



RESEARCH ARTICLE

Faster MCMC for Gaussian latent position network models

Neil A. Spencer^{1,*}, Brian W. Junker² and Tracy M. Sweet³

¹Harvard University, Boston, MA 02115, USA ²Carnegie Mellon University, Pittsburgh, PA 15213, USA and ³University of Maryland College Park, College Park, MD 20742, USA

*Corresponding author. Email: nspencer@hpsch.harvard.edu

Action Editor: Stanley Wasserman

Abstract

Latent position network models are a versatile tool in network science; applications include clustering entities, controlling for causal confounders, and defining priors over unobserved graphs. Estimating each node's latent position is typically framed as a Bayesian inference problem, with Metropolis within Gibbs being the most popular tool for approximating the posterior distribution. However, it is well-known that Metropolis within Gibbs is inefficient for large networks; the acceptance ratios are expensive to compute, and the resultant posterior draws are highly correlated. In this article, we propose an alternative Markov chain Monte Carlo strategy—defined using a combination of split Hamiltonian Monte Carlo and Firefly Monte Carlo—that leverages the posterior distribution's functional form for more efficient posterior computation. We demonstrate that these strategies outperform Metropolis within Gibbs and other algorithms on synthetic networks, as well as on real information-sharing networks of teachers and staff in a school district.

Keywords: Hamiltonian Monte Carlo; network data; firefly Monte Carlo; latent space model; longitudinal network data; Bayesian computation

1. Introduction

Network data—measurements of relationships across sets of entities—are becoming increasingly common across science and industry, largely due to technological advances in data collection and storage. Common sources of network data include social networks (Carrington et al., 2005), citation networks (Ji & Jin, 2016), gene regulatory networks (Hecker et al., 2009), disease transmission networks (Newman, 2002), neural connectomes (Chen et al., 2016), transportation networks (Xie & Levinson, 2009), and food webs (Chiu & Westveld, 2011). A broad range of statistical tools based on stochastic graphs (Goldenberg et al., 2010; Crane, 2018) are available for probabilistically modeling networks, ranging from the simple Erdős-Rényi model (Erdős & Rényi, 1960) to sophisticated latent variable models (Airolidi et al., 2008; Clauset et al., 2008; Fosdick et al., 2018; Dabbs et al., 2020). Latent variable models can be defined to capture common network properties such as community structure, hierarchical structure, and degree heterogeneity.

Evaluating the likelihood of a latent variable model has a computational complexity that is quadratic in the number of nodes. These models are thus costly to fit to large networks, especially if one wishes to quantify uncertainty in a Bayesian modeling and inference framework (Gelman et al., 2013). For instance, traditional Markov chain Monte Carlo algorithms (Gamerman & Lopes, 2006) such as Gibbs sampling or random walk Metropolis can require tens of thousands of likelihood evaluations to accurately quantify expectations and uncertainties. This computational

burden is even larger when the chains are slow-mixing, which is often the case for Bayesian hierarchical models.

In this work, we develop a faster Markov chain Monte Carlo algorithm for a class of latent variable network models called the latent position network model (LPM). LPMs—originally proposed by Hoff et al. (2002)—have been applied to a variety of statistical problems, including modeling network interventions (Sweet et al., 2013), clustering entities (Handcock et al., 2007), modeling social influence (Sweet & Adhikari, 2020), controlling for causal confounders (McFowland III & Shalizi, 2021), and defining priors on unobserved graphs (Linderman et al., 2016). Each node in a LPM possesses a real-valued latent variable (its *position*), with each edge treated as an independent Bernoulli random draw depending on the participating nodes' latent positions. These probabilities are modeled as a decreasing function of the nodes' latent distance, thus promoting homophily and triadic closure (i.e. a friend of a friend is more likely to be a friend). Edge probabilities may also depend on covariates, such as whether the entities share a commonly observed trait.

The principal task in fitting a LPM is to infer the latent positions, as well as the parameters of the link function. In a Bayesian modeling and inference framework, the posterior distribution of these parameters quantifies uncertainty in the corresponding estimates. Evaluating and summarizing this posterior distribution requires intensive computation, namely because of an intractable normalization constant.

The standard tool for computing posterior summaries has been Markov chain Monte Carlo (MCMC) via Metropolis within Gibbs (Handcock et al., 2007; Raftery et al., 2012). This technique side-steps explicit computation of the normalization constants and can approximate posterior expectations arbitrarily well if run long enough. However, accurate inference via Metropolis within Gibbs can be computationally infeasible for large networks, largely due to two problems: (1) The random walk step size required to obtain high acceptance rates shrinks as the number of nodes grows, resulting in slowly mixing chains for large networks, and (2) the computational complexity of performing a full sweep of position updates is quadratic in the number of nodes, so each iteration for a large network is expensive to compute. We address these challenges in this article through the development of a more efficient MCMC algorithm.

We are not the first to recognize these limitations of Metropolis within Gibbs for LPMs. In recent years, multiple approaches for approximating the likelihood have been proposed to scale up Bayesian inference of LPMs to large networks. Raftery et al. (2012) proposed a case-control-based approach, subsampling the non-edge dyads to approximate each acceptance ratio in Metropolis within Gibbs. Rastelli et al. (2018) proposed a discrete-grid approximation of the latent positions, simplifying each likelihood evaluation. Salter-Townshend & Murphy (2013) proposed variational inference as an alternative to MCMC. Though each of these approaches speeds up posterior inference, the improvements come at the cost of biasing the results with the likelihood approximations.

Our approach is instead based on Hamiltonian Monte Carlo (HMC). HMC (Duane et al., 1987; Neal, 2011; Betancourt, 2017) and its variants (Girolami & Calderhead, 2011; Hoffman & Gelman, 2014; Betancourt, 2016) are a class of MCMC algorithms that leverage Hamiltonian dynamics to construct gradient-informed proposals for differentiable posterior distributions. Well-tuned HMC proposals produce large moves while maintaining high Metropolis-Hastings acceptance rates. HMC can thus be much more efficient than traditional random walk-based methods, especially in high dimensions, without introducing any bias in the likelihood.

In recent years, the use of HMC algorithms has been democratized in the software Stan (Carpenter et al., 2017). Stan implements a specialized version of HMC that generally applies across a broad class of Bayesian models, with built-in tools for diagnosing Markov chain mixing problems. However, Stan's generality requires limiting its flexibility, such as requiring all discrete latent variables to be marginalized, and the rest to be updated simultaneously. Usually, these are small sacrifices for easy HMC implementation with built-in mixing diagnostics. However, MCMC

for large LPMs often stretches one's computational resources to their limit. We thus need all tools at our disposal to optimize our inference strategy, including sampling discrete random variables and block updates of variables.

The specialized HMC-based sampling strategy we develop in this article is specifically intended for Gaussian LPMs (Rastelli et al., 2016), a class of LPMs for which the link probability function decays like a half-Gaussian probability density function. This class of LPMs was originally studied because they are easy to work with analytically. We show here that the Gaussian-inspired link function also provides computational advantages—the log posterior can be split into Gaussian and non-Gaussian components, thus facilitating efficient integration of HMC via split HMC (Shahbaba et al., 2014). Moreover, we further increase the efficiency for sparse networks by developing an exact dyad subsampling scheme based on Firefly Monte Carlo (FlyMC; Maclaurin & Adams (2015)). This scheme allows us to subsample the non-edge dyads, decreasing the complexity of the non-Gaussian component of the posterior while maintaining an exact MCMC strategy. To complete the LPM fitting algorithm, we also include Markov chain updates for the parameters of the link function. We also extend our sampling strategy to accommodate categorical covariates in the link function, as well as prior dependence between latent positions in the network as in longitudinal latent position models (Kim et al., 2018).

The remainder of the article is organized as follows. Section 2 establishes notation and provides the necessary background information pertaining to LPMs, Gaussian LPMs, and HMC. Section 3 outlines the ingredients of our new computation methodology for Gaussian LPMs: split HMC and firefly Monte Carlo, then combines them with updates to the link function parameters to define a new Markov chain Monte Carlo strategy. Section 4 presents two empirical studies to demonstrate the superiority of our algorithm. Study 1 uses synthetically generated examples to demonstrate the superior performance of our method compared to a variety of existing approaches in the literature such as Metropolis within Gibbs, elliptical slice sampling, Stan, and the No-U-turn sampler. Study 2 demonstrates the extent to which our algorithms outperform Metropolis within Gibbs for fitting information-sharing models among teachers and staff in a school district. Section 5 contains some concluding remarks.

2. Preliminaries

The following notation will be used throughout the paper. We use \mathbb{R} to denote the set of real numbers, \mathbb{R}_+ to denote the set of non-negative real numbers, \mathbb{N} to denote the set of natural numbers, and $[n]$ to denote the set $\{1, \dots, n\}$ of natural numbers less than or equal to n . For a set S , we use S^d to denote the collection of all d -length vectors with entries from S and $S^{n \times d}$ to denote collection of possible $n \times d$ matrices with entries from S . For two sets S_1, S_2 , $S_1 \times S_2$ denotes their Cartesian product.

For a vector $z \in \mathbb{R}^d$, z_i denotes its i th entry and $\|z\|$ denote its Euclidean norm. For a matrix $B \in \mathbb{R}^{n \times d}$, B_i denotes its i th row, $B_{\cdot i}$ denotes its i th column, B_{ij} denote its (i, j) th entry, $B^T \in \mathbb{R}^{d \times n}$ denotes its transpose, and B^{-1} denotes its inverse. We use I_n to denote the $n \times n$ identity matrix.

We represent networks among n entities as undirected binary graphs on n nodes. We use $A \in \{0, 1\}^{n \times n}$ to denote the adjacency matrix of the graph, with $A_{ij} = 1$ indicating the presence of an edge between nodes i and j , and $A_{ij} = 0$ indicating its absence. Our focus is on undirected graphs, so $A_{ij} = A_{ji}$ for all dyads $(i, j) \in [n]^2$. For simplicity, we use A to refer to both a graph and its adjacency matrix interchangeably, using $[n]$ index the nodes according to the order of their rows in the adjacency matrix. We use the shorthand $E_A \subseteq [n]^2$ to denote the set of edges associated with A , and $\{(i, j) \notin E_A\}$ to denote the set of edges absent from E_A . The combinatorial Laplacian of A is denoted as $L^A \in \mathbb{R}^{n \times n}$. Specifically, $L^A = D^A - A$ where D^A is a diagonal matrix of the node degrees $D_{ii}^A = \sum_{j=1}^n A_{ij}$.

