



HAL
open science

Zero-inflation in the Multivariate Poisson Lognormal Family

Bastien Batardière, Julien Chiquet, François Gindraud, Mahendra Mariadassou

► **To cite this version:**

Bastien Batardière, Julien Chiquet, François Gindraud, Mahendra Mariadassou. Zero-inflation in the Multivariate Poisson Lognormal Family. *Annals of Applied Statistics*, In press. hal-04586927

HAL Id: hal-04586927

<https://hal.science/hal-04586927>

Submitted on 24 May 2024



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 - NoDerivatives 4.0 International License

ZERO-INFLATION IN THE MULTIVARIATE POISSON LOGNORMAL FAMILY

BY BASTIEN BATARDIÈRE^{1,a} JULIEN CHIQUET^{1,b} ,
FRANÇOIS GINDRAUD^{2,c} AND MAHENDRA MARIADASSOU^{3,d} 

¹Université Paris-Saclay, AgroParisTech, INRAE, UMR MIA Paris-Saclay, 91120, Palaiseau, France.,
^abastien.batardiere@inrae.fr; ^bjulien.chiquet@inrae.fr

²Université Lyon 1, CNRS, Laboratoire de Biométrie et Biologie Evolutive UMR 5558, F-69622, Villeurbanne, France.,
^cfrancois.gindraud@inrae.fr

³Université Paris-Saclay, INRAE, MaIAGE, 78350, Jouy-en-Josas, France, ^dmahendra.mariadassou@inrae.fr

Analyzing high-dimensional count data is a challenge and statistical model-based approaches provide an adequate and efficient framework that preserves explainability. The (multivariate) Poisson-Log-Normal (PLN) model is one such model: it assumes count data are driven by an underlying structured latent Gaussian variable, so that the dependencies between counts solely stems from the latent dependencies. However PLN doesn't account for zero-inflation, a feature frequently observed in real-world datasets. Here we introduce the Zero-Inflated PLN (ZIPLN) model, adding a multivariate zero-inflated component to the model, as an additional Bernoulli latent variable. The Zero-Inflation can be fixed, site-specific, feature-specific or depends on covariates. We estimate model parameters using variational inference that scales up to datasets with a few thousands variables and compare two approximations: (i) independent Gaussian and Bernoulli variational distributions or (ii) Gaussian variational distribution conditioned on the Bernoulli one. The method is assessed on synthetic data and the efficiency of ZIPLN is established even when zero-inflation concerns up to 90% of the observed counts. We then apply both ZIPLN and PLN to a cow microbiome dataset, containing 90.6% of zeroes. Accounting for zero-inflation significantly increases log-likelihood and reduces dispersion in the latent space, thus leading to improved group discrimination.

CONTENTS

1	Introduction	2
2	Model	4
3	Estimation by Variational Inference	6
	3.1 Choice of the variational family	6
	3.2 Expected lower bounds	7
4	Optimization	9
	4.1 Optimization of $J^{(1)}$	9
	4.2 Optimization of $J^{(2)}$	10
	4.3 Optimization using analytic law of $W_{ij} Y_{ij}$	10
	4.4 Implementation details	11
5	Simulation Study	11
	5.1 Experimental details	11
	5.2 Simulations when π^* fluctuate	12
	5.3 Simulations when XB^* fluctuate	13

Keywords and phrases: Count data, Poisson lognormal model, Zero inflated model, Variational Inference, Alternate Optimisation.

5.4	Simulations when n fluctuate	13
6	Application to cow microbiome data	17
6.1	Model selection	17
6.2	Modeling of counts	17
6.3	Latent means	18
7	Conclusion and discussion	19
	Funding	20
	Acknowledgments	20
	Implementation	20
	References	20
A	Technical results	22
A.1	ELBO derivation	22
A.2	Proofs	25
B	Additional simulations	26

1. Introduction. Count data appears in many different fields such as ecology, accidents analysis, single-cell RNA (scRNA) sequencing and metagenomics. For example, researchers may be interested in estimating the correlation between abundances of different species in a biome or expressions of different genes in a cell. More specifically, the model introduced in this paper is motivated by the increasing importance of microbiome studies. Broadly speaking, a microbiome is a collection of microbes, together with their genomes, found in a given habitat (*e.g.* plant leaves, human gut, waste water, etc.). The most widespread way of studying microbiomes is to amplify and sequence a marker gene, which acts as a molecular barcode. The sequences are processed through bioinformatics pipelines ([Escudié et al., 2017](#)) to produce Operational Taxonomic Units (OTUs) / Amplicon Sequence Variants (ASVs), a proxy for microbial species in microbial ecology, and enumerated to create count tables, recording the abundance of each OTU/ASV in each sample. Those tables are characterized by a very high fraction (ranging from 80 to 95%) of zero counts and a high number of variables.

Count data are hard to analyse as is and transformations must be performed beforehand in order to extract meaningful statistics. While log-transformation followed by Gaussian analyses is fast and widely used, it lacks sound statistical grounding and model-based approaches are much more suitable ([O’Hara and Kotze, 2010](#)). In particular, Negative-Binomial (NB) and Poisson-based models are popular choices for modeling count data and have been extensively used in RNAseq studies (see *e.g.* [Love, Huber and Anders, 2014](#)). The NB distribution, which uses Poisson emission law with Gamma-distributed parameter to induce overdispersion, is generally preferred to the standard Poisson distribution to satisfy the overdispersion property (higher variance than mean value) repeatedly observed in sequencing-based count data (including scRNA-seq, see [Choudhary and Satija, 2022](#)). However, these two distributions are ill-suited to model the fraction of zero counts independently from the mean count. The method described in [Lambert \(1992\)](#) incorporates a Dirac mass centered at zero to simulate zero-inflation, facilitating efficient execution of univariate differential analyses but disregards inter-variable dependencies. In order to address this problem, [Li et al. \(1999\)](#) extends the univariate Zero-Inflated Poisson (ZIP) model to Multivariate ZIP (MZIP) using a mixture of Poisson, a technique which is also used to extend univariate Zero-Inflated Negative-Binomial (ZINB) to multivariate ZINB ([Dong et al., 2014](#)). However, in both cases, the correlation between variables is very constrained. A bivariate zero-inflated negative-binomial model is investigated in [Cho et al. \(2023\)](#) to measure correlation between two genes in scRNAseq data but fails to scale to higher order dependencies.

The (multivariate) Poisson-Log-Normal (in short PLN, see [Aitchison and Ho, 1989](#)) model offers a general framework to multivariate count data by offering flexibility to describe dependencies between counts by means of a latent Gaussian variable. As a mixture of Poisson

with Log-Normal distributed parameters, PLN models results in overdispersion, just like the NB. However, the underlying Gaussian structure inherent to the PLN model makes correlation between variables natural, unlike its NB counterpart. More generally, the PLN model falls in the family of latent variable models (LVMs), and more specifically of multivariate generalized linear mixed models (mGLMMs) sometimes also called generalized linear latent variable models (GLLVMs). In those models, the distribution of observed responses usually belongs either to the exponential family (Bernoulli, Binomial, Poisson, Negative-Binomial, with or without Zero-Inflation, etc.) or the exponential dispersion model (Tweedie, etc.). Model parameters are related to linear combinations of latent variables (and possibly covariates) through a simple link function. Parameter estimation for common GLLVMs is efficiently implemented in some packages (Niku et al., 2019; Seabold and Perktold, 2010), making it a popular option for multivariate count data. However, while some models allows for dependency between variables in the latent space and other accounts for zero-inflation, no model accounts, to the best of our knowledge, for both at the same time.

We introduce here the Zero-Inflated Poisson Log-Normal (ZIPLN) model, based on the PLN model. ZIPLN benefits from the Gaussian structure of the PLN model, with an extra zero-inflated component. This extra layer adds flexibility to the model as its parameters can be chosen to be shared across the individuals, across the features or even to depend on its own set of covariates. As exact inference of (ZI)PLN is intractable and conditional laws are only partially known, we cannot rely on the Expectation-Maximization (EM) algorithm (Dempster, Laird and Rubin, 1977), as done for optimizing classical latent variable models. We instead rely on variational inference (Jaakkola and Jordan, 2000; Wainwright and Jordan, 2008; Hui et al., 2017; Blei, Kucukelbir and McAuliffe, 2017). Other approaches based on Monte Carlo techniques have been proposed (Jacquier, Johannes and Polson, 2007; Cappé et al., 2002; Stoehr and Robin, 2024) to infer the maximum likelihood estimator, but it does not scale with the dimension of the observations. Numerical integration can be performed (Aitchison and Ho, 1989) as an alternative to the variational approximation used here but becomes prohibitive when the number of dimensions exceeds 5. Here, we develop a Variational-EM algorithm where we propose two different variational approximations. The first assumes conditional independence between both components, resulting in a fast M step. By contrast, the second is slightly slower but leverages the dependence between components to use a more complex variational approximation.

Related works. ZINBWaVE, proposed by Risso et al. (2018) is the closest work to ours, modelling zero-inflation (resp. counts) as a logistic (resp. log-linear) regression involving sample-level, feature-level and (unobserved) sample-level covariates, where the unobserved covariates are presumed to be unwanted variations and captured through latent factors. This model however suffers from a lack of identifiability and is mostly interested in estimating the probability that a null count arises from zero-inflation. We distinguish ourselves from ZINBWaVE via identifiability of parameters and most importantly via the inherent and explicit dependency structure between variables.

Recently, deep neural networks have been proposed to model count data. Variational autoencoders (VAE) (Kingma and Welling, 2022) are particularly efficient, performing dimension reduction via a latent variable framework. Zhao et al. (2020) proposes a VAE to model high dimensional overdispersed count data based on the NB distribution and Jin et al. (2020) models sparse and imbalanced count data with VAE. Many VAE models have been proposed for the sole purpose of studying scRNA-seq data (see e.g. Lopez et al. (2018); Choi, Li and Quon (2022); Xu et al. (2023); Wang and Gu (2018)). Although model-based and highly effective for predictions, VAE remain significantly harder to interpret in terms of coefficients, outputs and impact of structuring factors of interest than their statistical latent variable models counterparts.

Our paper is organized as follows: in Section 2, we introduce the PLN model followed by the ZIPLN model. In Section 3, we discuss the variational inference and choices made on the variational families. In Section 4, we discuss the optimization strategy. We study the model performances on simulated data in Section 5, and apply it to a cow microbiome dataset in Section 6. We conclude in Section 7 with discussions and possible improvements.

2. Model.

Background: Multivariate Poisson lognormal-model. The multivariate Poisson lognormal model relates a p -dimensional observation count vector $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{ip}) \in \mathbb{N}^p$ to a p -dimensional vector of Gaussian latent variables $\mathbf{Z}_i \in \mathbb{N}^p$ with precision matrix $\mathbf{\Omega}$ (that is, covariance matrix $\mathbf{\Sigma} \triangleq \mathbf{\Omega}^{-1}$). We adopt a formulation of PLN close to a multivariate generalized linear model, where the main effect is due to a linear combination of d covariates $\mathbf{x}_i \in \mathbb{R}^d$ (including an intercept). We also let the possibility to add some offsets for the p variables in each sample, that is $\mathbf{o}_i \in \mathbb{R}^p$:

$$(1) \quad \begin{array}{l} \text{latent space } \mathbf{Z}_i \sim \mathcal{N}(\mathbf{x}_i^\top \mathbf{B}, \mathbf{\Omega}^{-1}), \\ \text{observation space } Y_{ij} | Z_{ij} \text{ indep. } \mathbf{Y}_i | \mathbf{Z}_i \sim \mathcal{P}(\exp\{\mathbf{o}_i + \mathbf{Z}_i\}). \end{array}$$

The $d \times p$ matrix \mathbf{B} is the latent matrix of regression parameters. The latent covariance matrix $\mathbf{\Sigma}$ describes the underlying residual structure of dependence between the p variables, once the covariates are accounted for. We denote by $\mathbf{Y}, \mathbf{O}, \mathbf{X}$ the observed matrices with respective sizes $n \times p, n \times p$ and $n \times d$ stacking row-wise the vectors of counts, offsets and covariates (respectively $\mathbf{Y}_i, \mathbf{o}_i$ and \mathbf{x}_i). We also denote by \mathbf{Z} the $n \times p$ matrix of unobserved latent Gaussian vectors \mathbf{Z}_i .

