ until convergence **do**

    Sample a permutation σ of $\{1, \ldots, n\}$ uniformly

    for $i = 1, \ldots, n$ **do**

        [1] Sampling Construct the random set $\mathcal{J}_i^\star = \mathcal{J}_{i1} \cup \mathcal{J}_{i0}^\star$ using uniform or adaptive subsampling

        [2] Local optimization Compute the locally optimized densities for $z_{ij}, j \in \mathcal{J}_i^\star$ leading to

        Truncated Normal distributions with natural parameters

$$\psi_{ij}\left(\boldsymbol{\lambda}_i^{(t-1)}, \boldsymbol{\lambda}_j^{(t-1)}\right) = \left[\left(-2\boldsymbol{\lambda}_{i2}^{(t-1)}\right)^{-1}\boldsymbol{\lambda}_{i1}^{(t-1)}\right]^{\mathsf{T}}\left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)}\right)^{-1}\boldsymbol{\lambda}_{j1}^{(t-1)}\right]$$

        [2] Global optimization. Update the global parameters leveraging a stochastic optimization

$$\boldsymbol{\lambda}_{i1}^{(t)} = (1 - \rho_t)\boldsymbol{\lambda}_{i1}^{(t-1)} + \rho_t \left\{ \sum_{j \in \mathcal{J}_{i1}} \left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)}\right)^{-1}\boldsymbol{\lambda}_{j1}^{(t-1)}\right]\tilde{z}_{ij} + \right.$$

$$\left. r_{i0}\sum_{j \in \mathcal{J}_{i0}^\star} \left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)}\right)^{-1}\boldsymbol{\lambda}_{j1}^{(t-1)}\right]\tilde{z}_{ij} + \mathbf{I}_H \mathbf{a}_0 \right\}$$

$$\boldsymbol{\lambda}_{i2}^{(t)} = (1 - \rho_t)\boldsymbol{\lambda}_{i2}^{(t-1)} - \rho_t 0.5 \left(\sum_{j \in \mathcal{J}_{i1}} \mathbf{S}_j^{(t-1)} + r_{i0}\sum_{j \in \mathcal{J}_{i0}^\star} \mathbf{S}_j^{(t-1)} + \mathbf{I}_H \right),$$

      with

$$\mathbf{S}_j^{(t-1)} = \left(-2\boldsymbol{\lambda}_{j2}^{(t-1)}\right)^{-1} + \left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)}\right)^{-1}\boldsymbol{\lambda}_{j1}^{(t-1)}\right]^{\mathsf{T}}\left[\left(-2\boldsymbol{\lambda}_{j2}^{(t-1)}\right)^{-1}\boldsymbol{\lambda}_{j1}^{(t-1)}\right]$$

      and

$$\tilde{z}_{ij}^{(t-1)} = \psi_{ij}^{(t-1)}\left(\boldsymbol{\lambda}_i^{(t-1)}, \boldsymbol{\lambda}_j^{(t-1)}\right) +$$

$$(2y_{ij} - 1)\phi\left(\psi_{ij}^{(t-1)}\left(\boldsymbol{\lambda}_i^{(t-1)}, \boldsymbol{\lambda}_j^{(t-1)}\right)\right)\Phi\left[(2y_{ij} - 1)\psi_{ij}^{(t-1)}\left(\boldsymbol{\lambda}_i^{(t-1)}, \boldsymbol{\lambda}_j^{(t-1)}\right)\right]^{-1},$$

        where $\phi(x)$ and $\Phi(x)$ represent the density and the cumulative distribution function of a standard

        Gaussian evaluated in x, respectively, and with $r_{i0} = n_{i0}/n_{i0}^\star$ in case of uniform sampling and

        $r_i = m_{i0}/m_{i0}^\star$ for the adaptive version, as outlined in Equation (20).

    Therefore, the approximating density for $\mathbf{w}_i$ is Gaussian with mean $\boldsymbol{\mu}_i^{(t)}$ and covariance $\boldsymbol{\Sigma}_i^{(t)}$, where

$$\boldsymbol{\mu}_i^{(t)} = \left(-2\boldsymbol{\lambda}_{i2}^{(t)}\right)^{-1}\boldsymbol{\lambda}_{i1}^{(t)}, \quad \boldsymbol{\Sigma}_i^{(t)} = \left(-2\boldsymbol{\lambda}_{i2}^{(t)}\right)^{-1}.$$

Output $q^\star(\mathbf{W}) = \prod_{i=1}^n q^\star(\mathbf{w}_i)$.

---

4 SIMULATION STUDIES

Table 1: Simulation studies. Sampling distributions for the three considered scenarios.