2.1 Latent position network models

In the distance-based latent position network model (LPM) of Hoff et al. (2002), each node $i \in [n]$ is modeled as having a d -dimensional latent position $z_i \in \mathbb{R}^d$ for some positive integer d (typically $d = 2$ or $d = 3$ to facilitate visualization). It is convenient to arrange these latent positions in a matrix $Z \in \mathbb{R}^{n \times d}$, where $Z_i \cdot = z_i$. The edges A_{ij} are modeled as being generated according to

$$\mathbb{P}(A_{ij} = 1 | Z) = K(\|z_i - z_j\|, x_{ij})$$

where $\|z_i - z_j\|$ denotes the distance between nodes i and j , x_{ij} represents any relevant edge-specific covariates for nodes i and j , and K is the *link function*—a non-increasing function from $\mathbb{R}_+ \times \mathcal{X}$ to $[0, 1]$. Here, \mathcal{X} denotes the range of possible values of the covariate x_{ij} for each dyad $(i, j) \in [n]^2$. In this article, we assume each covariate x_{ij} is categorical taking on C distinct values. For notational convenience if there are no covariates, we will take $C = 1$ and let x be an $n \times n$ matrix of ones.

In their original version of the LPM, Hoff et al. (2002) proposed modeling K as a logistic function of the latent distance and the covariate according to

$$K(\|z_i - z_j\|, x_{ij}) = (1 + \exp(\alpha + \beta_{x_{ij}} + \|z_i - z_j\|))^{-1},$$

where the parameters $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}^C$ control the total number of edges and the effect of the covariates, respectively. Recently, Rastelli et al. (2016) proposed an alternative form for K inspired by the functional form of the Gaussian probability density function—aptly named the Gaussian Latent Position Model (GLPM). Their original exposition did not consider covariates, taking the form

$$K(\|z_i - z_j\|) = \tau \exp\left(-\frac{1}{2\gamma^2} \|z_i - z_j\|^2\right),$$

where $\tau \in [0, 1]$ controls the number of edges (i.e. sparsity level) and $\gamma^2 > 0$ controls the decay of the link probabilities.

Thus far, two advantages of GLPMs over logistic LPMs have been identified in the literature. The Gaussian-like choice of K yields closed-form expressions for various network statistics of GLPMs which makes them easier to theoretically analyze than logistic LPMs (Rastelli et al., 2016). The lighter tails of the Gaussian link function are also conducive to proving consistency of the maximum likelihood estimator of the latent positions (Spencer & Shalizi, 2019). In this paper, we identify and explore yet another advantage of GLPMs—the Gaussian shape of K facilitates faster posterior inference techniques.

Our work considers an extension of the GLPM to accommodate categorical covariates. Specifically, we consider

$$K(\|z_i - z_j\|, x_{ij}) = \tau_{x_{ij}} \exp\left(-\frac{1}{2\gamma^2} \|z_i - z_j\|^2\right), \tag{1}$$

with parameters $\tau \in [0, 1]^C$ and $\gamma^2 > 0$. Here, the effect of the covariate is encoded in the vector τ , allowing for subnetworks corresponding to certain covariate categories to be sparser than others. This single covariate formulation can be extended without loss of generality to multiple discrete covariates, with or without interactions, by a suitable mapping of the joint range space of covariates into $[C]$. For notational conciseness, we occasionally omit the dependence of K on x_{ij} in this article. Together, the Gaussian shape and factorizability of K can be exploited to speed up Bayesian inference.

Before proceeding, it is important to acknowledge that the parameters γ and Z in a GLPM are together identified only up to a multiplicative constant. That is, for any viable estimates $\gamma = \hat{\gamma}$ and $Z = \hat{Z}$, a model defined by

$$\gamma = 1, \quad Z = \hat{\gamma}^{-1} \hat{Z} \tag{2}$$

has an equivalent likelihood. This second parameterization is known as a centered parameterization (Papaspiliopoulos et al., 2007), which can be computationally advantageous for inferring γ^2 . We exploit it in Section 3.3.2.

2.2 Existing computational strategies for Bayesian inference of LPMs

Fitting a LPM to a graph A can be separated into two interdependent tasks: (1) inferring the latent positions Z and (2) inferring the parameters of the link function. Depending on the modeling objective of the problem at hand, either (1) or (2) could be the primary inferential target. For instance, Z is the primary inferential target when controlling for causal confounders (McFowland III & Shalizi, 2021), but τ (as in 1) is the primary target when estimating the effect of a covariate on edge probabilities. Regardless, both inference tasks are typically carried out in a single Monte Carlo algorithm; independent priors are placed on both the parameters and the latent positions, and their posterior distribution is approximated with samples drawn according to Markov chain Monte Carlo.

To simplify exposition, we will present our strategies for the two inference tasks separately. Here, we review existing computational strategies from the literature for inferring the latent positions Z conditional on K . All discussions of inference on the parameters of K for the GLPM are deferred until Section 3.3.

Bayesian inference of the latent positions Z depends on the link function K , the observed covariates x , and two additional inputs: an observed network (encoded by an adjacency matrix $A \in \{0, 1\}^{n \times n}$) and a prior on $Z \in \mathbb{R}^{n \times d}$. The standard prior choice for Z in the literature has been an independent isotropic d -dimensional Gaussian on each row (i.e. latent position) of Z . Here, we generalize this prior to $Z_{\cdot k} \sim N(0, \Omega^{-1})$ —defined independently with a shared Ω for each $k \in [d]$. That is, the nodes' positions are independent and identically distributed across dimensions, but can be dependent across nodes within each dimension. Without loss of generality, any Gaussian prior exhibiting dependence of a nodes' position across dimensions can be transformed into an equivalent prior with independence across dimensions via a rotation¹.

This more general set-up for the prior allows for known structural information—such as feature-informed node clustering or temporal dependence—to be included as non-zero entries in the precision matrix $\Omega \in \mathbb{R}^{n \times n}$. Other priors, such as a mixture of Multivariate Gaussians (Handcock et al., 2007; Krivitsky et al., 2009), are beyond the scope of this paper, but would involve a straightforward extension of the methods presented here.

Given the prior $Z_{\cdot k} \sim N(0, \Omega^{-1})$ for $k \in [d]$, the posterior distribution on Z is given by

$$\mathbb{P}(Z|A) \propto \prod_{(i,j) \in E_A} K(\|z_i - z_j\|) \prod_{(i,j) \notin E_A} (1 - K(\|z_i - z_j\|)) \exp\left(-\frac{1}{2} \sum_{k=1}^d Z_{\cdot k}^T \Omega Z_{\cdot k}\right). \quad (3)$$

The normalization constant for this density is a $(n \times d)$ -dimensional integral that cannot be computed analytically. Instead, we must rely on approximate methods for calculating expectations with respect to the posterior.

In their seminal LPM paper, Hoff et al. (2002) proposed for the posterior computation of $\mathbb{P}(Z|A)$ to be carried out via Markov chain Monte Carlo (MCMC). They obtained a Markov Chain Z^1, Z^2, \dots, Z^T with stationary distribution $\mathbb{P}(Z|A)$ by repeatedly applying a random walk Metropolis update to all latent positions Z simultaneously. As is the case for most MCMC algorithms, ensuring an adequate Metropolis-Hastings acceptance rate requires that the standard deviations of these random walk updates be appropriately tuned using a series of short pilot runs. However, these joint random walk proposals are known to be inefficient when the posterior is high-dimensional (e.g. for networks with many nodes) because the random walk standard deviation required to obtain reasonable acceptance rates is simply too small to explore the space efficiently.

In an effort to alleviate this slow-mixing, the subsequent LPM literature (e.g. Handcock et al. 2007; Raftery et al. 2012) uses a Metropolis within Gibbs strategy for updating the latent positions instead. In *Metropolis within Gibbs*, the latent positions $(z_i)_{i \in [n]}$ are updated one at a time in sequence according to a random walk via a symmetric kernel q centered at its current position (e.g. a scaled isotropic Gaussian or multivariate uniform). This approach is implemented in the popular R package `latentnet` (Krivitsky & Handcock, 2008); it is still widely used today (Fosdick et al., 2018; Aliverti & Durante, 2019; Sweet & Adhikari, 2020).

A sweep of the Metropolis within Gibbs algorithm can be summarized as follows. For each $i \in [n]$,

- (1) Propose $z'_i \sim q_\delta(z_i, z'_i)$.
- (2) Accept this proposal with probability equal to

$$\frac{\mathbb{P}(Z'|A)}{\mathbb{P}(Z|A)} = \frac{\exp\left(\sum_{k=1}^d Z'^T_k \Omega Z_k\right)}{\exp\left(\sum_{k=1}^d (Z'_k)^T \Omega Z_k\right)} \prod_{j:(i,j) \in E_A} \frac{K(\|z'_i - z_j\|)}{K(\|z_i - z_j\|)} \prod_{j:(i,j) \notin E_A} \frac{1 - K(\|z'_i - z_j\|)}{1 - K(\|z_i - z_j\|)} \quad (4)$$

Otherwise reject and keep z_i .

The matrix Z' in (4) is constructed such that $Z'_i = z'_i$ and $Z'_j = z_j$ for all other $i \neq j$. The notation $z'_i \sim q_\delta(z_i, z'_i)$ denotes drawing z'_i from a symmetric distribution centered at z_i with $\delta > 0$ denoting a tuning parameter for the width or *step size* of the proposal. Computing (4) involves only the prior for z_i and the (at most n) likelihood terms corresponding to dyads containing i —all other terms are equivalent for Z and Z' . For a fully observed network A , each full sweep updating Z thus requires $O(n^2)$ computations.

As with random walk Metropolis, it is standard practice to tune δ using preliminary tuning runs to achieve a desired acceptance rate². The required value of δ typically shrinks as n grows, meaning that chains must be run much longer to achieve mixing when fitting larger networks. For example, Figure 7 in Section 6.6 of the Appendix demonstrates the decreasing relationship between the number of nodes and the tuned Metropolis within Gibbs step size δ for the variety of different synthetic networks considered in Section 4.1.

For large enough n , approximating the posterior using Metropolis within Gibbs thus also becomes intractable (Raftery et al., 2012)—the step size is too small to efficiently explore the space given the complexity of computing the Metropolis-Hastings acceptance ratios. There have been multiple recent proposals that approximate the LPM likelihood (Raftery et al., 2012; Rastelli et al., 2018) to alleviate the computational burden of the accept-reject step. But as noted in the introduction, these approximations introduce non-vanishing bias in the subsequent inference.

Our goal in this article is to avoid such bias completely by developing an MCMC algorithm that outperforms Metropolis within Gibbs without sacrificing exactness. To accomplish this, we use a Monte Carlo algorithm known as HMC.

2.3 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) is an auxiliary variable MCMC algorithm that uses the gradient of the log posterior to inform an efficient Markov proposal kernel. Inspired by Hamiltonian dynamics, HMC augments the posterior distribution with a “momentum” variable for each target parameter, framing the task of proposing the next state as that of simulating Hamiltonian motion of an object along a high-dimensional surface.

An HMC chain consists of a sequence of snapshots of an object sliding along the frictionless surface. The object’s momentum is randomly refreshed after each snapshot, thus resulting in a sequence of stochastic draws. By using the negative log posterior as the energy function to inform

its motion, HMC provides large step sizes that nevertheless maintain high Metropolis-Hastings acceptance rates. Moreover, these ratios have closed forms due to the properties of Hamiltonian dynamics (namely reversibility and volume preservation). The algorithm is thus efficient for exploring high-dimensional posteriors.

