



HAL
open science

Importance sampling-based gradient method for dimension reduction in Poisson log-normal model

Bastien Batardière, Julien Chiquet, Joon Kwon, Julien Stoehr

► **To cite this version:**

Bastien Batardière, Julien Chiquet, Joon Kwon, Julien Stoehr. Importance sampling-based gradient method for dimension reduction in Poisson log-normal model. 2024. hal-04712554v2

HAL Id: hal-04712554

<https://hal.science/hal-04712554v2>

Preprint submitted on 15 Nov 2024 (v2), last revised 18 Nov 2024 (v3)




HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - ShareAlike 4.0 International License

Importance sampling-based gradient method for dimension reduction in Poisson log-normal model

Bastien Batardière¹ , Julien Chiquet¹ ,
Joon Kwon¹ , and Julien Stoehr^{2,1} 

¹ *Université Paris-Saclay, AgroParisTech, INRAE, UMR MIA Paris-Saclay, 91120, Palaiseau, France, e-mail:*

bastien.batardiere@inrae.fr; julien.chiquet@inrae.fr; joon.kwon@inrae.fr

² *Université Paris-Dauphine, Université PSL, CNRS, CEREMADE, 75016 Paris, France, e-mail: stoehr@ceremade.dauphine.fr*

Abstract: High-dimensional count data poses significant challenges for statistical analysis, necessitating effective methods that also preserve explainability. We focus on a low rank constrained variant of the Poisson log-normal model, which relates the observed data to a latent low-dimensional multivariate Gaussian variable via a Poisson distribution. Variational inference methods have become a golden standard solution to infer such a model. While computationally efficient, they usually lack theoretical statistical properties with respect to the model. To address this issue we propose a projected stochastic gradient scheme that directly maximizes the log-likelihood. We prove the convergence of the proposed method when using importance sampling for estimating the gradient. Specifically, we obtain a rate of convergence of $O(T^{-1/2} + N^{-1})$ with T the number of iterations and N the number of Monte Carlo draws. The latter follows from a novel descent lemma for non convex L -smooth objective functions, and random biased gradient estimate. We also demonstrate numerically the efficiency of our solution compared to its variational competitor. Our method not only scales with respect to the number of observed samples but also provides access to the desirable properties of the maximum likelihood estimator.

MSC2020 subject classifications: Primary 00X00, 00X00; secondary 00X00.

Keywords and phrases: Dimension reduction, importance sampling, multivariate count data, Poisson log-normal model, projected stochastic gradient descent.

1. Introduction

Multivariate count data are prevalent in a widening range of applications such as ecology, genomics, microbiology, astronomy, and economy, just to name a few. This ubiquity has prompted the development of numerous statistical models, as unlike continuous multivariate distributions, a generic universal multivariate distribution for count data does not exist [27]. Most of the successful proposals are latent variable models belonging to the family of generalized multivariate mixed

models (GMMM). The latter offers the strength of model-based approaches, enabling the incorporation of external covariates and allowing the latent variables to be constrained in various ways to perform a specific task — regression, variable selection, and dimension reduction — while controlling the complexity of the model. These strengths contribute to the unwavering popularity of these models in the aforementioned fields of application, where both modeling and interpretability are essential prerequisites.

Significant milestones in the literature include a few generic frameworks initially developed through applications in ecology, where count and abundance tables have long been the norm. The generalized linear latent variable models (GLLVM) of [33] are instances of generalized multivariate mixed models with low-dimensional latent variables, where the distribution of observed responses usually belongs to the exponential family. The Hierarchical Modeling of Species Communities (HMSC), presented in [34], also falls within the class of generalized linear latent variable models with additional layers in the modeling of the latent variables. Multivariate Poisson Log-Normal models (PLN), as presented in [14], are yet another instance of the expansive family of generalized linear latent variable models. The latter confers the advantage of developing a generic and versatile framework capable of addressing various tasks, including dimension reduction, regression, clustering, discriminant analysis.

These modeling frameworks encounter the usual inference issues inherent to latent variable models. Specifically, direct and exact likelihood maximization is difficult since it requires evaluating an integral over a space of the latent variable dimension. Hence, direct numerical integration approaches [2] are limited to small-scale problems involving solely a few variables. Approaches based on Markov chain Monte Carlo (MCMC) techniques can handle medium-size problems but are computationally expensive [25, 42]. Alternatively, methods based on Laplace approximations exhibit greater computational efficiency but can potentially be inaccurate [21]. An indirect approach relies on the Expectation-Maximization (EM) algorithm, a well-established method for inference in incomplete data models since [16]. However, the M-step is practically intractable in GMMM, as it requires computing an expectation with respect to the distribution of the latent variable conditional on the observations. MCMC techniques can obviously be used in such a setting [*e.g.*, 30], though displaying the same shortcomings. More recently, the growing size of datasets and the porting of these methods to other fields of application where the number of variables expands drastically, such as genomics, have sparked interest in variational approaches [23, 5, 26, 33] as they provide a good compromise between accuracy and computational efficiency.

In the context of PLN models, which is the focus of this paper, [14] have extensively used this variational approach in conjunction with the EM algorithm and adapted it to several contexts, including dimension reduction, clustering, sparse covariance. The implementation provided in [15] is efficient and can deal with problems with thousands of observations and hundreds — even a couple of thousands — of variables. The resulting estimator can be shown to converge to the maximum of the surrogate likelihood function and enjoy asymptotical

normality [44]. However, these results pertain to the surrogate model and generally differ from the natural properties of an M-estimator associated with the likelihood [43]. In particular, while the maximum reached by the variational estimator seems at least empirically to coincide with the maximum likelihood estimator (MLE), there is no genuine estimator of the variance of the variational estimator that can be used to measure uncertainty properly: although the bootstrap method or the jackknife estimator could be used to build an estimator of the variance of the estimate, the variational solution is marred by the lack of relevant statistical guarantees. Consequently, the design of efficient algorithms that can directly maximize the likelihood and inherit the desirable properties of MLEs is still a key research issue for GMMM, particularly for PLN models. Such an algorithm allows a more direct assessment of estimator uncertainty by means of asymptotic variance estimates. It is in this spirit that [40] propose a variant of the Monte Carlo EM scheme that combines composite likelihood and importance sampling methods with a focus on applications in synecology. While the approach benefits from the properties of the maximum composite likelihood estimator, it necessitates splitting the data into overlapping blocks containing a small number of variables. The minimum number of blocks required grows quadratically in the number of variables (or species). As the computational complexity of their Monte Carlo EM increases linearly with the number of blocks, the solution is primarily suited for problems involving a few dozen variables but does not scale up efficiently to larger problems from a computational perspective.

Contributions This paper introduces a projected stochastic gradient descent (SGD) scheme based on self-normalized importance sampling to obtain gradient estimates for optimizing the marginal likelihood of the observed data in the Poisson log-normal model, subject to a rank constraint on the latent space. This model, introduced by [11], can be seen as a probabilistic version of Principal Component Analysis (PCA) with Poisson emission law, and its standard inference solution is a variational Expectation-Maximization (VEM) algorithm. Estimating parameters according to the maximum likelihood principle with Monte Carlo simulations is a long-standing idea for an unnormalized statistical model — a class of challenging models due to their intractable partition function which is a highly multidimensional integral depending on the parameters. For instance, Monte Carlo maximum likelihood estimation [19] uses importance sampling to estimate the partition function while contrastive divergence [24] estimates the gradient of the log partition function via Monte Carlo methods. More recently, the noise-contrastive estimation [22] reformulates the initial problem to avoid estimating the partition function or its gradient. Here, we rely on the fact that incomplete data models share similarities with the unnormalized models in that, under mild regularity conditions, the inference resumes to deal with an intractable integral, namely the score function for the observed likelihood [31]. We show that the PLN-PCA model falls within the set of incomplete data models for which the score function is written as an expected value with respect to the conditional distribution of the latent given the observed and can

thereby be estimated by simulation methods. The rank constraint ensures that the importance sampling estimator can handle problems with up to thousands of variables in the emission space, provided that the dimension of the latent space is controlled and limited to some tens. The projection step onto a convex compact set specifically guarantees that the objective function is L -smooth. We also show that it ensures a bounded mean squared error and bias for the gradient estimator. Such properties are common in the literature [20, 32, 38].

Our major contribution is a novel convergence theorem for the gradient method presented. To establish the result within the context of a self-normalized importance sampling estimator, we first present a general descent lemma applicable under minimal assumptions — specifically, L -smoothness, and bounded bias and quadratic error for the gradient estimator. To the best of our knowledge, it is the first result on projected stochastic gradient schemes for potentially both non convex objective functions and random biased gradient estimators. Given T iterations of our method and N Monte Carlo draws, we obtain a theoretical rate of $O(T^{-1/2} + N^{-1})$ for the gradient mapping norm. This convergence rate is consistent with those demonstrated in the literature, albeit in different contexts. For a non-convex setting with an unbiased gradient estimator, Ghadimi and Lan [20] derive an $O(T^{-1/2})$ convergence rate for the gradient’s norm objective, while Mai and Johansson [32] achieve the same rate for the gradient mapping norm but when a projection step is added. In situations where the gradient estimator is biased, with a bound b on the bias, Ajalloeian and Stich [3] retrieve a rate of $O(T^{-1/2} + b)$, but only when no projection step is performed.

The paper also includes an efficient implementation of our algorithm using JAX library [7], and GPU computing. As a by-product, we provide a PyTorch [36] version of the VEM solution, enabling the analysis of large-scale datasets with hundreds of thousands of observations and tens of thousands of variables.

Outline The paper begins with an introduction to the standard multivariate PLN model and its PCA version in Section 3. We then present our stochastic gradient scheme with convergence guarantees in Section 4. In Section 5, we propose a sequentially adapted Gaussian mixture distribution to serve as a valid importance sampling proposal distribution within our algorithm. Finally, Section 6 details a simulation study on synthetic data and an application to genomic data, where we deal with the problem of dimension reduction and visualization of a transcriptomic single-cell dataset. Technical details are postponed till Appendices A–D.

2. Notations and conventions

Let p and q be positive integers. The vector space of all $p \times q$ -matrices over a ring \mathbb{A} is denoted by $\mathcal{M}_{p \times q}(\mathbb{A})$. The subset of all symmetric, positive and definite $p \times p$ -matrices over \mathbb{R} is denoted by \mathcal{S}_{++}^p . We denote by $\langle \cdot, \cdot \rangle$ the scalar product on a real p -space \mathbb{R}^p , and $\|\cdot\|$ its associated norm. The matrix norm induced by $\|\cdot\|$ on $\mathcal{M}_{p \times p}(\mathbb{R})$ is also denoted by $\|\cdot\|$. $\text{Diag}(x)$ is a diagonal matrix with

diagonal equal to x for x a vector. When applied to matrices or vectors, simple functions like log, exp or square apply element-wise. The vector full of zeros of size p is denoted $\mathbf{0}_p$.

We denote by $\mathbb{M}_1(\mathbb{R}^p)$ the set of probability measures on \mathbb{R}^p . Given a probability measure $\pi \in \mathbb{M}_1(\mathbb{R}^p)$, \mathbb{M}_π is the set of probability measures that dominates π . The product measure $\prod_{i=1}^n \pi$ on $\mathbb{R}^{d \times n}$ is denoted by $\pi^{\otimes n}$. We use the same notation to refer to a measure and its associated density, meaning that if π is absolutely continuous with respect to the Lebesgue measure λ , $\pi(dx) = \pi(x)\lambda(dx)$. The expectation with respect to π is denoted by \mathbb{E}_π . When there is no ambiguity regarding the integration measure, we simply use the notation \mathbb{E} .

We denote by $\mathcal{N}(\boldsymbol{\mu}, \mathbf{S})$ a p -dimensional Gaussian variable with mean $\boldsymbol{\mu} \in \mathbb{R}^p$ and variance $\mathbf{S} \in \mathcal{S}_{++}^p$ and $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{S})$ its density evaluated at $\mathbf{x} \in \mathbb{R}^p$. We denote by $\mathcal{P}(\lambda)$ a Poisson variable with rate $\lambda > 0$.

The Kullback–Leibler divergence between $\pi \in \mathbb{M}_1(\mathbb{R}^p)$ and $\mu \in \mathbb{M}_\pi$ is defined by

$$\text{KL}(\pi \parallel \mu) = \int \log \frac{\pi(x)}{\mu(x)} \pi(dx).$$

The entropy of a random variable X distributed according to $\pi \in \mathbb{M}_1(\mathbb{R}^p)$ is defined by

$$\mathcal{H}_\pi(X) = - \int \log \pi(x) \pi(dx).$$

Given $\mathcal{X} \subseteq \mathbb{R}^p$, a differentiable function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is said to be L -smooth on \mathcal{X} with $L \geq 0$ if its gradient is L -Lipschitz on \mathcal{X} , namely, for any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathcal{X}$,

$$\|\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) - \nabla_{\boldsymbol{\theta}'} f(\boldsymbol{\theta}')\| \leq L \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|.$$

3. Dimension reduction in multivariate Poisson log-normal models

Background: Multivariate Poisson log-normal model Consider a data matrix $\mathbf{Y} = (Y_{ij}) \in \mathcal{M}_{n \times p}(\mathbb{N})$ storing n i.i.d. observations of a p -dimensional random vector $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{ip}) \in \mathbb{N}^p$, $i = 1, \dots, n$. The multivariate Poisson lognormal model [see 2, for its original formulation] relates each of the observed vector \mathbf{Y}_i to a latent (or unobserved) p -dimensional Gaussian vector \mathbf{Z}_i , whose covariance matrix $\boldsymbol{\Sigma}$ describes the underlying structure of dependence between the p variables. Following a formalism similar to that of GMMM, the model can also encompass a main effect due to a linear combination of d observed covariates $\mathbf{x}_i \in \mathbb{R}^d$ (including a vector of intercepts), and some possible user-specified offsets $\mathbf{o}_i = (o_{ij}) \in \mathbb{R}^p$ to take into account the sampling efforts between the samples. The model assumes that the observations Y_{ij} are independent conditionally on $\mathbf{Z}_i = (Z_{ij})$, and that the conditional distribution $p(Y_{ij} \mid Z_{ij})$ is a Poisson distribution, namely,

$$\begin{aligned} \{\mathbf{Z}_i\}_{1 \leq i \leq n} \text{ i.i.d. : } & \quad \mathbf{Z}_i \sim \mathcal{N}(\mathbf{0}_q, \boldsymbol{\Sigma}); \\ \{Y_{ij}\}_{\substack{1 \leq i \leq n \\ 1 \leq j \leq p}} \text{ ind. } \mid \{\mathbf{Z}_i\}_{1 \leq i \leq n} : & \quad Y_{ij} \mid Z_{ij} \sim \mathcal{P}(\exp(Z_{ij} + \mathbf{B}_j^\top \mathbf{x}_i + o_{ij})), \end{aligned} \quad (1)$$

where $\mathbf{B} = [\mathbf{B}_1, \dots, \mathbf{B}_p] \in \mathcal{M}_{d \times p}(\mathbb{R})$ is a latent matrix of regression parameters. In this framework, the main goal is to estimate the vector of parameters $\boldsymbol{\theta} = (\mathbf{B}, \boldsymbol{\Sigma}) \in \mathbb{R}^D$, with $D = dp + p(p+1)/2$, from the data matrices \mathbf{Y} and \mathbf{X} .

Poisson lognormal-model with low-rank constraint Throughout this paper, we focus on the PCA variant of model (1) as introduced in [11]. The latter is derived by adding a rank constraint on the latent covariance matrix, such that $\text{rank}(\boldsymbol{\Sigma}) = q < p$. The constraint alleviates the number of parameters to estimate, which can become prohibitive when the number of variables p is large in (1). This key feature is particularly relevant in the perspective of importance sampling. Indeed, this allows us to deal with problems where the dimension of the observation space is potentially large, in contrast with the number of parameters in the model and the dimension of the latent space, where the particles are sampled. The PCA version can be written in a hierarchical framework by adding a layer with, for each individual, a q -dimensional standard Gaussian vector \mathbf{W}_i , and introducing an individual-specific linear function $f_i : \mathbb{R}^q \rightarrow \mathbb{R}^p$ defined for all $\mathbf{w} \in \mathbb{R}^q$ as

$$f_i(\mathbf{w}; \mathbf{B}, \mathbf{C}) = \mathbf{C}\mathbf{w} + \mathbf{B}^\top \mathbf{x}_i + \mathbf{o}_i,$$

where $\mathbf{C} \in \mathcal{M}_{p \times q}(\mathbb{R})$ encodes the embedding of the observations into a space of lower dimension. The model is then written as

$$\begin{aligned} \{\mathbf{W}_i\}_{1 \leq i \leq n} \text{ i.i.d.} : & \quad \mathbf{W}_i \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q); \\ \{\mathbf{Z}_i\}_{1 \leq i \leq n} \text{ i.i.d.} : & \quad \mathbf{Z}_i = \mathbf{C}\mathbf{W}_i + \mathbf{B}^\top \mathbf{x}_i + \mathbf{o}_i; \\ \{Y_{ij}\}_{\substack{1 \leq i \leq n \\ 1 \leq j \leq p}} \text{ ind.} \mid \{\mathbf{Z}_i\}_{1 \leq i \leq n} : & \quad p(Y_{ij} \mid Z_{ij}) = \frac{\exp\{Y_{ij}Z_{ij} - \exp(Z_{ij})\}}{Y_{ij}!}. \end{aligned} \quad (2)$$

We refer to Model (2) as PLN-PCA. Adopting the PCA terminology, \mathbf{C} is the $p \times q$ matrix of loadings, and \mathbf{W}_i represents the vector of scores of the i -th observation in the low-dimensional latent subspace, whose dimension q corresponds to the number of components.

The vector of unknown parameters to be estimated is now $\boldsymbol{\theta} = (\mathbf{B}, \mathbf{C}) \in \mathbb{R}^d$, with $d = p(q+m)$ that is significantly smaller than D when $q \ll p$, the typical case at play in a context of dimension reduction. The complete log-likelihood of Model (2) can be written, up to additive constants with respect to model parameters, as

$$\sum_{i=1}^n \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{W}_i) = \sum_{i=1}^n \left\{ \langle \mathbf{Y}_i, \mathbf{Z}_i \rangle - \sum_{j=1}^p \exp(Z_{ij}) - \frac{1}{2} \|\mathbf{W}_i\|^2 \right\}.$$

Without any further assumption on \mathbf{C} , remark that $\boldsymbol{\theta}$ is not identifiable since the distribution of \mathbf{Z}_i is invariant when multiplying \mathbf{C} by any orthogonal matrix. However, since \mathbf{Y}_i depends on \mathbf{C} solely through the covariance $\boldsymbol{\Sigma} = \mathbf{C}\mathbf{C}^\top$, it is enough to have the identifiability for the model parametrized by $(\mathbf{B}, \boldsymbol{\Sigma})$. We will discuss this point further when introducing the gradient ascent algorithm.

4. Biased stochastic gradient descent

4.1. Solution to maximum likelihood principle

Estimation of θ is achieved by maximizing the likelihood of the observed data $p_{\theta}(\mathbf{Y})$, or equivalently by solving the optimization problem

$$\theta^{\text{MLE}} = \arg \min_{\theta \in \mathcal{X}} \ell(\theta), \quad \ell(\theta) = -\frac{1}{n} \sum_{i=1}^n \log p_{\theta}(\mathbf{Y}_i), \quad (3)$$

where ℓ is referred to as the loss function. Optimizing such a function is not straightforward because the marginal $p_{\theta}(\mathbf{Y}_i)$ requires integrating out the latent variable \mathbf{W}_i . We refer to θ^{MLE} as the maximum likelihood estimator.