Zero-inflated PLN regression model. We now aim to model an excess of zeroes in the data by adding zero-inflation to the standard PLN model (1), so that the zeroes in \mathbf{Y}_i arise from two different sources: either from a component where zero is the only possible value, or from a standard PLN component like in Equation (1). This two-component mixture is defined thanks to an additional latent vector $\mathbf{W}_i = (W_{i1}, \dots, W_{ip}) \in \mathbb{R}^p$ of Bernoulli random variables, parametrized by probabilities $\boldsymbol{\pi}_i = (\pi_{i1}, \dots, \pi_{ip})$ describing the probability that variable j in sample i belongs to the pure zero component:

$$(2) \quad \begin{array}{l} \text{PLN latent space } \mathbf{Z}_i = (Z_{ij})_{j=1\dots p} \sim \mathcal{N}(\mathbf{x}_i^\top \mathbf{B}, \mathbf{\Omega}^{-1}), \\ \text{excess of zero } \mathbf{W}_i = (W_{ij})_{j=1\dots p} \sim \otimes_{j=1}^p \mathcal{B}(\pi_{ij}), \\ \text{observation space } Y_{ij} | W_{ij}, Z_{ij} \sim^{\text{indep}} W_{ij} \delta_0 + (1 - W_{ij}) \mathcal{P}(\exp\{o_{ij} + Z_{ij}\}), \end{array}$$

where δ_0 is the Dirac distribution and we note $\boldsymbol{\pi}$ the matrix obtained by stacking the vectors $(\boldsymbol{\pi}_1^\top, \dots, \boldsymbol{\pi}_n^\top)$. Our model is flexible enough to accommodate different parametrizations for π_{ij} based on the availability of covariates and/or modeling choices made by the user. We consider three variants:

$$(3a) \quad \pi_{ij} = \pi \in [0, 1] \quad (\text{non-dependent - ND})$$

$$(3b) \quad \pi_{ij} = \text{logit}^{-1}(\mathbf{X}^0 \mathbf{B}^0)_{ij}, \quad \mathbf{X}^0 \in \mathbb{R}^{n \times d_0}, \quad \mathbf{B}^0 \in \mathbb{R}^{d_0 \times p} \quad (\text{column-wise dependence - CD})$$

$$(3c) \quad \pi_{ij} = \text{logit}^{-1}(\bar{\mathbf{B}}^0 \bar{\mathbf{X}}^0)_{ij}, \quad \bar{\mathbf{B}}^0 \in \mathbb{R}^{n \times d_0}, \quad \bar{\mathbf{X}}^0 \in \mathbb{R}^{d_0 \times p} \quad (\text{row-wise dependence - RD})$$

where $\text{logit}^{-1}(\cdot)$ is the logistic (or inverse logit) function, $d_0 \geq 1$, \mathbf{B}^0 (resp. $\bar{\mathbf{B}}^0$) are regression coefficients associated with row-wise matrix of covariates \mathbf{X}^0 (resp. column-wise covariates $\bar{\mathbf{X}}^0$), obtained by stacking the vectors $((\mathbf{x}_1^0)^\top, \dots, (\mathbf{x}_n^0)^\top)$, which may or may not be the same as in \mathbf{X} , the matrix of covariates in the PLN component. The ND inflation setting

allows the zero-inflation component to be shared along all individuals and variables, a simple but slightly unrealistic assumption, while RD (resp. CD) allows zero-inflation to be shared along all variables (resp. individuals) as it depends only on the individual (resp. variable) covariates. CD is useful when some variables are prone to zero-inflation across individuals (*e.g.* taxa for which the marker gene fails to amplify) whereas RD is useful when the set of zero-inflated variables rather depends on the individual's characteristics (*e.g.* acidophile taxa in soils with high pH). Note that when the column-wise (resp. row-wise) covariates reduce to a vector of one $\mathbf{1}_n$ (resp. $\mathbf{1}_p^\top$), the model corresponds to an inflation towards zero shared across rows (resp. columns) with vector of parameters $\boldsymbol{\pi} = (\pi_j) \in [0, 1]^p$ (resp. $\boldsymbol{\pi} = (\pi_i) \in [0, 1]^n$). We refer to Model 2 as the ZIPLN regression model.

Using standard results on Poisson and Gaussian distribution, we easily derive the expectation and variance of the ZIPLN regression model. Letting $A_{ij} \triangleq \exp(o_{ij} + \mu_{ij} + \sigma_{jj}/2)$ with $\mu_{ij} = \mathbf{x}_i^\top \mathbf{B}_j$, then

$$\mathbb{E}(Y_{ij}) = (1 - \pi_{ij})A_{ij} \geq 0,$$

$$\mathbb{V}(Y_{ij}) = (1 - \pi_{ij})A_{ij} + (1 - \pi_{ij})A_{ij}^2(e^{\sigma_{jj}} - (1 - \pi_{ij})).$$

In the following, we are interested in inferring the vector of parameters $\boldsymbol{\theta}$ where $\boldsymbol{\theta} = (\boldsymbol{\Omega}, \mathbf{B}, \boldsymbol{\pi})$, $\boldsymbol{\theta} = (\boldsymbol{\Omega}, \mathbf{B}, \mathbf{B}^0)$ and $\boldsymbol{\theta} = (\boldsymbol{\Omega}, \mathbf{B}, \bar{\mathbf{B}}^0)$ for Models 3a, 3b and 3c respectively, where $\boldsymbol{\Omega} \in \mathbb{S}_p^{++}$, $\mathbf{B} \in \mathbb{R}^{d \times p}$, $\boldsymbol{\pi} \in [0, 1]$, $\mathbf{B}^0 \in \mathbb{R}^{d_0 \times p}$ and $\bar{\mathbf{B}}^0 \in \mathbb{R}^{n \times d_0}$ with \mathbb{S}_p^{++} the set of $p \times p$ positive-definite matrices. We first show that Model 2 is identifiable.

Identifiability of ZIPLN models. Identifiability results are available for the ZI Poisson model (Li, 2012) and can be generalized to the ZIPLN regression model. To this end, we first consider the simple ZIPLN model, (*i.e.* a ZIPLN model without covariate), with a single sample, in order to drop the index i :

$$\begin{aligned} \mathbf{W} &= (W_j)_{j=1\dots p} \sim \mathcal{B}^{\otimes}(\boldsymbol{\pi}) = (\pi_1) \otimes \dots \otimes (\pi_p) \\ \mathbf{Z} &= (Z_j)_{j=1\dots p} \sim \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Omega}^{-1}) \\ Y_j | W_j, Z_j &\sim W_j \delta_0 + (1 - W_j) \mathcal{P}(e^{Z_j}), \quad Y_j \perp Y_k | \mathbf{W}, \mathbf{Z} \end{aligned} \tag{4}$$

PROPOSITION 1. *The simple ZIPLN model defined in (4) with parameter $\boldsymbol{\theta} = (\boldsymbol{\Omega}, \boldsymbol{\mu}, \boldsymbol{\pi})$ and parameter space $\mathbb{S}_p^{++} \times \mathbb{R}^p \times (0, 1)^p$ is identifiable.*

The proof relies on the method of moments and is postponed to the appendix. We now use this result to prove identifiability of the ZIPLN regression Model 3b (proof for 3c is achieved by replacing $\mathbf{X}^0 \mathbf{B}$ with $\bar{\mathbf{B}}^0 \bar{\mathbf{X}}$).

PROPOSITION 2. *The ZIPLN regression Model 2 with zero-inflation defined as in Equation (3b) and parameter $\boldsymbol{\theta} = (\boldsymbol{\Omega}, \mathbf{B}, \mathbf{B}^0)$ and parameter space $\mathbb{S}_p^{++} \times \mathcal{M}_{p,d}(\mathbb{R}) \times \mathcal{M}_{p,d}(\mathbb{R})$ is identifiable if and only if both $n \times d$ and $n \times d_0$ matrices of covariates \mathbf{X} and \mathbf{X}^0 are full rank.*

PROOF. If \mathbf{X} (resp. \mathbf{X}^0) is not full rank, there exists $\mathbf{B} \neq \mathbf{B}'$ such that $\mathbf{X}\mathbf{B} = \mathbf{X}\mathbf{B}'$ (resp. $\mathbf{X}^0\mathbf{B} = \mathbf{X}^0\mathbf{B}'$) and therefore the map $\boldsymbol{\theta} \mapsto p_{\boldsymbol{\theta}}$ is not one-to-one. We know from Proposition 1 that $\mathbf{O} + \mathbf{X}\mathbf{B}$ and $\boldsymbol{\pi} = \text{logit}^{-1}(\mathbf{X}^0\mathbf{B}^0)$ are identifiable. Since the affine function and the logit^{-1} are both one-to-one, parameters \mathbf{B} and \mathbf{B}^0 are identifiable as soon as the maps $\mathbf{B} \mapsto \mathbf{X}\mathbf{B}$ and $\mathbf{B}^0 \mapsto \mathbf{X}^0\mathbf{B}^0$ are injective, which is the case as soon as \mathbf{X} and \mathbf{X}^0 are full rank. \square

3. Estimation by Variational Inference. Our goal is to maximize the marginal likelihood. In the framework of latent models, a standard approach (e.g. with *Expectation-Maximization* algorithms) uses the following decomposition by integrating over the latent variables \mathbf{W}, \mathbf{Z}

$$(5) \quad \log p_\theta(\mathbf{Y}) = \log \frac{p_\theta(\mathbf{Z}, \mathbf{W}, \mathbf{Y})}{p_\theta(\mathbf{Z}, \mathbf{W}|\mathbf{Y})} = \int_{\mathbf{W}, \mathbf{Z}} \log \frac{p_\theta(\mathbf{Z}, \mathbf{W}, \mathbf{Y})}{p_\theta(\mathbf{Z}, \mathbf{W}|\mathbf{Y})} p_\theta(\mathbf{Z}, \mathbf{W}|\mathbf{Y}) d\mathbf{W} d\mathbf{Z}.$$

However, for the ZIPLN model, it is untractable since the conditional distribution $p_\theta(\mathbf{Z}, \mathbf{W}|\mathbf{Y})$ has no closed-form. To overcome this issue, we rely on a variational approximation of this distribution which will yield a lower bound of $\log p_\theta(\cdot)$: for observation i , we denote by $\tilde{p}_\psi(\mathbf{Z}_i, \mathbf{W}_i)$ the approximation of $p_\theta(\mathbf{Z}_i, \mathbf{W}_i|\mathbf{Y}_i)$ where ψ is a set of variational parameters to be optimized. Subtracting to the untractable Expression (5) of the log-likelihood the positive (and also untractable) quantity (known as the Kullback-Leibler divergence)

$$KL(\tilde{p}_\psi(\cdot)\|p_\theta(\cdot|\mathbf{Y})) = \int_{\mathbf{W}, \mathbf{Z}} \log \frac{\tilde{p}_\psi(\mathbf{Z}, \mathbf{W})}{p_\theta(\mathbf{Z}, \mathbf{W}|\mathbf{Y})} \tilde{p}_\psi(\mathbf{Z}, \mathbf{W}) d\mathbf{W} d\mathbf{Z}$$

results after some rearrangements in the following Evidence Lower BOund (ELBO):

$$\begin{aligned} J(\theta, \psi) &= \log p_\theta(\mathbf{Y}) - KL(\tilde{p}_\psi(\cdot)\|p_\theta(\cdot|\mathbf{Y})) \\ &= \int_{\mathbf{W}, \mathbf{Z}} \log \frac{p_\theta(\mathbf{Z}, \mathbf{W}, \mathbf{Y})}{\tilde{p}_\psi(\mathbf{Z}, \mathbf{W})} \tilde{p}_\psi(\mathbf{Z}, \mathbf{W}) d\mathbf{W} d\mathbf{Z} \\ (6) \quad &= \tilde{\mathbb{E}}[\log p_\theta(\mathbf{Z}, \mathbf{W}, \mathbf{Y})] - \tilde{\mathbb{E}}[\log \tilde{p}_\psi(\mathbf{Z}, \mathbf{W})]. \end{aligned}$$

This also looks like a plugin of integral Equation (5) with $p_\theta(\mathbf{Z}, \mathbf{W}|\mathbf{Y})$ replaced with $\tilde{p}_\psi(\mathbf{Z}, \mathbf{W})$. An appropriate choice of variational approximation will make the integral calculation tractable, while leading to an acceptable approximation of the log-likelihood (Blei, Kucukelbir and McAuliffe, 2017). The choice of the variational family is crucial as an inappropriate or too simplistic family can lead to bias and inconsistency in the resulting estimator (Westling and McCormick, 2019) whereas a too complex family would lead to an untractable optimization criterion.

3.1. Choice of the variational family.

Standard variational approximation. A straightforward, yet efficient, approach is to consider the mean field approximation, which breaks all dependencies between the vectors \mathbf{Z}_i and \mathbf{W}_i and their respective coordinates and approximates the conditional distribution as the product of its coordinate-wise marginals:

$$\tilde{p}_{\psi_i}^{(1)}(\mathbf{Z}_i, \mathbf{W}_i) \triangleq \tilde{p}_{\psi_i}(\mathbf{Z}_i) \tilde{p}_{\psi_i}(\mathbf{W}_i) = \otimes_{j=1}^p \tilde{p}_{\psi_i}(\mathbf{Z}_{ij}) \tilde{p}_{\psi_i}(\mathbf{W}_{ij}).$$

On top of that, we assume Gaussian and Bernoulli distribution for $\tilde{p}_{\psi_i}(\mathbf{Z}_{ij})$ and $\tilde{p}_{\psi_i}(\mathbf{W}_{ij})$ respectively, giving rise to the following variational approximation

$$(7) \quad \tilde{p}_{\psi_i}^{(1)}(\mathbf{Z}_i, \mathbf{W}_i) = \otimes_{j=1}^p \mathcal{N}(M_{ij}, S_{ij}^2) \mathcal{B}(P_{ij})$$

with $0 \leq P_{ij} \leq 1$ and $\psi_i = (M_{ij}, S_{ij}, P_{ij})_{1 \leq j \leq p}$. We denote \mathbf{M}, \mathbf{S} and \mathbf{P} the $n \times p$ matrices with respective entries M_{ij}, S_{ij} and P_{ij} ($1 \leq i \leq n, 1 \leq j \leq p$). This approximation therefore requires the estimation of $3np$ additional variational parameters on top of θ .