General implementations of HMC have recently gained traction in the literature for fitting LPMs within large hierarchical models (Linderman et al., 2016; Salter-Townshend & McCormick, 2017). In this work, we will develop an HMC algorithm that is specifically tooled for inference in LPMs. We now provide a description of HMC, placing emphasis on the components relevant to the development of our algorithm for the LPM. For more detailed reviews of the theory and practice of MCMC, see Neal (2011) or Betancourt (2017).

Consider a target density $p(Z)$ that is differentiable with respect to its real-valued arguments $Z \in \mathbb{R}^d$. HMC targets an augmented version of this density $p(Z, U) = p(Z)q(U)$ where $U \in \mathbb{R}^d$ is a vector of auxiliary momentum variables—each corresponding to an entry in Z . Note that because the density $p(Z, U)$ admits $p(Z)$ as a marginal, discarding the U 's from a Markov chain targeting $p(Z, U)$ yields draws from $p(Z)$.

Let $q(U)$ be a zero mean multivariate Gaussian density with covariance matrix $M \in \mathbb{R}^{d \times d}$, and let $H(Z, U) = -\log(p(Z)) - \log(q(U))$. This function $H(Z, U)$ plays the role of energy in the Hamiltonian dynamics of HMC, with the covariance M —referred to as the *Mass matrix* (Neal, 2011) or *Euclidean metric* (Betancourt, 2017)—controlling the effect of the momentum on the dynamics.

Hamiltonian motion over $H(Z, U)$ is governed by the following differential equations:

$$\frac{dZ_i}{dt} = \frac{\partial H(Z, U)}{\partial U_i} = (M^{-1}U)_i \quad (5)$$

$$\frac{dU_i}{dt} = \frac{-\partial H(Z, U)}{\partial Z_i} = \frac{\partial \log(p(Z))}{\partial Z_i}. \quad (6)$$

Here, $(M^{-1}U)_i$ denotes the i th coordinate in the vector $M^{-1}U$, and t represents the artificial “time” for which the Hamiltonian trajectory is computed. That is, the derivatives of Z and U with respect to t reflect the rate of change in these quantities along Hamiltonian trajectory. Given an initial state (Z^0, U^0) , HMC generates a Markov chain of snapshots $(Z^j, U^j)_{j \in \mathbb{N}}$ with stationary distribution $p(Z, U)$ by iterating between simulating Hamiltonian motion for a fixed integration time $T > 0$, then refreshing the momentum according to its conditional distribution. The integration time T is typically specified by the user to control the length of time between momentum updates.

Given (Z^j, U^j) , the next draw (Z^{j+1}, U^{j+1}) is obtained using the following steps

- (1) Update the momentum variables via Gibbs $U^j \sim \text{MVN}(0, M)$.
- (2) Simulate Hamiltonian motion $(Z^j, U^j) \rightarrow (Z^{j''}, U^{j''})$ for T time units.
- (3) Accept the move $(Z^{j+1}, U^{j+1}) = (Z^{j''}, -U^{j''})$ with probability

$$\min \left(\frac{p(Z^{j''}, U^{j''})}{p(Z^j, U^j)}, 1 \right). \quad (7)$$

Otherwise reject the move, letting $(Z^{j+1}, U^{j+1}) = (Z^j, U^j)$.

The negation of $U^{j''}$ in Step 3 ensures the proposal is reversible³.

The performance of HMC as described above depends on two user-specified parameters: the mass matrix M and the integration time T . Before we discuss choosing these parameters, we must first address how to simulate Hamiltonian motion.

If the Hamiltonian motion in Step 2 above was to be simulated exactly, the Metropolis-Hastings correction in Step 3 would be unnecessary because the ratio would be exactly one (Neal, 2011).

This property is guaranteed by the conservation of energy in Hamiltonian motion—Step 2 simply moves along a density contour of the augmented distribution. Unfortunately, exact simulation of the Hamiltonian motion is not possible for most posterior densities that arise in Bayesian inference—there is no known analytic way to move along the contours.

In practice, simulation of the trajectory of Hamiltonian motion is typically carried out using approximate numerical integrators of the differential equations, the most popular of which is the *leapfrog integrator* (Neal, 2011). The leapfrog integrator discretizes the Hamiltonian motion via alternating linear updates of U and Z until the trajectory of length T has been simulated. The following steps are iterated $L \in \mathbb{N}$ times:

$$\begin{aligned} U &\leftarrow U - \frac{\varepsilon}{2} \frac{\partial H}{\partial Z}(Z, U) \\ Z &\leftarrow Z + \varepsilon M^{-1} U \\ U &\leftarrow U - \frac{\varepsilon}{2} \frac{\partial H}{\partial Z}(Z, U) \end{aligned}$$

Together, the user-specified parameters $\varepsilon > 0$ (step size) and $L \in \mathbb{N}$ (the number of steps) define the integration time $T = L\varepsilon$. Smaller values of ε provide more accurate approximations of the Hamiltonian motion and thus a higher Metropolis-Hastings acceptance rate. However, they also require correspondingly larger values of L —and thus more computation—to simulate a given integration time T . It is thus important to strike a balance between the two to obtain adequately high acceptance rates for reasonably correlated draws without wasting computational resources.

Choosing the user-specified parameters M and T for HMC via leapfrog amounts to choosing three parameters: the mass matrix M , the step size ε , and the number of steps L . The integration time $T = L\varepsilon$ is a product of the choices of ε and L .

Like tuning the step size for traditional Metropolis algorithms, standard practice for choosing ε and L is to conduct preliminary tuning runs at various parameter levels, looking for values that maximize the chain's efficiency. The matrix M can also be chosen this way. However, a more theoretically motivated heuristic (Betancourt, 2017) is to set M to the precision matrix of the posterior as estimated from the preliminary chains. In practice, the true precision matrix may be dense (and thus expensive to compute), making a diagonal or low-rank approximations (Carpenter et al., 2017; Bales et al., 2019) more suitable.

Unfortunately, efficiently tuning all three of L , ε , and M can itself be a computationally burdensome because the three parameters are interdependent (the optimal choice of ε relies particularly heavily on the choice of M). Indeed, under the standard settings, it is typical for Stan to devote more time to adapting L , ε , and M than running the final Markov chain. When the computational problem is already straining the computational budget at hand—such as in the case of large LPMs—these tuning costs can be prohibitive. We thus seek a simpler, more easily tunable algorithm that is specialized to large LPMs.

Before proceeding, it is worth noting that strategies exist for which L , ε , and M are not necessarily held constant through all regions of the posterior. For instance, Riemannian HMC (Girolami & Calderhead, 2011) adapts M based on the current state of the chain, and the NUTS algorithm (Hoffman & Gelman, 2014) adaptively chooses L on-the-fly to avoid wasted computation. However, these algorithms tend to be computationally expensive, requiring many additional density, gradient, or Hessian evaluations. This will be evident when we compare Stan and NUTS to our strategy in Section 4.1.

3. New sampling methodology

Our new MCMC algorithm for the GLPM is composed of three novel components: a split HMC (Shahbaba et al., 2014) integrator to update the latent positions (Section 3.1), a Firefly Monte

Carlo (FlyMC; Maclaurin & Adams (2015)) auxiliary variable scheme to subsample non-edge dyads (Section 3.2), and Gibbs sampling strategies to update the parameters τ and γ^2 of the link probability function K (Section 3.3). We present each of these contributions in sequence.

3.1 Split Hamiltonian Monte Carlo

Though it is certainly the most popular for HMC, the leapfrog integrator is just one of many options for integrating Hamiltonian dynamics (Leimkuhler & Reich, 2004; Chao et al., 2015; Mannseth et al., 2016). Here, we consider an alternative called *split Hamiltonian Monte Carlo* (Shahbaba et al., 2014). Split HMC is a variant on the leapfrog strategy that efficiently simulates Hamiltonian motion by exploiting a Gaussian component of the posterior. It works best when the Gaussian component is a good approximation for the entire posterior. Split HMC also provides a natural choice for the mass matrix M that is locally adaptive without having to evaluate the Hessian.

The standard leapfrog update described in Section 2.3 is equivalent to decomposing the energy into three terms:

$$H(Z, U) = -\frac{1}{2} \log(p(Z)) - \log(q(U)) - \frac{1}{2} \log(p(Z)). \tag{8}$$

then cycling through isolated updates according (5) and (6) for each of the components individually. This “split” of the energy ensures that only one of Z or U is being updated at any given time, causing each isolated operation to be straightforward. In split HMC, we consider a different split of the energy function, decomposing it to exploit partial analytic solutions of Hamiltonian equations.

Hamilton’s equations usually lack an analytic solution, but a notable exception is when the energy is defined as the negative logarithm of a multivariate Gaussian density (Pakman & Paninski, 2014); the exponential and uniform distributions are others (Bloem-Reddy & Cunningham, 2016). For the Gaussian case, the motion and momentum updates can be simulated exactly along an ellipse (i.e. a contour of the Multivariate Gaussian distribution). Alone, this fact would have limited utility for Bayesian computation; exact algorithms for inference involving Gaussian posteriors are readily available. However, these analytic solutions can be remarkably useful as part of a splitting strategy.

The split Hamiltonian integrator (Shahbaba et al., 2014) alternates between joint position momentum updates based on the analytical solution of the Gaussian component of the posterior and updates to the momentum to correct for the remaining portion of the posterior. When the exact part is a good approximation for the entire posterior, this allows for a coarser ε to maintain a high acceptance rate. We now present the decomposition of the GLPM posterior into its Gaussian and non-Gaussian components. Specifically, the likelihood of the edges, the prior density, and the momentum forming the Gaussian component, and the likelihood of non-edges forms the remainder.

Treating the τ and γ^2 as known, the posterior (3) takes the form

$$\mathbb{P}(Z|A, \tau, \gamma^2) \propto \left(\prod_{(ij) \in E_A} \tau_{x_{ij}} \right) \exp\left(\frac{\sum_{\ell=1}^d Z_{\cdot\ell}^T \left(\Omega + \frac{L^A}{\gamma^2} \right) Z_{\cdot\ell}}{-2} \right) \prod_{(ij) \notin E_A} \left(1 - \tau_{x_{ij}} \exp\left(\frac{\|z_i - z_j\|^2}{-2\gamma^2} \right) \right) \tag{9}$$

where L^A denotes the Laplacian of A . This posterior is proportional to the product of the two components $\mathbb{P}(Z|A, \tau, \gamma^2) \propto \mathbb{P}_1(Z|A, \tau, \gamma^2)\mathbb{P}_0(Z|A, \tau, \gamma^2)$ where

$$\mathbb{P}_1(Z|A, \tau, \gamma^2) = \left(\prod_{(ij) \in E_A} \tau_{x_{ij}} \right) \exp\left(-\frac{1}{2} \sum_{\ell=1}^d Z_{\cdot\ell}^T \left(\Omega + \frac{1}{\gamma^2} L^A \right) Z_{\cdot\ell} \right)$$

corresponds to the contribution of the prior and likelihood of the observed edges and

$$\mathbb{P}_0(Z|A, \tau, \gamma^2) = \prod_{(ij) \notin E_A} \left(1 - \tau_{x_{ij}} \exp \left(-\frac{1}{2\gamma^2} \|z_i - z_j\|^2 \right) \right)$$

corresponds to the contribution to the likelihood of the non-edges. Using the shorthand

$$\Sigma = \left(\Omega + \frac{1}{\gamma^2} L^A \right), \tag{10}$$

we can now split the corresponding energy as

$$H(Z, U) = \left[-\frac{1}{2} \log (\mathbb{P}_0(Z|A, \tau, \gamma^2)) \right] + \left[\frac{1}{2} \sum_{\ell=1}^d \left(Z_{\cdot \ell}^T \Sigma Z_{\cdot \ell} + U_{\cdot \ell}^T M^{-1} U_{\cdot \ell} \right) \right] - \left[\frac{1}{2} \log (\mathbb{P}_0(Z|A, \tau, \gamma^2)) \right], \tag{11}$$

ignoring additive constants. The center term in this split is Gaussian.