	Model	Sampling distribution
s1	Latent factor model	$(y_{ij} \mid \pi_{ij}) \sim \mathrm{Ber}(\pi_{ij}), \quad i < j$ $\pi_{ij} = [1 + \exp(-\mathbf{w}_i^{\mathsf{T}}\mathbf{w}_j)]^{-1},$ $\mathbf{w}_i \sim \mathrm{N}_2(\mathbf{0}, 9\,\mathbf{I}_2), \quad i = 1, \dots, n$
s2	Latent distance model	$(y_{ij} \mid \pi_{ij}) \sim \mathrm{Ber}(\pi_{ij}), \quad i < j$ $\pi_{ij} = [1 + \exp(-||\mathbf{w}_i - \mathbf{w}_j||_2)]^{-1}, \mathbf{w}_i \in \mathbb{R}^2$ $\mathbf{w}_i \sim \mathrm{N}_2(\mathbf{0}, \mathbf{I}_2), \quad i = 1, \dots, n$
s3	Stochastic block model	$(y_{ij} \mid \pi_{ij}) \sim \mathrm{Ber}(\pi_{ij}), \quad i < j$ $\pi_{ij} = 0.6\,\mathrm{I}\{w_i = w_j\} + 0.2\,\mathrm{I}\{w_i \neq w_j\}$ $w_i \sim \mathrm{Ber}(0.5), \quad i = 1, \dots, n$

We conduct a simulation study to evaluate the performances of the proposed algorithms. To compare our methods with a similar alternative—in terms of variational approximation under the same statistical model—we used the Automated Differentiation Variational Inference algorithm (ADVI, Kucukelbir et al., 2015) to approximate the posterior distribution of the LFM specified in Equation 1. This algorithm is available in the software STAN (Stan Development Team, 2019), and it relies on a fully-factorized Gaussian approximation whose parameters are obtained via stochastic optimization, leveraging an adaptive step-size sequence; see Kucukelbir et al. (2015) for more details. We also compare SVILF with a Latent distance model (Hoff et al., 2002) estimated via Variational-EM (VBLPCM, Salter-Townshend and Murphy, 2013).

In this section, we focus on the logit link for the LFM; results for the probit link are reported in the Supplementary Materials. Computational performances are evaluated in terms of memory usage, elapsed time and goodness of fit. Memory is measured in MB of used RAM, while elapsed time is evaluated in seconds to run the VB routine until convergence. For SVILF and VBLPCM, we use as convergence criterion the mean squared difference between consecutive parameter values below 10^{-5}; for ADVI, we follow the default implementation, relying on a median ELBO difference below 10^{-2}. Predictions are evaluated in terms of adequacy in recovering the probability of observing an edge, according to the Area Under the Roc curve (AUC).

The simulations focus on three different data generating processes, summarized in Table 1.

Table 2: Simulation studies. Average RAM usage (in MB) across 100 simulation replicates, measured with the R function `gc`. Standard deviations are not reported since memory allocation is essentially constant across simulation replicates.

	n	100	200	300	500	1000	2000	3000
S1	SVILF	101	101	101	101	103	108	116
	SVILF (ADA)	101	101	101	101	103	108	116
	ADVI	112	115	119	128	156	227	308
	VBLPCM	117	122	131	158	285	792	1635
S2	SVILF	101	101	101	101	103	109	119
	SVILF (ADA)	101	101	101	101	103	109	119
	ADVI	112	115	119	128	156	226	308
	VBLPCM	117	123	133	164	310	891	1861
S3	SVILF	101	101	101	102	104	113	129
	SVILF (ADA)	101	101	101	102	104	113	129
	ADVI	112	115	119	128	156	224	308
	VBLPCM	118	127	141	186	395	1234	2631

Networks in the first scenario (S1) are generated according to a LFM model, with 2 latent factors randomly generated from multivariate Gaussian with diagonal covariance and standard deviation equal to 3. In the second simulation scenario (S2), data are generated from a latent distance model (Hoff et al., 2002) having 2-dimensional latent positions, generated from standard Gaussian distributions, while in the third and last scenario (S3), data are generated from a stochastic block-model with 2 latent groups, equal weights, within-group probability of connection equal to 0.6 and between-group probability equal to 0.2.

We focus on networks with a number of nodes $n \in \{100, 200, 300, 500, 1000, 2000, 3000\}$, and replicate the data generation process and posterior estimation with 100 different random seeds. We fix the number of latent factors and latent coordinates to $H = 4$, set $\mathbf{a}_0 = \text{logit}(2 \sum_{i<j} y_{ij}/n(n-1)) \cdot \mathbf{I}_H$ to account for the overall network sparsity, and set the SVILF step-size parameters $\alpha = 1, \beta = 0.75, \gamma = 2$. In the Supplementary Materials, we extend the simulations to different values of latent factors H and of the parameter γ. Posterior inference for ADVI and VBLPCM is performed with default parameters configuration. The ADVI algorithm often failed to converge with random initialization of the parameters. Instead, convergence criteria are met when the parameters are initialized from the eigenvalues of the adjacency matrix. For the SVILF algorithms parameters are initialized randomly, and in most settings convergence is reached in less than 50 iterations. We

Table 3: Simulation studies. Average elapsed time (standard deviation) in seconds across 100 simulation replicates.