Enhanced variational approximation. As W_{ij} can take only two values, the dependence between \mathbf{Z}_{ij} and \mathbf{W}_{ij} can easily be made explicit by noting that

$$(8) \quad Z_{ij}|W_{ij}, Y_{ij} = (Z_{ij}|Y_{ij}, W_{ij} = 1)^{W_{ij}} (Z_{ij}|Y_{ij}, W_{ij} = 0)^{1-W_{ij}}.$$

The conditional distribution of $Z_{ij}|Y_{ij}, W_{ij} = 1$ simplifies to $Z_{ij}|W_{ij} = 1$ and is thus known as Z_{ij} and W_{ij} are independent: it follows a Gaussian distribution with mean $\mathbf{x}_i^\top \mathbf{B}_j$ and variance Σ_{jj} . By contrast, $Z_{ij}|Y_{ij}, W_{ij} = 0$ is untractable and approximated by a Gaussian distribution, giving rise to an alternative and slightly more involved variational approximation:

$$(9) \quad \tilde{p}_{\psi_i}^{(2)}(\mathbf{Z}_i, \mathbf{W}_i) = \otimes_{j=1}^p \mathcal{N}(\mathbf{x}_i^\top \mathbf{B}_j, \Sigma_{jj})^{W_{ij}} \mathcal{N}(M_{ij}, S_{ij}^2)^{1-W_{ij}} W_{ij}, \quad W_{ij} \sim^{\text{indep}} \mathcal{B}(P_{ij}).$$

The resulting ELBOs are tractable and detailed below.

REMARK 1. *It can be shown that $W_{ij}|Y_{ij}, Z_{ij} \sim \mathcal{B}\left(\sigma\left(\log\left(\frac{\pi_{ij}}{1-\pi_{ij}}\right) + Z_{ij}\right)\right) \mathbf{1}_{Y_{ij}=0}$ where $\sigma(\cdot) = \text{logit}^{-1}$, so that one could condition the other way around. While this conditional law is intuitive and appears simpler than Equation (8) at first glance, as it only involves a Bernoulli variable, it turns out to be untractable. Indeed, the resulting ELBO involves the entropy term $\tilde{\mathbb{E}}[\log \tilde{p}_\psi(\mathbf{W}|\mathbf{Z})]$, the computation of which requires computing expectations of the form $\mathbb{E}[\log(\sigma(U))\sigma(U)]$ for arbitrary univariate Gaussians $U \sim \mathcal{N}(\mu, \sigma^2)$, which are untractable when U is non-degenerated.*

3.2. Expected lower bounds. We set $\psi = (\psi_i)_{1 \leq i \leq n}$ the variational parameters of the variational distribution $\tilde{p}_\psi^{(1)} = \prod_{i=1}^n \tilde{p}_{\psi_i}^{(1)}$ (resp. $\tilde{p}_\psi^{(2)} = \prod_{i=1}^n \tilde{p}_{\psi_i}^{(2)}$) defined in Equation (7) (resp. Equation (9)). We denote by $\tilde{\mathbb{E}}^{(1)}$ (resp. $\tilde{\mathbb{E}}^{(2)}$) its expectation and by $J^{(1)}(\psi, \theta)$ (resp. $J^{(2)}(\psi, \theta)$) the corresponding ELBO, the expression of which is detailed in the next proposition.

PROPOSITION 3. *The ELBO defined in Equation (6) with variational approximation $\tilde{p}_\psi^{(1)}$ can be written in matrix form as*

$$(10) \quad \begin{aligned} J^{(1)}(\psi, \theta) = & \tilde{\mathbb{E}}^{(1)}[\log p_\theta(\mathbf{Y}|\mathbf{Z}, \mathbf{W})] + \tilde{\mathbb{E}}^{(1)}[\log p_\theta(\mathbf{W})] + H(\mathbf{P}) + \frac{1}{2} \text{Tr}(\mathbf{1}_{n,p}^\top \log(\mathbf{S}^2)) \\ & + \frac{n}{2} \log |\boldsymbol{\Omega}| - \frac{1}{2} \text{Tr}\left(\boldsymbol{\Omega} \left(\text{Diag}(\bar{\mathbf{S}}^2) + g(\mathbf{M} - \mathbf{X}\mathbf{B})\right)\right) + \frac{np}{2} \end{aligned}$$

and with variational approximation $\tilde{p}_\psi^{(2)}$ we get

$$(11) \quad \begin{aligned} J^{(2)}(\psi, \theta) = & \tilde{\mathbb{E}}^{(2)}[\log p_\theta(\mathbf{Y}|\mathbf{Z}, \mathbf{W})] + \tilde{\mathbb{E}}^{(2)}[\log p_\theta(\mathbf{W})] + H(\mathbf{P}) + \frac{1}{2} \text{Tr}(\mathbf{Q}^\top \log(\mathbf{S}^2)) \\ & + \frac{n}{2} \log |\boldsymbol{\Omega}| - \frac{1}{2} \text{Tr}\left(\boldsymbol{\Omega} \left(\text{Diag}(\mathbf{1}_n^\top (\mathbf{Q} \odot \mathbf{S}^2)) + g(\mathbf{Q} \odot (\mathbf{M} - \mathbf{X}\mathbf{B}))\right)\right) \\ & - \frac{1}{2} \text{Tr}\left(\text{diag}(\boldsymbol{\Omega}) \mathbf{1}_n^\top \left((\mathbf{1}_n \text{diag}(\boldsymbol{\Sigma})^\top) \odot \mathbf{P} + \mathbf{P} \odot \mathbf{Q} \odot (\mathbf{M} - \mathbf{X}\mathbf{B})^2\right)\right) \\ & - \frac{1}{2} \mathbf{1}_n^\top \mathbf{P} \log(\text{diag}(\boldsymbol{\Sigma})) + \frac{np}{2}, \end{aligned}$$

where \odot denotes the Hadamard product, diag returns a vector constituted of the diagonal of the input squared matrix, $\mathbf{1}_n$ (resp. $\mathbf{1}_{n,p}$) is a column-vector (resp. matrix) of size n (resp. $n \times p$) filled with ones, Diag takes a vector x and returns a diagonal matrix with diagonal

x , logarithm and squared functions are applied component-wise, $\mathbf{Q} = \mathbf{1}_{n,p} - \mathbf{P}$ and $g(\mathbf{D}) = \mathbf{D}^\top \mathbf{D}$ for $\mathbf{D} \in \mathbb{R}^{n \times p}$. We denoted $\bar{\mathbf{S}}^2 = \mathbf{1}_n^\top \mathbf{S}^2$ and

$$\delta_{0,\infty}(x) = \begin{cases} 0 & \text{if } x = 0 \\ -\infty & \text{else} \end{cases}$$

with the convention that $0 \times \delta_{0,\infty}(x) = 0$ for all x . Note that both ELBOs share the following terms ($\tilde{\mathbb{E}}^{(1)}$ and $\tilde{\mathbb{E}}^{(2)}$ coincides for the following terms so that we drop the index):

$$\tilde{\mathbb{E}}[\log p_\theta(\mathbf{Y}|\mathbf{Z}, \mathbf{W})] = \text{Tr}(\mathbf{Q}^\top (\mathbf{Y} \odot (\mathbf{O} + \mathbf{M}) - \mathbf{A} - \log(\mathbf{Y}!)) + \mathbf{P}^\top \delta_{0,\infty}(\mathbf{Y})),$$

$$\tilde{\mathbb{E}}[\log p_\theta(\mathbf{W})] = \text{Tr}(\mathbf{P}^\top \boldsymbol{\mu}_0 - \mathbf{1}_{n,p}^\top \log(\mathbf{1}_{n,p} + e^{\boldsymbol{\mu}_0})),$$

$$H(\mathbf{P}) = -\text{Tr}(\mathbf{P}^\top \log(\mathbf{P}) + \mathbf{Q}^\top \log(\mathbf{Q})),$$

where factorial and exponential are applied component-wise and the matrix \mathbf{A} denotes $\exp(\mathbf{O} + \mathbf{M} + \mathbf{S}^2/2)$ where \exp is applied component-wise and $\boldsymbol{\mu}_0 = \mathbf{1}_{n,p} \times \text{logit}(\pi)$ in the ND case, $\boldsymbol{\mu}_0 = \mathbf{X}^0 \mathbf{B}^0$ in the CD case and $\boldsymbol{\mu}_0 = \bar{\mathbf{B}}^0 \bar{\mathbf{X}}^0$ in the RD case. We used the convention $0 \times \log(0) = 0$ for all x .

REMARK 2. The main and only goal of $\delta_{0,\infty}$ is to ensure that $P_{ij} = 0$ whenever $Y_{ij} \neq 0$, i.e. that Y_{ij} doesn't originate from the null component when it's positive.

Model selection criterion. When the modelling choice of zero-inflation is unclear, we consider two classical criteria: BIC (Schwarz, 1978) and ICL (Biernacki, Celeux and Govaert, 2000) to choose between Models 3a, 3b and 3c. The log-likelihood is replaced by its lower bound J . We recall that ICL uses the conditional entropy of the latent variables given the observations as an additional penalty with respect to BIC. The difference between BIC and ICL measures the uncertainty of the representation of the observations in the latent space. Because the true conditional distribution $p_\theta(\mathbf{Z}, \mathbf{W}|\mathbf{Y})$ is intractable, we replace it with its variational approximation $\tilde{p}_\psi(\mathbf{W}, \mathbf{Z})$ to evaluate this entropy. Recalling that $\psi = (\mathbf{M}, \mathbf{S}, \mathbf{P})$, the entropy for variational approximations $\tilde{p}_\psi^{(1)}$ and $\tilde{p}_\psi^{(2)}$ are respectively given by

$$H^{(1)}(\psi) = \frac{1}{2} \mathbf{1}_n^\top \log(\mathbf{S}^2) \mathbf{1}_p + \frac{np}{2} \log(2\pi e) + H(\mathbf{P})$$

$$H^{(2)}(\psi) = \frac{1}{2} \mathbf{1}_n^\top \text{Tr}((\mathbf{1}_{n,p} - \mathbf{P})^\top \log(\mathbf{S}^2)) \mathbf{1}_p - \frac{1}{2} \mathbf{1}_n^\top \mathbf{P} \log(\text{diag}(\Sigma)) + H(\mathbf{P})$$

where $H(\mathbf{P})$ is defined in 3. The BIC and ICL criterion for variational approximation $\tilde{p}_\psi^{(i)}$ ($i = \{1, 2\}$) are thus given by

$$\text{BIC}^{(i)} = J^{(i)} - K \log(n)$$

$$\text{ICL}^{(i)} = \text{BIC}^{(i)} - H^{(i)}(\psi)$$

where $K = p(p+1)/2 + pd + c$ is the number of parameters and c depends on the modelling choice for the zero-inflation component: 1 for Model 3a, pd_0 for Model 3b and nd_0 for Model 3c.

REMARK 3. We note that 3c is not a parametric model, for which BIC and ICL have a theoretical grounding, but a semi-parametric one. We nevertheless use those criteria to compare the three models to each other.

The following section discusses the optimization of both ELBOs to estimate the model parameters θ .

4. Optimization. Estimating θ is equivalent to solving the optimization problem

$$(12) \quad \arg \max_{\psi, \theta} J(\psi, \theta).$$

where J can be either $J^{(1)}$ (standard approximation) or $J^{(2)}$ (enhanced approximation).

4.1. *Optimization of $J^{(1)}$.* Past experience for standard PLN models (Chiquet, Mariadassou and Robin, 2017; Chiquet, Robin and Mariadassou, 2019; Chiquet, Mariadassou and Robin, 2021) (and analytical properties of $J^{(1)}$ derived in this section) suggests solving the above problem using alternated gradient descent.

Consider the ELBO $J^{(1)}$ defined in Proposition 3. The following technical propositions will serve to update some parameters in an alternate optimization scheme.

PROPOSITION 4. [Updates of $\mathbf{B}, \mathbf{\Omega}, \mathbf{P}$ and \mathbf{B}^0] For fixed ψ , the values of $\mathbf{\Omega}, \mathbf{B}$ maximizing $J^{(1)}$ are

$$\hat{\mathbf{\Omega}} = n \left[g(\mathbf{M} - \mathbf{X}\mathbf{B}) + \bar{\mathbf{S}}^2 \right]^{-1}, \quad \hat{\mathbf{B}} = [\mathbf{X}^\top \mathbf{X}]^{-1} \mathbf{X}^\top \mathbf{M}.$$

where $g(\mathbf{D}) = \mathbf{D}^\top \mathbf{D}$ as in Proposition 3. Furthermore, if $\mathbf{X}^0 = \mathbf{1}_n$, $J^{(1)}$ is maximized at

$$\hat{\mathbf{B}}^0 = \frac{1}{n} \mathbf{1}_n^\top \mathbf{P}.$$

When θ is fixed, $J^{(1)}$ is concave with respect to \mathbf{P} and maximized at

$$\hat{\mathbf{P}} = \text{logit}^{-1} (\mathbf{A} + \mathbf{X}^0 \mathbf{B}^0) \times \delta_0(\mathbf{Y}).$$

PROOF. Proofs for $\hat{\mathbf{\Omega}}$ and $\hat{\mathbf{B}}$ and $\hat{\mathbf{B}}^0$ are straightforward using null-gradient conditions given in the appendix (Proposition 6) and left to the reader. We only prove the concavity with respect to \mathbf{P} . As $J^{(1)}$ is separable in each P_{ij} , we only need to prove concavity with respect to each P_{ij} . If $Y_{ij} > 0$, $P_{ij} \delta_{0, \infty}(Y_{ij}) = -\infty$ as soon as $P_{ij} > 0$ and $J^{(1)}$ is therefore concave in P_{ij} . If $Y_{ij} = 0$, $J^{(1)}$ depends on P_{ij} only through $P_{ij} A_{ij} + P_{ij} [\mathbf{x}_i^{0\top} \mathbf{B}_j^0 - \text{logit}(P_{ij})] - \log(1 - P_{ij})$ which is concave in P_{ij} . \square

An alternated gradient descent optimizing $J^{(1)}$ is proposed in Algorithm 1 and convergence to a stationary point is a direct consequence of the following lemma.