In the above, we have departed from the typical notation in our definition of the mass matrix M and momentum variables U . In standard presentations of HMC (including our Section 2.3), the target parameters and momentum variables are naturally represented as vectors. However, for a LPM, the parameters Z are more suitably represented as a $n \times d$ matrix. We have thus chosen to also represent the momentum variables U as a $n \times d$ matrix. Since there are $n \times d$ momentum variables, the standard notation/definition of the mass matrix would require that $M \in \mathbb{R}^{nd \times nd}$. We have opted to instead define the full mass matrix block diagonally, using d repetitions of the same matrix $M \in \mathbb{R}^{n \times n}$. This choice facilitates the more compact representation in (11) without altering the validity of the algorithm.

On top of being notationally and computationally convenient, the use of an identical mass matrix across all dimensions is justified by symmetry in the target posterior—the marginal distribution of each column of Z is the same (Shortreed et al., 2006).

The above decomposition thus suggests a natural choice of M . Recall from Section 2.3 that the precision matrix of the posterior is an efficient choice for M . Accordingly, we suggest that Σ is a reasonable choice for M , as it should be a good approximation of the posterior precision matrix provided that \mathbb{P}_0 is a good approximation of the full posterior. Moreover, the choice $M = \Sigma$ is also particularly amenable to simulating the split HMC trajectories because it leads to arithmetic cancelations that simplify computation. Finally, the mass matrix M depends on the parameter γ^2 —when combined with a Monte Carlo strategy for inferring γ^2 (such as the one we present in Section 3.3), setting $M = \Sigma$ allows for the mass matrix to evolve adaptively with the state of γ^2 in the chain.

The following is a complete recipe for split HMC for LPMs, using the block diagonal mass matrix we have just defined. Note that the intermediate variable $V \in \mathbb{R}^{n \times d}$ introduced in Step 2 is a change of variable for efficiently parameterizing the contour of the multivariate Gaussian, and Step 4 inverts the change of variable to recover U . For more details on the exact simulation of HMC for Multivariate Gaussians, see Pakman & Paninski (2014).

Suppose that A denotes an observed adjacency matrix, $\tau \in [0, 1]^C$ and $\gamma^2 > 0$ denote the values of the parameters of the Gaussian link function, and $\varepsilon > 0, L \in \mathbb{N}$ and Σ (as defined in 10) denote the user-specified tuning parameters for split HMC. Given (Z^j, U^j) , the next split HMC draw (Z^{j+1}, U^{j+1}) is obtained via the following steps:

Algorithm 1: Split HMC Update

- (1) Update the momentum variables via Gibbs $U^j \sim \text{MVN}(0, \Sigma)$.
- (2) Define intermediate variables $V^j = \Sigma^{-1}U^j$ and $Z^{j'} = Z^j$.
- (3) Integrate Hamiltonian motion $(Z^j, U^j) \rightarrow (Z^{j'}, U^{j'})$ for $T = L\varepsilon$ time units by iterating the following updates L times:

$$\begin{aligned}
 V^j &\leftarrow V^j + \frac{\varepsilon}{2} \Sigma^{-1} \frac{\partial \log(\mathbb{P}_0(Z|A, \tau, \gamma^2))}{\partial Z}(Z^{j'}) \\
 (Z^{j'}, V^j) &\leftarrow \left(\sin(\varepsilon)V^j + \cos(\varepsilon)Z^{j'}, \cos(\varepsilon)V^j - \sin(\varepsilon)Z^{j'} \right) \\
 V^j &\leftarrow V^j + \frac{\varepsilon}{2} \Sigma^{-1} \frac{\partial \log(\mathbb{P}_0(Z|A, \tau, \gamma^2))}{\partial Z}(Z^{j'})
 \end{aligned}$$

Finish by setting $U^{j'} = \Sigma V^j$.

- (4) Accept the move $(Z^{j+1}, U^{j+1}) = (Z^{j'}, -U^{j'})$ with probability

$$\min \left(\frac{\mathbb{P}(Z^{j'}|A, \tau, \gamma^2)}{\mathbb{P}(Z^j|A, \tau, \gamma^2)} \exp \left(\frac{1}{2} \sum_{\ell=1}^d (U'_{\ell} - U''_{\ell})^T \Sigma^{-1} (U'_{\ell} - U''_{\ell}) \right), 1 \right).$$

Otherwise, the move is rejected and $(Z^{j+1}, U^{j+1}) = (Z^j, U^j)$.

In Step 3 of Algorithm 1, the gradient functions return $n \times d$ matrices defined by

$$\left(\frac{\partial \log(\mathbb{P}_0(Z|A, \tau, \gamma^2))}{\partial Z}(Z) \right)_{ik} = \frac{\partial \log(\mathbb{P}_0(Z|A, \tau, \gamma^2))}{\partial Z_{ik}} \tag{12}$$

$$= \sum_{j:(i,j) \notin E_A} \frac{(Z_{ik} - Z_{jk})}{\gamma^2} \frac{\tau_{x_{ij}} \exp \left(-\frac{\|z_i - z_j\|^2}{2\gamma^2} \right)}{1 - \tau_{x_{ij}} \exp \left(-\frac{\|z_i - z_j\|^2}{2\gamma^2} \right)}. \tag{13}$$

3.2 Firefly sampling of non-edges

Recall from Section 3.1 that exact simulation of the Hamiltonian motion is thwarted by the non-edge terms in the likelihood. Furthermore, the computational bottlenecks for running Algorithm 1 are the gradient evaluations in Step 3 and the acceptance ratio evaluation in Step 4. These computations each require an operation to be performed for each non-edge, making them especially expensive for large sparse networks due to the large number of non-edges. It can thus be beneficial to eliminate some non-edge terms from the likelihood at each iteration of split HMC. Here, we propose such a strategy.

Consider the following data augmentation scheme inspired by the Firefly Monte Carlo (FlyMC; Maclaurin & Adams (2015)). For each $(i, j) \in [n]^2$, we define auxiliary independent binary random variables θ_{ij} such that $\mathbb{P}(\theta_{ij} = 1 | \tau_{x_{ij}}) = \tau_{x_{ij}}$. Using these auxiliary variables, we can re-express the edge probabilities as

$$\begin{aligned}
 \mathbb{P}(A_{ij} = 1 | \theta_{ij} = 1, \tau_{x_{ij}}, \gamma^2, z_i, z_j) &= \exp \left(-\frac{1}{2\gamma^2} \|z_i - z_j\|^2 \right), \\
 \mathbb{P}(A_{ij} = 1 | \theta_{ij} = 0, \tau_{x_{ij}}, \gamma^2, z_i, z_j) &= 0,
 \end{aligned}$$

while maintaining the same marginal likelihood. Now,

$$\mathbb{P}(\theta_{ij} = 0 | A_{ij} = 0, Z, \tau_{x_{ij}}, \gamma^2) = \frac{1 - \tau_{x_{ij}}}{1 - \tau_{x_{ij}} \exp\left(-\frac{1}{2\gamma^2} \|z_i - z_j\|^2\right)},$$

$$\mathbb{P}(\theta_{ij} = 0 | A_{ij} = 1, Z, \tau_{x_{ij}}, \gamma^2) = 0.$$

Note that for all $(i, j) \in E_A$, $\theta_{ij} = 1$ must hold. Thus,

$$\mathbb{P}(Z|A, \theta, \tau, \gamma^2) = \mathbb{P}_1(Z|A, \tau, \gamma^2) \prod_{ij:\theta_{ij}=1, A_{ij}=0} \left(1 - \exp\left(-\frac{1}{2\gamma^2} \|z_i - z_j\|^2\right)\right),$$

meaning that

$$\mathbb{P}_0^*(Z|A, \theta, \tau, \gamma^2) = \prod_{ij:\theta_{ij}=1, A_{ij}=0} \left(1 - \exp\left(-\frac{1}{2\gamma^2} \|z_i - z_j\|^2\right)\right)$$

can replace $\mathbb{P}_0(Z|A, \tau, \gamma^2)$ in Split HMC once the θ variables are instantiated. If many of the θ_{ij} are 0, computing $\mathbb{P}_0(Z|A, \tau, \gamma^2)$ —and its gradients—is far cheaper than computing the marginal $\mathbb{P}_0(Z|A, \tau, \gamma^2)$ analogs. Combining this data augmentation strategy with split HMC can be a major computational improvement, provided that we can update the θ_{ij} values efficiently.

To do so, we propose a Metropolis-Hastings step using proposal $q(\theta_{ij} = 1) = \tau_{ij}$. Let $\text{MH}(\theta_{ij} = 0 \rightarrow \theta_{ij} = 1)$ denote the Metropolis-Hastings ratio associated with a proposed move from $\theta_{ij} = 0$ to $\theta_{ij} = 1$, and let $\text{MH}(\theta_{ij} = 1 \rightarrow \theta_{ij} = 0)$ denote the Metropolis-Hastings ratio associated with a proposed move from $\theta_{ij} = 1$ to $\theta_{ij} = 0$. The values of these ratios are given by

$$\text{MH}(\theta_{ij} = 0 \rightarrow \theta_{ij} = 1) = \left(1 - \exp\left(-\frac{1}{2\gamma_{ij}^2} \|z_i - z_j\|^2\right)\right), \tag{14}$$

$$\text{MH}(\theta_{ij} = 1 \rightarrow \theta_{ij} = 0) = \frac{1}{1 - \exp\left(-\frac{1}{2\gamma_{ij}^2} \|z_i - z_j\|^2\right)} \geq 1. \tag{15}$$

Thus, out of the four possible moves $0 \rightarrow 0, 0 \rightarrow 1, 1 \rightarrow 0, 1 \rightarrow 1$, the accept-reject step only needs to be performed for $0 \rightarrow 1$. In contrast, a Gibbs update of θ_{ij} from its full conditional would require that the posterior density be evaluated for any of the four moves. As such, this Metropolis update involves less computation than a full Gibbs update, especially for small values of τ .