	n	100	200	300	500	1000	2000	3000
S1	SVILF	1.2 (0.1)	2.0 (0.3)	4.2 (0.8)	14.8 (1.7)	112 (12)	893 (78)	3362 (337)
	SVILF (ADA)	2.5 (0.2)	7.6 (0.9)	15.2 (1.6)	37.6 (4.6)	188 (22)	1212 (126)	4270 (542)
	ADVI	7.4 (1.2)	27.8 (2.5)	62.4 (7.0)	176.3 (14.7)	711 (49)	2815 (229)	6543 (451)
	VBLPCM	5.4 (0.5)	19.5 (1.4)	48.5 (8.8)	159.9 (34.7)	738 (207)	3761 (901)	10247 (2334)
S2	SVILF	1.2 (0.1)	1.8 (0.2)	2.9 (0.4)	7.2 (1.1)	37 (4)	256 (42)	1017 (206)
	SVILF (ADA)	2.5 (0.2)	7.7 (0.5)	15.9 (1.1)	47.4 (3.7)	244 (19)	1445 (137)	4698 (629)
	ADVI	6.8 (1.6)	25.1 (2.8)	58.8 (9.7)	165.8 (16.2)	658 (46)	2578 (177)	6043 (427)
	VBLPCM	5.0 (0.5)	18.9 (1.5)	47.4 (9.0)	105.1 (11.3)	744 (210)	3977 (1013)	11711 (2622)
S3	SVILF	1.2 (0.1)	1.9 (0.3)	3.3 (0.3)	9.0 (1.0)	48 (4)	332 (33)	1306 (90)
	SVILF (ADA)	2.9 (0.2)	9.1 (0.6)	19.8 (1.2)	56.1 (3.6)	276 (15)	1622 (167)	5785 (597)
	ADVI	6.7 (0.6)	25.3 (2.2)	57.1 (5.3)	163.2 (19.1)	647 (40)	2542 (154)	5926 (375)
	VBLPCM	4.8 (0.6)	18.5 (1.5)	48.3 (9.3)	163.1 (37.8)	960 (174)	4586 (935)	14848 (2878)

found no difference when initialing the parameters from the eigenvalues of the adjacency matrix, and since eigen-decomposition can be demanding in high-dimensional settings, we recommend random initialization as a default option for the SVILF. For VBLPCM we used the build-in function `vblpcmstart` that generates sensible parameters' initialization.

Table 2 reports the average memory usage for posterior inference in all the considered simulation scenarios, across 100 simulation replicates. When the number of nodes is in the order of few hundred, the considered algorithms require a similar amounts of RAM. In contrast, when the number of nodes is in the order of few thousand, the proposed SVILF algorithms requires significantly less memory in all the scenarios. For example, in the case $n = 3000$, ADVI and VBLPCM require approximately 3 and 20 times as much memory as SVILF; this factor is increasing with n, showing the efficiency of the SVILF algorithms for high-dimensional networks. As expected, the two sampling schemes implemented under the SVILF algorithm perform identically in terms of memory usage, since no additional storage capacity is needed.

Table 3 compares the elapsed time required by the different algorithms to reach convergence. Results suggest that SVILF with uniform subsampling provides the fastest routine in all the considered settings, with an elapsed 2 to 15 times faster than ADVI and VBLPCM. For example, with $n = 2000$ and data generated under a latent-distance model (S2), SVILF is 13 times faster than the competitors, requiring in 4 minutes on average versus 42 of ADVI and 66 of VBLPCM. As expected,

Table 4: Simulation studies. Average (standard deviation) AUC in percentage values across 100 simulation replicates.

	n	100	200	300	500	1000	2000	3000
S1	SVILF	86 (1.6)	86 (1.3)	86 (0.9)	85 (0.7)	85 (0.5)	85 (0.4)	85 (0.4)
	SVILF (ADA)	86 (1.3)	85 (1.2)	85 (1.0)	85 (0.9)	84 (0.6)	84 (0.4)	84 (0.4)
	ADVI	85 (2.1)	86 (1.2)	86 (0.9)	86 (0.7)	86 (0.5)	86 (0.4)	86 (0.3)
	VBLPCM	87 (1.3)	88 (1.0)	87 (0.9)	87 (0.9)	84 (3.3)	82 (4.1)	76 (7.2)
S2	SVILF	72 (0.9)	68 (0.7)	66 (0.5)	64 (0.4)	63 (0.4)	62 (0.3)	62 (0.3)
	SVILF (ADA)	73 (0.7)	69 (0.5)	67 (0.5)	66 (0.4)	65 (0.4)	65 (0.3)	65 (0.2)
	ADVI	66 (2.8)	65 (1.1)	66 (1.0)	66 (0.6)	64 (0.5)	63 (0.3)	62 (0.2)
	VBLPCM	78 (0.8)	73 (0.6)	71 (0.5)	69 (0.5)	67 (0.6)	67 (0.4)	61 (2.0)
S3	SVILF	80 (0.6)	77 (0.4)	76 (0.2)	75 (0.2)	74 (0.1)	73 (0.1)	72 (0.1)
	SVILF (ADA)	80 (0.6)	77 (0.3)	76 (0.2)	75 (0.1)	74 (0.1)	73 (0.1)	72 (0.1)
	ADVI	75 (1.5)	74 (0.4)	73 (0.3)	73 (0.2)	72 (0.2)	72 (0.1)	72 (0.1)
	VBLPCM	80 (0.6)	76 (0.3)	75 (0.2)	73 (0.2)	72 (0.1)	71 (0.2)	64 (6.7)

the adaptive sampling scheme affects the execution time of the SVILF algorithm, because of the additional loop of order $\mathcal{O}(n_{i0})$ to compute the sampling weights at each iteration. However, even if this adaptive strategy is more expensive than the standard SVILF, it still provides significant computational advantages over ADVI and VBLPCM, reducing the elapsed time and memory usage.