LEMMA 1 (Convergence properties). $J^{(1)}$ is (separately) concave in θ and ψ .

PROOF. For $\psi = (\psi_1, \mathbf{P})$ with $\psi_1 = (\mathbf{M}, \mathbf{S})$, note first that $J^{(1)}$ is separable in ψ_1 and \mathbf{P} so we can prove it independently for each parameter. For ψ_1 , it follows from the same result in the standard PLN model (lemma 1 of Chiquet, Mariadassou and Robin (2017)). For \mathbf{P} , it follows from Proposition 4. For θ , note first that $J^{(1)}$ is separable in $(\mathbf{B}, \mathbf{\Omega})$ and \mathbf{B}^0 so we can prove it independently for each parameter. For the former, it follows from the same result in the standard PLN model (for $(\mathbf{B}, \mathbf{\Omega})$). For the latter, it follows from the concavity (for all $\mathbf{a} \in \mathbb{R}^d$ fixed) of the following function $f: \mathbb{R}^d \mapsto \mathbb{R}$:

$$f(\boldsymbol{\beta}) = \mathbf{a}^\top \boldsymbol{\beta} - \mathbf{1}_n^\top \log \left(1 + e^{\mathbf{X}^0 \boldsymbol{\beta}} \right).$$

\square

Algorithm 1: VEM

Input : $\theta^{(0)}, \psi^{(0)}$ initial point, $T \geq 1$ number of iterations.
for $s = 0, \dots, T - 1$ **do**

M-step

$$\boldsymbol{\Omega}^{(s+1)} = n \left[g \left(\mathbf{M}^{(s)} - \mathbf{X} \mathbf{B}^{(s)} \right) + \bar{\mathbf{S}}^2{}^{(s)} \right]^{-1}$$

$$\mathbf{B}^{(s+1)} = [\mathbf{X}^\top \mathbf{X}]^{-1} \mathbf{X}^\top \mathbf{M}^{(s)}$$

$$\mathbf{B}^{0,(s+1)} = \arg \max_{\mathbf{B}^0} \text{Tr} \left[\left(\mathbf{P}^{(s)} \right)^\top \mathbf{X}^0 \mathbf{B}^0 \right] - \text{Tr} \left[\mathbf{1}_{n,p}^\top \log \left(1 + e^{\mathbf{X}^0 \mathbf{B}^0} \right) \right]$$

VE-step

$$\mathbf{P}^{(s+1)} = \text{logit}^{-1} \left(\mathbf{A}^{(s)} + \mathbf{X}^0 \mathbf{B}^{0,(s+1)} \right) \times \delta_0(\mathbf{Y}), \quad \mathbf{Q}^{(s+1)} = \mathbf{1}_{n,p} - \mathbf{P}^{(s+1)}$$

$$\mathbf{M}^{(s+1)} = \arg \max_{\mathbf{M}} \left(\text{Tr} \left(\mathbf{Q}^{(s+1)\top} (\mathbf{Y} \odot \mathbf{M} - \mathbf{A}) \right) - \frac{1}{2} \text{Tr} \left(\boldsymbol{\Omega}^{(s+1)} g \left(\mathbf{M} - \mathbf{X} \mathbf{B}^{(s+1)} \right) \right) \right)$$

$$\mathbf{S}^{(s+1)} = \arg \max_{\mathbf{S}} \left(-\text{Tr} \left(\mathbf{Q}^{(s+1)\top} \mathbf{A} \right) - \frac{1}{2} \text{Tr} \left(\mathbf{1}_{n,p}^\top \log \left(\mathbf{S}^2 \right) \right) - \frac{1}{2} \text{Tr} \left(\boldsymbol{\Omega}^{(s+1)} \bar{\mathbf{S}}^2 \right) \right)$$

end

Output : $\theta^{(T)}, \psi^{(T)}$

4.2. *Optimization of $J^{(2)}$.* While optimization of $J^{(1)}$ is easily manageable using closed forms and benefits from a bi-concavity property, optimization of $J^{(2)}$ is more challenging. Indeed, the concavity in $\boldsymbol{\Omega}$ is lost and no closed form can be used for any parameter update.

We do not maximize the ELBO with respect to each parameter in an alternate coordinate-wise fashion but instead compute the gradient with respect to (ψ, θ) as if it were a single parameter and perform a gradient update. Formally, given $\psi^{(0)}, \theta^{(0)}$ and a learning rate $\eta > 0$, we perform the update step

$$(13) \quad (\psi^{(s+1)}, \theta^{(s+1)}) = (\psi^{(s)}, \theta^{(s)}) + \eta \nabla_{\psi, \theta} J^{(2)}(\psi^{(s)}, \theta^{(s)})$$

until a convergence criteria or a maximum number of iterations is reached.

4.3. *Optimization using analytic law of $W_{ij}|Y_{ij}$.* The exact conditional law $W_{ij}|Y_{ij}$ can be derived and is detailed in the next proposition.

PROPOSITION 5. *Let $1 \leq j \leq p$. The conditional law of $W_{ij}|Y_{ij}$ is given by*

$$W_{ij}|Y_{ij} \sim \mathcal{B} \left(\frac{\pi_{ij}}{\varphi(o_{ij} + \mathbf{x}_i^\top \mathbf{B}_j, \Sigma_{jj}) (1 - \pi_{ij}) + \pi_{ij}} \right) \mathbf{1}_{Y_{ij}=0}$$

with $\varphi(\mu, \sigma^2) = \mathbb{E}[\exp(-X)]$, $X \sim \mathcal{LN}(\mu, \sigma^2)$.

In Section 3 we made the variational approximation $\tilde{p}(W_{ij}) \sim \mathcal{B}(P_{ij})$, considered P_{ij} as free and optimized the ELBO with respect to P_{ij} . The above proposition suggests that P_{ij} can instead be derived directly from θ and not be considered as a free variational parameter. We consider $\tilde{J}^{(1)}$ (resp. $\tilde{J}^{(2)}$) the ELBO $J^{(1)}$ (resp. $J^{(2)}$) with $P_{ij} = \Psi(\theta)_{ij}$ with

$$\Psi(\theta) \triangleq \frac{\boldsymbol{\pi}}{\varphi(\mathbf{O} + \mathbf{X}^\top \mathbf{B}, \mathbf{1}_n \text{diag}(\boldsymbol{\Sigma})^\top) \odot (1 - \boldsymbol{\pi}) + \boldsymbol{\pi}} \odot \mathbf{1}_{\mathbf{Y}=0},$$

where φ and the division are applied component-wise and $\mathbf{1}_{\mathbf{Y}=\mathbf{0}}$ is a $n \times p$ matrix such that $(\mathbf{1}_{\mathbf{Y}=\mathbf{0}})_{ij} = 0$ if and only if $Y_{ij} = 0$. Formally, we have

$$\tilde{J}^{(1)}(M, S, \Omega, B, B^0) = J^{(1)}(M, S, \Psi(\Omega, B, B^0), \Omega, B, B^0),$$

and the same formula applies to $\tilde{J}^{(2)}$. Note that both $\tilde{J}^{(1)}$ and $\tilde{J}^{(2)}$ have np fewer variational parameters compared to $J^{(1)}$ and $J^{(2)}$ ($2np$ compared to $3np$) since \mathbf{P} is now completely determined by θ . The function φ is intractable but a sharp (derivable) approximation $\tilde{\varphi}$ is available and detailed in the next section. For the Standard approximation, a major drawback of this approach compared to optimizing $J^{(1)}$ is the lack of any closed form update as stationary points of $\tilde{\varphi}$ are intractable. For the optimization, we consider the gradient scheme defined in Equation (13) where ψ is replaced with $\psi_1 = (M, S)$.

4.4. Implementation details.

Gradient with respect to Ω . When no closed form is available, optimization with respect to Ω must be adapted to ensure that Ω remains symmetric and positive definite. Instead of maximizing directly over Ω , we introduce a $p \times p$ unconstrained matrix C and use the following parametrization for Ω

$$(14) \quad \Omega = (CC^\top)^{-1}$$

and compute the gradient with respect to C .

Approximation of φ . The function φ defined in Proposition 5 is intractable but an approximation (Rojas-Nandayapa, 2008) can be computed:

$$\varphi(\mu, \sigma^2) \approx \tilde{\varphi}(\mu, \sigma^2) = \frac{\exp\left(-\frac{W^2(\sigma^2 e^\mu) + 2W(\sigma^2 e^\mu)}{2\sigma^2}\right)}{\sqrt{1 + W(\sigma^2 e^\mu)}},$$

where $W(\cdot)$ is the Lambert function (i.e. $z = x \exp(x) \Leftrightarrow x = W(z)$, $x, z \in \mathbb{R}$). An analysis of its sharpness is performed in Asmussen, Jensen and Rojas-Nandayapa (2014). Derivability of $\tilde{\varphi}$ is ensured as $W(\cdot)$ is derivable.

Stochastic Gradient Ascent. When n is large, computing the whole gradient $\nabla_\psi J$ is time-consuming. As the ELBOs defined in Proposition 3 are additive in the variational parameters ψ , Stochastic Gradient Ascent (Robbins and Monro, 1951) can be applied to scale the algorithm to large datasets.

5. Simulation Study. In this section, we evaluate the statistical and computational performances of Models 3a, 3b and 3c on simulated data.

5.1. *Experimental details.* We set $n = 500, p = 150, d = 3$ (and $d_0 = 4$ for 3b, 3c) to mimick typical sizes observed in microbiome studies. Given $\theta^* = (\Sigma^*, B^*, \pi^*)$ for Model 3a, $\theta^* = (\Sigma^*, B^*, B^{0*})$ for Model 3b and $\theta^* = (\Sigma^*, B^*, \bar{B}^{0*})$ for Model 3c, we simulate n independent observations \mathbf{Y}_i for each model, and consider the following estimation strategies:

- Standard ($J^{(1)}$),
- Enhanced ($J^{(2)}$),
- Standard Analytic ($\tilde{J}^{(1)}$),
- Enhanced Analytic ($\tilde{J}^{(2)}$),
- PLN,

- Oracle PLN,

where PLN competitor corresponds to a PLN model fitted directly on the zero-inflated data matrix \mathbf{Y} whereas Oracle PLN is fitted on the non-inflated data \mathbf{T} , where $T_{ij}|Z_{ij} \sim \mathcal{P}(\exp(O_{ij} + Z_{ij}))$. While the former evaluates model performance without considering zero-inflation during modeling, the latter serves as a reference for the Poisson component, as it is unaffected by the signal degradation caused by zero-inflation. To assess inference quality, we report the following metrics:

- Root Mean Squared Error (RMSE) between true parameters θ^* and estimates $\hat{\theta}$, as well as between true zero-inflation probabilities π^* and estimated probabilities $\hat{\pi}$,
- ELBO,
- Reconstruction error, computed as the RMSE between the original data matrix \mathbf{Y} and the reconstructed data matrix $\hat{\mathbf{Y}}$.
- Computation time.

The results pertaining to the reconstruction error and computation are deferred to Appendix B for detailed analysis.

We investigate how the models respond to fluctuations in the zero-inflation probability π^* (Section 5.2) and fluctuations in the mean \mathbf{XB}^* of the Gaussian component \mathbf{Z} (Section 5.3). In the former, we explore whether an increase in zero-inflation probability results in better estimation of the zero-inflation parameter π^* (or B^{0*} , \bar{B}^{0*}) at the expense of degraded estimation for the PLN parameters (Σ^* , \mathbf{B}^*). In the latter, we assess the model’s accuracy in challenging scenarios where \mathbf{XB}^* is small, leading to numerous zeros in both the ZI and PLN components. Additionally, in Section 5.4, we examine the performance enhancement as the sample size n increases.

In all simulations, the covariance matrix Σ^* is structured as a block-diagonal matrix comprising five blocks, each with a size of $p/5$ and filled with values of 1. To ensure invertibility, we add the identity matrix. The $n \times d$ parameters \mathbf{X} , \mathbf{X}^0 , and $\bar{\mathbf{B}}^{0*}$ are composed of independent entries taking values of ± 1 with equal probability, except for the first column (intercept), where entries are fixed to 1. Similarly, all entries of \mathbf{B}^* are independently sampled from a Gaussian distribution with a mean $\gamma \in \mathbb{R}$ and a variance of $1/\sqrt{d}$, ensuring that \mathbf{XB}^* exhibits independent Gaussian entries centered on γ with unit variance. Parameters B^{0*} and $\bar{\mathbf{X}}^0$ follow the same generation process, except that the Gaussian mean is determined by $\text{logit}(\rho)$ for some $0 < \rho < 1$. A larger ρ corresponds to increased zero-inflation, while a larger γ indicates a larger Poisson Log-Normal (PLN) component. Specific values for ρ and γ are provided in each subsection. Notably, offsets (\mathbf{O}) are not considered in these simulations, and are set to a zero matrix of dimensions $n \times p$.