Going forward, we refer to the parameter augmentation and update strategy described above as FlyMC. The reduction in computational cost of evaluating the posterior density and gradients under FlyMC is most prevalent when most of the θ_{ij} are 0. Because $\mathbb{P}(\theta_{ij} = 0 | \tau_{x_{ij}}) = 1 - \tau_{x_{ij}}$, the computational gains from FlyMC are largest when $\tau_{x_{ij}}$ is small. On the other hand, when $\tau_{x_{ij}}$ are relatively large (close to 1), most values of the θ_{ij} will be one, meaning the computational improvements in evaluating the gradient may not justify the computational expense of instantiating and updating the θ variables. In the extreme case of $\tau_{x_{ij}} = 1$, no subsampling will occur at all, so FlyMC should not be included.

For sparse networks, however, $\tau_{x_{ij}}$ may be very small for some values of x_{ij} , leading to substantial computational gains. In addition to providing a computational speed-up, the new FlyMC posterior facilitates the inference of τ via Gibbs steps.

3.3 Bayesian inference of the parameters of the link function

Thus far, our posterior computation strategy for Z has held τ and γ^2 —the parameters of the link function—at fixed values. In most applications, τ and γ^2 are unknown—they need to be inferred

along with Z . As we noted in Section 2.2, τ may even be the primary inferential target. It is thus important that our posterior computation strategy compute the full joint posterior of τ , γ^2 , and Z . Here, we describe efficient Gibbs updates for both τ (Section 3.3.1) and γ^2 (Section 3.3.2) to be alternated with our split HMC + FlyMC.

The updates we describe here apply to specific families of priors on τ and γ^2 . In particular, we use independent Beta priors for each entry in τ , along with an inverse Gamma prior γ^2 . Moreover, the update for τ is applicable only in conjunction with the FlyMC strategy outlined in Section 3.2. If FlyMC is not used, we recommend a simple random walk Metropolis-Hastings update for τ instead. Finally, our update strategy for γ^2 depends on the centered re-parameterization mentioned in Section 2.1 where γ is treated as a scaling factor for the latent positions. It is applicable whether or not FlyMC is used. An detailed expression of the full posterior and relevant conditionals are available in Section 6.4 of the Appendix.

3.3.1 Updating τ given Z, θ, γ^2

We assume independent $\text{Beta}(\alpha_c, \beta_c)$ priors on the entries in τ , with $c \in [C]$ indexing the possible levels of the covariate, and $\alpha_c, \beta_c \in \mathbb{R}_+$. If Section 3.2's FlyMC strategy is used, each τ_c can then be updated according to its conditional posterior distribution given the FlyMC variables θ and the covariates x . Indeed, this posterior distribution is conjugate to the Beta prior because inferring τ_c given θ , and the covariate values x is equivalent to inferring the probability parameter of a sequence of Bernoulli trials (specifically the θ_{ij} for which $x_{ij} = c$). Thus,

$$\tau_c \mid \theta, x \sim \text{Beta}(\alpha_c + \Theta_c^1, \beta_c + \Theta_c^0), \quad (16)$$

where $\Theta^0, \Theta^1 \in (\{0\} \cup \mathbb{N})^C$ are defined according to

$$\begin{aligned} \Theta_c^0 &= \left| \left\{ (i, j) \in [n]^2 : \theta_{ij} = 0 \text{ and } x_{ij} = c \right\} \right| \\ \Theta_c^1 &= \left| \left\{ (i, j) \in [n]^2 : \theta_{ij} = 1 \text{ and } x_{ij} = c \right\} \right| \end{aligned}$$

This update can be made efficient by keeping track of Θ^0 and Θ^1 during FlyMC updates.

3.3.2 Updating γ^2 given Z, τ, θ

We assume a Inverse Gamma $\text{IG}(a, b)$ prior on γ^2 , where $a, b > 0$. Then, the posterior density of $\gamma^2 \mid A, Z, \theta$ is proportional to

$$p(\gamma^2 \mid A, Z, \theta) \propto \text{IG}(\gamma^2 \mid a, b) \exp \left(-\frac{\sum_{\ell=1}^d Z_{\cdot \ell}^T L^A Z_{\cdot \ell}}{2\gamma^2} \right) \prod_{ij: \theta_{ij}=1, A_{ij}=0} \left(1 - \exp \left(-\frac{\|z_i - z_j\|^2}{2\gamma^2} \right) \right)$$

where $\text{IG}(\gamma^2 \mid a, b)$ denotes the probability density function of the Inverse Gamma distribution. This density provides no closed-form Gibbs update and is expensive to evaluate.

Alternatively, we propose using the centered parameterization (2) in Section 2.1. By modifying the prior on Z to depend on γ through $Z_{\cdot \ell} \sim N(0, \gamma^{-2}\Omega)$ for $\ell \in [d]$, we can instead treat γ^2 as known to be 1 in the link function (1). Instead, we infer a scale parameter for the variance of the latent positions.

Accordingly, the conditional distribution of γ^2 given A, Z, θ, τ, x is given by

$$\begin{aligned} p(\gamma^2 \mid A, Z, \theta, x) &= \text{IG}(\gamma^2 \mid a, b) \exp \left(-\frac{\gamma^2}{2} \sum_{\ell=1}^d (Z_{\cdot \ell})^T \Omega^{-1} Z_{\cdot \ell} \right) \\ &\propto \text{IG} \left(\gamma^2 \mid a + \frac{nd}{2}, b + \frac{1}{2} \sum_{\ell=1}^d (Z_{\cdot \ell})^T \Omega^{-1} Z_{\cdot \ell} \right). \end{aligned}$$

The posterior dependence of γ^2 on θ and A has been eliminated and the prior on γ^2 is now conjugate, allowing for a simple Gibbs update.

This re-parameterization and the corresponding Gibbs update is straightforward to incorporate with the other updates of Z , θ , and τ described in Sections 3.1, Section 3.2, and Section 3.3.1, respectively. We simply fix the γ^2 in the link function at 1 and replace Ω with $\gamma^{-2}\Omega$ for the prior covariance in those sections. After performing MCMC sampling in this re-parametrized setting, one can recover the original parameterization with γ^2 in the link function by scaling the Z draws.

Sections 3.1, 3.2, and 3.3 together define a full MCMC strategy for posterior computation of Z , τ , γ^2 that also introduces and updates auxiliary FlyMC variables θ . Each sweep of the chain iterates through a split HMC update of Z , a Metropolis update of each θ , a Gibbs update for each entry in τ , and a Gibbs update of γ^2 . Alternatively, if FlyMC is not incorporated, posterior computation of Z , τ , and γ^2 can be performed by alternating the split HMC update of Z , a Metropolis update of each entry in τ , then a Gibbs update of γ^2 . For completeness, the functional form of the posterior being computed in both the split HMC and split HMC + FlyMC algorithms is outlined in Section 6.4 of the Appendix.

4. Empirical studies

To explore and understand the relative strengths and weaknesses of our new posterior inference algorithm for GLPMs, we conduct two empirical studies. Study 1 consists of a “bake-off” between various inference algorithms, comparing the efficiency of our method to that of plausible competitors from the literature. These comparisons span a variety of synthetically generated networks of different sizes and sparsity levels. Study 2 is a real data example, demonstrating efficiency of split HMC and split HMC + FlyMC compared to traditional Metropolis within Gibbs for modeling information-sharing among elementary school teachers and staff in a school district. Multiple model set-ups are considered in Study 2, including the use of categorical covariates and longitudinal network data.

Given our interest in unbiased methods for posterior computation, we focus exclusively on MCMC algorithms in both experiments. These MCMC algorithms all approximate posterior distributions equally well if run for a sufficiently long time, so our comparisons are based on their relative efficiency. Descriptions of the software and hardware we use for the experiments, details of how we tune the algorithms, and details of how we compare their performance are available in Sections 6.1, 6.2, and 6.3 of the Appendix, respectively.

4.1 Study 1: Synthetic data

In this empirical study, we investigate the efficiency of split HMC (Section 3.1) and split HMC + FlyMC (Sections 3.1 and 3.2) compared to nine other exact MCMC algorithms from the literature. We are especially interested in fitting LPMs for large sparse networks, so we have tooled the study design to compare performance as networks get larger and more sparse. Our investigation involves fitting Gaussian LPMs to sixteen different synthetically generated networks. These networks—stochastically generated according to a GLPMs with pre-specified values of τ , γ^2 , and n —demonstrate a variety of different sizes and sparsity. Specifically, we consider the full factorial design outlined by the first three columns of Table 1. No covariates are included, and all latent positions are drawn from a two-dimensional isotropic Gaussian. Note that the type of sparsity driven by small τ has different structure than that driven by small γ^2 , so our design considers both.

Figures 1 and 2 report the relative performance of eleven different MCMC posterior inference algorithms across the 16 networks described in Table 1. The algorithms vary along two criteria: the proposal used to update Z , and whether or not FlyMC (Section 3.2) is used to subsample the non-edges. We consider five strategies for updating Z : Metropolis within Gibbs (Section 2.2),




Figure 1. A depiction of the relationship between the number of nodes in the synthetically generated networks ($\tau = 0.2, 0.8, \gamma^2 = 0.2, 1.0$) for Empirical Study 1 and the relative efficiency of the five posterior computation algorithms. For each algorithm, relative efficiency (vertical axis) is quantified as the median across 500 dyads in the synthetic network of the relative Markov chain efficiency compared to Metropolis within Gibbs as a baseline. Note that the vertical axis is on the log scale.

elliptical slice sampling (Murray et al., 2010), elliptical slice sampling within Gibbs (Hahn et al., 2019), split HMC (Section 3.1) with $T \approx 2$, and an alternative implementation of Split HMC that uses NUTS (Hoffman & Gelman, 2014) to adaptively choose the integration time T . For each of these strategies, we consider implementations both with and without FlyMC. Finally, we include HMC as implemented in Stan version 2.18.2 (Carpenter et al., 2017) as an additional competitor algorithm, bringing the total number of algorithms to eleven. A FlyMC version of Stan is not possible. Metropolis within Gibbs serves as a standard baseline, the elliptical slice sampling algorithms represent alternative ways to exploit the Gaussian component in the posterior, the NUTS version of Split HMC illustrates the extra computational cost of adaptively choosing T , and Stan represents existing software.

In addition to sampling the latent positions, we use each algorithm to sample τ using a uniform prior and γ^2 using an inverse gamma $IG(1,1)$ prior. For the five FlyMC algorithms, we alternate between updates of Z , updates of the FlyMC variables θ according to the Metropolis strategy outlined in Section 3.2, updates of τ according to the Gibbs strategy outlined in Section 3.3.1, and an update of γ^2 to the Gibbs strategy outlined in Section 3.3.2. Where FlyMC is not used, τ is updated using a random walk Metropolis algorithm instead of Gibbs.