In order to evaluate the quality of predictions, Table 4 compares the goodness of fit of the posterior predicted probabilities in terms of AUC, reported in percentage values. The considered approaches share similar performances in most simulation scenarios. In the second Scenario, VBLPCM is the correctly specified model and it performs better than other methods on average. However, the performance gap is decreasing with the size of n, and when $n = 3000$ the AUC are not significantly different. Therefore, the main advantage of SVILF is to provide massive savings in terms of computational time and memory usage, achieving accurate performance in reasonable time also in high-dimensional settings. In the Supplementary Materials, we include additional simulations to illustrate the performance of SVILF with different specification of the number of latent factors H and the parameter γ, suggesting that the performance can be improved increasing H without affecting the computational performance.

5 APPLICATION

We analyze three high-dimensional datasets provided by Rozemberczki et al. (2021) and freely available on the repository `http://snap.stanford.edu/data/index.html`. The main aim of this section is to show how SVILF algorithms provide accurate representations of high-dimensional network data in terms of link prediction.

The first dataset involves mutual interconnections among Facebook public pages related to politicians, governmental organizations, television shows and companies; the resulting network consists of $n = 22700$ nodes corresponding to verified Facebook pages, while edges represent the mutual "likes" between these pages. The second dataset includes the mutual following relationships among $n = 37700$ users of the popular software developing platform GitHub. The third dataset includes $n = 11631$ articles on crocodiles from the English version of Wikipedia; an edge indicates mutual links between articles. In all the examples, we estimate an LFM using SVILF algorithms with the same settings as in the simulations studies, increasing $\gamma = 3$ to handle the massive sparsity of the networks (the average density is 0.001). We did not compare SVILF with the competitors as the number of nodes is excessively large.

Table 5 and Figure 1 illustrate, respectively, the AUC and the ROC curves for the considered approaches. Results suggest a satisfactory performance for the proposed methods, with an AUC above 0.75 in most settings. Coherently with the simulations, our empirical findings do not indicate that the adaptive subsampling is systematically preferable over uniform sampling. Results from Table 5 suggest that the best-performing approach is different across applications; for example, uniform subsampling achieves lower AUC than the adaptive approach in the Facebook example,

Table 5: Application. AUC for the proposed methods.

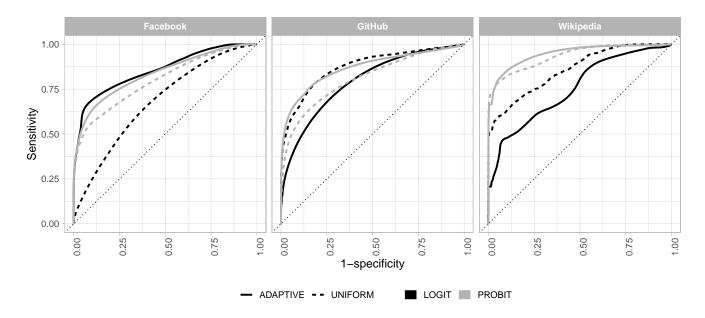|		LOGIT	PROBIT
Facebook	SVILF (ADA)	0.855	0.835
	SVILF	0.678	0.798
GitHub	SVILF (ADA)	0.791	0.867
	SVILF	0.874	0.809
Wikipedia	SVILF (ADA)	0.757	0.950
	SVILF	0.867	0.940

Figure 1: Application. Smoothed ROC curves for the proposed approaches over the three examples.

and higher in the GitHub dataset. Such differences might be related to the structure of network, and it will be of interest in future developments to investigate under which network structure one particular sampling approach is preferable.

6 DISCUSSION

Motivated by the abundance of large network-valued data, we have proposed a novel computational approach for approximate posterior inference under the latent factor model for networks. Specifically, we developed a stochastic Variational Bayes routine that explicitly leverages the sparsity of high-dimensional networks to perform efficient computation. The empirical evaluations suggest that the proposed algorithms lead to significant benefits in terms of computational efficiency, without affecting accuracy in modeling the network connectivity structure.

Although the LFM is routinely used to model undirected binary networks, several applications involve more complicated structures. Some examples includes multiple networks, presence of weighted edges or the desire to include additional covariates into the analysis. The SVILF algorithms for the LFM can be directly extended to such settings considering different conditional distributions for the elements y_{ij} (e.g., Poisson or Gaussian), and allowing the probability of observing an edge to change according to some edge-specific covariates. Additionally, the subsampling

strategy proposed in Section 3 can be potentially adapted to approximate the gradients in algorithms for latent distance models (e.g., Salter-Townshend and Murphy, 2013; Gollini and Murphy, 2016). However, the implementation requires ad-hoc steps, since these models do not fall under the same conditionally conjugate structure used by SVILF.