Reminder on the EM approach The Expectation-Maximization algorithm [16] circumvent this issue by using a decomposition of the log-likelihood of the observed data \mathbf{Y}_i into

$$\begin{aligned} \log p_{\theta}(\mathbf{Y}_i) &= \int_{\mathbb{R}^q} \log p_{\theta}(\mathbf{Y}_i, \mathbf{w}) p_{\theta}(\mathrm{d}\mathbf{w} \mid \mathbf{Y}_i) - \int_{\mathbb{R}^q} \log p_{\theta}(\mathbf{w} \mid \mathbf{Y}_i) p_{\theta}(\mathrm{d}\mathbf{w} \mid \mathbf{Y}_i) \\ &= \mathbb{E} [\log p_{\theta}(\mathbf{Y}_i, \mathbf{W}_i) \mid \mathbf{Y}_i, \theta] + \mathcal{H}_{p_{\theta}(\cdot \mid \mathbf{Y}_i)}(\mathbf{W}_i), \end{aligned}$$

The algorithm proceeds by evaluating the conditional expectation of the complete log-likelihood using the current estimates $\theta^{(t)}$ of the model parameters:

$$Q(\theta \mid \theta^{(t)}) = \sum_{i=1}^n \mathbb{E} [\log p_{\theta}(\mathbf{Y}_i, \mathbf{W}_i) \mid \mathbf{Y}_i, \theta^{(t)}].$$

By iteratively maximizing this quantity, the algorithm generates a sequence that converges under suitable regularity conditions to the maximum likelihood estimator θ^* [45, 6]. However, the conditional $p_{\theta}(\mathbf{W}_i \mid \mathbf{Y}_i)$ is intractable for the PLN model and its PCA extension. To address this challenge, variational inference approximates $p_{\theta}(\mathbf{W}_i \mid \mathbf{Y}_i)$ with a surrogate distribution $\phi_{\psi}(\mathbf{V})$ from a parametric family \mathcal{P} . For instance, [12, 13] resorted to multivariate Gaussian distributions $\mathcal{N}(\mathbf{m}_i^{\text{VEM}}, \mathbf{S}_i^{\text{VEM}})$ with diagonal covariance matrix. The Variational EM method alternates between updates of the variational parameter ψ and the model parameter θ , aiming to maximize a lower bound of the log-likelihood, defined as

$$\log p_{\theta}(\mathbf{Y}_i) - \text{KL} [\phi_{\psi}(\mathbf{W}_i) \parallel p_{\theta}(\mathbf{W}_i \mid \mathbf{Y}_i)] = \mathbb{E}_{\phi_{\psi}} [\log p_{\theta}(\mathbf{Y}_i, \mathbf{V})] + \mathcal{H}_{\phi_{\psi}}(\mathbf{V}).$$

At convergence, the variational solution is associated with

$$\phi_{\psi} = \arg \min_{\phi \in \mathcal{P}} \text{KL} [\phi \parallel p_{\theta}(\cdot \mid \mathbf{Y}_i)]. \quad (4)$$

Stochastic gradient scheme In contrast to existing methods, we propose to address the optimization problem (3) directly with an SGD scheme. Our approach leverages that the Louis principle applies to Model (2), as outlined below (see Appendix A for a proof).

Proposition 4.1. *For all individual $i = 1, \dots, n$, the incomplete log-likelihood $\theta \mapsto \log p_\theta(\mathbf{Y}_i)$ of Model (2) is twice continuously differentiable on \mathbb{R}^d and its score function can be written as*

$$s_i(\theta) = \int_{\mathbb{R}^q} \nabla_\theta \log p_\theta(\mathbf{Y}_i, \mathbf{w}) p_\theta(d\mathbf{w} \mid \mathbf{Y}_i) = \mathbb{E}[\nabla_\theta \log p_\theta(\mathbf{Y}_i, \mathbf{W}_i) \mid \mathbf{Y}_i, \theta]. \quad (5)$$

While the intractability of the conditional distribution renders exact calculation infeasible even for a small value of q , identity (5) is instrumental in estimating the score function with Monte Carlo methods and designing an SGD scheme. Given a learning rate $\gamma \in \mathbb{R}_+^*$ and an initial point $\theta^{(1)} \in \mathbb{R}^d$, the SGD scheme recursively defines a sequence $\{\theta^{(t)}\}_{t \in \mathbb{N}^*}$ through the equation

$$\theta^{(t+1)} = \theta^{(t)} - \gamma \hat{g}^{(t)}, \quad (6)$$

where $\hat{g}^{(t)}$ is a possibly biased estimator of $\nabla_\theta \ell(\theta^{(t)})$. Here we explore the opportunity of importance sampling methods [28, 29] to define $\hat{g}^{(t)}$. Indeed, the lack of closed-form for $p_\theta(\mathbf{W}_i \mid \mathbf{Y}_i)$ hinders Monte Carlo methods that rely on exact samples from it. However, importance sampling overcomes this difficulty by changing the integration measure. For any $\theta \in \mathbb{R}^d$, it approximates $p_\theta(\cdot \mid \mathbf{Y}_i)$ with a random probability measure based on weighted samples from a probability density function $\nu_i(\cdot; \theta)$, possibly depending on θ and referred to as proposal distribution, such that $p_\theta(\cdot \mid \mathbf{Y}_i)$ is absolutely continuous with respect to $\nu_i(\cdot; \theta)$. The importance sampling method is then based on the following identity

$$s_i(\theta) = \int_{\mathbb{R}^q} \frac{p_\theta(\mathbf{v} \mid \mathbf{Y}_i)}{\nu_i(\mathbf{v}; \theta)} \nabla_\theta \log p_\theta(\mathbf{Y}_i, \mathbf{v}) \nu_i(d\mathbf{v}; \theta).$$

To circumvent the issue of evaluating the intractable distribution in the above, the method leverages a tractable non-normalized version of the conditional distribution, namely the joint distribution. Let us introduce the Radon-Nikodym derivative of $p_\theta(\mathbf{Y}_i, \cdot)$ with respect to $\nu_i(\cdot; \theta)$:

$$\rho_{\theta,i}(\mathbf{v}) = \frac{p_\theta(\mathbf{Y}_i, \mathbf{v})}{\nu_i(\mathbf{v}; \theta)}, \quad \mathbf{v} \in \mathbb{R}^q.$$

The score can be written as

$$s_i(\theta) = \int_{\mathbb{R}^q} \rho_{\theta,i}(\mathbf{v}) \nabla_\theta \log p_\theta(\mathbf{Y}_i, \mathbf{v}) \nu_i(d\mathbf{v}; \theta) \Big/ \int_{\mathbb{R}^q} \rho_{\theta,i}(\mathbf{v}) \nu_i(d\mathbf{v}; \theta).$$

Thereby, given $N \in \mathbb{N}^*$ independent samples $\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,N}$ from $\nu_i(\cdot; \theta)$, the self-normalized importance sampling (SNIS) estimator of $s_i(\theta)$ is

$$\hat{s}_i^N(\theta) = \sum_{r=1}^N \rho_{\theta,i}(\mathbf{v}_{i,r}) \nabla_\theta \log p_\theta(\mathbf{Y}_i, \mathbf{v}_{i,r}) \Big/ \sum_{s=1}^N \rho_{\theta,i}(\mathbf{v}_{i,s}). \quad (7)$$

Algorithm 1: Importance Sampling based Gradient Descent (ISGD)

Input: initial point $\theta^{(1)}$, learning rate γ , number of iterations T , number of Monte Carlo draws N .
Output: the sequence $\theta^{(1)}, \dots, \theta^{(T+1)}$
for $t = 1$ **to** T **do**
 Sample i uniformly in $\{1, \dots, n\}$;
 Sample $(\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,N})$ from $\nu_i^{\otimes N}(\cdot; \theta^{(t)})$;
 Update $\theta^{(t+1)} = P_{\Theta}(\theta^{(t)} - \gamma \hat{g}^{(t)})$ with $\hat{g}^{(t)} = -\hat{s}_i^N(\theta^{(t)})$ as in Equation (7);
end

A possible solution is then to define $\hat{g}^{(t)}$ by averaging the estimator (7) across individuals. In what follows, we rely on a mini-batch strategy where we use a single individual at a time. Specifically, at iteration $t \in \mathbb{N}^*$, we draw an individual $i(t)$ uniformly in $\{1, \dots, n\}$ and the gradient estimator within the update Equation (6) is

$$\hat{g}^{(t)} = -\hat{s}_{i(t)}^N(\theta^{(t)}). \quad (8)$$

4.2. Convergence guarantees

The SNIS estimator (7) is strongly consistent, but exhibits bias for a fixed sample size N . As shown in this section, controlling this bias is essential to ensure the convergence of our algorithm. This requires imposing further constraints on the optimization problem. In the remainder of the paper, let $\Theta \subset \mathbb{R}^d$ be a nonempty, compact, and convex set. In place of the standard update (6), we employ a projected SGD algorithm, which is defined by

$$\theta^{(t+1)} = P_{\Theta}(\theta^{(t)} - \gamma \hat{g}^{(t)}), \quad (9)$$

where P_{Θ} stands for the orthogonal projection on Θ . Algorithm 1 summarizes the overall scheme. Formally, the projection step ensures that all iterates remain bounded, which in turn guarantees the L -smoothness and the bounded gradient conditions commonly assumed in the literature to establish convergence properties.

Preliminary result for L -smooth functions In the following, we delve into the analysis of an SGD algorithm (i) for non-convex and constrained optimization with L -smooth objective, (ii) with a gradient estimator that is both biased and random. To the best of our knowledge, no established convergence theorem exists for these conditions. Mai and Johansson [32, Lemma 3.2] achieve a similar result but for unbiased gradient estimator. Let $F : \mathbb{R}^d \rightarrow \mathbb{R}$ be defined by

$$F(\theta) = \ell(\theta) + I_{\Theta}(\theta), \quad I_{\Theta}(\theta) = \begin{cases} 0 & \text{if } \theta \in \Theta, \\ \infty & \text{if } \theta \notin \Theta. \end{cases} \quad (10)$$

For a real $\eta > 0$, its Moreau envelope is

$$F_\eta(\boldsymbol{\theta}) = \inf_{\boldsymbol{\theta}' \in \mathbb{R}^d} \left\{ F(\boldsymbol{\theta}') + \frac{1}{2\eta} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 \right\}.$$

The Moreau envelope is useful from an optimization perspective, as it has the same set of minimizers as F , while being differentiable, unlike F . This characteristic is instrumental in analyzing the convergence of proximal gradient methods such as the one from Equation (9). Our following result yields control over the Moreau envelope, providing we have a L -smooth loss function and a gradient estimator $\widehat{g}^{(t)}$ with bounded mean squared error and bias (see Appendix B for proof).

Lemma 4.2 (Descent lemma). *Let consider the gradient scheme as defined by Equation (9). Assume that*

- (i) *the function ℓ is L -smooth on Θ , and denote $\Gamma = \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta})\|$;*
- (ii) *for $t \in \mathbb{N}^*$, one has*

$$\begin{aligned} \sigma^{(t)} &= \mathbb{E} \left[\left\| \widehat{g}^{(t)} - \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \right\|^2 \mid \boldsymbol{\theta}^{(t)} \right] < \infty, \\ \xi^{(t)} &= \left\| \mathbb{E} \left[\widehat{g}^{(t)} \mid \boldsymbol{\theta}^{(t)} \right] - \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \right\| < \infty. \end{aligned}$$

Then, for any real constant $\eta \in (0, 1/\max\{2\Gamma + L, 2L\}]$,

$$\begin{aligned} \mathbb{E} \left[F_\eta(\boldsymbol{\theta}^{(t+1)}) \mid \boldsymbol{\theta}^{(t)} \right] &\leq F_\eta(\boldsymbol{\theta}^{(t)}) - \frac{\gamma}{2} \|\nabla F_\eta(\boldsymbol{\theta}^{(t)})\|^2 \\ &\quad + \frac{\gamma + \gamma^2 \Gamma}{\eta} \xi^{(t)} + \frac{\gamma^2}{2\eta} (\sigma^{(t)} + \Gamma^2). \end{aligned} \quad (11)$$

Remark that Γ is finite as the gradient is L -Lipschitz continuous on a bounded set. Consequently the interval $(0, 1/\max\{2\Gamma + L, 2L\}]$ is nonempty.

Applicability to the PLN-PCA model We now demonstrate that the loss function for the PLN-PCA model and the importance sampling estimate (8) satisfy the assumptions from Lemma 4.2.

Proposition 4.3. *Under Model (2), for any nonempty compact subset $\Theta \subset \mathbb{R}^d$, we have that*

- (i) *for any individual $i = 1, \dots, n$, there exists a real*

$$0 < \zeta_i = \inf_{\boldsymbol{\theta} \in \Theta} p_{\boldsymbol{\theta}}(\mathbf{Y}_i);$$

- (ii) *there exists a real $L \geq 0$ such that the objective function ℓ , as defined in (3), is L -smooth on Θ .*

Regarding Assumption (ii) of Lemma 4.2, we should note that both $\sigma^{(t)}$ and $\xi^{(t)}$ are driven by, respectively, the mean squared error and the bias of the importance sampling estimate (7). Agapiou et al. [1, Theorem 2.3] yield

sufficient conditions to have control over the bias and error. The latter led to the subsequent assumptions on the proposal distribution for our optimization problem.

Assumption 1. For all individual $i \in \{1, \dots, n\}$, the proposal distribution ν_i is chosen such that

$$\lambda_i = \sup_{(\boldsymbol{\theta}, \mathbf{V}) \in \Theta \times \mathbb{R}^q} \rho_{\boldsymbol{\theta}, i}(\mathbf{V}) < \infty,$$

$$\beta_i = \sup_{\boldsymbol{\theta} \in \Theta} \mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[\|\nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{V})\|_1^4 \right] < \infty.$$

Following the work of [1], we obtain finite bound when integrating the mean squared error and the bias of the importance sampling estimate (7) with respect to the random mini-batch index $i(t)$ (see Appendix C for proof).

Proposition 4.4. Let $\{i(t)\}_{t \in \mathbb{N}^*}$ and $\{\boldsymbol{\theta}^{(t)}\}_{t \in \mathbb{N}^*}$ be the associated random sequences generated by Algorithm 1. Under Model (2), if Assumption 1 holds, then, for all $t \in \mathbb{N}$,

$$\bar{\sigma}_{\text{IS}}^{(t)} = \mathbb{E} \left[\left\| \widehat{\mathbf{s}}_{i(t)}^N(\boldsymbol{\theta}^{(t)}) - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\|^2 \middle| \boldsymbol{\theta}^{(t)} \right] \leq \frac{d}{N} M_{\sigma},$$

$$\bar{\xi}_{\text{IS}}^{(t)} = \mathbb{E} \left[\left\| \mathbb{E} \left[\widehat{\mathbf{s}}_{i(t)}^N(\boldsymbol{\theta}^{(t)}) - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \middle| \boldsymbol{\theta}^{(t)}, i(t) \right] \right\|^2 \middle| \boldsymbol{\theta}^{(t)} \right] \leq \frac{\sqrt{d}}{N} M_{\xi},$$

where M_{σ} and M_{ξ} are two finite and positive constants given by

$$M_{\sigma} = \frac{12}{n} \sum_{i=1}^n \frac{\lambda_i^2 \sqrt{\beta_i}}{\zeta_i^2} \left(1 + \frac{250\lambda_i}{\zeta_i} + \frac{9\lambda_i^2}{\zeta_i^2} \right),$$

$$M_{\xi} = \frac{4}{n} \sum_{i=1}^n \frac{\lambda_i^2 \beta_i^{1/4}}{\zeta_i^2} \left\{ 2 + \sqrt{3 \left(1 + \frac{250\lambda_i}{\zeta_i} + \frac{9\lambda_i^2}{\zeta_i^2} \right)} \right\},$$

with λ_i and β_i as in Assumption 1.

As detailed in the proof of Theorem 4.5 in Appendix C, this result implies the conditions required on $\sigma^{(t)}$ and $\xi^{(t)}$, since there is a constant A such that

$$\sigma^{(t)} \leq A \left(\bar{\sigma}_{\text{IS}}^{(t)} + \bar{\xi}_{\text{IS}}^{(t)} + 1 \right), \quad \text{and} \quad \xi^{(t)} \leq \bar{\xi}_{\text{IS}}^{(t)}.$$

Convergence of the gradient mapping As the local minimum of the loss function may lie outside the compact set Θ , we cannot prove that the norm of $\nabla \ell(\boldsymbol{\theta})$ becomes arbitrarily small within Algorithm 1. Instead, in the context of gradient methods incorporating a projection step, the convergence rate is characterized in terms of the norm of the gradient mapping $G_{\eta}^{(t)}$ [18] defined for any real $\eta > 0$ by

$$G_{\eta}^{(t)} = \frac{\boldsymbol{\theta}^{(t)} - P_{\Theta}(\boldsymbol{\theta}^{(t)} - \eta \nabla \ell(\boldsymbol{\theta}^{(t)}))}{\eta}. \quad (12)$$

This mapping is a tailored gradient objective, specifically modified to handle the projection step, whose norm is equivalent to the gradient norm of the Moreau envelope (see Lemma B.3). Our next result shows that the norm of the gradient mapping for Algorithm 1 can be rendered arbitrarily small, provided we use a sufficiently large number of iterations T and particles N .

Theorem 4.5. *Let $\boldsymbol{\theta}^{(1)} \in \Theta$ be an initial value and $\gamma_0 \in \mathbb{R}_+^*$ a user-specified initial learning rate. Under Model (2), if Assumption 1 holds, then for any $T \in \mathbb{N}^*$ and any real constant $\eta \in (0, 1/\max\{2\Gamma, L\}]$, the sequence $\{\boldsymbol{\theta}^{(t)}\}_{1 \leq t \leq T}$ defined by Algorithm 1 with $\gamma = \gamma_0/\sqrt{T}$ satisfies*

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\|G_\eta^{(t)}\|^2 \right] &\leq \frac{2\tau}{\gamma_0(L\eta + 1)\sqrt{T}} \left(\ell(\boldsymbol{\theta}^{(1)}) - \ell(\boldsymbol{\theta}^{MLE}) \right. \\ &\quad \left. + \frac{\gamma_0^2(L\eta + 1)}{2\eta} \left[\Delta^2 + \Gamma^2 + \frac{d}{N} \left\{ M_\sigma + \frac{2(\Delta + \Gamma)}{\sqrt{d}} M_\xi \right\} \right] \right) \\ &\quad + \frac{2\tau\sqrt{d}}{\eta N} M_\xi, \end{aligned} \tag{13}$$

with L the smoothness constant of $\ell(\cdot)$, $\Gamma = \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla \ell(\boldsymbol{\theta})\|$, constants M_σ and M_ξ as defined in Proposition 4.4, and

$$\tau = \frac{(2L\eta + 1)^2}{L\eta + 1} \left(1 + \sqrt{\frac{L\eta}{L\eta + 1}} \right)^2, \quad \Delta = \max_{i=1, \dots, n} \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i)\|.$$

Remark 1. While the upper bound $1/\max\{2\Gamma, L\}$ is unknown, Algorithm 1 does not depend on the choice of η in practice, and therefore neither Γ nor L need to be estimated.