5.2. Simulations when π^ fluctuate.* The parameter γ is assigned a value of $\gamma = 2$ to introduce a moderately large Poisson Log-Normal (PLN) component, characterized by a low probability (6.5%) of generating zeroes. To regulate the degree of zero-inflation, we systematically increase the probability of zero-inflation from 0.2 to 0.9, with increments of 0.1. This results in a range of values for π^* within $\Pi \triangleq 0.2, 0.3, \dots, 0.8, 0.9$. For Model 3a, this adjustment is straightforward, as π^* is directly set to values within Π . In the case of Model 3b (and Model 3c), the control of zero-inflation is achieved by simulating B^{0*} (or $\bar{\mathbf{X}}^0$), as described previously, with ρ taking values from Π . This process results in the generation of eight distinct parameter sets θ for each model, corresponding to each value in Π . Subsequently, for each θ , we simulate 30 unique datasets \mathbf{Y} according to Equation (2). The obtained results are presented in Figure 1.

It is observed that the RMSE concerning Σ^* remains consistent across different model choices. Notably, the Standard variational approximation (VA) outperforms other VA, exhibiting expectedly poorer performance compared to the Oracle PLN. The Standard VA is notably affected by high values of π^* , whereas other VA methods demonstrate relatively stable performance with respect to π^* , albeit yielding inferior results. All variants of the ZIPLN VA exhibit significantly better performance compared to PLN, as expected.

Regarding π , observations from Models 3b and 3c indicate a progressive improvement in RMSE until π^* reaches a specific threshold, after which it begins to decline. This trend reflects the delicate balance in model performance: when π^* is too low, the scarcity of observed zeroes leads to an unreasonable RMSE. Conversely, excessively high values of π^* result in an abundance of zeroes, adversely impacting model performance. An optimal balance is achieved, dependent on the choice of VA and model, where an intermediate value of π^* ensures the best performance. In the case of Model 3a, both Enhanced VA methods exhibit improved performance as π^* increases. However, the Enhanced Analytic VA reaches a plateau when $\pi^* \geq 0.5$. This plateau corresponds to a scenario characterized by a high incidence of observed zeroes and a low signal from the Poisson component, thus justifying the observed stabilization in performance.

The RMSE concerning B increases with π^* up to $\pi^* \simeq 0.5$, above which the RMSE unexpectedly starts to decrease.

5.3. Simulations when XB^ fluctuate.* In Model 3a, we fix π^* at 0.3, a value chosen based on its propensity to produce moderately challenging models and yield contrasting results in previous experiments. For Model 3b (and Model 3c), we simulate the parameters B^{0^*} (and \bar{X}^0) as previously described, setting $\rho = 0.3$. To systematically enhance the signal of the Poisson component, we increment γ from 0 to 3 in steps of 0.5, thereby covering values in $\Gamma \triangleq 0, 0.5, \dots, 2.5, 3$. Following this methodology, we generate seven distinct parameter sets θ for each model, each corresponding to a value in Γ . Subsequently, for each θ , we simulate 30 unique datasets Y according to Equation (2). The obtained results are depicted in Figure 2.

Regarding the RMSE with respect to Σ , consistency is observed across all three models. Both Enhanced VA methods demonstrate stability, independent of XB , unlike the Standard VA, where the RMSE decreases with XB similar to the behavior observed with Oracle PLN. This last behavior is anticipated due to the exponential shrinkage towards zero for low values, making it challenging to differentiate between low and very low values. Notably, the Standard Analytic VA does not benefit from high magnitudes of XB , as indicated by an increase in RMSE with higher magnitudes.

Regarding π , all VA methods (except Enhanced VA) exhibit a consistent negative correlation between the RMSE and the magnitude of XB . With larger magnitudes of XB , the Poisson component generates fewer zeros, leading to clearer identification of the origin of zeros. Only the Enhanced VA demonstrates challenges with large magnitudes of XB for Models 3a and 3b, yet performs better when magnitudes are low or with Model 3c.

Regarding B , both the Standard and Enhanced Analytic VA methods exhibit consistency across all three models, performing similarly and benefiting from high magnitudes of XB , akin to Oracle PLN. Conversely, the Standard Analytic and Enhanced VA methods encounter difficulties when magnitudes are high.

5.4. Simulations when n fluctuate. The parameter γ is maintained at $\gamma = 2$, as discussed in Section 5.2, while the zero-inflation parameter π^* remains set at 0.3, consistent with the conditions outlined in Section 5.3. Furthermore, the number of variables is held constant at $p = 150$. To investigate the impact of the sample size, we incrementally vary n from 100 to




FIG 1. Simulations when the probability of zero-inflation π varies. To improve visibility, the scale is **not** shared on the second row.

600, with increments of 100, forming the set $N = 100, 200, \dots, 600$. For each value in N , we generate 30 distinct parameter sets θ , resulting in a total of 30×6 different parameter combinations. Subsequently, we simulate datasets \mathbf{Y} for each parameter θ . The obtained results are presented in detail in Figure 3.

For the Enhanced VA, the addition of more samples does not yield a substantial improvement in RMSE. Hence, further exploration into this aspect is deemed unnecessary. Notably, it performs comparably to other VA methods when n is low.

Regarding the RMSE with respect to Σ^* , a decrease is observed with increasing sample size n , albeit at a minimal rate for the Enhanced Analytic VA. As for π , the RMSE consistently decreases with n across all VA methods, particularly noticeable in Model 3b. In Model 3a, the behavior of the Standard Analytic VA exhibits a plateau phenomenon, observed when n exceeds 250, while the Standard VA shows a gradual deterioration once n surpasses 300, eventually stabilizing when $n \geq 500$. Conversely, the Enhanced Analytic VA demonstrates a consistent decrease in RMSE, aligning with expectations. Analyzing Model 3c with respect

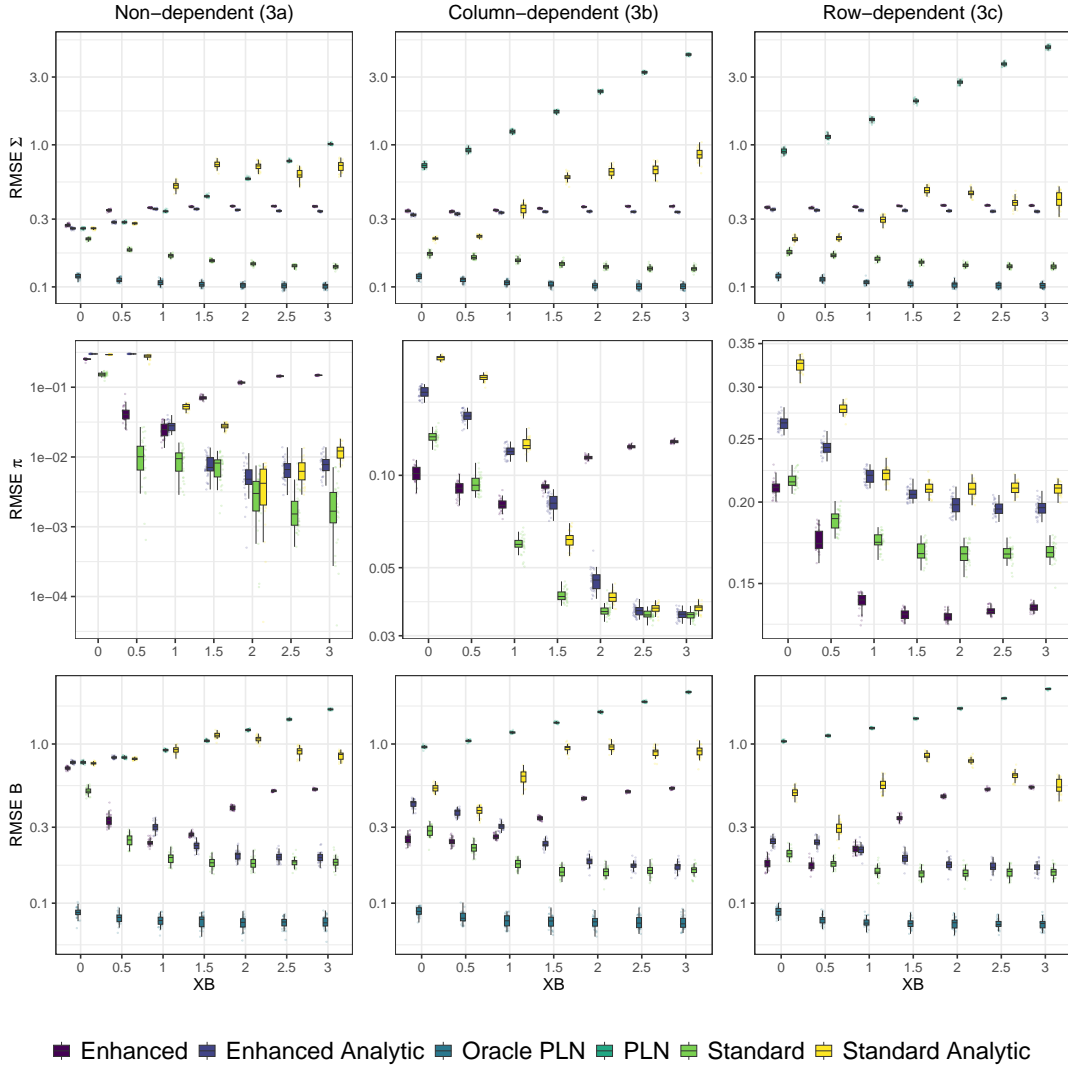


FIG 2. Simulations when the magnitude of the gaussian component \mathbf{XB} varies. To improve visibility, the scale is *not* shared on the second row.

to π presents challenges due to the increasing number of parameters of \mathbf{B}^0 with n , influencing π . The RMSE with respect to \mathbf{B} displays a consistent decrease for all VA methods, except the Enhanced one. Both the Standard and Enhanced Analytic VA methods exhibit similar performances, outperforming the Standard Analytic VA by a considerable margin.

Figure 4 presents the ELBOs for a set of sample sizes $n = 100, 300, 500$, showcasing only the highest-performing VA methods identified from Figures 1 to 3: Enhanced Analytic and Standard VA. The parameter γ is fixed at $\gamma = 2$, following the conditions outlined in Section 5.2, while the zero-inflation parameter π^* remains set at 0.3, consistent with the specifications in Section 5.3. Additionally, the dimensionality is maintained at $p = 150$. For each sample size n , 30 distinct parameter sets θ are generated, resulting in a total of 30×3 different parameter combinations. Subsequently, datasets \mathbf{Y} are simulated for each parameter θ . The ELBO is computed as the mean across the 30 different runs, accompanied by a 95% confidence interval.

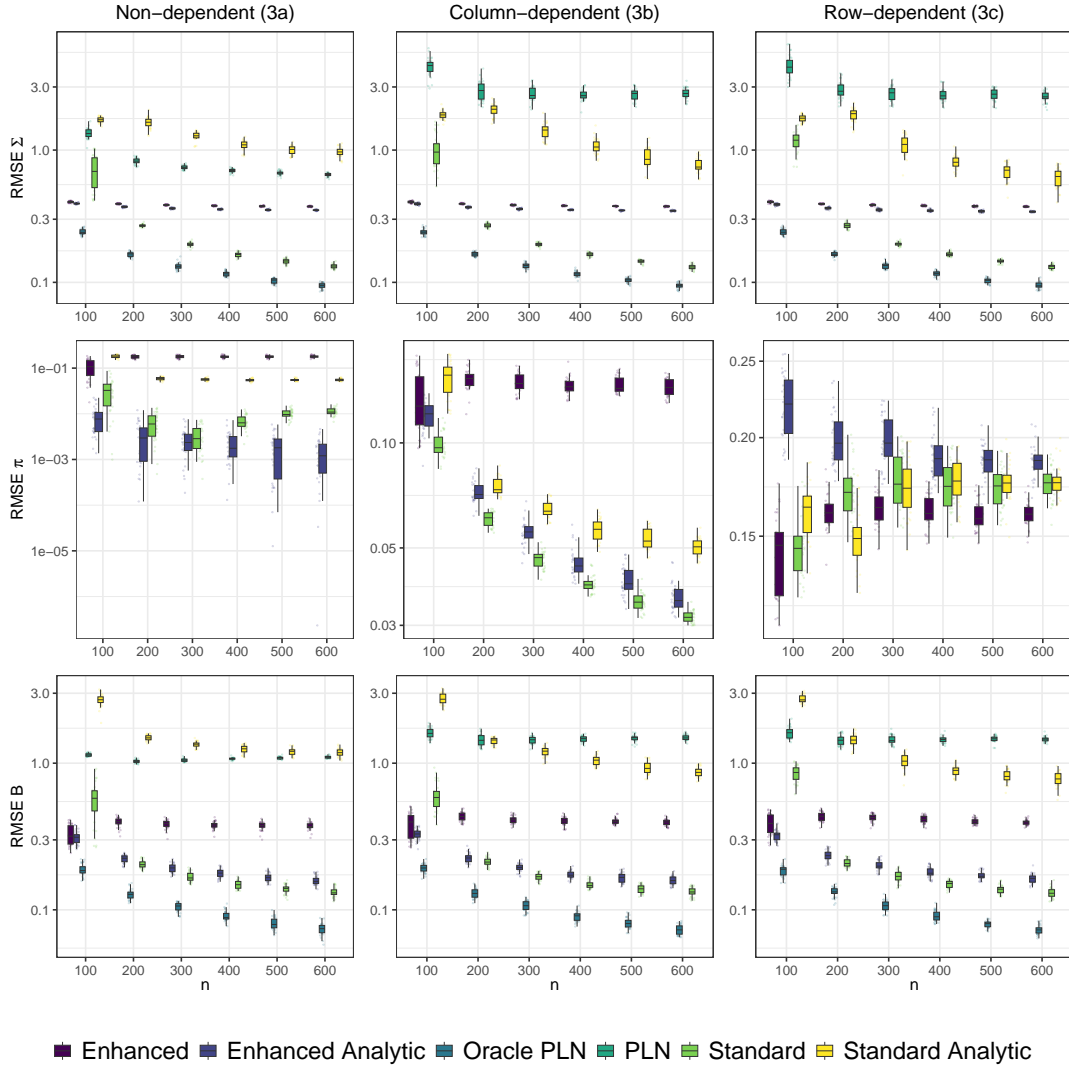


FIG 3. Simulations when the number of samples n grows. To improve visibility, the scale is *not* shared on the second row.

	n = 100			n = 300			n = 500		
	ND	CD	RD	ND	CD	RD	ND	CD	RD
Enhanced Analytic	-40086±1240	-44143±764	-40211±1045	-124765±2049	-136998±1413	-123418±2229	-208690±1929	-230619±2122	-209290±2728
Standard	-37737±1134	-41507±781	-37931±939	-121803±1979	-133404±1353	-120589±2156	-205011±1899	-226171±2055	-205689±2672

FIG 4. Comparison of ELBOs (higher the better) when the number of samples grows (ND:non-dependent 3a, CD:column-dependent 3b, RD:row-dependent 3c).