7 ACKNOWLEDGEMENTS

REFERENCES

Agresti, A. (2015). *Foundations of linear and generalized linear models*. John Wiley & Sons.

Albert, J. H. and Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–679.

Aliverti, E. and Durante, D. (2019). Spatial modeling of brain connectivity data via latent distance models with nodes clustering. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 12(3):185–196.

Athreya, A., Fishkind, D. E., Tang, M., Priebe, C. E., Park, Y., Vogelstein, J. T., Levin, K., Lyzinski, V., Qin, Y., and Sussman, D. L. (2018). Statistical inference on random dot product graphs: a survey. *Journal of Machine Learning Research*, 18(226):1–92.

Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.

Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877.

Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198.

Consonni, G. and Marin, J.-M. (2007). Mean-field variational approximate Bayesian inference for latent variable models. *Computational Statistics & Data Analysis*, 52(2):790–798.

Csardi, G. and Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal*, Complex Systems:1695.

Durante, D. and Dunson, D. B. (2014). Nonparametric Bayes dynamic modelling of relational data. *Biometrika*, 101(4):883–898.

Durante, D., Dunson, D. B., and Vogelstein, J. T. (2017). Nonparametric Bayes modeling of populations of networks. *Journal of the American Statistical Association*, 112(520):1516–1530.

Durante, D. and Rigon, T. (2019). Conditionally conjugate mean-field variational Bayes for logistic models. *Statistical Science*, 34(3):472–485.

D'Angelo, S., Murphy, T. B., Alfò, M., et al. (2019). Latent space modelling of multidimensional networks with application to the exchange of votes in eurovision song contest. *The Annals of Applied Statistics*, 13(2):900–930.

Gelfand, A. E. and Smith, A. F. (1990). Sampling-based approaches to calculating marginal densities. *Journal of The American Statistical Association*, 85(410):398–409.

Gollini, I. and Murphy, T. B. (2016). Joint modeling of multiple network views. *Journal of Computational and Graphical Statistics*, 25(1):246–265.

Ho, Q., Yin, J., and Xing, E. P. (2016). Latent space inference of internet-scale networks. *The Journal of Machine Learning Research*, 17(1):2756–2796.

Hoff, P. D. (2005). Bilinear mixed-effects models for dyadic data. *Journal of the American Statistical Association*, 100(469):286–295.

Hoff, P. D. (2021). Additive and multiplicative effects network models. *Statistical Science*, 36(1):34–50.

Hoff, P. D., Raftery, A. E., and Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098.

Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. (2013). Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347.

Keeling, M. J. and Eames, K. T. (2005). Networks and epidemic models. *Journal of the Royal Society Interface*, 2(4):295–307.

Kucukelbir, A., Ranganath, R., Gelman, A., and Blei, D. (2015). Automatic variational inference in stan. In *Advances in neural information processing systems*, pages 568–576.

Newman, M. (2018). *Networks*. Oxford University Press.

Polson, N. G., Scott, J. G., and Windle, J. (2013). Bayesian inference for logistic models using Pólya–Gamma latent variables. *Journal of the American Statistical Association*, 108(504):1339–1349.

Raftery, A. E., Niu, X., Hoff, P. D., and Yeung, K. Y. (2012). Fast inference for the latent space network model using a case-control approximate likelihood. *Journal of Computational and Graphical Statistics*, 21(4):901–919.

Robbins, H. and Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407.

Rozemberczki, B., Allen, C., and Sarkar, R. (2021). Multi-scale attributed node embedding. *Journal of Complex Networks*, 9.

Salter-Townshend, M. and Murphy, T. B. (2013). Variational Bayesian inference for the latent position cluster model. *Computational Statistics and Data Analysis*, 57(1):661–671.

Sewell, D. K. and Chen, Y. (2015). Latent space models for dynamic networks. *Journal of the American Statistical Association*, 110(512):1646–1657.

Stan Development Team (2019). RStan: the R interface to Stan. R package version 2.19.2.

Tang, M., Sussman, D. L., Priebe, C. E., et al. (2013). Universally consistent vertex classification for latent positions graphs. *Annals of Statistics*, 41(3):1406–1430.

Wu, J., Vallenius, T., Ovaska, K., Westermarck, J., Mäkelä, T. P., and Hautaniemi, S. (2009). Integrated network analysis platform for protein-protein interactions. *Nature Methods*, 6(1):75–77.

Young, S. J. and Scheinerman, E. R. (2007). Random dot product graph models for social networks. In *International Workshop on Algorithms and Models for the Web-Graph*, pages 138–149. Springer.

Stratified stochastic variational inference for high-dimensional network factor model

Supplementary materials

Emanuele Aliverti and Massimiliano Russo

S1 Simulations with different latent dimension H and parameter γ

We evaluate the performance of the SVILF algorithm described in Section 4 of Aliverti and Russo (2022) with all the combination $(H, \gamma) \in \{2, \ldots, 10\} \times \{1, \ldots 5\}$. The simulations focus on the same data generating processes outlined in Section 4 and Table 1 of the manuscript.