5. Importance sampling proposal choice

In this section, we propose a specific choice of proposal distribution that satisfies Assumption 1, namely a mixture distribution. Mixture distributions are often chosen for this task [e.g., 9] because of their flexibility as parametric models. In what follows, we leverage the model structure and focus on two-component Gaussian mixture distributions. Given two real constants $\alpha \in [0, 1]$ and $\delta > 0$, denote the two-component Gaussian mixture with mean $\boldsymbol{\mu} \in \mathbb{R}^q$ and covariance $\mathbf{S} \in \mathcal{S}_{++}^q$ by

$$\mathcal{GM}(\cdot; \boldsymbol{\mu}, \mathbf{S}, \alpha, \delta) = (1 - \alpha)\mathcal{N}(\cdot; \boldsymbol{\mu}, \mathbf{S}) + \alpha\mathcal{N}(\cdot; \boldsymbol{\mu}, \delta\mathbf{I}_q). \tag{14}$$

For each individual $i = 1, \dots, n$, the proposal distribution $\nu_i(\cdot; \boldsymbol{\theta})$ is set to such a mixture, and $\boldsymbol{\mu}$ and \mathbf{S} are iteratively adapted according to the current estimate $\boldsymbol{\theta}^{(t)}$ (see Algorithm 2). However, to ensure convergence, it is necessary to impose conditions on (14) (see Appendix D for proof).

Lemma 5.1. *Let $\alpha \in (0, 1]$ and $\delta > 1$. If for any $i = 1, \dots, n$, $\boldsymbol{\theta} \mapsto \boldsymbol{\mu}_i(\boldsymbol{\theta}) \in \mathbb{R}^q$ and $\boldsymbol{\theta} \mapsto \mathbf{S}_i(\boldsymbol{\theta}) \in \mathcal{S}_{++}^q$ are continuous on Θ , then the proposal distribution defined by*

$$\nu_i(\cdot; \boldsymbol{\theta}) = \mathcal{GM}(\cdot; \boldsymbol{\mu}_i(\boldsymbol{\theta}), \mathbf{S}_i(\boldsymbol{\theta}), \alpha, \delta) \quad (15)$$

fulfils Assumption 1.

Constraining the parameter δ to live in $(1, +\infty)$ is sufficient to guarantee that the Radon–Nikodym derivative $\rho_{\boldsymbol{\theta}, i}$ is uniformly bounded with respect to $\boldsymbol{\theta} \in \Theta$. However, this does not ensure an efficient gradient estimator in terms of error or bias. The bias and the error are both related to the Kullback–Leibler divergence $\text{KL}[p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i) \| \nu_i(\cdot; \boldsymbol{\theta})]$ [1, 10]. Specifically, these results show that the estimator exhibits enhanced efficiency for a fixed computational budget, as the Kullback–Leibler divergence decreases. Adaptive importance sampling addresses the minimization of $\nu \mapsto \text{KL}[p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i) \| \nu]$ over a given class of probability measure. For instance, the Population Monte Carlo proposed by [9] provides a solution when the proposal is a mixture distribution. Nevertheless, implementing such methods within an SGD scheme can be computationally intensive. Indeed, the target distribution of the adaptive scheme changes at each iteration of the gradient scheme, necessitating a full run of the adaptive method at each iteration.

Practical implementation In the following, we present a simpler heuristic that is efficient for the class of problems presented in this paper, albeit not optimal in terms of the Kullback–Leibler divergence. We consider using the mean and covariance of the conditional distribution $p_{\boldsymbol{\theta}}(\mathbf{W}_i | \mathbf{Y}_i)$, namely

$$\boldsymbol{\mu}_i(\boldsymbol{\theta}) = \mathbb{E}[\mathbf{W}_i | \mathbf{Y}_i, \boldsymbol{\theta}], \quad \mathbf{S}_i(\boldsymbol{\theta}) = \mathbb{V}[\mathbf{W}_i | \mathbf{Y}_i, \boldsymbol{\theta}]. \quad (16)$$

Both functions are continuous on Θ (see Appendix D). Moreover, the parameter α can be interpreted as a regularization parameter. Indeed, in the limiting case $\alpha = 0$, the proposal distribution as defined in (15) resumes to the optimal Gaussian proposal distribution, that is

$$\nu_i(\cdot; \boldsymbol{\theta}) = \arg \min_{\nu \in \mathcal{F}} \text{KL}[p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i) \| \nu], \quad \mathcal{F} = \{\mathcal{N}(\boldsymbol{\mu}, \mathbf{S}); \boldsymbol{\mu} \in \mathbb{R}^q, \mathbf{S} \in \mathcal{S}_{++}^q\}. \quad (17)$$

While the Radon–Nikodym derivative $\rho_{\boldsymbol{\theta}, i}$ with respect to such a proposal may not necessarily be bounded for any $\boldsymbol{\theta} \in \Theta$, it points out the effect of α . The mixture distribution (15) balances a component that informs on the intractable conditional distribution and a regularization or defensive component that plays a similar role to that of [9]. Practically speaking, we should opt for a small value of α to improve the efficiency of the importance sampling.

Both $\boldsymbol{\mu}_i(\boldsymbol{\theta})$ and $\mathbf{S}_i(\boldsymbol{\theta})$ are unknown and must be estimated. Obviously, we can use the importance sampling method, since we could simply recycle the particles simulated to estimate the gradient. However, such a solution may lead to poorly conditioned and non-positive definite matrix estimates for $\mathbf{S}_i(\boldsymbol{\theta})$. A more robust alternative can be achieved using the Hessian of the log-complete

Algorithm 2: Adaptive Importance Sampling based Gradient Descent

Input: initial point $\theta^{(1)}$, learning rate γ , number of iterations T , number of Monte Carlo draws N , mixture parameter $0 < \alpha \leq 1$, parameter $\delta > 0$.
Output: the sequence $\theta^{(1)}, \dots, \theta^{(T+1)}$

for $t = 1$ **to** T **do**

- Sample i uniformly in $\{1, \dots, n\}$;
- Compute the estimate $\hat{\mu}_i$ of $\mu_i(\theta^{(t)})$ as defined in Equation (16) ;
- Compute the estimate $\hat{\mathbf{S}}_i$ of $\mathbf{S}_i(\theta^{(t)})$ as defined in Equations (16) or (18) ;
- Set $\nu_i = \mathcal{GM}(\cdot; \hat{\mu}_i, \hat{\mathbf{S}}_i, \alpha, \delta)$;
- Sample $(\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,N})$ from $\nu_i^{\otimes N}(\cdot; \theta^{(t)})$;
- Compute $\theta^{(t+1)} = P_{\Theta}(\theta^{(t)} - \gamma \hat{g}^{(t)})$ with $\hat{g}^{(t)} = -\hat{s}_i^N(\theta^{(t)})$ as in Equation (7);

end

likelihood:

$$\begin{aligned} \mathbf{S}_i^H(\boldsymbol{\theta}) &= - \left[\nabla_{\mathbf{w}}^2 \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{w}) \Big|_{\mathbf{w}=\boldsymbol{\mu}_i(\boldsymbol{\theta})} \right]^{-1} \\ &= [\mathbf{I}_q + \mathbf{C}^\top \text{Diag}[\exp\{f_i(\boldsymbol{\mu}_i(\boldsymbol{\theta}); \mathbf{B}, \mathbf{C})\}] \mathbf{C}]^{-1}. \end{aligned} \quad (18)$$

This alternative stems from the second order Taylor expansion of the complete log-likelihood, and has been used in various contexts, such as posterior approximation [41] and importance sampling [35, Chapter 9]. Interestingly, for a Gaussian distribution, the Taylor expansion exactly relates the curvature of the scalar field at its mode to the variance, namely $\mathbf{S}_i^H(\boldsymbol{\theta})$ corresponds to the variance. In contrast to a Monte Carlo estimate of the covariance of the conditional distribution, it directly follows from Equation (18) that the Hessian of the log-complete likelihood, and consequently its inverse, is definite positive. Moreover, Lemma 5.1 also applies to this choice of covariance matrix, since the function $\boldsymbol{\theta} \mapsto \mathbf{S}_i^H(\boldsymbol{\theta})$ is continuous as a composition of continuous functions.

6. Simulation study

Competitors We compare different variants of our algorithm corresponding to specific choices of the proposal distribution $\nu_i(\cdot; \boldsymbol{\theta})$:

- **ISGD-VEM:** we set the proposal distribution to the variational distribution ϕ_ψ , as defined in Equation (4). Although it differs from the optimal proposal distribution (17), it represents a natural choice as it serves as the optimal Gaussian surrogate for the conditional distribution in terms of the Kullback–Leibler divergence $\phi \mapsto \text{KL}[\phi \| p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i)]$.
- **ISGD-VEMmix:** there is no guarantee that the Radon–Nikodym derivative with respect to the VEM proposal is bounded, and thus ensures the convergence of Algorithm 1. To address this, we introduce a defensive component in this version and consider

$$\nu_i(\cdot; \boldsymbol{\theta}) = \mathcal{GM}(\cdot; \mathbf{m}_i^{\text{VEM}}, \mathbf{S}_i^{\text{VEM}}, \alpha, \delta),$$

where $\mathbf{m}_i^{\text{VEM}}$ and $\mathbf{S}_i^{\text{VEM}}$ are the mean and covariance of the variational distribution ϕ_ψ .

- **AISGD-SNIS:** it corresponds to Algorithm 2 and the choice of the mixture distribution

$$\nu_i(\cdot; \boldsymbol{\theta}) = \mathcal{GM}(\cdot; \widehat{\mathbf{m}}_i(\boldsymbol{\theta}), \widehat{\mathbf{S}}_i(\boldsymbol{\theta}), \alpha, \delta),$$

where $\widehat{\mathbf{m}}_i(\boldsymbol{\theta})$ and $\widehat{\mathbf{S}}_i(\boldsymbol{\theta})$ are SNIS estimators of the mean and the covariance of the conditional distribution $p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i)$.

- **AISGD-Hessian:** it corresponds to Algorithm 2 and the choice of the mixture distribution

$$\nu_i(\cdot; \boldsymbol{\theta}) = \mathcal{GM}(\cdot; \widehat{\mathbf{m}}_i(\boldsymbol{\theta}), \mathbf{S}_i^H(\boldsymbol{\theta}), \alpha, \delta),$$

with $\mathbf{S}_i^H(\boldsymbol{\theta})$, as defined in Equation (18).

All these methods are initialized with the variational estimator $\boldsymbol{\theta}^{\text{VEM}}$ fitted with the standard VEM algorithm implemented in `pyPLNmodels`¹. Throughout the numerical study, the mixture hyperparameters are set to $\alpha = 0.001$ and $\delta = 1.1$.

6.1. Synthetic data

Data generation We consider simulation settings with $n = 300$ individuals, $p = 150$ variables, $d = 1$ covariate (that is, one intercept), and rank constraints $q = 3, 5$ and 15 . The offset term \mathbf{o} is set to zero. For each value of q , we sample $M = 100$ datasets $\mathbf{Y}^{(q,m)}$, $m = 1, \dots, M$, according to the PLN-PCA model (2) with the following regression parameters

$$B_{kj}^* \sim \mathcal{N}(2, 1), \quad k = 1, \dots, m, j = 1, \dots, p.$$

The covariance matrix $\boldsymbol{\Sigma}^*$ is set to the closest rank- q approximation using singular value decomposition of the Toeplitz matrix (t_{jk}) , $j, k = 1, \dots, p$, defined as

$$t_{jk} = 3 \times \mathbf{1}_{j=k} + u^{|j-k|},$$

where u is drawn uniformly between 0.6 and 0.8.

Experimental design Each algorithm runs for a total of 1010 epochs with two regimes of batch sizes, $\mathfrak{B} = n$ and $\mathfrak{B} = 1$, used along iterations. We resort to the full dataset for the first 1000 epochs. During this initial phase, the gradient estimator is computed as the average of SNIS estimators (7) over all individuals, namely

$$\widehat{\mathbf{g}}^{(t)} = -\frac{1}{n} \sum_{i=1}^n s_i^N(\boldsymbol{\theta}^{(t)}).$$

The latter is based on `Rprop` updates [37]. We refer the reader to Appendix E.1 for additional details on the use of `Rprop`. This practical approach is motivated

¹<https://github.com/PLN-team/pyPLNmodels>

by the empirical observations that **Rprop** updates tend to converge much faster than the original updates using a single learning rate and could serve as a warm-up phase. Our study focuses on the last 10 epochs with a batch size of 1. This setting corresponds to $T = 3000$ iterations of Algorithm 1. The learning rate γ is determined via a grid search. We present results corresponding to $N = 500, 1000$, or 2000 samples for the SNIS estimator (7).

Quality of the importance sampling proposal distribution The proposal distributions $\nu_i(\cdot; \boldsymbol{\theta})$ of the four competitors are compared at initialization with the three following metrics:

- the Kullback–Leibler divergence $\text{KL}[\nu_i(\cdot; \boldsymbol{\theta}) \| p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i)]$, which gives a discrepancy measure in terms of the variational objective function, and thus provides a comparison between the proposal distribution and the variational distribution. Given a N -sample $\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,N}$ from $\nu_i(\cdot; \boldsymbol{\theta})$, its Monte Carlo estimator is

$$-\log(N) - \frac{1}{N} \sum_{r=1}^N \log(\omega_{\boldsymbol{\theta},i,r}), \quad \omega_{\boldsymbol{\theta},i,r} = \frac{\rho_{\boldsymbol{\theta},i}(\mathbf{v}_{i,r})}{\sum_{s=1}^N \rho_{\boldsymbol{\theta},i}(\mathbf{v}_{i,s})}.$$

- the Kullback–Leibler divergence $\text{KL}[p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i) \| \nu_i(\cdot; \boldsymbol{\theta})]$, which relates to the efficiency of the importance sampling scheme in terms of bias and quadratic error. Its Monte Carlo estimator is

$$\log(N) + \sum_{r=1}^N \omega_{\boldsymbol{\theta},i,r} \log(\omega_{\boldsymbol{\theta},i,r}).$$

- The Effective Sample Size (ESS), which assesses how accurately the weighted samples from the importance sampling method approximates the target distribution $p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i)$: a higher effective sample size indicates a better empirical approximation of the target distribution. It is estimated by

$$\left(\sum_{r=1}^N \omega_{\boldsymbol{\theta},i,r}^2 \right)^{-1}.$$

The results are displayed in Figure 1. A key distinction between the ISGD-VEM proposal and other methods is the use of a Gaussian mixture distribution. A comparison between ISGD-VEM and ISGD-VEMmix highlights the impact of adding a defensive component to the variational solution. As shown in the top row of Figure 1, the defensive component does not yield improvement in terms of the variational objective function. However, the ISGD-VEMmix proposal offers equivalent performances regarding the importance sampling method (middle and bottom rows of Figure 1).

Furthermore, the comparison of ISGD-VEMmix with the two AISGD variants (SNIS and Hessian) enables the evaluation of the impact of replacing the diagonal covariance from the variational distribution with either the conditional

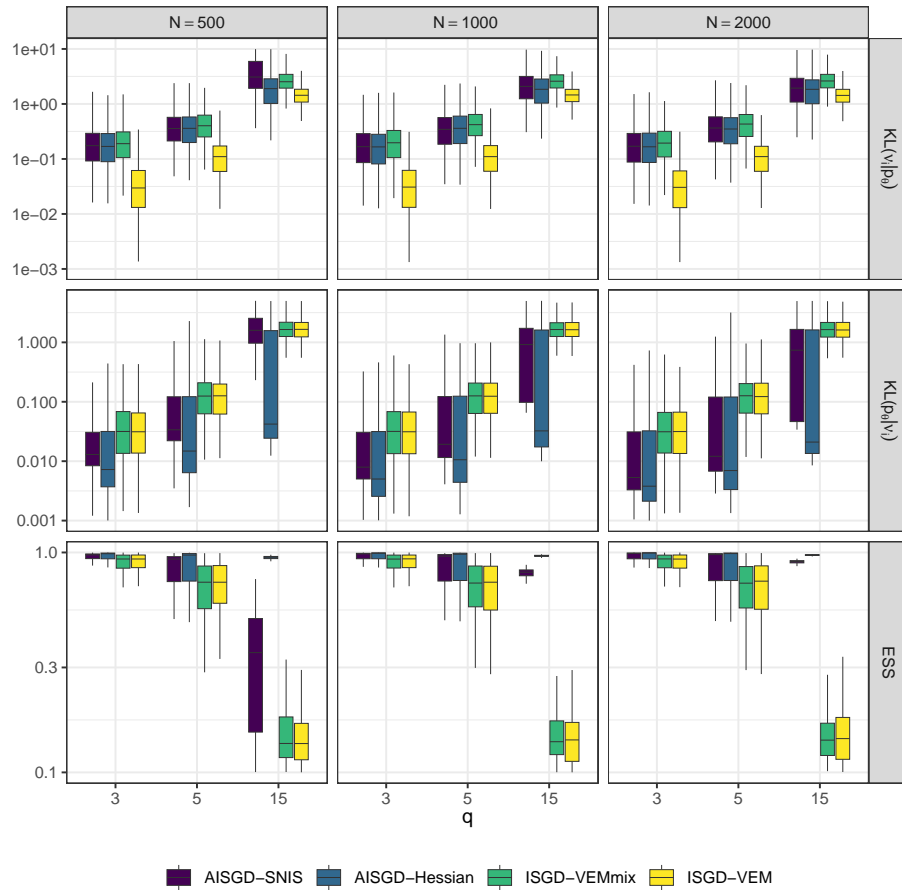


FIGURE 1. Distribution of the Kullback–Leibler divergence $KL[v_i(\cdot; \theta) \| p_{\theta}(\cdot | \mathbf{Y}_i)]$ (top row), the Kullback–Leibler divergence $KL[p_{\theta}(\cdot | \mathbf{Y}_i) \| v_i(\cdot; \theta)]$ (middle row), and the effective sample size (bottom row) as a function of the number of Monte Carlo draws N , the rank constraint q , and the inference algorithms (AISGD-SNIS, AISGD-Hessian, ISGD-VEMmix, ISGD-VEM) at initialization, that is $\theta^{(1)}$ is the VEM estimate. Each boxplot is based on $M = 100$ synthetic datasets, with each metric estimated using the specified N Monte Carlo draws.

covariance estimate or the curvature estimate. Unsurprisingly, AISGD-SNIS becomes numerically unstable when we do not have enough Monte Carlo draws for a given rank constraint q . It is well known that estimating a covariance matrix is much more demanding when the dimension of the sample space increases. From the variational point of view (top row of Figure 1), AISGD-SNIS and AISGD-Hessian do not exhibit significant improvement compared to ISGD-VEMmix. However, providing sufficient Monte Carlo draws are available, both AISGD-SNIS and AISGD-Hessian consistently demonstrate superior

performance for the importance sampling scheme in comparison to variational-based proposals (middle and bottom row of Figure 1). The lower values for the $\text{KL}[p_{\boldsymbol{\theta}}(\cdot | \mathbf{Y}_i) \|\nu_i(\cdot; \boldsymbol{\theta})]$ indicate a better control over bias and variance, while the performances in terms of ESS demonstrates that AISGD solutions are robust and efficient when the dimension q increases. Overall, AISGD-Hessian offers the most favorable practical performances. In what follows, we thus focus solely on AISGD-Hessian.