Conclusion of the study. The Standard VA method consistently exhibited superior performance across diverse scenarios, showing better RMSEs and demonstrating higher ELBO values. While Enhanced Analytic VA method showed promise in certain scenarios, such as fluctuations in π^* and \mathbf{XB}^* , it did not outperform the Standard VA consistently. Therefore, based on the results of this simulation study, the Standard VA method appears to be the preferred choice.

		Model		Criteria	
		ZI component	PLN component	BIC	ICL
PLN	No zero inflation		~1	-335 526.4	-856 215.2
			~site	-335 814.7	-856 263.3
			~time	-336 467.6	-854 845.5
			~time*site	-341 653.1	-859 927.5
ZI		~1		-331 848.1	-594 390.5
		~site	~1	-319 518.0	-531 080.2
		~time	~1	-321 122.8	-578 202.4
		~time*site	~1	-310 751.7	-546 597.5
ZIPLN		~site	~site	-323 527.1	-563 262.0
		~time	~time	-328 061.9	-534 570.5
		~time*site	~time*site	-332 835.8	-500 573.7

TABLE 1

Model selection criteria (BIC and ICL) and model details for the 11 ZIPLN models with (ZI and ZIPLN) or without (PLN) zero inflation fitted to the cow microbiota dataset. The best value for each criteria is in bold.

6. Application to cow microbiome data. We now consider a study on the structure and evolution of the microbiota of 45 lactating cows before and after calving (Mariadassou et al., 2023). In this experiment, three body sites in addition to the milk of the 4 teats were sampled (vagina, mouth, rumen) at 4 times points: 1 week before calving (except for the milk), 1 month, 3 months and 7 months after calving. The data include $n = 921$ samples with sequencing depths ranging from 1,003 to 81,591 reads. After preprocessing, as described in Mariadassou et al. (2023), a total of 1209 Amplicon Sequence Variants (ASV) were identified using the FROGS pipeline (Escudié et al., 2017) based on DADA2 (Callahan et al., 2016) and taxonomy was assigned using reference databases. We filtered out ASV with prevalence lower than 5% and removed samples for which the total count was zero, resulting in a count table of $n = 899$ samples (see) with $p = 259$ ASV and a mean proportion of zeroes of 90.3%.

Previous analyses have already shown that body site, sampling time and their interaction are strong structuring factors of the microbiota. We are interested in comparing the results of PLN and ZIPLN on this dataset, in particular how the explicit modeling of the Zero-Inflation changes (i) the fit of low value counts and (ii) the position and clustering of samples in the latent space.

6.1. Model selection. We considered a total of 11 models, without (PLN, 4 models) or with (ZIPLN, 7 models) zero inflation. For PLN, the four models correspond to different sets of covariates: no covariate, site, time and a full model with both covariates and their interaction. For ZIPLN, we considered seven models: one with no covariate in either the PLN or ZI components, three with covariates only in the ZI component (and no covariate for the PLN component) and the last with the same covariates in both the ZI and PLN components. Based on the results of section 5, each ZIPLN model was fitted with the Standard VA ($J^{(1)}$). We compared the 11 models in terms of BIC and ICL in Table 1. Both criteria show that accounting for zero inflation and more specifically the effect of site, time and their interaction on it has a bigger effect than modeling how those covariates affect the counts and leads to massive improvement in the penalized likelihood values. For the rest of the analyses, we focus on the full models: PLN with the site \times time interaction and ZIPLN with the site \times time interaction in both components.

6.2. Modeling of counts. Comparing the model fits of PLN and ZIPLN (Figure 5) show that that both models are good at fitting the counts along the observed range (left panel) with a small upward bias for ZIPLN for low counts. However, when focusing on observed zeroes,

we find out that ZIPLN predicts lower values than PLN (middle panel). The bimodality observed for ZIPLN fitted values corresponds to a mixture of two populations (right panel): on the one hand, counts for which both the variational probability of zero-inflation P_{ij} is close to 1 and the latent mean M_{ij} is very negative and on the other hand counts for which M_{ij} is much higher but compensated by a high value of P_{ij} . The first population corresponds to zeroes that could be well-captured by a PLN model without zero inflation whereas the second one corresponds to zeroes which are only fitted well thanks to the zero inflation component.




FIG 5. Model fits of PLN and ZIPLN in terms of fitted versus observed counts (left panel), fitted values for null counts (middle panel) and comparison of P_{ij} and M_{ij} estimated for null counts by ZIPLN (right panel).

Looking at the variational parameters estimated by ZIPLN (Figure 6), we find that the different microbiota are highly structured in blocks of species that are only prevalent and abundant in specific sites at specific times. In particular, few ASVs are detected in the oral microbiota (negative M and P close to 1, corresponding to mostly null counts). The blue bands in M indicate ASVs likely to be structurally absent from this microbiota. In contrast, the bright yellow blocks in P , and the corresponding ones in M , highlight ASVs that are systematically present and abundant in the oral microbiota 1 week after calving (P close to 0 and high values of M across all samples in that category). Finally, many ASVs have large positive M values in the milk microbiota but a heterogeneous pattern of zero inflation across samples, corresponding to ASVs that are not systematically found in the milk but abundant when present. This is in line with the documented (Mariadassou et al., 2023) high diversity and large biological diversity observed in the milk compared to microbiota from other body sites.

6.3. *Latent means.* Finally, we focus on the similarity of samples in the latent space. We do so with a PCA of the latent means M (Figure 7) inferred by PLN (left panel) and ZIPLN (right panel). The results of both models are quite similar with a strong stratification of microbiota according to body site along the diagonal and according to time along the antidiagonal. Likewise, the site \times time groups are in the same positions in both panels. The striking difference between both panels lies in the scale of the within-group dispersion of the samples: large for PLN and much smaller for ZIPLN, leading to tighter and better separated groups. This is coherent with the observations from Figure 5 (right panel): observed zeroes can be captured by values of $P_{ij} \simeq 1$, without affecting M_{ij} , unlike PLN where observed zeroes systematically lead to highly negative values M_{ij} (mean value of M_{ij} for observed zeroes: -1.12 with ZIPLN and -13.6 with PLN) therefore translating the samples in the latent space and leading to higher variability. This is also reflected in the estimate of Σ which has a much smaller volume for ZIPLN ($\log |\Sigma| = -206.24$, $\text{Tr}(\Sigma) = 216.8$) than for PLN ($\log |\Sigma| = 689.11$, $\text{Tr}(\Sigma) = 15407$) corresponding to narrower distributions in the latent space.




FIG 6. Variational latent means M (left) and zero inflation probability P (right) estimated by ZIPLN. The ASVs are ordered using a hierarchical clustering (Ward linkage) on the centered-log-ratio-transformed observed counts. The first grouping factor is sampling time: 1 week before (-1W), 1 month (1M), 3 months (3M) and 7 months (7M) after calving and the second is the type of microbiota: oral (O), nasal (N), vaginal (V) or from the milk (M).




FIG 7. PCA of the variational latent means M inferred by PLN (left) and ZIPLN (right). The samples are colored according to the site \times time categories, using the same abbreviations for site and time as in Figure 6.

7. Conclusion and discussion. In the context of analysis of high-dimensional count data, we introduce the ZIPLN model, driven by a latent Gaussian variable managing the structure between counts and a zero-inflated component explaining the zeroes that fails to be explained by a standard PLN model. The zero-inflation is flexible as it can be fixed, site-

specific, feature-specific or depends on covariates. We use two variational approximations, one breaking all dependencies between features and one relying on conditional law of counts given the observed zeroes. We compare and assess the quality of both variational approximations on synthetic data and show the efficiency of ZIPLN even when 90% of the counts are corrupted by zero-inflation. Our results show that the standard VA is faster than the enhanced VA, as expected, but also performs better or almost equivalently in terms of log-likelihood and parameter estimation, a surprise for us. The model is motivated by an application on the structure of the microbiota of 45 lactating cows where 90.3% of counts are null. We show the zero-inflation leads to massive improvements in terms of log likelihood and point out that ZIPLN is indeed capturing zeroes that PLN can not, leading to better and more tightly separated groups inside the microbiota dataset.

A natural follow-up of this work would consider a low-rank approximation of the covariance matrix (Chiquet, Mariadassou and Robin, 2017) in order to scale to even higher dimensional data, as the complexity is for now quadratic in the number of variables p . Another approach would consider a ℓ_1 penalty on the inverse covariance matrix to lower the number of variables correlated with each other.

Funding. Bastien Bartardière, Julien Chiquet and François Gindraud are supported by the French ANR grant ANR-18-CE45-0023 Statistics and Machine Learning for Single Cell Genomics (SingleStatOmics).

Acknowledgments. We would like to thank Jean-Benoist Léger for the implementation of the Lambert function and the precious advices to build the `pyPLNmodels` package.

Implementation. All the algorithms are available in the python package `pyPLNmodels`¹ and the R package `PLNmodels`².

REFERENCES

- AITCHISON, J. and HO, C. H. (1989). The multivariate Poisson-log normal distribution. *Biometrika* **76** 643-653. <https://doi.org/10.1093/biomet/76.4.643>
- ASMUSSEN, S., JENSEN, J. L. and ROJAS-NANDAYAPA, L. (2014). On the Laplace Transform of the Lognormal Distribution. *Methodology and Computing in Applied Probability* **18** 441 - 458.
- BIERNACKI, C., CELEUX, G. and GOVAERT, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE transactions on pattern analysis and machine intelligence* **22** 719–725.
- BLEI, D. M., KUCUKELBIR, A. and MCAULIFFE, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American statistical Association* **112** 859–877.
- CALLAHAN, B. J., MCMURDIE, P. J., ROSEN, M. J., HAN, A. W., JOHNSON, A. J. A. and HOLMES, S. P. (2016). DADA2: High-resolution sample inference from Illumina amplicon data. *Nature Methods* **13** 581–583. <https://doi.org/10.1038/nmeth.3869>
- CAPPÉ, O., DOUC, R., MOULINES, E. and ROBERT, C. (2002). On the Convergence of the Monte Carlo Maximum Likelihood Method for Latent Variable Models. *Scandinavian Journal of Statistics* **29** 615–635.
- CHIQUET, J., MARIADASSOU, M. and ROBIN, S. (2017). Variational inference for probabilistic Poisson PCA. *Ann. Appl. Stat.*
- 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.
- CHIQUET, J., ROBIN, S. and MARIADASSOU, M. (2019). Variational inference for sparse network reconstruction from count data. In *International Conference on Machine Learning* 1162–1171. PMLR.
- CHO, H., LIU, C., PREISSER, J. S. and WU, D. (2023). A bivariate zero-inflated negative binomial model and its applications to biomedical settings. *Statistical Methods in Medical Research* **32** 1300-1317. PMID: 37167422. <https://doi.org/10.1177/09622802231172028>