Figure S1 reports the elapsed time for SVILF. The dimension of the latent space and the magnitude of γ do not affect the performance of the algorithm in terms of timing. For a fixed n the elapsed time is essentially constant for all the considered (H, γ) values. This is expected since, as outlined in Section 4 of the manuscript, the SVILF computational cost is dominated by n. Figure S2 shows the memory usage, suggesting that also this quantity is constant for different values of (H, γ). The highest inflation in terms of memory usage is observed for $n = 3000$, where the used RAM ranges from 128 MB with $H = 2$ to 132 MB with $H = 10$, irrespectively of the specific value of γ. Lastly, Figure S3 depicts the values for the AUC. In Scenario 1, the LFM is the correctly specified model, and we do not observe performance improvements for $H \geq 3$. In Scenarios 2 and 3 we observe a moderate increasing trend in H, with performance that improves for larger number of latent factors. Results are consistent for different values of γ, suggesting that this parameter does not influence much model performance when selected in a reasonable range.
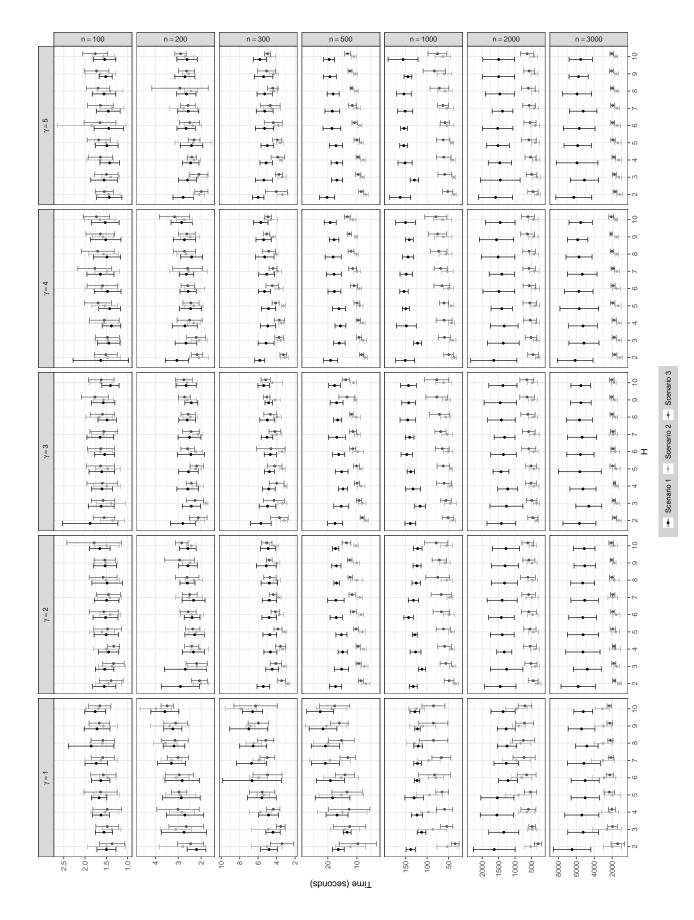
1