Asymptotic normality of the regression coefficient A key property of the MLE is the asymptotic normality that provides tests and confidence intervals on the model parameters. This property is not guaranteed by the variational estimator, maximizing a *surrogate* log-likelihood (often called ELBO). In order to assess the ability of our method to provide valid tests and confidence intervals for each, say, regression parameter B_{kj} , we examine the standardized estimates

$$\tilde{B}_{kj} = \left(\hat{B}_{kj} - B_{kj}^* \right) / \sqrt{\hat{\mathbb{V}}[\hat{B}_{kj}]} \quad (19)$$

where B_{kj}^* denotes the true value, and $\hat{\mathbb{V}}[\hat{B}_{kj}]$ stands for the estimated variance of \hat{B}_{kj} . Recall that the Fisher information matrix (FIM) is defined as

$$\mathcal{I}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}} \left[\nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_1) \{ \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_1) \}^{\top} \right]. \quad (20)$$

According to the M-estimator theory, the (asymptotic) variance of \hat{B}_{kj} is given by the diagonal term corresponding to B_{kj} of the inverse of the FIM, and the distribution of the $\tilde{B}_{kj}^{(m)}$ across simulations $m = 1, \dots, M$ should be close to a standard normal.

We illustrate the fit of the test statistics, as defined in (19), to the standard normal distribution at three different stages of the estimation procedure. The first stage corresponds to the VEM initialization, that is $\hat{B}_{kj} = \hat{B}_{kj}^{\text{VEM}}$, and a variational proxy for the unknown (asymptotic) variance of the variational estimator. This solution assumes that the evidence lower bound can be used in place of the log-likelihood in the definition (20). We refer to this solution as variational FIM. The second stage corresponds to the first iteration of Algorithm 2. The test statistics are also computed for $\hat{B}_{kj} = \hat{B}_{kj}^{\text{VEM}}$, but instead of using the variational FIM, we compute a Monte Carlo estimator of the FIM, referred to as SNIS FIM, given by

$$\hat{\mathcal{I}}_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \hat{s}_i^N(\boldsymbol{\theta}) \{ \hat{s}_i^N(\boldsymbol{\theta}) \}^{\top}. \quad (21)$$

The third stage corresponds to the end of the optimization scheme, that is, test statistics computed for $\hat{B}_{kj} = \hat{B}_{kj}^{(T)}$. The FIM at that point is also estimated using the Monte Carlo estimator (21).

We first perform a Kolmogorov-Smirnov test for the three aforementioned stages with $q = 5$ and $N = 1000$. Figure 2 represents Kolmogorov-Smirnov

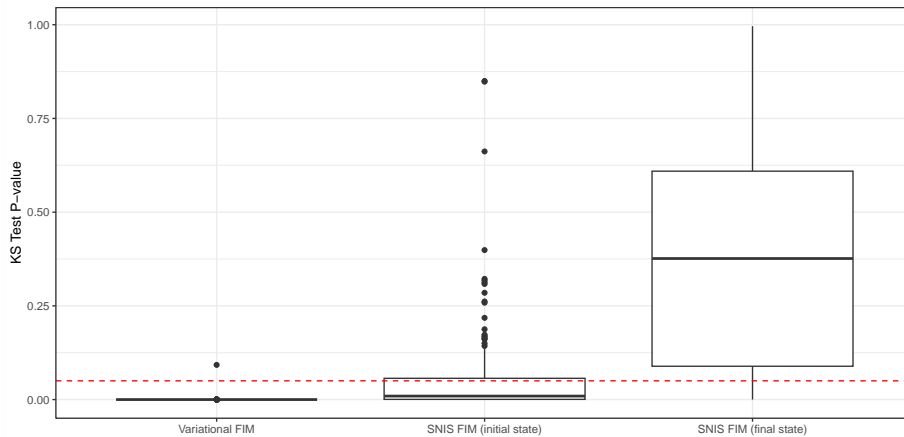


FIGURE 2. Distribution of the p -values of the Kolmogorov-Smirnov test for the distribution of the standardized estimates \hat{B}_{kj} , as defined in (19), over $M = 100$ simulations ($n = 300$, $p = 150$, $q = 5$, and $d = 1$) at three different stages: initial state $\theta^{(1)}$ (VEM estimate) with variational FIM (left) and with SNIS FIM (middle), and final state $\theta^{(T)}$ of AISGD-Hessian with SNIS FIM (right). AISGD-Hessian is run with $N = 1000$ Monte Carlo draws. Each boxplot is built across the $d \times p$ normalized coefficients \hat{B}_{kj} . Red dashed line [---]: the $\alpha = 5\%$ significance threshold.

p -values associated with the $d \times p$ regression parameter B_{kj} . The figure shows that the normality hypothesis is not rejected for the AISGD-Hessian. However, as we could expect, it is rejected for the VEM standardized estimates.

Figure 3 first illustrates why a departure from normality is observed for the standardized variational estimates on certain regression coefficients. The variational FIM solution significantly underestimates the variance (left column), implying, for instance, that using a full-based variational approach would lead to too narrow confidence intervals. Similar results were reported on the PLN model in [40]. Conversely, the SNIS FIM offers a satisfying and more accurate estimate of the variance (middle and right columns). The VEM estimator also exhibits bias (middle column). Further examination of additional qq-plots (not displayed here) indicates that VEM tends to overestimate the regression coefficients. Finally, the right column of Figure 3 shows that AISGD-Hessian corrects this bias and thus provides a good fit of the test statistics to the standard normal.

Conclusion AISGD-Hessian is a compelling method for inferring the parameters of a PLN PCA model. The algorithm offers an estimation procedure with valid uncertainty measures (or statistical tests). In contrast, the variational solution leads to unreliable outputs due to the approximation used throughout the VEM scheme and the crude variational FIM estimate.

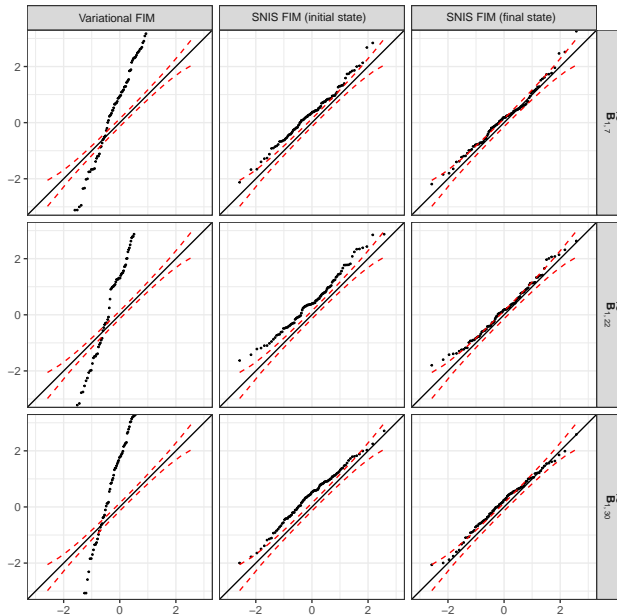


FIGURE 3. qq-plots of the standardized regression coefficients \tilde{B}_{1j} ($j = 7, 22, 30$), as defined in (19), over $M = 100$ simulations ($n = 300$, $p = 150$, $q = 5$, and $d = 1$) at three different stages: initial state $\theta^{(1)}$ (VEM estimate) with variational FIM (left) and with SNIS FIM (middle), and final state $\theta^{(T)}$ of AISGD-Hessian with SNIS FIM (right). AISGD-Hessian is run with $N = 1000$ Monte Carlo draws. x -axis: standard normal quantiles, y -axis: quantiles of \tilde{B}_{1j} (black dot $[\bullet]$), red dashed lines $[- -]$: 95% bounds for the standard normal qq-plot, black solid line $[-]$: perfect fit.

6.2. Inference on real data

Dataset Real data analysis is based on the scMARK dataset [17]. The latter is a benchmark for single-cell Ribonucleic acid (scRNA) data designed to serve as an RNA-seq equivalent of the MNIST dataset — each cell being labeled by one of the 28 possible cell types. It corresponds to $n = 19998$ samples (cells) and $p = 14059$ features (gene expression).

Due to memory limitations, we cannot address the problem with the original $p = 14059$ features. Therefore, we perform inference on a dataset reduced to $n = 300$ randomly chosen samples from the two most prevalent cell types (*T cells CD4+* and *T cells CD8+*), focusing on the $p = 100$ features with the largest variance. The offset term \mathbf{o} is set to zero. The two cell types are used as covariates, such that for each cell i , the covariates are given by

$$x_{i1} = \begin{cases} 1 & \text{if cell } i \text{ is a T cells CD4+} \\ 0 & \text{else} \end{cases}, \quad x_{i2} = 1 - x_{i1}.$$

Illustration of Theorem 4.5 For each rank constraint $q = 3, 5, 15$, we perform 10 runs of the AISGD-Hessian version of Algorithm 2 initialized at $\theta^{(1)} = \theta^{\text{VEM}}$ for 100 epochs and $N = 5000$ Monte Carlo draws. In Figure 4, we monitor the negative marginal log-likelihood and the gradient norm, averaged over the 10 runs, over the iterations. We observe a significant gain in terms of the marginal log-likelihood, suggesting that VEM has not converged to the maximum likelihood estimator. Interestingly, AISGD-Hessian proves to be increasingly more efficient as q increases. This behavior could result from using proposal distributions with a full covariance structure, which provides additional insights on the underlying geometry as compared to the diagonal approximation used in VEM. On the other hand, the norm of the gradient of the objective function decreases in average with T as stated by Theorem 4.5.

Parameters significance A major interest of our method is to achieve maximum likelihood estimation, but also to gain in interpretability thanks to an accurate estimation of the variance of PLN PCA parameters estimates. This allows practitioners to perform statistical tests and determine if a covariate has a significant effect or not. Figure 5 presents the confidence intervals for regression coefficients obtained at the initial and final iterations of a run of the AISGD-Hessian method on the reduced scMARK dataset with $q = 5$. As detailed in Section 6.1, we used here a warm-up phase of 2000 epochs, a batch size of $n = 300$ and $N = 1000$ Monte Carlo draws. We solely represent the coefficients whose estimates lie between -1 and 2 , leaving 34 coefficients out of the 200 available.

As previously observed in Figure 3, the VEM method seems to overestimate the regression coefficients consistently. By correcting this bias, the AISGD-Hessian method leads to differences in the interpretation for a few coefficients. For instance, we could reject the null hypothesis $B_{kj} = 0$ for all CD4+ coefficients at the VEM initialization, implying a positive impact of all the covariates. Conversely, at the final iteration, the decision changes for CD4+;63, CD4+;45 which are no longer significant. The same observation stands for CD8+;71 and CD8+;60. Additionally, CD8+;83, which initially has no impact, is found to have a negative impact at the final iteration.

7. Discussion

The paper explores the opportunity of a projected stochastic gradient scheme for inferring the parameters of a PLN-PCA model. Unlike competing variational approaches, our method allows to retrieve statistical guarantees on the resulting estimate, representing a significant step forward for the practitioners. This enables the construction of uncertainty measures and statistical tests. We have illustrated the benefits of our method on synthetic and real datasets using mixture distributions as proposals. Obviously we could resort to any distribution provided it satisfies Assumption 1. A future work could be to investigate the potential of normalizing flows as proposal distribution, especially for larger rank constraints.

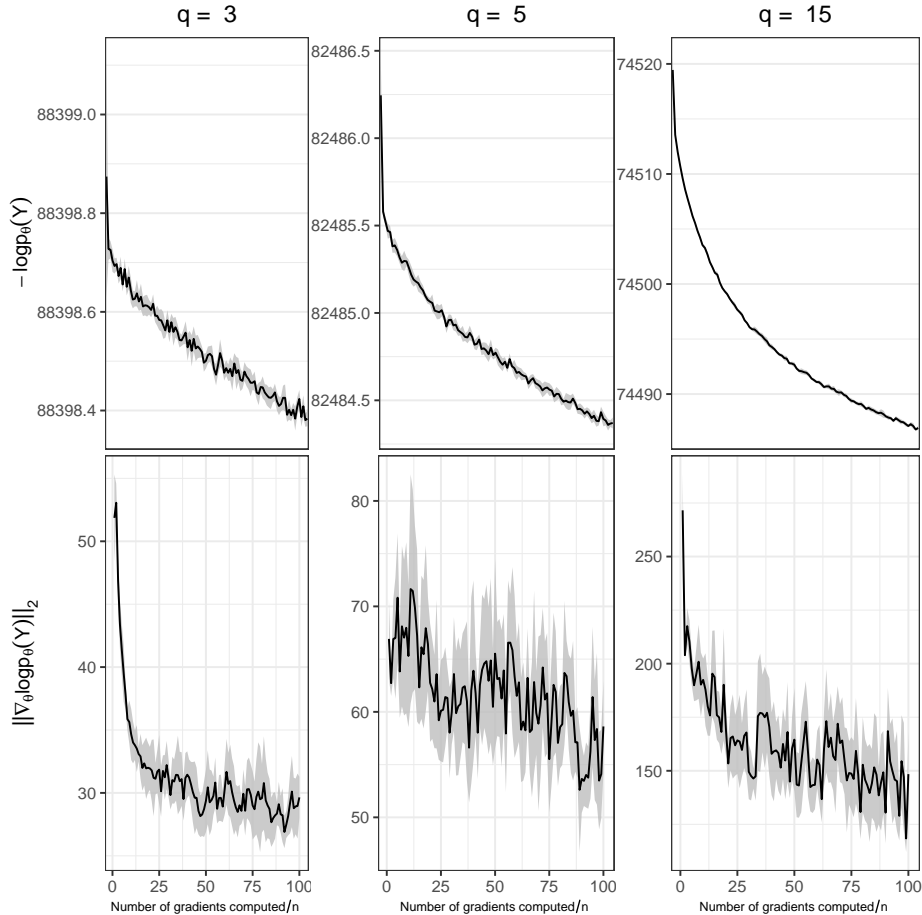


FIGURE 4. Negative marginal log-likelihood (top row) and norm of the gradient of the objective function (bottom row) for AISGD-Hessian as functions of the rank constraint q and the number of epochs for the scMARK dataset reduced to $n = 300$ samples and $p = 100$ variables. Black solid line $[-]$: averages over 10 runs of the AISGD-Hessian method initialized with the VEM estimate and $N = 5000$ Monte Carlo draws, grey area: 95% confidence regions.

While we have proven the convergence for the PLN-PCA model using an importance sampling-based gradient estimator, the descent lemma 4.2 opens new vista for latent variable models. Indeed, its minimal assumptions offers broader applicability, and it can serve as the corner stone for extending Theorem 4.5 to models with L -smooth loss function and arbitrary random and biased gradient estimator.

A natural extension is to consider models where the emission distribution belongs to a natural (or canonical) exponential family. While deriving regularity of the log likelihood, and hence a score estimator, assessing the L -smoothness of

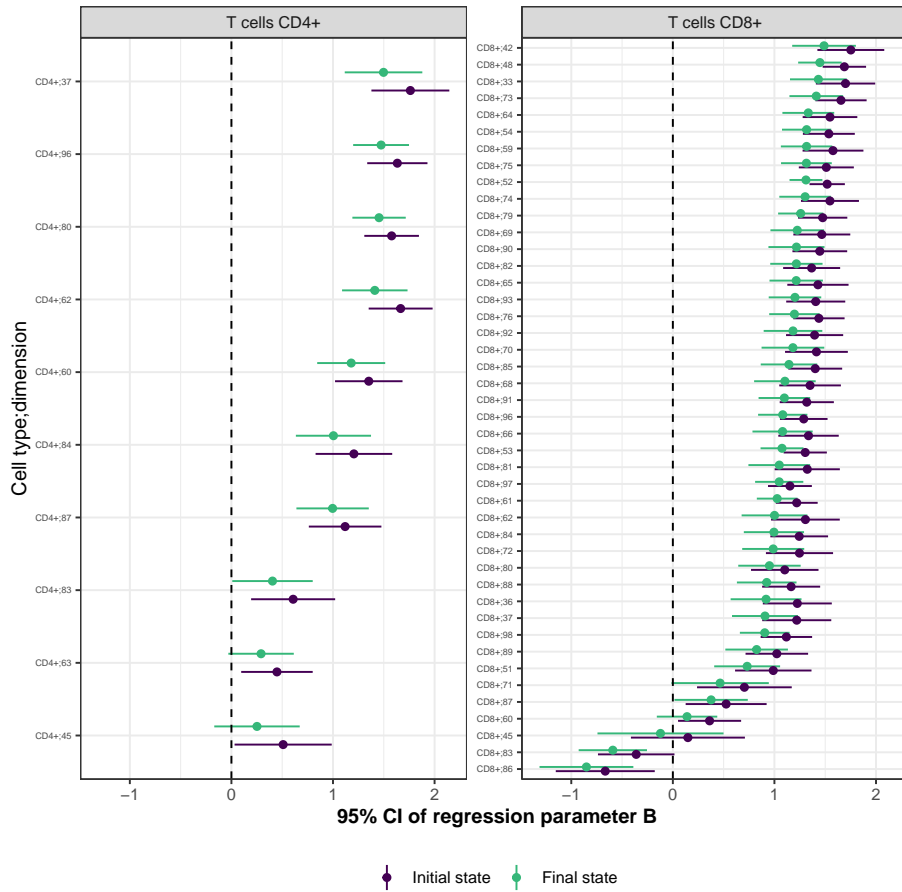


FIGURE 5. Confidence intervals for regression coefficients of $m = 2$ cell groups ($CD8+$ and $CD4+$ T cells) from the scMARK dataset, reduced to $n = 300$ samples and $p = 100$ variables. The intervals are computed at the initial state corresponding to the VEM estimate (purple), and the final state of a run of AISGD-Hessian with $N = 1000$ Monte Carlo draws. The variance was computed using the SNIS FIM estimator (21). x-axis: regression coefficient value, y-axis: cell type and feature number.

the loss function is relatively straightforward for some families — typically when the natural parameter space is the entire vector space —, such as the Binomial distribution. Additional technical conditions may however be required for families with constrained natural parameter spaces, like the negative-Binomial or Gamma distributions. These conditions might involve integrability constraints on the moments of the emission distribution with respect to the latent distribution or modifications to the link function f_i to account for the parameter constraints.

Beside addressing the inference for other models, the assumption of lemma

4.2 on the gradient estimator allows to consider other, possibly more elaborated, simulation based method to estimate the gradient, such as diffusion models.

Fundings

Bastien Bartardière and Julien Chiquet are supported by the French ANR grant ANR-18-CE45-0023 Statistics and Machine Learning for Single Cell Genomics (SingleStatOmics). Julien Stoehr has been partly funded by the European Union (ERC-2022-SYG-OCEAN-101071601). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] AGAPIOU, S., PAPASPILIOPOULOS, O., SANZ-ALONSO, D. and STUART, A. M. (2017). Importance Sampling: Intrinsic Dimension and Computational Cost. *Statistical Science* **32** 405–431.
- [2] AITCHISON, J. and HO, C. H. (1989). The multivariate Poisson-log normal distribution. *Biometrika* **76** 643–653.
- [3] AJALLOEIAN, A. and STICH, S. U. (2021). On the Convergence of SGD with Biased Gradients.
- [4] BERTSEKAS, D. P. (1999). *Nonlinear Programming*. Athena Scientific.
- [5] BLEI, D. M., KUCUKELBIR, A. and MCAULIFFE, J. D. (2017). Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association* 859–877.
- [6] BOYLES, R. A. (1983). On the Convergence of the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* **45** 47–50.
- [7] BRADBURY, J., FROSTIG, R., HAWKINS, P., JOHNSON, M. J., LEARY, C., MACLAURIN, D., NECULA, G., PASZKE, A., VANDERPLAS, J., WANDERMAN-MILNE, S. and ZHANG, Q. (2018). JAX: composable transformations of Python+NumPy programs.
- [8] BUBECK, S. (2015). *Convex Optimization: Algorithms and Complexity*.
- [9] CAPPÉ, O., DOUC, R., GUILLIN, A., MARIN, J. M. and ROBERT, C. P. (2008). Adaptive importance sampling in general mixture classes. *Statistics and Computing* **18** 447–459.
- [10] CHATTERJEE, S. and DIACONIS, P. (2018). The sample size required in importance sampling. *The Annals of Applied Probability* **28** 1099–1135.
- [11] CHIQUET, J., MARIADASSOU, M. and ROBIN, S. (2018a). Variational inference for probabilistic Poisson PCA. *Ann. Appl. Statist.* **12** 2674–2698.
- [12] CHIQUET, J., MARIADASSOU, M. and ROBIN, S. (2018b). Variational inference for probabilistic Poisson PCA. *The Annals of Applied Statistics* **12** 2674 – 2698.
- [13] CHIQUET, J., MARIADASSOU, M. and ROBIN, S. (2019). Variational Inference for sparse network reconstruction from count data. In *International Conference on Machine Learning*.