Each algorithm was initialized identically using the maximum likelihood estimate of Z and the true values of τ and γ^2 . The FlyMC and non-FlyMC versions of each algorithm were each simulated for 10,000 iterations. The Stan chain was run for just 2000 iterations due to its much longer runtime. The various hyperparameters (step size ϵ for HMC and random walk standard deviation δ for Metropolis) are shown in Table 1; these values were tuned according to the strategies outlined in Section 6.2 of the Appendix.

To evaluate the relative performances of each algorithm on each network, we consider the effective number of samples generated per second of runtime. The effective sample size ESS_f of a




Figure 2. A depiction of the relationship between the number of nodes in the synthetically generated networks ($\tau = 0.2, 0.8, \gamma^2 = 0.2, 1.0$) for Empirical Study 1 and the relative efficiency (compared to Metropolis within Gibbs) of the five FlyMC posterior computation algorithms. The solid black baseline is included for easy comparison to Metropolis within Gibbs. Note that the vertical axis is on the log scale.

Markov chain $\theta^1, \theta^2, \dots, \theta^N$ is defined as

$$ESS_f(\theta^1, \dots, \theta^N) = \frac{N}{1 + 2 \sum_{t=1}^{\infty} \rho_{t,f}}$$

where $\rho_{t,f}$ denotes the t -lag autocorrelation of the function $f(\theta)$ in the chain. For each network, we compute the ESS_f of each algorithm for 500 distinct choices of f . Each choice corresponds to the log probability of an edge at a randomly selected dyad. For each dyad, we calculate

$$\frac{ESS_f(\theta^1, \dots, \theta^N)}{ESS_f(\theta_m^1, \dots, \theta_m^N)} \times \frac{\text{time (in seconds) taken to compute the chain } \theta_m^1, \dots, \theta_m^N}{\text{time (in seconds) taken to compute the chain } \theta^1, \dots, \theta^N} \quad (17)$$

where $\theta^1, \theta^2, \dots, \theta^N$ denotes draws according to the algorithm being evaluated, $\theta_m^1, \dots, \theta_m^N$ denotes draws according to a well-tuned Metropolis within Gibbs algorithm exploring the same posterior, and f denotes the log probability of the given dyad. Figure 1 reports the median of (17) across 500 dyads for each of the non-FlyMC algorithms. For readability, the results from the analogous FlyMC algorithms are presented separately as Figure 2, using the same colors (but dashed instead of solid lines). For more explanation and justification of the evaluation metric in (17), see Section 6.3 of the Appendix.

The results shown in Figures 1 and 2 demonstrate several phenomena. Notably, Split HMC and Split HMC + FlyMC are the standout performers across all networks considered. Split HMC clearly outperforms Metropolis within Gibbs for all networks, and Split HMC + FlyMC outperforms Metropolis within Gibbs for all networks except the smaller networks in the sparsest regime. Notably, both implementations of Split HMC with $T \approx 2$ outperform their NUTS counterparts and Stan, demonstrating that the extra computational cost of using NUTS or Stan to adaptively update L may be unwarranted for LPMs.

Table 1. A summary of the simulation design for Study 1. The first three columns specify the number of nodes, the link function parameter γ^2 , and the link function parameter τ used to generate each of the synthetic networks. The edge density column reports the fraction of dyads exhibiting edges in each network. The final two columns report the tuned values of hyperparameters for the algorithms

Number of nodes	τ	γ^2	Edge density	Step size ε for HMC (+ FlyMC)	Step size δ for metropolis (+ FlyMC)
50	0.20	0.20	0.022	0.14 (0.08)	2.40 (2.40)
50	0.80	0.20	0.087	0.29 (0.16)	1.20 (1.00)
50	0.20	1.00	0.079	0.57 (0.21)	1.20 (1.20)
50	0.80	1.00	0.304	0.42 (0.28)	0.60 (0.60)
100	0.20	0.20	0.022	0.20 (0.06)	2.40 (1.60)
100	0.80	0.20	0.083	0.23 (0.15)	0.80 (0.80)
100	0.20	1.00	0.075	0.40 (0.20)	1.00 (0.80)
100	0.80	1.00	0.285	0.31 (0.28)	0.40 (0.40)
200	0.20	0.20	0.02	0.27 (0.08)	1.60 (1.20)
200	0.80	0.20	0.08	0.25 (0.18)	0.60 (0.60)
200	0.20	1.00	0.07	0.28 (0.19)	0.80 (0.60)
200	0.80	1.00	0.28	0.28 (0.25)	0.30 (0.30)
500	0.20	0.20	0.016	0.27 (0.08)	0.80 (0.80)
500	0.80	0.20	0.066	0.20 (0.13)	0.40 (0.40)
500	0.20	1.00	0.062	0.18 (0.18)	0.50 (0.40)
500	0.80	1.00	0.254	0.16 (0.17)	0.20 (0.20)

All methods based on elliptical slice sampling perform poorly, demonstrating that HMC is a better method for exploiting the near-Gaussianity of the posterior than elliptical slice sampling. Indeed, the elliptical slice sampling algorithms performed worse than Metropolis within Gibbs. The poor performance of the elliptical slice within Gibbs algorithms was due to its runtime—the conditional means and variances of each latent position at each iteration are very expensive to compute. The joint update elliptical slice algorithms performed poorly for the opposite reason. They had much faster runtimes, but the corresponding chains mixed very slowly because the draws exhibited very high autocorrelation.

The dominance of the Split HMC methods appears to be more pronounced for larger networks. For the denser $\tau = 0.8$ networks, Split HMC performs remarkably well, showing a distinct upward trend, indicating that its dominance over Metropolis within Gibbs would be even more pronounced for larger networks. For the sparser $\tau = 0.2$ networks, Split HMC + FlyMC demonstrates a similar upward trend. The tuned values of ε and δ shown in Table 1 demonstrate that split HMC is more robust to large network sizes: while the tuned step sizes for the Metropolis within Gibbs decay as the number of nodes increases, the tuned values of ε for split HMC remain more stable.

To facilitate comparison between the FlyMC and non-FlyMC implementations of split HMC, we have also included Figure 3. Instead of just reporting the median, it summarizes the entire distribution of (17) across the 500 sampled dyads. From the side-by-side boxplots, we can see that split HMC clearly outperforms split HMC + FlyMC for the $\tau = 0.8$ networks of all sizes. For the $\tau = 0.2$ networks, the comparison is less clear cut. For the very sparse network $\tau = 0.2$, $\gamma^2 = 0.2$, split HMC outperforms the FlyMC version for the smaller networks, but FlyMC edges it out for the 500 node network. It is worth noting that in this very sparse regime for smaller networks, the extreme lack of edges can lead to ambiguity in whether the extreme sparsity is driven by small τ ,




Figure 3. Boxplots showing the relative efficiency of Split HMC + FlyMC and Split HMC relative to Metropolis within Gibbs across 500 dyads in each network.

small γ^2 , or both. The joint posteriors of τ , γ^2 for different synthetic networks shown in Figure 8 in Section 6.6 of the Appendix demonstrates this phenomenon—there is a remarkable amount of uncertainty in the posterior of τ for the smaller sparse networks.

We have noticed that the FlyMC updates of τ and θ tend to mix slowly in these uncertain situations, thus leading to slow exploration of the joint distribution of (τ, γ^2) . Unencumbered by slow-mixing θ variables, split HMC tends to perform better in these under-identified settings, suggesting that FlyMC should only be used when the τ variable is better identified. This seems to be the case in the $\tau = 0.2, \gamma^2 = 1.0$ networks, where the FlyMC clearly outperforms the non-FlyMC version.

The full distribution of the relative efficiencies highlights another observation—although the split HMC algorithms tend to outperform Metropolis within Gibbs for the vast majority of dyads, there is often a small minority of dyads for which Metropolis within Gibbs performs better (seen as the lower tails of the boxplots in Figure 3 sometimes extending below 1). A thorough investigation into these dyads revealed no apparent pattern for which dyads tend to perform relatively poorly in a given network, suggesting that the primary explanation is simply the high-dimensionality of the posterior—with so many dimensions along which to mix, there will often be a small minority that mix more slowly. Regardless, the underperformance is not drastic for the large networks in which we are interested—split HMC still performs at the same order of magnitude as Metropolis within Gibbs.

From this empirical study, we have demonstrated that split HMC and split HMC + FlyMC tend to outperform competitors in the literature on synthetically generated data. For denser networks, or networks for which τ and γ^2 are poorly identified (i.e. smaller sparse networks), split HMC tends to be the better choice. For larger sparse networks, split HMC + FlyMC seems to be the top performer. In all cases, these strategies perform far better than simple Metropolis within Gibbs.

Table 2. A summary of the data configurations for Study 2

Dataset	Survey years	School IDs	Number of nodes	Number of edges
One year, one school	1	ID 4	32	150
One year, all schools	1	ID 1 to ID 14	326	1363
All years, one school	1-5	ID 4	14	79
All years, all schools	1-5	ID 1 to ID 14	129	1038

4.2 Study 2: Network of information-sharing in a school district

To demonstrate the efficacy of our split HMC strategies on real data, we now showcase several applications of LPMs to information-sharing networks of teachers and staff in a school district. These applications involve several model/network configurations commonly encountered in practice: networks with covariates encoding group memberships of the nodes, longitudinally observed networks with models promoting serial dependence of the latent positions, and combinations of the two. The data we use were collected in a mid-sized suburban school district in the Midwestern United States as part of the Distributed Leadership Studies at Northwestern University, a comprehensive program of research involving several longitudinal studies of workplace and social interactions among school staff and school systems. For more details about this particular dataset, see Spillane & Hopkins (2013); Spillane et al. (2018).

In five separate years, elementary school teachers and staff within this district were surveyed about who in the district they went to for advice, as well as the school in which they worked and other relevant covariates. Over these years, 661 distinct individuals responded to the survey in at least one year. 129 of them were present for all five surveys.

For the purposes of this empirical study, we have compiled the survey responses into a series of five undirected *information-sharing* networks—one for each year of data. These undirected information-sharing relationships were obtained by symmetrizing the information in the advice-seeking survey. That is, for each network, an edge is present between two individuals if either of them reported going to the other for advice in that year. In addition to the edge information for each dyad (i, j) , we have access to indicators x_{ij1} and x_{ij2} defined as:

$$x_{ij1} = \begin{cases} 1 & \text{if individuals } i \text{ and } j \text{ work at the same school} \\ 0 & \text{otherwise} \end{cases}$$

$$x_{ij2} = \begin{cases} 1 & \text{if individuals } i \text{ and } j \text{ shared advice in the previous year} \\ 0 & \text{otherwise} \end{cases}$$

which can be used to form covariates. From this sequence of five networks, we have extracted four different datasets, summarized in Table 2.

Using the datasets in Table 2, we fit six different models: one model for each of the *one school, one year* and *one school, all years* networks, and two models with different covariate configurations for each of the *all schools, one year* and *all schools, all years* networks. Table 2 summarizes these six models, their covariate formulations, and their resultant model fits. For all models, we assume a uniform prior on the entries in τ , a $\text{IG}(1,1)$ prior on γ^2 , and a latent dimension $d = 2$ for the latent positions. For models 1-3, we assume independent isotropic Gaussian priors on the latent positions. Models 4-6 follow the same nodes across multiple years, so we use an autoregressive Gaussian prior; each nodes' sequence of latent positions is assumed to have an autocorrelation of 0.95 a priori, and distinct nodes are assumed to be independent.