Figure S1: Average elapsed time (in seconds) for SVILF across 10 replicates of the simulations described in Section 4 with $(H, \gamma) \in \{2, \ldots, 10\} \times \{1, \ldots, 5\}$. The error-bars represent the average elapsed time plus/minus two standard deviations.
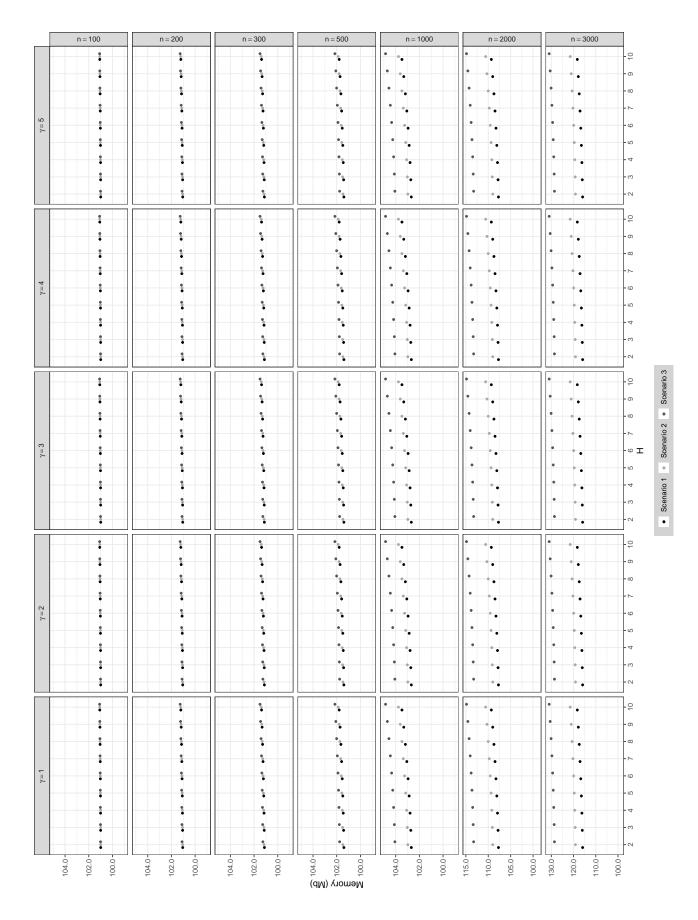
Figure S2: Average RAM usage (in MB) for SVILF across 10 replicates of the simulations described in Section 4 with $(H, \gamma) \in \{2, \ldots, 10\} \times \{1, \ldots, 5\}$. Memory usage is measured with the R function gc. Variability is not reported since memory allocation is essentially constant across simulation replicates.
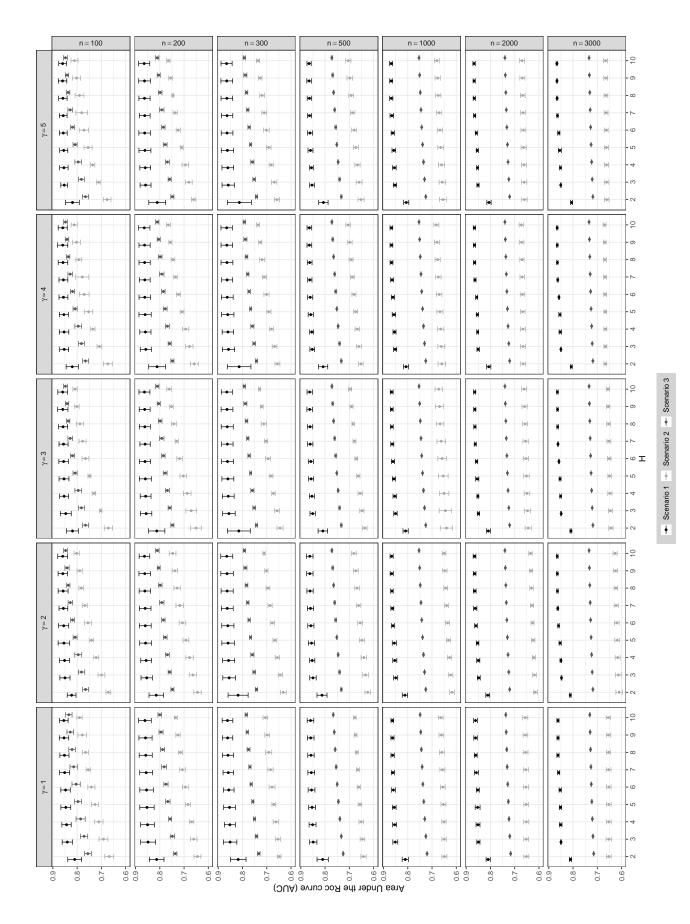
Figure S3: Average AUC for SVILF across 10 replicates of the simulations described in Section 4 with $(H, \gamma) \in \{2, \ldots, 10\} \times \{1, \ldots 5\}$. The error-bars represent the average AUC plus/minus two standard deviations.

4

Table S1: Simulation studies. Average RAM usage (in MB) measured with the R function `gc` across 100 replicates of the simulation study. Standard deviations are not reported since memory allocation is essentially constant across simulation replicates.

	n	100	200	300	500	1000	2000	3000
s1	SVILF	101	101	101	101	103	108	116
	SVILF (ADA)	101	101	101	101	102	108	116
s2	SVILF	101	101	101	102	103	109	119
	SVILF (ADA)	101	101	101	102	103	109	119
s3	SVILF	101	101	101	102	104	113	129
	SVILF (ADA)	101	101	101	102	104	113	129

Table S2: Simulation studies. Average elapsed time (standard deviation) in seconds across 100 replicates of the simulation study.

	n	100	200	300	500	1000	2000	3000
s1	SVILF	1.8 (0.5)	2.1 (0.6)	2.8 (0.8)	5.8 (1.3)	36.1 (6.3)	397 (71)	1460 (188)
	SVILF (ADA)	1.9 (0.4)	3.3 (0.9)	5.7 (1.1)	17.1 (2.3)	85.3 (9.0)	575 (85)	1857 (257)
s2	SVILF	1.7 (0.4)	2.1 (0.50)	2.8 (0.7)	6.1 (1.2)	31.6 (5.1)	321 (46)	1203 (137)
	SVILF (ADA)	2.0 (0.5)	3.8 (0.58)	6.7 (1.2)	17.0 (1.60)	85.2 (8.4)	619 (87)	1862 (228)
s3	SVILF	1.7 (0.5)	2.4 (0.6)	3.9 (1.0)	9.2 (1.3)	55 (9.2)	554 (64)	1829 (155)
	SVILF (ADA)	2.2 (0.5)	4.7 (0.7)	9.6 (0.9)	26.2 (2.3)	152.4 (17.3)	1326 (178)	4270 (430)

Table S3: Simulation studies. Average (standard deviation) AUC in percentage values across 100 replicates of the simulation study.