-
- [14] CHIQUET, J., MARIADASSOU, M. and ROBIN, S. (2021). The Poisson-Lognormal Model as a Versatile Framework for the Joint Analysis of Species Abundances. *Frontiers in Ecology and Evolution* **9** 188.
 - [15] CHIQUET, J., MARIADASSOU, M., ROBIN, S., POGGIATO, G. and GINDRAUD, F. (2023). PLNmodels: Poisson Lognormal Models PLN teams.
 - [16] DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum Likelihood from Incomplete Data via the EM algorithm. *Journal of the Royal Statistical Society: Series B* **39** 1–38.
 - [17] DIAZ-MEJIA, J. (2021). scMARK an 'MNIST' like benchmark to evaluate and optimize models for unifying scRNA data.
 - [18] DRUSVYATSKIY, D. and PAQUETTE, C. (2019). Efficiency of minimizing compositions of convex functions and smooth maps. *Math. Program.* **178** 503–558.
 - [19] GEYER, C. J. (1994). On the Convergence of Monte Carlo Maximum Likelihood Calculations. *Journal of the Royal Statistical Society. Series B (Methodological)* **56** 261–274.
 - [20] GHADIMI, S. and LAN, G. (2013). Stochastic First- and Zeroth-order Methods for Nonconvex Stochastic Programming.
 - [21] GÓMEZ-RUBIO, V. (2020). *Bayesian inference with INLA*. CRC Press.
 - [22] GUTMANN, M. U. and HYVÄRINEN, A. (2012). Noise-Contrastive Estimation of Unnormalized Statistical Models, with Applications to Natural Image Statistics. *Journal of Machine Learning Research* **13** 307–361.
 - [23] HALL, P., ORMEROD, J. T. and WAND, M. (2011). Theory of Gaussian variational approximation for a Poisson mixed model. *Statistica Sinica* 369–389.
 - [24] HINTON, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural computation* **14** 1771–1800.
 - [25] HUI, F. K. C. (2016). boral - Bayesian Ordination and Regression Analysis of Multivariate Abundance Data in r. *Methods in Ecology and Evolution* **7** 744–750.
 - [26] HUI, F. K. C., WARTON, D. I., ORMEROD, J. T., HAAPANIEMI, V. and TASKINEN, S. (2017). Variational Approximations for Generalized Linear Latent Variable Models. *Journal of Computational and Graphical Statistics* **26** 35–43.
 - [27] INOUE, D. I., YANG, E., ALLEN, G. I. and RAVIKUMAR, P. (2017). A review of multivariate distributions for count data derived from the Poisson distribution. *Wiley Interdisciplinary Reviews: Computational Statistics* **9** e1398.
 - [28] KAHN, H. (1949). *Stochastic (Monte Carlo) Attenuation Analysis*. Rand Corporation.
 - [29] KAHN, H. and HARRIS, T. E. (1951). Estimation of particle transmission by random sampling. *National Bureau of Standards applied mathematics series* **12** 27–30.
 - [30] KARLIS, D. (2005). EM algorithm for mixed Poisson and other discrete distributions. *ASTIN Bulletin: The Journal of the IAA* **35** 3–24.
 - [31] LOUIS, T. A. (1982). Finding the Observed Information Matrix When
-

- Using the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* **44** 226-233.
- [32] MAI, V. V. and JOHANSSON, M. (2021). Convergence of a Stochastic Gradient Method with Momentum for Non-Smooth Non-Convex Optimization.
- [33] NIKU, J., HUI, F. K. C., TASKINEN, S. and WARTON, D. I. (2019). gllvm: Fast analysis of multivariate abundance data with generalized linear latent variable models in R. *Methods in Ecology and Evolution* **10** 2173-2182.
- [34] OVASKAINEN, O., TIKHONOV, G., NORBERG, A., BLANCHET, F. G., DUAN, L., DUNSON, D., ROSLIN, T. and ABREGO, N. (2017). How to make more out of community data? A conceptual framework and its implementation as models and software. *Ecology Letters* **20** 561-576.
- [35] OWEN, A. B. (2013). *Monte Carlo theory, methods and examples*.
- [36] PASZKE, A., GROSS, S., MASSA, F., LERER, A., BRADBURY, J., CHANAN, G., KILLEEN, T., LIN, Z., GIMELSHEIN, N., ANTIGA, L. et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32**.
- [37] RIEDMILLER, M. and BRAUN, H. (1993). A direct adaptive method for faster backpropagation learning: the RPROP algorithm. In *IEEE International Conference on Neural Networks* **1** 586-591.
- [38] SCAMAN, K., MALHERBE, C. and SANTOS, L. D. (2022). Convergence Rates of Non-Convex Stochastic Gradient Descent Under a Generic Łojasiewicz Condition and Local Smoothness. In *Proceedings of the 39th International Conference on Machine Learning* **162** 19310-19327. PMLR.
- [39] SIMON, M. K. (2002). *Probability distributions involving Gaussian random variables: A handbook for engineers and scientists*. Springer.
- [40] STOEHR, J. and ROBIN, S. S. (2024). Composite likelihood inference for the Poisson log-normal model.
- [41] TIERNEY, L. and KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the american statistical association* **81** 82-86.
- [42] TIKHONOV, G., OPEDAL, O. H., ABREGO, N., LEHIKONEN, A., DE JONGE, M. M. J., OKSANEN, J. and OVASKAINEN, O. (2020). Joint species distribution modelling with the r-package Hmsc. *Methods in Ecology and Evolution* **11** 442-447.
- [43] VAN DER VAART, A. (1998). *Asymptotic statistics. Cambridge Series in Statistical and Probabilistic Mathematics* **27**. Cambridge Univ. Press, New York.
- [44] WESTLING, T. and MCCORMICK, T. H. (2019). Beyond prediction: A framework for inference with variational approximations in mixture models. *Journal of Computational and Graphical Statistics* **28** 778-789.
- [45] WU, C. F. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics* **11** 95-103.

Appendix A: Properties of the PLN-PCA model

Lemma A.1. *Under Model (2), for any compact set $\Theta \subset \mathbb{R}^d$, for all individual $i \in \{1, \dots, n\}$, there exists two constants $K_i^\Theta > 0$ and $\kappa_i^\Theta \in \mathbb{R}$, such that for any $\boldsymbol{\theta} \in \Theta$*

$$\log p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) \leq K_i^\Theta \|\mathbf{W}_i\| + \kappa_i^\Theta.$$

Proof. Given $i \in \{1, \dots, n\}$, $\boldsymbol{\theta} = (\mathbf{B}, \mathbf{C}) \in \mathbb{R}^d$, due to the conditional independence, we have for $\mathbf{Z}_i = (Z_{ij})$ as defined by Equation (2),

$$\log p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) = \langle \mathbf{Y}_i, \mathbf{Z}_i \rangle - \sum_{j=1}^p \exp(Z_{ij})$$

Since \exp is convex and differentiable on \mathbb{R} , for any $z, z_0 \in \mathbb{R}$, we have

$$\exp(z) \geq \exp(z_0) + (z - z_0) \exp(z_0).$$

It follows

$$\log p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) \leq \langle \mathbf{Y}_i, \mathbf{Z}_i \rangle - \exp(z_0) \sum_{j=1}^p Z_{ij} - \sum_{j=1}^p (1 - z_0) \exp(z_0).$$

Using the definition of \mathbf{Z}_i from Equation (2), we set

$$\begin{aligned} K_i(\boldsymbol{\theta}) &= \mathbf{C}^\top \{\mathbf{Y}_i - \exp(z_0) \mathbf{1}_p\}, \\ \kappa_i(\boldsymbol{\theta}) &= \langle \mathbf{Y}_i - \exp(z_0) \mathbf{1}_p, \mathbf{B}^\top \mathbf{x}_i + \mathbf{o}_i \rangle - \sum_{j=1}^p (1 - z_0) \exp(z_0). \end{aligned}$$

Then,

$$\log p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) \leq \langle K_i(\boldsymbol{\theta}), \mathbf{W}_i \rangle + \kappa_i(\boldsymbol{\theta}),$$

and using the Cauchy–Schwarz inequality we get

$$\log p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) \leq \|K_i(\boldsymbol{\theta})\| \|\mathbf{W}_i\| + \kappa_i(\boldsymbol{\theta}).$$

The functions K_i and κ_i are linear with respect to the parameter coordinates and, consequently, continuous on \mathbb{R}^d , and hence on Θ . Therefore, we can define the upper bounds $K_i^\Theta = \sup_{\boldsymbol{\theta} \in \Theta} \|K_i(\boldsymbol{\theta})\|$ and $\kappa_i^\Theta = \sup_{\boldsymbol{\theta} \in \Theta} \kappa_i(\boldsymbol{\theta})$, which provides the result. \square

Proposition 4.1. *For all individual $i = 1, \dots, n$, the incomplete log-likelihood $\boldsymbol{\theta} \mapsto \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i)$ of Model (2) is twice continuously differentiable on \mathbb{R}^d and its score function can be written as*

$$s_i(\boldsymbol{\theta}) = \int_{\mathbb{R}^q} \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{w}) p_{\boldsymbol{\theta}}(d\mathbf{w} \mid \mathbf{Y}_i) = \mathbb{E} [\nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{W}_i) \mid \mathbf{Y}_i, \boldsymbol{\theta}]. \quad (5)$$

Proof. The result is a direct application of the dominated convergence theorem. Indeed, for all individual $i = 1, \dots, n$, the likelihood writes as

$$p_{\boldsymbol{\theta}}(\mathbf{Y}_i) = \int_{\mathbb{R}^q} p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{w}) \mathcal{N}(d\mathbf{w}; \mathbf{0}_q, \mathbf{I}_q).$$

The function $\boldsymbol{\theta} \mapsto p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i)$ is twice continuously differentiable on \mathbb{R}^d as a composition of such functions. Moreover, each component Z_{ij} , $j = 1, \dots, p$, is a linear function of the components of $\boldsymbol{\theta}$, and for $r, s = 1, \dots, d$,

$$\frac{\partial}{\partial \theta_r} Z_{ij} = \sum_{k=1}^m x_{ik} \mathbf{1}_{\theta_r = B_{kj}} + \sum_{k=1}^q W_{ik} \mathbf{1}_{\theta_r = C_{jk}}, \quad \frac{\partial^2}{\partial \theta_s \partial \theta_r} Z_{ij} = 0. \quad (22)$$

Consequently, given a component θ_r , there is a unique index in $\{1, \dots, p\}$ such that the first partial derivative is non-zero. Denote by j and k such indices for the partial derivatives with respect to θ_r and θ_s , respectively. We then have

$$\begin{aligned} \frac{\partial}{\partial \theta_r} p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) &= p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) \{Y_{ij} - \exp(Z_{ij})\} \frac{\partial}{\partial \theta_r} Z_{ij}, \\ \frac{\partial^2}{\partial \theta_s \partial \theta_r} p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) &= p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) \left[\{Y_{ij} - \exp(Z_{ij})\} \{Y_{ik} - \exp(Z_{ik})\} \frac{\partial}{\partial \theta_r} Z_{ij} \frac{\partial}{\partial \theta_s} Z_{ik} \right. \\ &\quad \left. - \exp(Z_{ij}) \frac{\partial}{\partial \theta_r} Z_{ij} \frac{\partial}{\partial \theta_s} Z_{ij} \right]. \end{aligned}$$

It follows from (22) that for any $r = 1, \dots, d$, and any $j = 1, \dots, p$

$$\left| \frac{\partial}{\partial \theta_r} Z_{ij} \right| \leq \|\mathbf{x}_i\| + \|\mathbf{W}_i\|.$$

Given a non-empty open and bounded set $\Theta \subset \mathbb{R}^d$, it follows from Lemma A.1 that it exists two constants $K_i^\Theta > 0$ and κ_i^Θ such that for any $\boldsymbol{\theta} \in \Theta$

$$p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W}_i) \leq \exp(K_i^\Theta \|\mathbf{W}_i\| + \kappa_i^\Theta).$$

On the other hand, for any $\boldsymbol{\theta} \in \Theta$, any $j = 1, \dots, p$

$$|Z_{ij}| \leq A^\Theta \|\mathbf{W}_i\| + R_i^\Theta, \quad \text{with } A^\Theta = \sup_{\boldsymbol{\theta} \in \Theta} \|\boldsymbol{\theta}\|, \quad R_i^\Theta = \|\mathbf{x}_i\| \sup_{\boldsymbol{\theta} \in \Theta} \|\boldsymbol{\theta}\| + \max_{j=1, \dots, p} |o_{ij}|.$$

Consequently, for any $j = 1, \dots, p$, since $\exp(Z_{ij}) \leq \exp(|Z_{ij}|)$,

$$|Y_{ij} - \exp(Z_{ij})| \leq \exp(|Z_{ij}|) \{ \|\mathbf{Y}_i\| \exp(-|Z_{ij}|) + 1 \} \leq \exp(A^\Theta \|\mathbf{W}_i\| + R_i^\Theta) (\|\mathbf{Y}_i\| + 1).$$

Thereby, for any $j = 1, \dots, p$,

$$\begin{aligned} \left| \{Y_{ij} - \exp(Z_{ij})\} \frac{\partial}{\partial \theta_r} Z_{ij} \right| &\leq (\|\mathbf{x}_i\| + \|\mathbf{W}_i\|) (\|\mathbf{Y}_i\| + 1) \exp(A^\Theta \|\mathbf{W}_i\| + R_i^\Theta) \\ \left| \exp(Z_{ij}) \frac{\partial}{\partial \theta_r} Z_{ij} \frac{\partial}{\partial \theta_s} Z_{ij} \right| &\leq (\|\mathbf{x}_i\| + \|\mathbf{W}_i\|)^2 \exp(A^\Theta \|\mathbf{W}_i\| + R_i^\Theta) \\ &\leq (\|\mathbf{x}_i\| + \|\mathbf{W}_i\|)^2 \exp(2A^\Theta \|\mathbf{W}_i\| + 2R_i^\Theta), \end{aligned}$$

where the last inequality follows because A^Θ and R_i^Θ are both positive. Finally, we get that on Θ

$$\begin{aligned} \left| \frac{\partial}{\partial \theta_r} p_\theta(\mathbf{Y}_i | \mathbf{W}_i) \right| \exp\left(-\frac{\|\mathbf{W}_i\|^2}{2}\right) &\leq (\|\mathbf{x}_i\| + \|\mathbf{W}_i\|) (\|\mathbf{Y}_i\| + 1) Q_i^\Theta(\mathbf{W}_i, 1) \\ \left| \frac{\partial}{\partial \theta_s \partial \theta_r} p_\theta(\mathbf{Y}_i | \mathbf{W}_i) \right| \exp\left(-\frac{\|\mathbf{W}_i\|^2}{2}\right) &\leq (\|\mathbf{x}_i\| + \|\mathbf{W}_i\|)^2 \left\{ (\|\mathbf{Y}_i\| + 1)^2 + 1 \right\} Q_i^\Theta(\mathbf{W}_i, 2), \end{aligned}$$

with

$$Q_i^\Theta(\mathbf{W}_i, k) = \exp\left\{-\frac{1}{2}\|\mathbf{W}_i\|^2 + (K_i^\Theta + kA^\Theta)\|\mathbf{W}_i\| + \kappa_i^\Theta + kR_i^\Theta\right\}.$$

Crucially, each of these upper bounds is a Lebesgue integrable function on \mathbb{R}^q that does not depend on $\theta \in \Theta$. Consequently, we can conclude that the likelihood is twice continuously differentiable on Θ with the dominated convergence theorem. Since the result holds for any open and bounded set $\Theta \subset \mathbb{R}^d$, for all $\theta \in \mathbb{R}^d$, we can apply it to an open d -ball with center θ . Therefore, the likelihood is twice continuously differentiable on \mathbb{R}^d . Moreover

$$\nabla_\theta p_\theta(\mathbf{Y}_i) = \int_{\mathbb{R}^q} \nabla_\theta p_\theta(\mathbf{Y}_i | \mathbf{w}) \mathcal{N}(d\mathbf{w}; \mathbf{0}_q, \mathbf{I}_q) = \int_{\mathbb{R}^q} \nabla_\theta p_\theta(\mathbf{Y}_i, \mathbf{w}) d\mathbf{w}.$$

For all $\mathbf{Y}_i \in \mathbb{N}^p$, the likelihood $\theta \mapsto p_\theta(\mathbf{Y}_i)$ is positive on \mathbb{R}^d . Indeed, the integrand is positive everywhere since, by definition

$$\mathbf{w} \mapsto p_\theta(\mathbf{Y}_i | \mathbf{w}) = \prod_{j=1}^p p(Y_{ij} | z_{ij}), \quad (z_{i1}, \dots, z_{ip})^\top = f_i(\mathbf{w}; \mathbf{B}, \mathbf{C}),$$

and for any $z \in \mathbb{R}$, $y \mapsto p(y | z)$ is positive on \mathbb{N} . The continuous differentiability of the log-likelihood follows directly by composition. Finally,

$$\begin{aligned} \nabla_\theta \log p_\theta(\mathbf{Y}_i) &= \frac{1}{p_\theta(\mathbf{Y}_i)} \int_{\mathbb{R}^q} \nabla_\theta p_\theta(\mathbf{Y}_i, \mathbf{w}) d\mathbf{w} = \int_{\mathbb{R}^q} \frac{p_\theta(\mathbf{Y}_i, \mathbf{w}) \nabla_\theta \log p_\theta(\mathbf{Y}_i, \mathbf{w})}{p_\theta(\mathbf{Y}_i)} d\mathbf{w} \\ &= \int_{\mathbb{R}^q} \nabla_\theta \log p_\theta(\mathbf{Y}_i, \mathbf{w}) p_\theta(d\mathbf{w} | \mathbf{Y}_i). \end{aligned}$$

□

Proposition 4.3. *Under Model (2), for any nonempty compact subset $\Theta \subset \mathbb{R}^d$, we have that*

(i) *for any individual $i = 1, \dots, n$, there exists a real*

$$0 < \zeta_i = \inf_{\theta \in \Theta} p_\theta(\mathbf{Y}_i);$$

(ii) *there exists a real $L \geq 0$ such that the objective function ℓ , as defined in (3), is L -smooth on Θ .*

Proof. Let $\Theta \subset \mathbb{R}^d$ be a nonempty compact set.