For all six models, we employed the same tuning strategy as in Study 1 (described in Section 6.2 of the Appendix) to choose the hyperparameters and random walk step sizes. Each Metropolis within Gibbs chain was run for 20,000 iterations, and each split HMC and split HMC + FlyMC




Figure 4. Boxplots depicting the relative efficiency (in terms of effective sample size per second) of split HMC (both with and without FlyMC) compared to Metropolis within Gibbs. The six different data/model configuration described above are considered. For each configuration, the relative speed-up in effective sample size per second is provided for computing the posterior log probability of a random subset of dyads in the network. Note that the horizontal axis is provided on the log scale.

chain was run for 10,000 iterations. For each algorithm and model, the relative speed-up is summarized using a boxplot for at most 500 randomly selected dyads in Figure 4.

For each of the six models, both the non-FlyMC and FlyMC versions of split HMC clearly outperform Metropolis within Gibbs, even more so than in Study 1. The speed-up is most pronounced in Model 6, with the algorithms being almost 1000 times more efficient than Metropolis within Gibbs. The two HMC algorithms perform comparably across the different models. The most noticeable difference is for Model 2 and Model 3. For Model 2, the version of split HMC without FlyMC performs better. For Model 3, the FlyMC version is the better performer.

Both of these models are fit to the same network, but they differ in that Model 2 does not use the same school indicator in its covariate. As such, their disparity provides a good case study for when FlyMC is most useful. Table 3 shows that for Model 3, $\hat{\tau}_{EAP} = 0.86$ for within-school dyads and $\hat{\tau}_{EAP} = 0.01$ for between-school dyads. This stark difference shows that the covariate in the link function has captured a sparsity of between-school edges. Meanwhile, Model 2 does not have access to the same school covariate, forcing it to use a single $\hat{\tau}_{EAP} = 0.43$ across all dyads. To compensate, Model 2’s fitted link function is much steeper than that of Model 3—the $\hat{\gamma}_{EAP}^2$ values are 0.15 and 1.18, respectively. Turning our attention to Figure 5, we can see how and why this compensation works.

Figure 5(a) and (b) summarize the estimated latent positions in Models 2 and 3, respectively, with Figure 5(c) and (d) illustrating the contours of the latent position posterior distributions for two individuals across the two models. Figure 5(a) shows that Model 2’s estimated latent positions are essentially acting as proxies for the absent same school covariate information, clustering the individuals in the latent space according to their school memberships. The steep decline of

Table 3. A summary of the six model configurations and fits for Study 2. For each model, the Dataset column reports the network data which the model is fit, the x_{ij1} and x_{ij2} columns report the indicators used to form the link function covariate, and the $\hat{\gamma}_{EAP}^2$ and $\hat{\tau}_{EAP}$ columns report the expectations a posteriori of link function parameters γ^2 and τ , respectively. Note that a different value of τ is fit for each level of the covariate

Model	Dataset	x_{ij1} (Same school)	x_{ij2} (Previous edge)	$\hat{\gamma}_{EAP}^2$	$\hat{\tau}_{EAP}$	
1	One year, one school	Excluded	Excluded	1.22	0.85	
2	One year, all schools	Excluded	Excluded	0.15	0.43	
3	One year, all schools	Included	Excluded	1.18	0.86	$x_{ij1} = 1$
					0.01	$x_{ij1} = 0$
4	all years, one school	Excluded	Included	1.00	0.96	$x_{ij2} = 1$
					0.47	$x_{ij2} = 0$
5	All years, all schools	Excluded	Included	0.18	0.95	$x_{ij2} = 1$
					0.27	$x_{ij2} = 0$
6	All years, all schools	Included	Included	1.09	0.86	$x_{ij1} = 1, x_{ij2} = 1$
					0.74	$x_{ij1} = 1, x_{ij2} = 0$
					0.32	$x_{ij1} = 0, x_{ij2} = 1$
					0.01	$x_{ij1} = 0, x_{ij2} = 0$

the link function thus allows it to capture the sparsity of between-school edges, with the few between-school ties in the network determining the relative positioning of the different school clusters. Meanwhile, the latent positions in Figure 5(b) show much less structure, with the posteriors in Figure 5(d) being far more diffuse than those in 5(d). This suggests that the majority of the network's structure is already captured by the covariate in Model 3. Accordingly, the edge probabilities decay more gradually in the latent space.

The difference in algorithm performances between Models 2 and 3 thus boils down to the difference between τ -driven sparsity versus γ^2 -driven sparsity. The superior performance of the FlyMC algorithm for Model 3 is due to its ability exploit the sparsity (captured by τ) between schools. The corresponding θ variables drastically downsample the number of dyads involved in each likelihood and gradient calculation, speeding up computation. This illustrates that FlyMC is especially useful when applied to networks containing mainly dyads for which the value of $\tau_{x_{ij}}$ is expected to be small. On the other hand, in models where the sparsity is driven γ^2 , such as Model 2, there is less benefit to applying FlyMC.

We now shift our focus to Models 5 and 6, which have been fit to large longitudinal networks. For these models, the particularly strong performance of the split HMC algorithms compared to Metropolis within Gibbs can be attributed to two factors: the size of the network and the extra structure imposed by the temporal autocorrelation in latent positions. The joint gradient-informed update of all nodes in split HMC is particularly well-suited to account for the autocorrelation in the prior, where the uninformed random walk update of Metropolis within Gibbs is not.

For Model 6, we see in Table 2 that the same school covariate once again captures most of the network's sparsity by allowing for $\hat{\tau}_{EAP} = 0.01$ in across-school dyads that did not previously have an edge. However, FlyMC does not appear to be as beneficial in this case. This is due partially to the slower mixing of the τ variables in the FlyMC implementations—although they facilitate a closed-form Gibbs update, the auxiliary θ variables can also lead to extra autocorrelation in the chains for τ as they slowly evolve. Another contributor to the smaller disparity is a larger gap between




Figure 5. The top panels depict point estimates of the teachers' latent positions for Models 2 and 3 (i.e. the one year, all schools models both with and without the same school covariate). These point estimates are obtained via multi-dimensional scaling on the posterior expectation of the matrix of squared latent distances. The shape/color combinations are assigned according to the teacher's school. The edges indicate information-sharing relationships. The bottom two panels depict posterior distribution contours for the latent positions for two teachers (one from school 1, one from school 14).

tuned values of ε . Our approach for tuning ε and L yielded $\varepsilon = 0.31$, $L = 7$ for the non-FlyMC implementation, and $\varepsilon = 0.14$ and $L = 15$ for the FlyMC implementation. In the corresponding one year models, the gap was not as large— $\varepsilon = 0.22$, $L = 10$ for the non-FlyMC implementation compared to $\varepsilon = 0.14$ and $L = 15$ for the FlyMC implementation.

5. Concluding remarks

In this article, we developed new algorithms for estimating Gaussian latent position models (GLPM, Equation (1)) using a variation of split HMC (Shahbaba et al., 2014) adapted to the link function of the model and a variation of Firefly Monte Carlo (Maclaurin & Adams, 2015) for subsampling non-edge dyads while keeping the Monte Carlo algorithm exact up to Monte Carlo error. We conducted two empirical studies to investigate the performance of the split HMC and split HMC + FlyMC approaches: one on synthetic data (Section 4.1) and one on real data concerning information-sharing among teachers and staff (Section 4.2).

Our synthetic data study demonstrated that when split HMC and split HMC + FlyMC are tuned to have an acceptance rate of around 0.8-0.85 (for integration time $T \approx 2$), these algorithms vastly outperform similarly well-tuned competitors in the literature, especially for large networks. The competitors we considered include standard algorithms such as Metropolis within Gibbs and elliptical slice sampling, as well as more sophisticated HMC algorithms such as the NUTS and Stan that adaptively set T . Across the board, our implementations of split HMC and split HMC + FlyMC outperformed all competitors, even the more sophisticated HMC algorithms. Of the two new algorithms proposed in this article, our implementation of split HMC seemed to be the better performer for denser networks, as well as smaller sparse networks for which the link function

parameters were poorly identified. Split HMC + FlyMC performs best on large, sparse networks that contain sufficient information to identify the link function parameters. Given these results, we would expect split HMC + FlyMC to perform especially well in settings for which τ , the maximum link probability in Equation (1), is very small, such as the sparse graphon version of GLPMs (Borgs et al., 2014; Spencer & Shalizi, 2019).

At first, we were surprised by the extent to which the HMC integration time T at approximately 2 outperformed the adaptive strategies for setting T used in NUTS and Stan. Adaptive strategies have been shown to outperform fixed integration strategies across a variety of models and perform at least comparably for others (Hoffman & Gelman, 2014; Betancourt, 2016). However, the relatively poor performance of these algorithms on LPMs is less surprising when one considers the criteria used by Stan and NUTS to choose T . Both algorithms are configured to optimize the mixing of the latent positions. However, as we elaborate on in Section 6.3 of the Appendix, the latent positions themselves are under-identified in the posterior—the edge probabilities or latent distances are more appropriate targets because they are identified in the model. This insight suggests that adapting the NUTS criterion to optimize the mixing of functions of the parameters could be a worthwhile extension.

Our real data study demonstrated that FlyMC also performs well when τ varies with a categorical covariate, taking on small values for some categories. Another potential avenue of future research would be to consider a hybrid of split HMC and split HMC + FlyMC. FlyMC would only be applied to the subset of dyads for which τ is expected to be small. FlyMC would thus be exploited only where it is most effective, limiting any potential mixing problems due to the extra auxiliary variables.

In our real data study, we also considered the application of our algorithm to fit Gaussian LPMs for longitudinal networks. Of the variety of different ways to configure longitudinal LPMs (Kim et al., 2018), we used a simple model structure that used a structured Gaussian prior to promote serial dependence of the nodes' latent positions across time points, as well as covariates to promote persistent edges across time points. For these models, our split HMC algorithm performed particularly well because gradient-informed proposals naturally accommodate the extra structure in the prior. Therefore, we anticipate that our insights could also be applied to achieve substantial computational gains in other more structured priors for LPMs, such as latent cluster model (Krivitsky et al., 2009) or the multi-resolution network model Fosdick et al. (2018). Recently, Turnbull (2020) explored the use of sequential Monte Carlo (Doucet & Johansen, 2009) for Bayesian inference of longitudinal latent position models. They found that the algorithm scaled poorly to large networks due to expensive likelihood evaluations. These findings suggest that a combination of their approach with the algorithms we present here could be fruitful.