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

²<https://github.com/PLN-team/PLNmodels>

- CHOI, Y., LI, R. and QUON, G. (2022). siVAE: interpretable deep generative models for single-cell transcriptomes. *Genome Biology*. <https://doi.org/10.1186/s13059-023-02850-y>
- CHOUDHARY, S. and SATIJA, R. (2022). Comparison and evaluation of statistical error models for scRNA-seq. *Genome biology* **23** 27. <https://doi.org/10.1186/s13059-021-02584-9>
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum Likelihood from Incomplete Data via the EM algorithm. **39** 1–38.
- DONG, C., CLARKE, D. B., YAN, X., KHATTAK, A. and HUANG, B. (2014). Multivariate random-parameters zero-inflated negative binomial regression model: An application to estimate crash frequencies at intersections. *Accident Analysis & Prevention* **70** 320-329. <https://doi.org/10.1016/j.aap.2014.04.018>
- ESCUDIÉ, F., AUER, L., BERNARD, M., MARIADASSOU, M., CAUQUIL, L., VIDAL, K., MAMAN, S., HERNANDEZ-RAQUET, G., COMBES, S. and PASCAL, G. (2017). FROGS: Find, Rapidly, OTUs with Galaxy Solution. *Bioinformatics* **34** 1287-1294. <https://doi.org/10.1093/bioinformatics/btx791>
- HUI, F. K., 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.
- JAAKKOLA, T. S. and JORDAN, M. I. (2000). Bayesian parameter estimation via variational methods. *Statistics and Computing* **10** 25–37.
- JACQUIER, E., JOHANNES, M. and POLSON, N. (2007). MCMC maximum likelihood for latent state models. *Journal of Econometrics* **137** 615-640. <https://doi.org/10.1016/j.jeconom.2005.11.017>
- JIN, Y., LIU, M., LI, Y., XU, R., DU, L., GAO, L. and XIANG, Y. (2020). Variational auto-encoder based Bayesian Poisson tensor factorization for sparse and imbalanced count data. *Data Mining and Knowledge Discovery* **35** 505–532. <https://doi.org/10.1007/s10618-020-00723-7>
- KINGMA, D. P. and WELING, M. (2022). Auto-Encoding Variational Bayes.
- LAMBERT, D. (1992). Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics* **34** 1–14.
- LI, C.-S. (2012). Identifiability of zero-inflated Poisson models. *Brazilian Journal of Probability and Statistics* **26** 306 – 312. <https://doi.org/10.1214/10-BJPS137>
- LI, C.-S., LU, J.-C., PARK, J., KIM, K., BRINKLEY, P. A. and PETERSON, J. P. (1999). Multivariate Zero-Inflated Poisson Models and Their Applications. *Technometrics* **41** 29–38.
- LOPEZ, R., REGIER, J., COLE, M., JORDAN, M. I. and YOSEF, N. (2018). Deep Generative Modeling for Single-cell Transcriptomics. *Nature methods* **15** 1053 - 1058.
- LOVE, M. I., HUBER, W. and ANDERS, S. (2014). Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome biology* **15** 1–21.
- MARIADASSOU, M., NOUVEL, L. X., CONSTANT, F., MORGAVI, D. P., RAULT, L., BARBEY, S., HELLOIN, E., RUÉ, O., SCHBATH, S., LAUNAY, F., SANDRA, O., LEFEBVRE, R., LE LOIR, Y., GERMON, P., CITTI, C. and EVEN, S. (2023). Microbiota members from body sites of dairy cows are largely shared within individual hosts throughout lactation but sharing is limited in the herd. *Animal Microbiome* **5** 1–17. <https://doi.org/10.1186/s42523-023-00252-w>
- NIKU, J., HUI, F. K., 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.
- O’HARA, R. and KOTZE, D. J. (2010). Do not log-transform count data. *Methods in Ecology and Evolution* **1** 118–122.
- RISSO, D., PERRAUDEAU, F., GRIBKOVA, S., DUDOIT, S. and VERT, J.-P. (2018). A general and flexible method for signal extraction from single-cell RNA-seq data. *Nature communications* **9** 1–17.
- ROBBINS, H. and MONRO, S. (1951). A stochastic approximation method. *The annals of mathematical statistics* 400–407.
- ROJAS-NANDAYAPA, L. (2008). Risk probabilities: asymptotics and simulation.
- SCHWARZ, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics* **6** 461 – 464. <https://doi.org/10.1214/aos/1176344136>
- SEABOLD, S. and PERKTOLD, J. (2010). statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*.
- STOEHR, J. and ROBIN, S. S. (2024). Composite likelihood inference for the Poisson log-normal model.
- WAINWRIGHT, M. J. and JORDAN, M. I. (2008). Graphical Models, Exponential Families, and Variational Inference. *Found. Trends Mach. Learn.* **1** 1–305.
- WANG, D. and GU, J. (2018). VASC: Dimension Reduction and Visualization of Single-cell RNA-seq Data by Deep Variational Autoencoder. *Genomics, Proteomics & Bioinformatics* **16** 320-331. Bioinformatics Commons (II). <https://doi.org/10.1016/j.gpb.2018.08.003>
- 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. PMID: 32713999. <https://doi.org/10.1080/10618600.2019.1609977>

- XU, J., XU, J., MENG, Y., LU, C., CAI, L., ZENG, X., NUSSINOV, R. and CHENG, F. (2023). Graph embedding and Gaussian mixture variational autoencoder network for end-to-end analysis of single-cell RNA sequencing data. *Cell Reports Methods* **3** 100382. <https://doi.org/10.1016/j.crmeth.2022.100382>
- ZHAO, H., RAI, P., DU, L., BUNTINE, W., PHUNG, D. and ZHOU, M. (2020). Variational Autoencoders for Sparse and Overdispersed Discrete Data. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics* (S. CHIAPPA and R. CALANDRA, eds.). *Proceedings of Machine Learning Research* **108** 1684–1694. PMLR.

APPENDIX A: TECHNICAL RESULTS

A.1. ELBO derivation.

Evidence Lower bound (ELBO). We recall that $\mathbf{M} = [\mathbf{M}_1, \dots, \mathbf{M}_n]^\top$, $\mathbf{P} = [\mathbf{P}_1, \dots, \mathbf{P}_n]^\top$, $\mathbf{S} = [\mathbf{S}_1, \dots, \mathbf{S}_n]^\top$ and $\tilde{\mathbf{S}}^2 = \mathbf{1}_n^\top \mathbf{S}^2$ and $\tilde{\mathbb{E}}^{(1)}$ (resp. $\tilde{\mathbb{E}}^{(2)}$) the expectation under $\tilde{p}^{(1)}$ (resp. $\tilde{p}^{(2)}$).

PROPOSITION 3. *The ELBO defined in Equation (6) with variational approximation $\tilde{p}_\psi^{(1)}$ can be written in matrix form as*

$$(10) \quad \begin{aligned} J^{(1)}(\psi, \theta) = & \tilde{\mathbb{E}}^{(1)} [\log p_\theta(\mathbf{Y} | \mathbf{Z}, \mathbf{W})] + \tilde{\mathbb{E}}^{(1)} [\log p_\theta(\mathbf{W})] + H(\mathbf{P}) + \frac{1}{2} \text{Tr}(\mathbf{1}_{n,p}^\top \log(\mathbf{S}^2)) \\ & + \frac{n}{2} \log |\boldsymbol{\Omega}| - \frac{1}{2} \text{Tr} \left(\boldsymbol{\Omega} \left(\text{Diag}(\tilde{\mathbf{S}}^2) + g(\mathbf{M} - \mathbf{X}\mathbf{B}) \right) \right) + \frac{np}{2} \end{aligned}$$

and with variational approximation $\tilde{p}_\psi^{(2)}$ we get

$$(11) \quad \begin{aligned} J^{(2)}(\psi, \theta) = & \tilde{\mathbb{E}}^{(2)} [\log p_\theta(\mathbf{Y} | \mathbf{Z}, \mathbf{W})] + \tilde{\mathbb{E}}^{(2)} [\log p_\theta(\mathbf{W})] + H(\mathbf{P}) + \frac{1}{2} \text{Tr}(\mathbf{Q}^\top \log(\mathbf{S}^2)) \\ & + \frac{n}{2} \log |\boldsymbol{\Omega}| - \frac{1}{2} \text{Tr} \left(\boldsymbol{\Omega} \left(\text{Diag}(\mathbf{1}_n^\top (\mathbf{Q} \odot \mathbf{S}^2)) + g(\mathbf{Q} \odot (\mathbf{M} - \mathbf{X}\mathbf{B})) \right) \right) \\ & - \frac{1}{2} \text{Tr} \left(\text{diag}(\boldsymbol{\Omega}) \mathbf{1}_n^\top \left((\mathbf{1}_n \text{diag}(\boldsymbol{\Sigma})^\top) \odot \mathbf{P} + \mathbf{P} \odot \mathbf{Q} \odot (\mathbf{M} - \mathbf{X}\mathbf{B})^2 \right) \right) \\ & - \frac{1}{2} \mathbf{1}_n^\top \mathbf{P} \log(\text{diag}(\boldsymbol{\Sigma})) + \frac{np}{2}, \end{aligned}$$

where \odot denotes the Hadamard product, diag returns a vector constituted of the diagonal of the input squared matrix, $\mathbf{1}_n$ (resp. $\mathbf{1}_{n,p}$) is a column-vector (resp. matrix) of size n (resp. $n \times p$) filled with ones, Diag takes a vector x and returns a diagonal matrix with diagonal x , logarithm and squared functions are applied component-wise, $\mathbf{Q} = \mathbf{1}_{n,p} - \mathbf{P}$ and $g(\mathbf{D}) = \mathbf{D}^\top \mathbf{D}$ for $\mathbf{D} \in \mathbb{R}^{n \times p}$. We denoted $\tilde{\mathbf{S}}^2 = \mathbf{1}_n^\top \mathbf{S}^2$ and

$$\delta_{0,\infty}(x) = \begin{cases} 0 & \text{if } x = 0 \\ -\infty & \text{else} \end{cases}$$

with the convention that $0 \times \delta_{0,\infty}(x) = 0$ for all x . Note that both ELBOs share the following terms ($\tilde{\mathbb{E}}^{(1)}$ and $\tilde{\mathbb{E}}^{(2)}$ coincides for the following terms so that we drop the index):

$$\begin{aligned} \tilde{\mathbb{E}} [\log p_\theta(\mathbf{Y} | \mathbf{Z}, \mathbf{W})] &= \text{Tr}(\mathbf{Q}^\top (\mathbf{Y} \odot (\mathbf{O} + \mathbf{M}) - \mathbf{A} - \log(\mathbf{Y}!)) + \mathbf{P}^\top \delta_{0,\infty}(\mathbf{Y})), \\ \tilde{\mathbb{E}} [\log p_\theta(\mathbf{W})] &= \text{Tr}(\mathbf{P}^\top \boldsymbol{\mu}_0 - \mathbf{1}_{n,p}^\top \log(\mathbf{1}_{n,p} + e^{\boldsymbol{\mu}_0})), \\ H(\mathbf{P}) &= -\text{Tr}(\mathbf{P}^\top \log(\mathbf{P}) + \mathbf{Q}^\top \log(\mathbf{Q})), \end{aligned}$$

where factorial and exponential are applied component-wise and the matrix \mathbf{A} denotes $\exp(\mathbf{O} + \mathbf{M} + \mathbf{S}^2/2)$ where \exp is applied component-wise and $\boldsymbol{\mu}_0 = \mathbf{1}_{n,p} \times \logit(\pi)$ in

the ND case, $\boldsymbol{\mu}_0 = \mathbf{X}^0 \mathbf{B}^0$ in the CD case and $\boldsymbol{\mu}_0 = \bar{\mathbf{B}}^0 \bar{\mathbf{X}}^0$ in the RD case. We used the convention $0 \times \log(0) = 0$ for all x .

PROOF. We recall that

$$J^{(k)}(\psi, \theta) = \tilde{\mathbb{E}}^{(k)}[\log p_\theta(\mathbf{Z}, \mathbf{W}, \mathbf{Y})] - \tilde{\mathbb{E}}^{(k)}[\log \tilde{p}_\psi^{(k)}(\mathbf{Z}, \mathbf{W})],$$

where $\tilde{\mathbb{E}}^{(k)}$ stands for the expectation with variational approximation $\tilde{p}_\psi^{(k)}$ ($k = \{1, 2\}$). We first compute the entropy term $\tilde{\mathbb{E}}^{(k)}[\log \tilde{p}_\psi^{(k)}(\mathbf{Z}, \mathbf{W})]$ for each variational distribution, starting with the standard one:

$$\begin{aligned} -\tilde{\mathbb{E}}^{(1)}[\log \tilde{p}_\psi(\mathbf{Z}, \mathbf{W})] &= \sum_i -\tilde{\mathbb{E}}^{(1)}[\log \tilde{p}_{\psi_1}(\mathbf{Z}_i)] - \tilde{\mathbb{E}}^{(1)}[\log \tilde{p}_{\psi_2}(\mathbf{W}_i)] \\ (15) \qquad \qquad \qquad &= \frac{1}{2} \sum_i (p + \mathbf{1}_p^\top \log \mathbf{S}_i^2 + p \log(2\pi)) - \sum_{i,j} \left[P_{ij} \log(P_{ij}) + Q_{ij} \log(Q_{ij}) \right], \end{aligned}$$

where $Q_{ij} = 1 - P_{ij}$. For the Enhanced approximation, denoting $\sum_{i_0, j_0} \triangleq \sum_{1 \leq i \leq n, 1 \leq j \leq p, Y_{ij}=0}$, $\sum_{i_1, j_1} \triangleq \sum_{1 \leq i \leq n, 1 \leq j \leq p, Y_{ij}>0}$ and $\mathcal{N}(x; \mu, \sigma^2)$ the density of a Gaussian with mean μ and variance σ^2 evaluated at $x \in \mathbb{R}$, we get

$$\begin{aligned} \tilde{\mathbb{E}}^{(2)}[\log \tilde{p}_\psi(\mathbf{Z}|\mathbf{W})] &= \\ &\sum_{i_0, j_0} \tilde{\mathbb{E}}^{(2)}[W_{ij} \log(\mathcal{N}(Z_{ij}; \mathbf{x}_i^\top \mathbf{B}_j, \Sigma_{jj})) + (1 - W_{ij}) \log(\mathcal{N}(Z_{ij}; M_{ij}, S_{ij}^2))] \\ &+ \sum_{i_1, j_1} \tilde{\mathbb{E}}^{(2)}[\log(\mathcal{N}(Z_{ij}; M_{ij}, S_{ij})) | W_{ij} = 0] \\ &= \sum_{i_0, j_0} P_{ij} \tilde{\mathbb{E}}^{(2)}[\log(\mathcal{N}(Z_{ij}; \mathbf{x}_i^\top \mathbf{B}_j, \Sigma_{jj})) | W_{ij} = 1] \\ &+ \sum_{i_0, j_0} (1 - P_{ij}) \tilde{\mathbb{E}}^{(2)}[\log(\mathcal{N}(Z_{ij}; M_{ij}, S_{ij})) | W_{ij} = 0] - \sum_{i_1, j_1} \log(|S_{ij}| \sqrt{2\pi e}) \\ &= \sum_{i_0, j_0} \left(-P_{ij} \log(\sqrt{\Sigma_{jj} 2\pi e}) + P_{ij} \log(|S_{ij}| \sqrt{2\pi e}) \right) - \sum_{i, j} \log(|S_{ij}| \sqrt{2\pi e}) \\ &= - \sum_{i_0, j_0} P_{ij} \left(\frac{\log(\Sigma_{jj})}{2} - \log(|S_{ij}|) \right) - \sum_{i, j} (\log(|S_{ij}|) +) - \frac{np}{2} \log(2\pi e) \\ &= - \sum_{i, j} \delta_{0, \infty}(Y_{ij}) P_{ij} \left(\frac{\log(\Sigma_{jj})}{2} - \log(|S_{ij}|) \right) - \sum_{i, j} (\log(|S_{ij}|) +) - \frac{np}{2} \log(2\pi e). \end{aligned}$$