- (i) From Proposition 4.1, the likelihood of the observed data $\boldsymbol{\theta} \mapsto p_{\boldsymbol{\theta}}(\mathbf{Y}_i)$ is continuous on the compact set Θ . Thus, it is bounded and attains its bounds on Θ . Moreover, the likelihood of the observed data is positive on \mathbb{R}^d , so is its lower bound on Θ .
- (ii) A direct consequence of Proposition 4.1, is that $\boldsymbol{\theta} \mapsto \nabla \ell(\boldsymbol{\theta})$ is continuous differentiable on the compact set Θ . Therefore, it is Lipschitz continuous. \square

Appendix B: Intermediate optimization results and proof of Lemma 4.2

In this Section, we consider results for minimizing a function ℓ on \mathbb{R}^d . We denote $\Theta \subset \mathbb{R}^d$ a compact and convex set.

Lemma B.1 (Non-expansiveness, Bertsekas [4, Prop 2.1.3]). *The projection P_{Θ} satisfies the non-expansiveness property, namely for all $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^d$*

$$\|P_{\Theta}(\boldsymbol{\theta}) - P_{\Theta}(\boldsymbol{\theta}')\| \leq \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|.$$

For a real $\eta > 0$, the proximal operator of F as defined in Equation (10) is given by

$$\text{prox}_F^{\eta}(\boldsymbol{\theta}) = \arg \min_{\boldsymbol{\theta}' \in \mathbb{R}^d} \left\{ F(\boldsymbol{\theta}') + \frac{1}{2\eta} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 \right\}.$$

Lemma B.2 (Drusvyatskiy and Paquette [18, Lemma 4.3]). *If the function ℓ is L -smooth, for any real constant $\eta < 1/L$, the Moreau envelope F_{η} is differentiable on \mathbb{R}^d and relates to the proximal operator through the following equation*

$$\nabla F_{\eta}(\boldsymbol{\theta}) = \frac{1}{\eta} \{\boldsymbol{\theta} - \text{prox}_F^{\eta}(\boldsymbol{\theta})\}.$$

Lemma B.3 (Drusvyatskiy and Paquette [18, Theorem 4.5]). *If the function ℓ is L -smooth, for any real constant $\eta > 0$, the gradient mapping $G_{\eta}^{(t)}$, as defined in Equation (12), satisfy the following inequality*

$$\|G_{\eta}^{(t)}\| \leq \left(1 + \frac{L\eta}{L\eta + 1}\right) \left(1 + \sqrt{\frac{L\eta}{L\eta + 1}}\right) \|\nabla F_{\frac{\eta}{L\eta + 1}}(\boldsymbol{\theta}^{(t)})\|.$$

A differentiable function $f : \mathbb{R}^d \mapsto \mathbb{R}$ is said to be μ -strongly-convex, $\mu \in \mathbb{R}_+^*$, if for all $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^d$,

$$f(\boldsymbol{\theta}') \geq f(\boldsymbol{\theta}) + \langle \nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle + \frac{\mu}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2.$$

Lemma B.4. *If the function ℓ is L -smooth, then, for any real constant $\eta \in (0, 1/L)$, and any $\bar{\boldsymbol{\theta}} \in \mathbb{R}^d$, the following function is $(\eta^{-1} - L)$ -strongly convex*

$$f : \boldsymbol{\theta} \mapsto \ell(\boldsymbol{\theta}) + \frac{1}{2\eta} \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}\|^2.$$

Proof. Let $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^d$. We have

$$f(\boldsymbol{\theta}) + \langle \nabla f(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle = \ell(\boldsymbol{\theta}) + \frac{1}{2\eta} \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}\|^2 + \langle \nabla \ell(\boldsymbol{\theta}) + \eta^{-1}(\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle. \quad (23)$$

According to Bubeck [8, Lemma 3.4], since ℓ is L -smooth, it satisfies for all $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^d$

$$|\ell(\boldsymbol{\theta}') - \ell(\boldsymbol{\theta}) - \langle \nabla \ell(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle| \leq \frac{L}{2} \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2. \quad (24)$$

It directly follows that

$$\ell(\boldsymbol{\theta}) + \langle \nabla \ell(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle - \frac{L}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 \leq \ell(\boldsymbol{\theta}').$$

Combining the latter inequality with Equation (23) yields

$$\begin{aligned} & f(\boldsymbol{\theta}) + \langle \nabla f(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle + \frac{\eta^{-1} - L}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 \\ & \leq \ell(\boldsymbol{\theta}') + \frac{1}{2\eta} (\|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}\|^2 + 2\langle \boldsymbol{\theta} - \bar{\boldsymbol{\theta}}, \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle + \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2) \end{aligned}$$

For a scalar product and its associated norm, the identity $\|a + b\|^2 = \|a\|^2 + 2\langle a, b \rangle + \|b\|^2$ gives

$$f(\boldsymbol{\theta}) + \langle \nabla f(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle + \frac{\eta^{-1} - L}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 \leq \ell(\boldsymbol{\theta}') + \frac{1}{2\eta} \|\boldsymbol{\theta}' - \bar{\boldsymbol{\theta}}\|^2 = f(\boldsymbol{\theta}').$$

Since $\eta^{-1} + L > 0$, we can conclude that f is $(\eta^{-1} + L)$ -strongly convex. \square

Lemma B.5. *If the function ℓ is L -smooth, then for any $\boldsymbol{\theta} \in \Theta$, and any real constant $\eta \in (0, (2\Gamma + L)^{-1})$, with $\Gamma = \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla \ell(\boldsymbol{\theta})\|$, we have*

$$\|\boldsymbol{\theta} - \text{prox}_F^\eta(\boldsymbol{\theta})\| \leq 1.$$

Proof. Given $\boldsymbol{\theta} \in \Theta$, we show that any point at a distance more than one from $\boldsymbol{\theta}$ is not $\text{prox}_F^\eta(\boldsymbol{\theta})$. This implies that necessarily

$$\|\boldsymbol{\theta} - \text{prox}_F^\eta(\boldsymbol{\theta})\| \leq 1.$$

Let $\boldsymbol{\theta}' \in \mathbb{R}^d$ be a point such that $\|\boldsymbol{\theta} - \boldsymbol{\theta}'\| > 1$. According to Bubeck [8, Lemma 3.4], since ℓ is L -smooth, it satisfies for all $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^d$

$$\ell(\boldsymbol{\theta}) \leq \ell(\boldsymbol{\theta}') + \langle \nabla \ell(\boldsymbol{\theta}'), \boldsymbol{\theta} - \boldsymbol{\theta}' \rangle + \frac{L}{2} \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2. \quad (25)$$

Let first assume that $\boldsymbol{\theta}' \in \Theta$. Using successively the Cauchy–Schwarz inequality, the positive bound Γ and the assumption $\|\boldsymbol{\theta} - \boldsymbol{\theta}'\| > 1$, we get

$$\langle \nabla \ell(\boldsymbol{\theta}'), \boldsymbol{\theta} - \boldsymbol{\theta}' \rangle \leq \|\nabla \ell(\boldsymbol{\theta}')\| \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| < \Gamma \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2. \quad (26)$$

It results from Equation (25) and Equation (26) that

$$\ell(\boldsymbol{\theta}) < \ell(\boldsymbol{\theta}') + \left(\Gamma + \frac{L}{2} \right) \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2.$$

For $\eta \leq (2\Gamma + L)^{-1}$, we have $(\Gamma + L/2) \leq 1/(2\eta)$. Moreover, the functions F and ℓ coincide on Θ , so that

$$F(\boldsymbol{\theta}) < F(\boldsymbol{\theta}') + \frac{1}{2\eta} \|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2.$$

We can therefore conclude that $\boldsymbol{\theta}' \neq \text{prox}_F^\eta(\boldsymbol{\theta})$. Conversely, if $\boldsymbol{\theta}' \notin \Theta$, by definition of F , we also have that $\boldsymbol{\theta}' \neq \text{prox}_F^\eta(\boldsymbol{\theta})$. \square

Lemma 4.2 (Descent lemma). *Let consider the gradient scheme as defined by Equation (9). Assume that*

- (i) *the function ℓ is L -smooth on Θ , and denote $\Gamma = \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta})\|$;*
- (ii) *for $t \in \mathbb{N}^*$, one has*

$$\begin{aligned} \sigma^{(t)} &= \mathbb{E} \left[\left\| \widehat{\boldsymbol{g}}^{(t)} - \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \right\|^2 \mid \boldsymbol{\theta}^{(t)} \right] < \infty, \\ \xi^{(t)} &= \left\| \mathbb{E} \left[\widehat{\boldsymbol{g}}^{(t)} \mid \boldsymbol{\theta}^{(t)} \right] - \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \right\| < \infty. \end{aligned}$$

Then, for any real constant $\eta \in (0, 1/\max\{2\Gamma + L, 2L\}]$,

$$\begin{aligned} \mathbb{E} \left[F_\eta(\boldsymbol{\theta}^{(t+1)}) \mid \boldsymbol{\theta}^{(t)} \right] &\leq F_\eta(\boldsymbol{\theta}^{(t)}) - \frac{\gamma}{2} \|\nabla F_\eta(\boldsymbol{\theta}^{(t)})\|^2 \\ &\quad + \frac{\gamma + \gamma^2 \Gamma}{\eta} \xi^{(t)} + \frac{\gamma^2}{2\eta} (\sigma^{(t)} + \Gamma^2). \end{aligned} \quad (11)$$

Proof. Given $t \in \mathbb{N}^*$ and $\eta \in (0, \max\{2\Gamma + L, 2L\}^{-1}]$, we define the virtual iterates $\bar{\boldsymbol{\theta}}^{(t)}$ as

$$\bar{\boldsymbol{\theta}}^{(t)} = \text{prox}_F^\eta(\boldsymbol{\theta}^{(t)}) = \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \left\{ F(\boldsymbol{\theta}) + \frac{1}{2\eta} \|\boldsymbol{\theta} - \boldsymbol{\theta}^{(t)}\|^2 \right\}.$$

Note that $\bar{\boldsymbol{\theta}}^{(t)} \in \Theta$. Indeed, according to the definition of the Moreau envelope

$$F(\bar{\boldsymbol{\theta}}^{(t)}) = F_\eta(\boldsymbol{\theta}^{(t)}) - \frac{1}{2\eta} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2 \leq F(\boldsymbol{\theta}^{(t)}).$$

The upper bound is finite, since $\boldsymbol{\theta}^{(t)}$ as defined in (9) is within the compact Θ . The operator F is thus finite at $\bar{\boldsymbol{\theta}}^{(t)}$, and consequently $\bar{\boldsymbol{\theta}}^{(t)} \in \Theta$ since F takes infinite value outside of Θ .

As the virtual iterates are within Θ , the definition of the Moreau envelope yields

$$F_\eta(\boldsymbol{\theta}^{(t+1)}) \leq \ell(\bar{\boldsymbol{\theta}}^{(t)}) + \frac{1}{2\eta} \|\boldsymbol{\theta}^{(t+1)} - \bar{\boldsymbol{\theta}}^{(t)}\|^2. \quad (27)$$

Using the definition of $\boldsymbol{\theta}^{(t+1)}$ and the fact that $\bar{\boldsymbol{\theta}}^{(t)} \in \Theta$, we obtain from Lemma B.1

$$\begin{aligned} \frac{1}{2\eta} \|\boldsymbol{\theta}^{(t+1)} - \bar{\boldsymbol{\theta}}^{(t)}\|^2 &= \frac{1}{2\eta} \|P_\Theta(\boldsymbol{\theta}^{(t)} - \gamma\hat{\boldsymbol{g}}^{(t)}) - P_\Theta(\bar{\boldsymbol{\theta}}^{(t)})\|^2 \\ &\leq \frac{1}{2\eta} \|\boldsymbol{\theta}^{(t)} - \gamma\hat{\boldsymbol{g}}^{(t)} - \bar{\boldsymbol{\theta}}^{(t)}\|^2. \end{aligned}$$

The bilinearity of the scalar product gives

$$\|\boldsymbol{\theta}^{(t)} - \gamma\hat{\boldsymbol{g}}^{(t)} - \bar{\boldsymbol{\theta}}^{(t)}\|^2 = \|\boldsymbol{\theta}^{(t)} - \bar{\boldsymbol{\theta}}^{(t)}\|^2 + 2\gamma \langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \hat{\boldsymbol{g}}^{(t)} \rangle + \gamma^2 \|\hat{\boldsymbol{g}}^{(t)}\|^2,$$

and the inequality (27) becomes

$$\begin{aligned} F_\eta(\boldsymbol{\theta}^{(t+1)}) &\leq \ell(\bar{\boldsymbol{\theta}}^{(t)}) + \frac{1}{2\eta} \|\boldsymbol{\theta}^{(t)} - \bar{\boldsymbol{\theta}}^{(t)}\|^2 + \frac{\gamma}{\eta} \langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \hat{\boldsymbol{g}}^{(t)} \rangle + \frac{\gamma^2}{2\eta} \|\hat{\boldsymbol{g}}^{(t)}\|^2 \\ &\leq F_\eta(\boldsymbol{\theta}^{(t)}) + \frac{\gamma}{\eta} \langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \hat{\boldsymbol{g}}^{(t)} \rangle + \frac{\gamma^2}{2\eta} \|\hat{\boldsymbol{g}}^{(t)}\|^2. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E} [F_\eta(\boldsymbol{\theta}^{(t+1)}) \mid \boldsymbol{\theta}^{(t)}] &\leq F_\eta(\boldsymbol{\theta}^{(t)}) + \frac{\gamma}{\eta} \langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \mathbb{E} [\hat{\boldsymbol{g}}^{(t)} \mid \boldsymbol{\theta}^{(t)}] \rangle \\ &\quad + \frac{\gamma^2}{2\eta} \mathbb{E} [\|\hat{\boldsymbol{g}}^{(t)}\|^2 \mid \boldsymbol{\theta}^{(t)}]. \end{aligned}$$

We now study more in detail the second and the third terms of the right-hand side.

Regarding the second term, we have

$$\langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \mathbb{E} [\hat{\boldsymbol{g}}^{(t)} \mid \boldsymbol{\theta}^{(t)}] \rangle = \langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \rangle + \langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, b^{(t)} \rangle,$$

where $b^{(t)}$ stands for the bias of the gradient estimate, namely

$$b^{(t)} = \mathbb{E} [\hat{\boldsymbol{g}}^{(t)} \mid \boldsymbol{\theta}^{(t)}] - \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}).$$

By combining the Cauchy–Schwarz inequality with Lemma B.5, which applies since $\eta \leq (2\Gamma + L)^{-1}$, we get

$$\langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, b^{(t)} \rangle \leq \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\| \|b^{(t)}\| \leq \|b^{(t)}\|.$$

Since the function ℓ is L -smooth, we can use the inequality (24), namely

$$\langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \rangle \leq \ell(\bar{\boldsymbol{\theta}}^{(t)}) - \ell(\boldsymbol{\theta}^{(t)}) + \frac{L}{2} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2.$$

In conclusion we obtain

$$\langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \mathbb{E} [\hat{g}^{(t)} \mid \boldsymbol{\theta}^{(t)}] \rangle \leq \ell(\bar{\boldsymbol{\theta}}^{(t)}) - \ell(\boldsymbol{\theta}^{(t)}) + \frac{L}{2} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2 + \|b^{(t)}\|. \quad (28)$$

We aim at obtaining a bound not depending on the virtual iterates. As $\eta^{-1} \geq 2L > L$, Lemma B.4 applies, and the function

$$g : \boldsymbol{\theta} \mapsto \ell(\boldsymbol{\theta}) + \frac{1}{2\eta} \|\boldsymbol{\theta} - \boldsymbol{\theta}^{(t)}\|^2$$

is $(\eta^{-1} - L)$ -strongly convex. By definition, this function achieves a minimum at $\bar{\boldsymbol{\theta}}^{(t)}$. Consequently its gradient is 0 at that point and the strong-convexity gives

$$\ell(\boldsymbol{\theta}^{(t)}) \geq \ell(\bar{\boldsymbol{\theta}}^{(t)}) + \frac{1}{2\eta} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2 + \frac{\eta^{-1} - L}{2} \|\boldsymbol{\theta}^{(t)} - \bar{\boldsymbol{\theta}}^{(t)}\|^2.$$

Using this last inequality, Equation (28) becomes

$$\begin{aligned} \langle \bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}, \mathbb{E} [\hat{g}^{(t)} \mid \boldsymbol{\theta}^{(t)}] \rangle &\leq -\frac{1}{2\eta} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2 - \frac{\eta^{-1} - 2L}{2} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2 \\ &\quad + \|b^{(t)}\| \\ &\leq -\frac{1}{2\eta} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2 + \|b^{(t)}\|, \end{aligned}$$

since $\eta^{-1} \geq 2L$. Finally, using Lemma B.2 leads to

$$\frac{1}{2\eta} \|\bar{\boldsymbol{\theta}}^{(t)} - \boldsymbol{\theta}^{(t)}\|^2 = \frac{1}{2\eta} \|\eta \nabla_{F_\eta}(\boldsymbol{\theta}^{(t)})\|^2 = \frac{\eta}{2} \|\nabla_{F_\eta}(\boldsymbol{\theta}^{(t)})\|^2.$$

So far, we have hence proven that

$$\begin{aligned} \mathbb{E} [F_\eta(\boldsymbol{\theta}^{(t+1)}) \mid \boldsymbol{\theta}^{(t)}] &\leq F_\eta(\boldsymbol{\theta}^{(t)}) - \frac{\gamma}{2} \|\nabla_{F_\eta}(\boldsymbol{\theta}^{(t)})\|^2 \\ &\quad + \frac{\gamma}{\eta} \|b^{(t)}\| + \frac{\gamma^2}{2\eta} \mathbb{E} [\|\hat{g}^{(t)}\|^2 \mid \boldsymbol{\theta}^{(t)}]. \end{aligned}$$

We now focus on the third term. The bilinearity of the scalar product gives

$$\begin{aligned} \mathbb{E} [\|\hat{g}^{(t)}\|^2 \mid \boldsymbol{\theta}^{(t)}] &= \sigma^{(t)} - \|\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)})\|^2 + 2 \langle \mathbb{E} [\hat{g}^{(t)} \mid \boldsymbol{\theta}^{(t)}], \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \rangle \\ &= \sigma^{(t)} + \|\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)})\|^2 + 2 \langle b^{(t)}, \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) \rangle. \end{aligned}$$

The Cauchy-Schwarz inequality along with the definition of the bound Γ yields

$$\mathbb{E} \left[\|\widehat{g}^{(t)}\|^2 \mid \boldsymbol{\theta}^{(t)} \right] \leq \sigma^{(t)} + \Gamma^2 + 2\Gamma \|b^{(t)}\|.$$

It follows that

$$\begin{aligned} \mathbb{E} \left[F_\eta(\boldsymbol{\theta}^{(t+1)}) \mid \boldsymbol{\theta}^{(t)} \right] &\leq F_\eta(\boldsymbol{\theta}^{(t)}) - \frac{\gamma}{2} \|\nabla F_\eta(\boldsymbol{\theta}^{(t)})\|^2 \\ &\quad + \frac{\gamma + \gamma^2 \Gamma}{\eta} \|b^{(t)}\| + \frac{\gamma^2}{2\eta} (\sigma^{(t)} + \Gamma^2). \end{aligned}$$

□

Appendix C: Proof of Theorem 4.5

Proposition 4.4. *Let $\{i(t)\}_{t \in \mathbb{N}^*}$ and $\{\boldsymbol{\theta}^{(t)}\}_{t \in \mathbb{N}^*}$ be the associated random sequences generated by Algorithm 1. Under Model (2), if Assumption 1 holds, then, for all $t \in \mathbb{N}$,*

$$\begin{aligned} \bar{\sigma}_{\text{IS}}^{(t)} &= \mathbb{E} \left[\left\| \widehat{s}_{i(t)}^N(\boldsymbol{\theta}^{(t)}) - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\|^2 \mid \boldsymbol{\theta}^{(t)} \right] \leq \frac{d}{N} M_\sigma, \\ \bar{\xi}_{\text{IS}}^{(t)} &= \mathbb{E} \left[\left\| \mathbb{E} \left[\widehat{s}_{i(t)}^N(\boldsymbol{\theta}^{(t)}) - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \mid \boldsymbol{\theta}^{(t)}, i(t) \right] \right\| \mid \boldsymbol{\theta}^{(t)} \right] \leq \frac{\sqrt{d}}{N} M_\xi, \end{aligned}$$

where M_σ and M_ξ are two finite and positive constants given by

$$\begin{aligned} M_\sigma &= \frac{12}{n} \sum_{i=1}^n \frac{\lambda_i^2 \sqrt{\beta_i}}{\zeta_i^2} \left(1 + \frac{250\lambda_i}{\zeta_i} + \frac{9\lambda_i^2}{\zeta_i^2} \right), \\ M_\xi &= \frac{4}{n} \sum_{i=1}^n \frac{\lambda_i^2 \beta_i^{1/4}}{\zeta_i^2} \left\{ 2 + \sqrt{3 \left(1 + \frac{250\lambda_i}{\zeta_i} + \frac{9\lambda_i^2}{\zeta_i^2} \right)} \right\}, \end{aligned}$$

with λ_i and β_i as in Assumption 1.