	n	100	200	300	500	1000	2000	3000
s1	SVILF	85 (1.7)	85 (1.2)	85 (0.9)	85 (0.7)	85 (0.5)	85 (0.3)	85 (0.3)
	SVILF (ADA)	86 (1.6)	86 (1.2)	86 (0.9)	86 (0.7)	86 (0.5)	86 (0.3)	86 (0.3)
s2	SVILF	69 (1.2)	66 (0.8)	64 (0.6)	63 (0.4)	62 (0.5)	61 (0.4)	61 (0.3)
	SVILF (ADA)	72 (1.2)	69 (0.6)	67 (0.6)	66 (0.4)	65 (0.4)	65 (0.3)	65 (0.3)
s3	SVILF	78 (0.8)	76 (0.4)	75 (0.3)	74 (0.2)	73 (0.1)	72 (0.1)	72 (0.1)
	SVILF (ADA)	79 (0.7)	77 (0.4)	76 (0.2)	75 (0.1)	74 (0.1)	73 (0.1)	72 (0.1)

S2 Simulation studies with probit link

Results from the Section 4 of Aliverti and Russo (2022) are reproduced relying on a LFM with probit link function, implemented within the R package `svilf` available at `https://github.com/emanuelealiverti/svilf`. Tables S1-S3 illustrate the average memory, elapsed time and AUC, confirming a modest use of computational resources also with large networks, as well as satisfactory performance in terms of AUC (Table S3). The results in Tables S1 to S3

for the probit link function are comparable with the logit link ones reported in Section 4 of the paper (Tables 2 to 4)

S3   CAVI algorithm for lfm with probit link

Algorithm S1 illustrates the pseudo-code for the CAVI algorithm under a LFM with probit link function. Leveraging the conditionally-conjugate representation outlined in the article, the implementation is similar with Algorithm 1 of Aliverti and Russo (2022).

Algorithm S1: CAVI for LFM with probit link.

Initialize $\left\{\mathbf{S}_1^{(1)}, \ldots, \mathbf{S}_n^{(1)}\right\}$ and $\left\{\boldsymbol{\mu}_1^{(1)}, \ldots, \boldsymbol{\mu}_n^{(1)}\right\}$.

for $t = 2$ until convergence **do**

> [1] $q^{(t)}(\mathbf{w}_i)$ is a density of a $\mathrm{N}_H(\boldsymbol{\mu}_i^{(t)}, \boldsymbol{\Sigma}_i^{(t)})$ with $\boldsymbol{\mu}_i^{(t)} = [-2\boldsymbol{\lambda}_{i2}^{(t)}]^{-1}\boldsymbol{\lambda}_{i1}^{(t)}$,
> $\boldsymbol{\Sigma}_i^{(t)} = [-2\boldsymbol{\lambda}_{i2}^{(t)}]^{-1}$ and
>
> $$\boldsymbol{\lambda}_{i1}^{(t)} = \sum_{j \neq i} \boldsymbol{\mu}_j^{(t-1)} \tilde{z}_{ij}^{(t-1)} + \mathbf{I}_H \mathbf{a}_0, \quad \boldsymbol{\lambda}_{i2}^{(t)} = -0.5 \left(\sum_{j \neq i} \left[\boldsymbol{\Sigma}_j^{(t-1)} + \left(\boldsymbol{\mu}_j^{(t-1)} \right) \left(\boldsymbol{\mu}_j^{(t-1)} \right)^{\mathsf{T}} \right] + \mathbf{I}_H \right),$$
>
> for $i = 1, \ldots, n$. In the above expression,
>
> $$\tilde{z}_{ij}^{(t-1)} = \gamma_{ij}^{(t-1)} + (2y_{ij} - 1)\phi(\gamma_{ij}^{(t-1)})\Phi[(2y_{ij} - 1)\gamma_{ij}^{(t-1)}]^{-1}, \quad \gamma_{ij}^{(t-1)} = \left(\boldsymbol{\mu}_i^{(t-1)} \right)^{\mathsf{T}} \left(\boldsymbol{\mu}_j^{(t-1)} \right)$$
>
> and $\phi(x)$ and $\Phi(x)$ represent the density and the cumulative distribution function of a standard Gaussian evaluated in x.
>
> [2] $q^{(t)}(z_{ij})$, is the density of a $\mathrm{TN}[\gamma_{ij}^{(t)}, 1, (-\infty, 0)]$ if $y_{ij} = 0$ and $\mathrm{TN}[\gamma_{ij}^{(t)}, 1, (0, +\infty)]$ if $y_{ij} = 1$, for $i = 2, \ldots, n$ and $j = 1, \ldots, i - 1$.

Output $q^\star(\mathbf{W}, \mathbf{z}) = \prod_{i=1}^n q^\star(\mathbf{w}_i) \prod_{i=2}^n \prod_{j=1}^{n-1} q^\star(z_{ij})$.

References

Aliverti, E. and Russo, M. (2022). Stratified stochastic variational inference for high-dimensional network factor model. *Journal of Computational and Graphical Statistics*, 31(2):502–511.