Finally, it is worth briefly commenting on our motivation for focusing on the Gaussian link function instead of the traditional logistic link function. When proposing the GLPM, Rastelli et al. (2016) argued that the Gaussian LPM yields results that are analogous and comparable to those of the logistic LPM, while providing interpretable link function parameters and making it easier to derive theoretical results. Since then, Spencer & Shalizi (2019) proved consistent estimation results for the Gaussian link function—analogue results under the logistic link function remain an open problem. In this article, we have demonstrated additional benefits of the Gaussian link function over the logistic link function. The closed-form Gaussian decomposition underlying our split HMC algorithm is not possible with a logistic link, and the FlyMC strategy cannot be applied because it relies on a factorization of the sparsity parameter τ from the link function. Moreover, the likelihood of the logistic link LPM is not differentiable when two nodes have the same position, which complicates the application of HMC. For all of these reasons, we propose that the Gaussian LPM would be a more suitable “default” choice of the link function—the logistic link should only be used when its particular functional form is justified by the application.

Acknowledgments. We would like to acknowledge helpful suggestions and feedback from Robert Kass, Jared Murray, Cosma Shalizi, participants of the Networkshop research group at Carnegie Mellon University, the audience at the Statistical Inference in Network Models satellite symposium of Netsci 2019, seminar attendees at Dalhousie University, and those who visited our poster at the OBayes17 conference.

We would like to thank Alex Reinhart, Carl Skipper, posters on The Stan Forum, and Roberto Gomez for their help and guidance regarding how to design and run the empirical studies in a supercomputing environment. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).

All three authors gratefully acknowledge funding from IES (US Dept of ED) Award 305D150045, NAS gratefully acknowledges an NSERC Postgraduate Scholarship, and funding from National Institute of Mental Health grant R01MH064537.

Finally, we would like to acknowledge Jim Spillane and the Distributed Leadership Studies (<http://www.distributedleadership.org>) at Northwestern University, including the NebraskaMATH study, for the use of their data in this work. We greatly appreciate the help provided by Jim and his colleagues.

Competing interests. None.

Supplementary material

For supplementary material for this article, please visit <http://dx.doi.org/10.1017/nws.2022.1>

Notes

1 In this sense, the dimensions of Z behave like principal components in principal component analysis

2 Empirically, we have found that for LPMs, a Metropolis within Gibbs acceptance rate somewhere between 20% and 30% gives optimal results—this is consistent with related optimal scaling theory (Roberts et al., 2001)

3 In practice, the marginal distribution of Z (and not the joint distribution of Z, U) is the target of HMC, so this negation step can be omitted because it is immediately changed by subsequent Gibbs update of U (Neal, 2011).

References

- Airoldi, E. M., Blei, D. M., Fienberg, S. E., & Xing, E. P. (2008). Mixed membership stochastic blockmodels. *Journal of Machine Learning Research*, 9(Sep), 1981–2014.
- Aliverti, E., & 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.
- Bales, B., Pourzanjani, A., Vehtari, A., & Petzold, L. (2019). Selecting the metric in Hamiltonian Monte Carlo. *arXiv preprint arXiv:1905.11916*.
- Betancourt, M. (2016). Identifying the optimal integration time in Hamiltonian Monte Carlo. *arXiv preprint arXiv:1601.00225*.
- Betancourt, M. (2017). A conceptual introduction to Hamiltonian Monte Carlo. *arXiv preprint arXiv:1701.02434*.
- Bloem-Reddy, B., & Cunningham, J. (2016). Slice sampling on Hamiltonian trajectories. In *International Conference on Machine Learning* (pp. 3050–3058).
- Borgs, C., Chayes, J., Cohn, H., & Zhao, Y. (2014). An ℓ^p theory of sparse graph convergence i: Limits, sparse random graph models, and power law distributions. *Transactions of the American Mathematical Society*.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., . . . Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).
- Carrington, P. J., Scott, J., & Wasserman, S. (2005). *Models and Methods in Social Network Analysis*, vol. 28. Cambridge university press.
- Chao, W.-L., Solomon, J., Michels, D., & Sha, F. (2015). Exponential integration for Hamiltonian Monte Carlo. In *International Conference on Machine Learning* (pp. 1142–1151).
- Chen, L., Vogelstein, J. T., Lyzinski, V., & Priebe, C. E. (2016). A joint graph inference case study: the *c. elegans* chemical and electrical connectomes. In *Worm*, vol. 5, (p. e1142041). Taylor & Francis.
- Chiu, G. S., & Westveld, A. H. (2011). A unifying approach for food webs, phylogeny, social networks, and statistics. *Proceedings of the National Academy of Sciences*, 108(38), 15881–15886.
- Clauset, A., Moore, C., & Newman, M. E. (2008). Hierarchical structure and the prediction of missing links in networks. *Nature*, 453(7191), 98.
- Crane, H. (2018). *Probabilistic Foundations of Statistical Network Analysis*. Chapman and Hall/CRC.

- Dabbs, B., Adhikari, S., & Sweet, T. (2020). Conditionally independent dyads (CID) network models: a latent variable approach to statistical social network analysis. *Social Networks, Revision Under Review*.
- Doucet, A., & Johansen, A. M. (2009). A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering*, 12(656–704), 3.
- Duane, S., Kennedy, A. D., Pendleton, B. J., & Roweth, D. (1987). Hybrid Monte Carlo. *Physics letters B*, 195(2), 216–222.
- Erdős, P., & Rényi, A. (1960). On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5(1), 17–60.
- Fosdick, B. K., McCormick, T. H., Murphy, T. B., Ng, T. L. J., & Westling, T. (2018). Multiresolution network models. *Journal of Computational and Graphical Statistics* (pp. 1–12).
- Gamerman, D., & Lopes, H. F. (2006). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. CRC Press.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian Data Analysis*. CRC Press.
- Girolami, M., & Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2), 123–214.
- Goldenberg, A., Zheng, A. X., Fienberg, S. E., Airoldi, E. M., (2010). A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2(2), 129–233.
- Hahn, P. R., He, J., & Lopes, H. F. (2019). Efficient sampling for gaussian linear regression with arbitrary priors. *Journal of Computational and Graphical Statistics*, 28(1), 142–154.
- Handcock, M. S., Raftery, A. E., & Tantrum, J. M. (2007). Model-based clustering for social networks. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 170(2), 301–354.
- Hecker, M., Lambeck, S., Toepfer, S., Van Someren, E., & Guthke, R. (2009). Gene regulatory network inference: Data integration in dynamic models—a review. *Biosystems*, 96(1), 86–103.
- Hoff, P. D., Raftery, A. E., & Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460), 1090–1098.
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1), 1593–1623.
- Ji, P., & Jin, J. (2016). Coauthorship and citation networks for statisticians. *The Annals of Applied Statistics*, 10(4), 1779–1812.
- Kim, B., Lee, K. H., Xue, L., & Niu, X. (2018). A review of dynamic network models with latent variables. *Statistics Surveys*, 12, 105.
- Krivitsky, P. N., & Handcock, M. S. (2008). Fitting position latent cluster models for social networks with latentnet. *Journal of Statistical Software*, 24.
- Krivitsky, P. N., Handcock, M. S., Raftery, A. E., & Hoff, P. D. (2009). Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Social Networks*, 31(3), 204–213.
- Leimkuhler, B., & Reich, S. (2004). *Simulating Hamiltonian Dynamics*, vol. 14. Cambridge University Press.
- Linderman, S., Adams, R. P., & Pillow, J. W. (2016). Bayesian latent structure discovery from multi-neuron recordings. In *Advances in Neural Information Processing Systems* (pp. 2002–2010).
- Maclaurin, D., & Adams, R. P. (2015). Firefly Monte Carlo: Exact MCMC with subsets of data. In *International Joint Conference on Artificial Intelligence*.
- Mannseth, J., Kleppe, T. S., & Skaug, H. J. (2016). On the application of higher order symplectic integrators in Hamiltonian Monte Carlo. *arXiv preprint arXiv:1608.07048*.
- McFowland III, E., & Shalizi, C. R. (2021). Estimating causal peer influence in homophilous social networks by inferring latent locations. *Journal of the American Statistical Association* (pp. 1–12).
- Murray, I., Adams, R., & MacKay, D. (2010). Elliptical slice sampling. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (pp. 541–548).
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. L. Jones, & X.-L. Meng (Eds.), *Handbook of Markov chain Monte Carlo*, chap. 5. New York: CRC Press, pp. 113–162.
- Newman, M. E. (2002). Spread of epidemic disease on networks. *Physical Review E*, 66(1), 016128.
- Pakman, A., & Paninski, L. (2014). Exact Hamiltonian Monte Carlo for truncated multivariate gaussians. *Journal of Computational and Graphical Statistics*, 23(2), 518–542.
- Papaspiliopoulos, O., Roberts, G. O., & Sköld, M. (2007). A general framework for the parametrization of hierarchical models. *Statistical Science*, (pp. 59–73).
- Raftery, A. E., Niu, X., Hoff, P. D., & 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.
- Rastelli, R., Friel, N., & Raftery, A. E. (2016). Properties of latent variable network models. *Network Science*, 4(4), 407–432.
- Rastelli, R., Maire, F., & Friel, N. (2018). Computationally efficient inference for latent position network models. *arXiv preprint arXiv:1804.02274*.
- Roberts, G. O., Rosenthal, J. S., et al. (2001). Optimal scaling for various metropolis-hastings algorithms. *Statistical Science*, 16(4), 351–367.
- Salter-Townshend, M., & McCormick, T. H. (2017). Latent space models for multiview network data. *The Annals of Applied Statistics*, 11(3), 1217.

- Salter-Townshend, M., & Murphy, T. B. (2013). Variational Bayesian inference for the latent position cluster model for network data. *Computational Statistics & Data Analysis*, 57(1), 661–671.
- Shahbaba, B., Lan, S., Johnson, W. O., & Neal, R. M. (2014). Split Hamiltonian Monte Carlo. *Statistics and Computing*, 24(3), 339–349.
- Shortreed, S., Handcock, M. S., & Hoff, P. (2006). Positional estimation within a latent space model for networks. *Methodology*, 2(1), 24–33.
- Spencer, N. A., & Shalizi, C. R. (2019). Projective, sparse, and learnable latent position network models. *arXiv preprint arXiv:1709.09702*.
- Spillane, J. P., & Hopkins, M. (2013). Organizing for instruction in education systems and school organizations: How the subject matters. *Journal of Curriculum Studies*, 45(6), 721–747.
- Spillane, J. P., Hopkins, M., & Sweet, T. M. (2018). School district educational infrastructure and change at scale: Teacher peer interactions and their beliefs about mathematics instruction. *American Educational Research Journal*, 55(3), 532–571.
- Sweet, T., & Adhikari, S. (2020). A latent space network model for social influence. *Psychometrika*, (pp. 1–24).
- Sweet, T. M., Thomas, A. C., & Junker, B. W. (2013). Hierarchical network models for education research: Hierarchical latent space models. *Journal of Educational and Behavioral Statistics*, 38(3), 295–318.
- Turnbull, K. (2020). *Advancements in Latent Space Network Modelling*. Ph.D. thesis, Lancaster University.
- Xie, F., & Levinson, D. (2009). Modeling the growth of transportation networks: A comprehensive review. *Networks and Spatial Economics*, 9(3), 291–307.