Proof. For any $t \in \mathbb{N}$, since the individual $i(t)$ is drawn uniformly independently of $\boldsymbol{\theta}^{(t)}$, we have

$$\begin{aligned} \bar{\sigma}_{\text{IS}}^{(t)} &= \mathbb{E} \left[\mathbb{E} \left[\left\| \widehat{s}_{i(t)}^N - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\|^2 \mid \boldsymbol{\theta}^{(t)}, i(t) \right] \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \widehat{s}_i^N - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i) \right\|^2 \mid \boldsymbol{\theta}^{(t)}, i(t) = i \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\nu_i^{\otimes N}(\cdot; \boldsymbol{\theta}^{(t)})} \left[\left\| \widehat{s}_i^N - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i) \right\|^2 \right] \end{aligned}$$

and, similarly

$$\begin{aligned}\bar{\xi}_{\text{IS}}^{(t)} &= \frac{1}{n} \sum_{i=1}^n \left\| \mathbb{E} \left[\widehat{s}_{i(t)}^N - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \mid \boldsymbol{\theta}^{(t)}, i(t) = i \right] \right\| \\ &= \frac{1}{n} \sum_{i=1}^n \left\| \mathbb{E}_{\nu_i^{\otimes N}(\cdot; \boldsymbol{\theta}^{(t)})} \left[\widehat{s}_i^N - \nabla \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i) \right] \right\|\end{aligned}$$

Consider an individual $i \in \{1, \dots, n\}$ and a parameter $\boldsymbol{\theta} \in \Theta$. As in [1], the r -th central moment, $r \in \mathbb{N}^*$, of a measurable function $h : \mathbb{R}^q \rightarrow \mathbb{R}$ with respect to the proposal distribution is denoted

$$m_r[h] = \mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[|h(\mathbf{V}) - \mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} [h(\mathbf{V})]|^r \right].$$

We also denote ϕ_k , $k \in \{1, \dots, d\}$, the application that returns the k -th component of the score, namely for all $\mathbf{V} \in \mathbb{R}^q$

$$\phi_k(\mathbf{V}) = \frac{\partial}{\partial \theta_k} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{V}).$$

For a N -sample $(\mathbf{V}_1, \dots, \mathbf{V}_N)$ from $\nu_i(\cdot; \boldsymbol{\theta})$, the importance sampling estimate $\widehat{s}_i^N(\boldsymbol{\theta}) = (\widehat{s}_{ik})_{1 \leq k \leq d}$ of the score g_i , as defined in Equation (5), can thereby be written as

$$\widehat{s}_{ik} = \frac{1}{\sum_{s=1}^N \rho_{\boldsymbol{\theta}, i}(\mathbf{V}_s)} \sum_{r=1}^N \rho_{\boldsymbol{\theta}, i}(\mathbf{V}_r) \phi_k(\mathbf{V}_r).$$

According to Theorem 2.3 in [1], we can guarantee control over the bias and mean squared error of the importance sampling estimate for the k -th component of the score, $k \in \{1, \dots, d\}$, provided that the following quantity is finite

$$\begin{aligned}C_{\text{MSE}, k}(\boldsymbol{\theta}) &= \frac{3}{p_{\boldsymbol{\theta}}(\mathbf{Y}_i)^2} \left\{ m_2[\phi_k \rho_{\boldsymbol{\theta}, i}] \right. \\ &\quad \left. + \frac{9}{p_{\boldsymbol{\theta}}(\mathbf{Y}_i)^2} \sqrt{m_4[\rho_{\boldsymbol{\theta}, i}] \mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[|\phi_k(\mathbf{V}) \rho_{\boldsymbol{\theta}, i}(\mathbf{V})|^4 \right]} \right. \\ &\quad \left. + \frac{125}{p_{\boldsymbol{\theta}}(\mathbf{Y}_i)} \sqrt{m_6[\rho_{\boldsymbol{\theta}, i}] \mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[|\phi_k(\mathbf{V})|^4 \right]} \right\}.\end{aligned}$$

To prove such a statement, first note that by definition

$$0 < \lambda_i, \quad \mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[|\phi_k(\mathbf{V})|^4 \right] \leq \beta_i. \quad (29)$$

The monotonicity of the expected value thus yields

$$\mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[|\phi_k(\mathbf{V}) \rho_{\boldsymbol{\theta}, i}(\mathbf{V})|^4 \right] \leq \lambda_i^4 \mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[|\phi_k(\mathbf{V})|^4 \right] \leq \lambda_i^4 \beta_i.$$

Furthermore, we use that for all $r \in \mathbb{N}^*$, $\mathbb{E}[|X - \mathbb{E}[X]|^r] \leq 2^r \mathbb{E}[|X|^r]$ (this is a direct consequence of the Minkowski and Jensen inequalities). Then, we have for all $r \in \mathbb{N}^*$

$$\begin{aligned} m_r[\rho_{\theta,i}] &\leq 2^r \mathbb{E}_{\nu_i(\cdot; \theta)} [\rho_{\theta,i}(\mathbf{V})^r] \leq 2^r \lambda_i^r, \\ m_2[\phi_k \rho_{\theta,i}] &\leq 4 \mathbb{E}_{\nu_i(\cdot; \theta)} [|\phi_k(\mathbf{V}) \rho_{\theta,i}(\mathbf{V})|^2] \leq 4 \lambda_i^2 \mathbb{E}_{\nu_i(\cdot; \theta)} [|\phi_k(\mathbf{V})|^2]. \end{aligned} \quad (30)$$

Finally, the Jensen inequality for the square root function provides that

$$\mathbb{E}_{\nu_i(\cdot; \theta)} [|\phi_k(\mathbf{V})|^2] \leq \sqrt{\mathbb{E}_{\nu_i(\cdot; \theta)} [|\phi_k(\mathbf{V})|^4]} \leq \sqrt{\beta_i} \quad (31)$$

Combining the upper-bounds from Equations (29)–(31) leads to

$$C_{\text{MSE},k}(\theta) \leq \frac{12\lambda_i^2 \sqrt{\beta_i}}{p_{\theta}(\mathbf{Y}_i)^2} \left\{ 1 + \frac{9\lambda_i^2}{p_{\theta}(\mathbf{Y}_i)^2} + \frac{250\lambda_i}{p_{\theta}(\mathbf{Y}_i)} \right\}.$$

Since λ_i and β_i are all finite constants, and for any $\theta \in \mathbb{R}^d$ and any $\mathbf{Y}_i \in \mathbb{N}$, $p_{\theta}(\mathbf{Y}_i)$ is positive, it follows that $C_{\text{MSE},k}(\theta)$ is finite. Consequently, Theorem 2.3 in [1] states that

$$\begin{aligned} \mathbb{E}_{\nu_i^{\otimes N}(\cdot; \theta)} \left[\left\{ \widehat{s}_{ik} - \frac{\partial}{\partial \theta_k} \log p_{\theta}(\mathbf{Y}_i) \right\}^2 \right] &\leq \frac{1}{N} C_{\text{MSE},k}(\theta), \\ \left| \mathbb{E}_{\nu_i^{\otimes N}(\cdot; \theta)} [\widehat{s}_{ik}] - \frac{\partial}{\partial \theta_k} \log p_{\theta}(\mathbf{Y}_i) \right| &\leq \frac{2}{N p_{\theta}(\mathbf{Y}_i)} \left\{ \frac{1}{p_{\theta}(\mathbf{Y}_i)} \sqrt{m_2[\rho_{\theta,i}] m_2[\overline{\phi}_k \rho_{\theta,i}]} \right. \\ &\quad \left. + \sqrt{C_{\text{MSE},k}(\theta) \mathbb{E}_{\nu_i(\cdot; \theta)} [\rho_{\theta,i}(\mathbf{V})^2]} \right\}, \end{aligned}$$

where $\overline{\phi}_k : \mathbf{V} \mapsto \phi_k(\mathbf{V}) - \mathbb{E}_{\nu_i(\cdot; \theta)} [\phi_k(\mathbf{U})]$. To conclude, we further need to eliminate the dependence in θ in these upper bounds. For any $\mathbf{Y}_i \in \mathbb{N}^p$, the function $\theta \mapsto p_{\theta}(\mathbf{Y}_i)$ is bounded below by $\zeta_i > 0$ on Θ (Proposition 4.3). Moreover, combining the argument of Equation (30) with the positivity of the variance and Equation (31), we have

$$m_2[\overline{\phi}_k \rho_{\theta,i}] \leq 4 \lambda_i^2 \mathbb{E}_{\nu_i(\cdot; \theta)} [|\overline{\phi}_k(\mathbf{V})|^2] \leq 4 \lambda_i^2 \mathbb{E}_{\nu_i(\cdot; \theta)} [|\phi_k(\mathbf{V})|^2] \leq 4 \lambda_i^2 \sqrt{\beta_i}.$$

Overall, these last two bounds in addition with Equations (29)–(31) yield the conclusion, namely

$$\begin{aligned} \overline{\sigma}_{\text{IS}}^{(t)} &= \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^d \mathbb{E}_{\nu_i^{\otimes N}(\cdot; \theta)} \left[\left\{ \widehat{s}_{ik} - \frac{\partial}{\partial \theta_k} \log p_{\theta}(\mathbf{Y}_i) \right\}^2 \right] \\ &\leq \frac{12d}{nN} \sum_{i=1}^n \frac{\lambda_i^2 \sqrt{\beta_i}}{\zeta_i^2} \left(1 + \frac{250\lambda_i}{\zeta_i} + \frac{9\lambda_i^2}{\zeta_i^2} \right), \end{aligned}$$

and

$$\begin{aligned}
\bar{\xi}_{\text{IS}}^{(t)} &= \frac{1}{n} \sum_{i=1}^n \sqrt{\sum_{k=1}^d \left| \mathbb{E}_{\nu_i^{\otimes N}(\cdot; \boldsymbol{\theta})} [\hat{s}_{ik}] - \frac{\partial}{\partial \theta_k} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i) \right|^2} \\
&\leq \frac{1}{n} \sum_{i=1}^n \frac{2\sqrt{d}}{N\zeta_i} \left\{ \frac{1}{\zeta_i} 4\lambda_i^2 \beta_i^{1/4} + \sqrt{\frac{12\lambda_i^4 \sqrt{\beta_i}}{\zeta_i^2} \left(1 + \frac{250\lambda_i}{\zeta_i} + \frac{9\lambda_i^2}{\zeta_i^2} \right)} \right\} \\
&\leq \frac{4\sqrt{d}}{nN} \sum_{i=1}^n \frac{\lambda_i^2 \beta_i^{1/4}}{\zeta_i^2} \left\{ 2 + \sqrt{3 \left(1 + \frac{250\lambda_i}{\zeta_i} + \frac{9\lambda_i^2}{\zeta_i^2} \right)} \right\}.
\end{aligned}$$

□

Theorem 4.5. Let $\boldsymbol{\theta}^{(1)} \in \Theta$ be an initial value and $\gamma_0 \in \mathbb{R}_+^*$ a user-specified initial learning rate. Under Model (2), if Assumption 1 holds, then for any $T \in \mathbb{N}^*$ and any real constant $\eta \in (0, 1/\max\{2\Gamma, L\}]$, the sequence $\{\boldsymbol{\theta}^{(t)}\}_{1 \leq t \leq T}$ defined by Algorithm 1 with $\gamma = \gamma_0/\sqrt{T}$ satisfies

$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\|G_{\eta}^{(t)}\|^2 \right] &\leq \frac{2\tau}{\gamma_0(L\eta + 1)\sqrt{T}} \left(\ell(\boldsymbol{\theta}^{(1)}) - \ell(\boldsymbol{\theta}^{\text{MLE}}) \right. \\
&\quad \left. + \frac{\gamma_0^2(L\eta + 1)}{2\eta} \left[\Delta^2 + \Gamma^2 + \frac{d}{N} \left\{ M_{\sigma} + \frac{2(\Delta + \Gamma)}{\sqrt{d}} M_{\xi} \right\} \right] \right) \\
&\quad + \frac{2\tau\sqrt{d}}{\eta N} M_{\xi},
\end{aligned} \tag{13}$$

with L the smoothness constant of $\ell(\cdot)$, $\Gamma = \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla \ell(\boldsymbol{\theta})\|$, constants M_{σ} and M_{ξ} as defined in Proposition 4.4, and

$$\tau = \frac{(2L\eta + 1)^2}{L\eta + 1} \left(1 + \sqrt{\frac{L\eta}{L\eta + 1}} \right)^2, \quad \Delta = \max_{i=1, \dots, n} \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i)\|.$$

Proof. According to Proposition 4.3, the function ℓ is L -smooth. Using successively that the square function and the expectation are increasing on \mathbb{R}_+ , for any real constant $\eta > 0$, Lemma B.3 yields

$$\mathbb{E} \left[\|G_{\eta}^{(t)}\|^2 \right] \leq \left(1 + \frac{L\eta}{L\eta + 1} \right)^2 \left(1 + \sqrt{\frac{L\eta}{L\eta + 1}} \right)^2 \mathbb{E} \left[\|\nabla F_{\frac{\eta}{L\eta + 1}}(\boldsymbol{\theta}^{(t)})\|^2 \right]. \tag{32}$$

To elaborate on the upper bound, we aim at using Lemma 4.2. Therefore, we first need to prove that we have control over the bias $\xi^{(t)}$ and the mean squared error $\sigma^{(t)}$ of the gradient estimates, as defined in Lemma 4.2.

Let $t \in \mathbb{N}^*$. Due to the bilinearity of the scalar product, the following identity holds:

$$\begin{aligned} \sigma^{(t)} &= \mathbb{E} \left[\left\| \widehat{g}^{(t)} + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\|^2 + \left\| \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\|^2 \mid \boldsymbol{\theta}^{(t)} \right] \\ &\quad - 2\mathbb{E} \left[\left\langle \widehat{g}^{(t)} + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}), \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\rangle \mid \boldsymbol{\theta}^{(t)} \right] \end{aligned}$$

Note that the first term of the right hand side corresponds to $\bar{\sigma}_{\text{IS}}^{(t)}$. Let bound the two remaining terms. According to Proposition 4.1, for all $i = 1, \dots, n$, the function

$$\boldsymbol{\theta} \mapsto \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i)$$

is continuous on Θ , and thereby bounded. Denote

$$\Delta = \max_{i=1, \dots, n} \sup_{\boldsymbol{\theta} \in \Theta} \|\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i)\|.$$

We have

$$\mathbb{E} \left[\left\| \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\|^2 \mid \boldsymbol{\theta}^{(t)} \right] \leq \Delta^2$$

Moreover, $i(t)$ being a uniform random variable on $\{1, \dots, n\}$ independent of the sequence $\boldsymbol{\theta}^{(t)}$,

$$\begin{aligned} &\mathbb{E} \left[\left\langle \widehat{g}^{(t)} + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}), \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\rangle \mid \boldsymbol{\theta}^{(t)} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\left\langle -\widehat{s}_{i(t)}^N(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}), \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\rangle \mid \boldsymbol{\theta}^{(t)}, i(t) \right] \right] \\ &= \mathbb{E} \left[\left\langle -\mathbb{E} \left[\widehat{s}_{i(t)}^N(\boldsymbol{\theta}^{(t)}) \mid \boldsymbol{\theta}^{(t)}, i(t) \right] + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}), \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\rangle \right] \\ &= -\frac{1}{n} \sum_{i=1}^n \left\langle \mathbb{E}_{\nu_i^{\otimes N}(\cdot, \boldsymbol{\theta}^{(t)})} \left[\widehat{s}_i^N(\boldsymbol{\theta}^{(t)}) \right] - \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i), \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i) \right\rangle. \end{aligned}$$

Providing Assumption 1 holds, the Cauchy-Schwarz inequality along with Proposition 4.4 yields

$$\begin{aligned} &-2\mathbb{E} \left[\left\langle \widehat{g}^{(t)} + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}), \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \right\rangle \mid \boldsymbol{\theta}^{(t)} \right] \\ &\leq \frac{2}{n} \sum_{i=1}^n \left\| \mathbb{E}_{\nu_i^{\otimes N}(\cdot, \boldsymbol{\theta}^{(t)})} \left[\widehat{s}_i^N(\boldsymbol{\theta}^{(t)}) \right] - \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i) \right\| \left\| \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}^{(t)}) + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i) \right\| \\ &\leq \frac{2\Delta\sqrt{d}}{N} M_{\xi}. \end{aligned}$$

In conclusion, for any $t \in \mathbb{N}^*$, we have a finite upper bound for the mean squared error $\sigma^{(t)}$, namely

$$\sigma^{(t)} \leq \frac{d}{N} M_{\sigma} + \Delta^2 + \frac{2\Delta\sqrt{d}}{N} M_{\xi} = \Delta^2 + \frac{d}{N} \left(M_{\sigma} + \frac{2\Delta}{\sqrt{d}} M_{\xi} \right).$$

Conversely, leveraging Assumption 1 once again, Proposition 4.4 shows that we have a finite upper bound for the bias $\xi^{(t)}$. Indeed

$$\begin{aligned}\xi^{(t)} &= \left\| \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[-\widehat{s}_{i(t)}^N \mid \boldsymbol{\theta}^{(t)}, i(t) = i \right] + \frac{1}{n} \sum_{i=1}^n \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_i) \right\| \\ &= \left\| \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[-\widehat{s}_{i(t)}^N + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \mid \boldsymbol{\theta}^{(t)}, i(t) = i \right] \right\| \\ &\leq \frac{1}{n} \sum_{i=1}^n \left\| \mathbb{E} \left[-\widehat{s}_{i(t)}^N + \nabla_{\boldsymbol{\theta}} \log p_{\boldsymbol{\theta}^{(t)}}(\mathbf{Y}_{i(t)}) \mid \boldsymbol{\theta}^{(t)}, i(t) = i \right] \right\| = \bar{\xi}_{\text{IS}}^{(t)} \leq \frac{\sqrt{d}}{N} M_{\xi}.\end{aligned}$$

Consequently, Lemma 4.2 applies for any real constant $\eta \in (0, 1/\max\{2\Gamma, L\}]$, as

$$\frac{\eta}{L\eta + 1} \leq \frac{1}{L + \max\{2\Gamma, L\}} = \frac{1}{\max\{2\Gamma + L, 2L\}}.$$

After integrating both sides of Equation (11) with respect to $\boldsymbol{\theta}^{(t)}$ and plugging in the aforementioned upper bounds, we get

$$\begin{aligned}\frac{\gamma}{2} \mathbb{E} \left[\left\| \nabla F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(t)}) \right\|^2 \right] &\leq \mathbb{E} \left[F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(t)}) \right] - \mathbb{E} \left[F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(t+1)}) \right] \\ &\quad + \frac{(\gamma + \gamma^2\Gamma)(L\eta + 1)\sqrt{d}}{\eta N} M_{\xi} \\ &\quad + \frac{\gamma^2(L\eta + 1)}{2\eta} \left\{ \Delta^2 + \frac{d}{N} \left(M_{\sigma} + \frac{2\Delta}{\sqrt{d}} M_{\xi} \right) + \Gamma^2 \right\}.\end{aligned}$$

Then, summing along the iterations $t = 1, \dots, T$ yields

$$\begin{aligned}\sum_{t=1}^T \mathbb{E} \left[\left\| \nabla F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(t)}) \right\|^2 \right] &\leq \frac{2}{\gamma} \left\{ F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(1)}) - \mathbb{E} \left[F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(T+1)}) \right] \right\} \\ &\quad + \frac{2(L\eta + 1)T\sqrt{d}}{\eta N} M_{\xi} + \frac{\gamma(L\eta + 1)T}{\eta} \left[\Delta^2 \right. \\ &\quad \left. + \frac{d}{N} \left\{ M_{\sigma} + \frac{2(\Delta + \Gamma)}{\sqrt{d}} M_{\xi} \right\} + \Gamma^2 \right].\end{aligned}$$

For $\boldsymbol{\theta}^{\text{MLE}} = \arg \min_{\boldsymbol{\theta} \in \Theta} \ell(\boldsymbol{\theta})$, according to the definition of the Moreau envelope, for any $\boldsymbol{\theta} \in \Theta$

$$\ell(\boldsymbol{\theta}^{\text{MLE}}) \leq F_{\eta/(L\eta+1)}(\boldsymbol{\theta}),$$

and thus,

$$\begin{aligned}\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\left\| \nabla F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(t)}) \right\|^2 \right] &\leq \frac{2}{\gamma T} \left\{ F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(1)}) - \ell(\boldsymbol{\theta}^{\text{MLE}}) \right\} \\ &\quad + \frac{2(L\eta + 1)\sqrt{d}}{\eta N} M_{\xi} + \frac{\gamma(L\eta + 1)}{\eta} \left[\Delta^2 \right. \\ &\quad \left. + \frac{d}{N} \left\{ M_{\sigma} + \frac{2(\Delta + \Gamma)}{\sqrt{d}} M_{\xi} \right\} + \Gamma^2 \right].\end{aligned}$$

Combining the last inequality with Equation (32) and using $F_{\eta/(L\eta+1)}(\boldsymbol{\theta}^{(1)}) \leq \ell(\boldsymbol{\theta}^{(1)})$ since $\boldsymbol{\theta}^{(1)} \in \Theta$, we get for $\gamma = \gamma_0/\sqrt{T}$

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\|G_\eta^{(t)}\|^2 \right] &\leq \frac{2\tau}{\gamma_0(L\eta+1)\sqrt{T}} \left\{ F_{\frac{\eta}{L\eta+1}}(\boldsymbol{\theta}^{(1)}) - \ell(\boldsymbol{\theta}^{\text{MLE}}) \right\} + \frac{2\tau\sqrt{d}}{\eta N} M_\xi \\ &\quad + \frac{\gamma_0\tau}{\eta\sqrt{T}} \left[\Delta^2 + \frac{d}{N} \left\{ M_\sigma + \frac{2(\Delta+\Gamma)}{\sqrt{d}} M_\xi \right\} + \Gamma^2 \right]. \end{aligned}$$

□

Appendix D: Results on the importance sampling proposal distribution

Lemma 5.1. *Let $\alpha \in (0, 1]$ and $\delta > 1$. If for any $i = 1, \dots, n$, $\boldsymbol{\theta} \mapsto \boldsymbol{\mu}_i(\boldsymbol{\theta}) \in \mathbb{R}^q$ and $\boldsymbol{\theta} \mapsto \mathbf{S}_i(\boldsymbol{\theta}) \in \mathcal{S}_{++}^q$ are continuous on Θ , then the proposal distribution defined by*

$$\nu_i(\cdot; \boldsymbol{\theta}) = \mathcal{GM}(\cdot; \boldsymbol{\mu}_i(\boldsymbol{\theta}), \mathbf{S}_i(\boldsymbol{\theta}), \alpha, \delta) \quad (15)$$

fulfils Assumption 1.

Proof. Let $i \in \{1, \dots, n\}$, $\alpha \in (0, 1]$ and $\delta > 1$.

(A1.1) Each component of a mixture distribution being a non-negative function, we have for any $\mathbf{V} \in \mathbb{R}^q$ and any $\boldsymbol{\theta} \in \mathbb{R}^d$,

$$\rho_{\boldsymbol{\theta}, i}(\mathbf{V}) = \frac{p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{V})}{\mathcal{GM}(\mathbf{V}; \boldsymbol{\mu}_i(\boldsymbol{\theta}), \mathbf{S}_i(\boldsymbol{\theta}), \alpha, \delta)} \leq \frac{1}{\alpha} \frac{p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{V})}{\mathcal{N}(\mathbf{V}; \boldsymbol{\mu}_i(\boldsymbol{\theta}), \delta \mathbf{I}_q)}.$$

On the compact set Θ , there exist real constants $K_i^\Theta > 0$ and κ_i^Θ (Lemma A.1) such that

$$\log \frac{p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{V})}{\mathcal{N}(\mathbf{V}; \boldsymbol{\mu}_i(\boldsymbol{\theta}), \delta \mathbf{I}_q)} \leq K_i^\Theta \|\mathbf{V}\| + \kappa_i^\Theta - \frac{1}{2} \|\mathbf{V}\|^2 + \frac{1}{2\delta} \|\mathbf{V} - \boldsymbol{\mu}_i(\boldsymbol{\theta})\|^2 + \frac{q}{2} \log(\delta).$$

Using the Cauchy-Schwarz inequality, we get

$$\|\mathbf{V} - \boldsymbol{\mu}_i(\boldsymbol{\theta})\|^2 = \|\mathbf{V}\|^2 + \|\boldsymbol{\mu}_i(\boldsymbol{\theta})\|^2 - 2\langle \boldsymbol{\mu}_i(\boldsymbol{\theta}), \mathbf{V} \rangle \leq \|\mathbf{V}\|^2 + \|\boldsymbol{\mu}_i(\boldsymbol{\theta})\|^2 + 2\|\mathbf{V}\| \|\boldsymbol{\mu}_i(\boldsymbol{\theta})\|.$$

$\boldsymbol{\theta} \mapsto \boldsymbol{\mu}_i(\boldsymbol{\theta})$ is continuous on the compact set Θ , and hence bounded, say by u_* . It follows that

$$\|\mathbf{V} - \boldsymbol{\mu}_i(\boldsymbol{\theta})\|^2 \leq \|\mathbf{V}\|^2 + u_*^2 + 2u_* \|\mathbf{V}\|.$$

Consequently, we can derive an upper bound independent of $\boldsymbol{\theta}$, namely

$$\frac{p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{V})}{\mathcal{N}(\mathbf{V}; \boldsymbol{\mu}_i(\boldsymbol{\theta}), \delta \mathbf{I}_q)} \leq \exp \left[\left(K_i^\Theta + \frac{u_*}{\delta} \right) \|\mathbf{V}\| - \frac{\delta-1}{2\delta} \|\mathbf{V}\|^2 + \kappa_i^\Theta + \frac{u_*^2}{2\delta} + \frac{q}{2} \log(\delta) \right].$$

Providing $\delta > 1$, the quadratic term in \mathbf{V} overweighs any other terms and the supremum in \mathbf{V} is finite:

$$\sup_{(\boldsymbol{\theta}, \mathbf{V}) \in \Theta \times \mathbb{R}^q} \rho_{\boldsymbol{\theta}, i}(\mathbf{V}) \leq \frac{1}{\alpha} \sup_{(\boldsymbol{\theta}, \mathbf{V}) \in \Theta \times \mathbb{R}^q} \frac{p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{V})}{\mathcal{N}(\mathbf{V}; \boldsymbol{\mu}_i(\boldsymbol{\theta}), \delta \mathbf{I}_q)} < \infty.$$

(A1.2) Following Equation (22), there is a unique j in $\{1, \dots, p\}$ such that

$$\left| \frac{\partial}{\partial \theta_r} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{W}_i) \right| \leq (\|\mathbf{x}_i\| + \|\mathbf{W}_i\|) \{ \|\mathbf{Y}_i\| + \exp(Z_{ij}) \}.$$

We then have,

$$\begin{aligned} \left| \frac{\partial}{\partial \theta_r} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{W}_i) \right|^4 &\leq (8\|\mathbf{x}_i\|^4 + 8\|\mathbf{W}_i\|^4) \{8\|\mathbf{Y}_i\|^4 + 8\exp(4Z_{ij})\} \\ &\leq 64\|\mathbf{Y}_i\|^4 (\|\mathbf{x}_i\|^4 + \|\mathbf{W}_i\|^4) \\ &\quad + 64 \exp(4\mathbf{B}_j^\top \mathbf{x}_i + 4o_{ij}) \{ \|\mathbf{x}_i\|^4 + \|\mathbf{W}_i\|^4 \} \exp(4\mathbf{C}_j^\top \mathbf{W}_i), \end{aligned}$$

where $\mathbf{B}_j = (B_{1j}, \dots, B_{pj})^\top$ and $\mathbf{C}_j = (C_{j1}, \dots, C_{jq})^\top$ stand for the j -th column and row of \mathbf{B} and \mathbf{C} , respectively.

Let show that the upper bound admits a finite expectation with respect to a multivariate Gaussian distribution $\mathcal{N}(\mu, S)$, for any $\mu \in \mathbb{R}^q$ and $S \in \mathbb{S}_{++}^q$. To demonstrate this, we use the following identities. For any $\mathbf{w} \in \mathbb{R}^q$, a straightforward rewriting yields

$$\mathcal{N}(\mathbf{w}; \mu, S) \exp(4\mathbf{C}_j^\top \mathbf{w}) = \mathcal{N}(\mathbf{w}; \mu + 4S\mathbf{C}_j, S) \exp(4\mathbf{C}_j^\top \mu + 8\mathbf{C}_j^\top S\mathbf{C}_j). \quad (33)$$

Moreover, we have (see for instance [39])

$$\mathbb{E}_{\mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)} \left[\|\mathbf{W}_i\|^4 \right] \leq q(q+2).$$

Using that

$$\mathbb{E}_{\mathcal{N}(\mu, S)} \left[\|\mathbf{W}_i\|^4 \right] = \mathbb{E}_{\mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)} \left[\left\| \mu + S^{1/2} \mathbf{W}_i \right\|^4 \right] \leq 8\|\mu\|^4 + 8\|S^{1/2}\|^4 \mathbb{E}_{\mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)} \left[\|\mathbf{W}_i\|^4 \right],$$

we get

$$\mathbb{E}_{\mathcal{N}(\mu, S)} \left[\|\mathbf{W}_i\|^4 \right] \leq 8\|\mu\|^4 + 8q(q+2)\|S^{1/2}\|^4 \quad (34)$$

Equations (33) and (34) thus lead to

$$\begin{aligned} \Phi_j(\boldsymbol{\theta}, \mu, S) &= \exp(4\mathbf{C}_j^\top \mu + 8\mathbf{C}_j^\top S\mathbf{C}_j) \\ &= \mathbb{E}_{\mathcal{N}(\mu, S)} \left[\exp(4\mathbf{C}_j^\top \mathbf{W}_i) \right], \\ \Psi_j(\boldsymbol{\theta}, \mu, S) &= \exp(4\mathbf{C}_j^\top \mu + 8\mathbf{C}_j^\top S\mathbf{C}_j) \left\{ 8\|\mu + 4S\mathbf{C}_j\|^4 + 8q(q+2)\|S^{1/2}\|^4 \right\} \\ &\geq \mathbb{E}_{\mathcal{N}(\mu, S)} \left[\|\mathbf{W}_i\|^4 \exp(4\mathbf{C}_j^\top \mathbf{W}_i) \right]. \end{aligned}$$

Therefore, for any $\mu \in \mathbb{R}^q$ and $S \in \mathbb{S}_{++}^q$

$$\begin{aligned} \mathbb{E}_{\mathcal{N}(\mu, S)} \left[\left| \frac{\partial}{\partial \theta_r} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{W}_i) \right|^4 \right] &\leq 64\|\mathbf{Y}_i\|^4 \left\{ \|\mathbf{x}_i\|^4 + 8\|\mu\|^4 + 8q(q+2)\|S^{1/2}\|^4 \right\} \\ &\quad + 64 \exp(4\mathbf{B}_j^\top \mathbf{x}_i + 4o_{ij}) \|\mathbf{x}_i\|^4 \Phi_j(\boldsymbol{\theta}, \mu, S) \\ &\quad + 64 \exp(4\mathbf{B}_j^\top \mathbf{x}_i + 4o_{ij}) \Psi_j(\boldsymbol{\theta}, \mu, S). \end{aligned}$$

If $\boldsymbol{\theta} \mapsto \boldsymbol{\mu}_i(\boldsymbol{\theta})$ and $\boldsymbol{\theta} \mapsto \mathbf{S}_i(\boldsymbol{\theta})$ are continuous on Θ , the functions

$$\begin{aligned}\boldsymbol{\theta} &\mapsto \mathbf{S}_i(\boldsymbol{\theta})^{1/2}, \\ \boldsymbol{\theta} &\mapsto (1 - \alpha)\Phi_j(\boldsymbol{\theta}, \boldsymbol{\mu}_i(\boldsymbol{\theta}), \mathbf{S}_i(\boldsymbol{\theta})) + \alpha\Phi_j(\boldsymbol{\theta}, \boldsymbol{\mu}_i(\boldsymbol{\theta}), \delta\mathbf{I}_q), \\ \boldsymbol{\theta} &\mapsto (1 - \alpha)\Psi_j(\boldsymbol{\theta}, \boldsymbol{\mu}_i(\boldsymbol{\theta}), \mathbf{S}_i(\boldsymbol{\theta})) + \alpha\Psi_j(\boldsymbol{\theta}, \boldsymbol{\mu}_i(\boldsymbol{\theta}), \delta\mathbf{I}_q),\end{aligned}$$

are also continuous on the compact set Θ (by a composition argument), and therefore bounded on Θ , say by c_* , ϕ_* and ψ_* respectively. On the other hand,

$$\exp(4\mathbf{B}_j^\top \mathbf{x}_i + 4o_{ij}) \leq \exp(4\|\mathbf{x}_i\|\|\mathbf{B}_j\| + 4\|\mathbf{o}_i\|) \leq \exp(4\|\mathbf{x}_i\| \sup_{\boldsymbol{\theta} \in \Theta} \|\boldsymbol{\theta}\| + 4\|\mathbf{o}_i\|).$$

Consequently, with u_* the bound of $\boldsymbol{\theta} \mapsto \boldsymbol{\mu}_i(\boldsymbol{\theta})$ on Θ , for any $r \in \{1, \dots, d\}$,

$$\begin{aligned}\mathbb{E}_{\nu_i(\cdot; \boldsymbol{\theta})} \left[\left| \frac{\partial}{\partial \theta_r} \log p_{\boldsymbol{\theta}}(\mathbf{Y}_i, \mathbf{W}_i) \right|^4 \right] &\leq 64\|\mathbf{Y}_i\|^4 [\|\mathbf{x}_i\|^4 + 8u_* + 8q(q+2) \{(1-\alpha)c_*^4 + \alpha\delta^4\}] \\ &\quad + 64 \exp(4\|\mathbf{x}_i\| \sup_{\boldsymbol{\theta} \in \Theta} \|\boldsymbol{\theta}\| + 4\|\mathbf{o}_i\|) (\|\mathbf{x}_i\|^4 \phi_* + \psi_*),\end{aligned}$$

which yields the conclusion. \square

Lemma D.1. *When considering Model (2), for any individual $i = 1, \dots, n$, and any $k \in \mathbb{N}^*$, the function*

$$\boldsymbol{\theta} \mapsto \int_{\mathbb{R}^q} \mathbf{w}^k p_{\boldsymbol{\theta}}(d\mathbf{w} \mid \mathbf{Y}_i)$$

is continuous on \mathbb{R}^d .

Proof. Let $i \in \{1, \dots, n\}$, $k \in \mathbb{N}^*$, and $\Theta \subset \mathbb{R}^d$ a non-empty bounded and open set. For any $\mathbf{w} \in \mathbb{R}^q$ and $\boldsymbol{\theta} \in \Theta$, the Bayes rule yields

$$p_{\boldsymbol{\theta}}(\mathbf{w} \mid \mathbf{Y}_i) = \frac{p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{w}_i) \mathcal{N}(\mathbf{w}; \mathbf{0}_q, \mathbf{I}_q)}{p_{\boldsymbol{\theta}}(\mathbf{Y}_i)}.$$

Given $\mathbf{w} \in \mathbb{R}^q$, $\boldsymbol{\theta} \mapsto p_{\boldsymbol{\theta}}(\mathbf{Y}_i \mid \mathbf{W})$ is continuous on Θ , and $\boldsymbol{\theta} \mapsto p_{\boldsymbol{\theta}}(\mathbf{Y}_i)$ is continuous and positive on Θ (see Proposition 4.1 and its proof). This proves the continuity of $\boldsymbol{\theta} \mapsto \mathbf{w}^k p_{\boldsymbol{\theta}}(\mathbf{w} \mid \mathbf{Y}_i)$ on Θ .

Additionally, Lemma A.1 and Proposition 4.3 state that there are real constants $K_i^\Theta > 0$, κ_i^Θ , and $\zeta_i > 0$, such that for any $\boldsymbol{\theta} \in \Theta$

$$\|\mathbf{w}^k p_{\boldsymbol{\theta}}(\mathbf{w} \mid \mathbf{Y}_i)\| \leq \frac{\|\mathbf{w}\|^k}{\zeta_i} \exp \left\{ K_i^\Theta \|\mathbf{w}\| - \frac{1}{2} \|\mathbf{w}\|^2 + \kappa_i^\Theta - \frac{q}{2} \log(2\pi) \right\}.$$

The upper bound is Lebesgue integrable on \mathbb{R}^q and independent of $\boldsymbol{\theta} \in \Theta$. Consequently, the dominated convergence theorem yields the continuity on any non-empty bounded and open set. \square

Appendix E: Simulation study
E.1. Note on Rprop updates

Rprop assigns an individual learning rate to each parameter, which is adjusted based on the gradient. Formally, given $A^{(1)} = \mathbf{I}_d$ and $(\eta_+, \eta_-) = (1.2, 0.5)$, Rprop updates correspond to

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - A^{(t)} \widehat{\mathbf{g}}^{(t)},$$

with $A_{jk}^{(t)} = 0$ for $j \neq k$, $j, k = 1, \dots, d$ and

$$A_{jk}^{(t)} = \begin{cases} \eta_+ A_{kk}^{(t-1)} & \text{if } \widehat{g}_k^{(t)} \text{ has same sign than } \widehat{g}_k^{(t-1)} \\ \eta_- A_{kk}^{(t-1)} & \text{else} \end{cases}, \quad k = 1, \dots, d$$