For the full entropy, we have

$$\begin{aligned} \tilde{\mathbb{E}}^{(2)}[\log \tilde{p}_\psi(\mathbf{Z}, \mathbf{W})] &= \tilde{\mathbb{E}}^{(2)}[\log \tilde{p}_\psi(\mathbf{W})] + \tilde{\mathbb{E}}^{(2)}[\log \tilde{p}_\psi(\mathbf{Z}|\mathbf{W})] \\ (16) \qquad \qquad \qquad &= \sum_{i, j} \left[P_{ij} \log(P_{ij}) + Q_{ij} \log(Q_{ij}) \right] + \tilde{\mathbb{E}}^{(2)}[\log \tilde{p}_\psi(\mathbf{Z}|\mathbf{W})]. \end{aligned}$$

Now, the complete data log-likelihood of the ZIPLN regression model is given by

$$\begin{aligned}
\log p_\theta(\mathbf{Z}, \mathbf{W}, \mathbf{Y}) &= \log p_\theta(\mathbf{Y}|\mathbf{Z}, \mathbf{W}) + \log p_\theta(\mathbf{Z}) + \log p_\theta(\mathbf{W}) \\
(17) \quad &= \sum_{i=1}^n \sum_{j=1}^p W_{ij} \delta_{0,\infty}(Y_{ij}) + (1 - W_{ij}) (Y_{ij}(o_{ij} + Z_{ij}) - e^{o_{ij}+Z_{ij}} - \log(Y_{ij}!)) \\
&\quad - \frac{1}{2} \sum_{i=1}^n (\|\mathbf{Z}_i - \mathbf{x}_i^\top \mathbf{B}\|_\Omega^2 - \log |\Omega| + p \log(2\pi)) \\
&\quad + \sum_{i=1}^n \sum_{j=1}^p W_{ij} \mathbf{x}_i^{0\top} \mathbf{B}_j^0 - \log \left(1 + e^{\mathbf{x}_i^{0\top} \mathbf{B}_j^0} \right).
\end{aligned}$$

We start with computing the expectation $\tilde{\mathbb{E}}^{(k)}[\log \tilde{p}(\mathbf{Z})]$, set $\mathbf{H}_i = \mathbf{Z}_i - \mathbf{x}_i^\top \mathbf{B}$ and denotes H_{ij} its j^{th} coordinate. Under the Enhanced approximation we have

$$H_{ij}|W_{ij} \sim \mathcal{N}(h_{ij}, s_{ij}^2)^{1-W_{ij}} \mathcal{N}(0, \Sigma_{jj})^{W_{ij}}$$

with $h_{ij} = M_{ij} - \mathbf{x}_i^\top \mathbf{B}_j$. Now we have

$$\tilde{\mathbb{E}}^{(2)}[\|\mathbf{Z}_i - \mathbf{x}_i^\top \mathbf{B}\|_\Omega^2] = \tilde{\mathbb{E}}^{(2)}[\|\mathbf{H}_i\|_\Omega^2] = \sum_{k=1}^p \sum_{l=1}^p \tilde{\mathbb{E}}^{(2)}[H_{ik} \Omega_{kl} H_{il}].$$

Since the $(H_{ik})_{1 \leq k \leq p}$ are independent, for all $1 \leq k, l \leq p$ we have

$$\tilde{\mathbb{E}}^{(2)}[H_{ik} H_{il}] = \begin{cases} \tilde{\mathbb{E}}^{(2)}[H_{ik}] \tilde{\mathbb{E}}^{(2)}[H_{il}] & \text{if } k \neq l \\ \tilde{\mathbb{E}}^{(2)}[H_{ik}]^2 + \tilde{\mathbb{V}}^{(2)}[H_{ik}] & \text{if } k = l, \end{cases}$$

where $\tilde{\mathbb{V}}^{(2)}$ denotes the variance under the Enhanced approximation. Furthermore, $\tilde{\mathbb{E}}^{(2)}[H_{ij}] = \tilde{\mathbb{E}}^{(2)}[\tilde{\mathbb{E}}^{(2)}[H_{ij}|W_{ij}]] = h_{ij}(1 - P_{ij})$. Using the law of total variance, we also have

$$\begin{aligned}
\tilde{\mathbb{V}}^{(2)}[H_{ij}] &= \tilde{\mathbb{V}}^{(2)}[\tilde{\mathbb{E}}^{(2)}[H_{ij}|W_{ij}]] + \tilde{\mathbb{E}}^{(2)}[\tilde{\mathbb{V}}^{(2)}[H_{ij}|W_{ij}]] \\
&= \tilde{\mathbb{V}}^{(2)}[h_{ij}(1 - W_{ij})] + \tilde{\mathbb{E}}^{(2)}[s_{ij}^2(1 - W_{ij}) + \Sigma_{jj} W_{ij}] \\
&= h_{ij}^2 P_{ij}(1 - P_{ij}) + s_{ij}^2(1 - P_{ij}) + \Sigma_{jj} P_{ij}.
\end{aligned}$$

This gives

$$\begin{aligned}
\tilde{\mathbb{E}}^{(2)}[\|\mathbf{H}_i\|_\Omega^2] &= \sum_{k=1}^p \sum_{l=1}^p \tilde{\mathbb{E}}^{(2)}[H_{ik} \Omega_{kl} H_{il}] = \sum_{k=1}^p \sum_{l=1}^p \Omega_{kl} \tilde{\mathbb{E}}^{(2)}[H_{ik} H_{il}] \\
&= \sum_{k=1}^p \sum_{l=1}^p \Omega_{kl} \tilde{\mathbb{E}}^{(2)}[H_{ik}] \tilde{\mathbb{E}}^{(2)}[H_{il}] + \sum_{j=1}^p \Omega_{jj} \tilde{\mathbb{V}}^{(2)}[H_{ij}] \\
&= \tilde{\mathbb{E}}^{(2)}[\mathbf{H}_i]^\top \Omega \tilde{\mathbb{E}}^{(2)}[\mathbf{H}_i] + \text{diag}(\Omega)^\top \text{diag}(\tilde{\mathbb{V}}^{(2)}[\mathbf{H}_i]) \\
(18) \quad &= [(\mathbf{M}_i - \mathbf{x}_i^\top \mathbf{B}) \odot (\mathbf{1}_p - \mathbf{P}_i)]^\top \Omega [(\mathbf{M}_i - \mathbf{x}_i^\top \mathbf{B}) \odot (\mathbf{1}_p - \mathbf{P}_i)] \\
&\quad + \text{diag}(\Omega)^\top (\mathbf{P}_i \odot \mathbf{Q}_i \odot (\mathbf{M}_i - \mathbf{x}_i^\top \mathbf{B})^2 + \mathbf{Q}_i \odot \mathbf{s}_i^2 + \text{diag}(\Sigma) \odot \mathbf{P}_i),
\end{aligned}$$

where $\mathbf{Q}_i = \mathbf{1}_p - \mathbf{P}_i$. For the Standard approximation, a similar argument applies but $\tilde{\mathbb{E}}^{(1)}[H_{ij}] = M_{ij}$ and $\tilde{\mathbb{V}}^{(1)}[H_{ij}] = s_{ij}^2$, giving the much simpler term

$$(19) \quad \tilde{\mathbb{E}}^{(1)}[\|\mathbf{Z}_i - \mathbf{x}_i^\top \mathbf{B}\|_\Omega^2] = (\mathbf{M}_i - \mathbf{x}_i^\top \mathbf{B})^\top \Omega (\mathbf{M}_i - \mathbf{x}_i^\top \mathbf{B}) + \text{diag}(\Omega)^\top \mathbf{s}_i^2.$$

As both $\tilde{\mathbb{E}}^{(1)}$ and $\tilde{\mathbb{E}}^{(2)}$ coincides on the remaining terms of the complete data log-likelihood, we drop the index and denote $\tilde{A}_{ij} = \tilde{\mathbb{E}}[\exp(o_{ij} + Z_{ij})] = \exp(o_{ij} + M_{ij} + S_{ij}^2/2)$. Taking the variational expectation gives

$$\begin{aligned}
& \tilde{\mathbb{E}}[\log p_{\theta}(\mathbf{Y}|\mathbf{Z}, \mathbf{W}) + \log p_{\theta}(\mathbf{W})] = \\
& \sum_{i,j} \tilde{\mathbb{E}}[(1 - W_{ij})(Y_{ij}(o_{ij} + Z_{ij}) - e^{o_{ij} + Z_{ij}} - \log(Y_{ij}!)) + W_{ij}\delta_{0,\infty}(Y_{ij})] \\
& + \sum_{i,j} \tilde{\mathbb{E}}[W_{ij}\mathbf{x}_i^{0\top} \mathbf{B}_j^0 - \log(1 + e^{\mathbf{x}_i^{0\top} \mathbf{B}_j^0})] \\
(20) \quad & = \sum_{i,j} (1 - P_{ij})(Y_{ij}(o_{ij} + M_{ij}) - \tilde{A}_{ij} - \log(Y_{ij}!)) + P_{ij}\delta_{0,\infty}(Y_{ij}) \\
& + \sum_{i,j} P_{ij}\mathbf{x}_i^{0\top} \mathbf{B}_j^0 - \log(1 + e^{\mathbf{x}_i^{0\top} \mathbf{B}_j^0})
\end{aligned}$$

where all operations (exp, log, logit, etc) are applied component-wise. Putting Equations (15), (17), (19) and (20) together gives $J^{(1)}$ and Equations (16) to (18) and (20) gives $J^{(2)}$. The writing in compact matrix form is left to the reader. \square

A.2. Proofs.

PROPOSITION 1. *The simple ZIPLN model defined in (4) with parameter $\theta = (\Omega, \mu, \pi)$ and parameter space $\mathbb{S}_p^{++} \times \mathbb{R}^p \times (0, 1)^p$ is identifiable.*

PROOF. We use the moments of \mathbf{Y} to prove identifiability. Letting $A_j = \exp(\mu_j + \sigma_{jj}/2) = \mathbb{E}[e^{Z_j}]$, and using results on moments of Poisson and Gaussian distributions,

- (i) If $U \sim \mathcal{N}(\mu, \sigma^2)$, then $\mathbb{E}[e^U] = \exp(\mu + \sigma^2/2)$,
- (ii) If $U \sim \mathcal{P}(\lambda)$ then $\mathbb{E}[U] = \lambda$, $\mathbb{E}[U^2] = \lambda(1 + \lambda)$, $\mathbb{E}[U^3] = \lambda(1 + 3\lambda + \lambda^2)$,

we have that $\mathbb{E}[(e^{Z_j})^2] = A_j^2 e^{\sigma_{jj}}$ and $\mathbb{E}[(e^{Z_j})^3] = A_j^3 e^{3\sigma_{jj}}$. By the law of total expectation, we get the first three moments of Y_j :

$$\begin{aligned}
\mathbb{E}[Y_j] &= (1 - \pi_j)A_j \\
\mathbb{E}[Y_j^2] &= (1 - \pi_j)A_j[1 + A_j e^{\sigma_{jj}}] \\
\mathbb{E}[Y_j^3] &= (1 - \pi_j)A_j[1 + 3A_j e^{\sigma_{jj}} + A_j^2 e^{3\sigma_{jj}}]
\end{aligned}$$

In parallel, using the conditional independence of Y_j and Y_k ($j \neq k$) knowing W_j, W_k , the independence of W_j, W_k and the law of total covariance, we have

$$\text{Cov}(Y_j, Y_k) = (1 - \pi_j)(1 - \pi_k)A_j A_k (e^{\sigma_{jk}} - 1).$$

Using arithmetic manipulations of the moments, we have

$$\begin{aligned}
A_j e^{\sigma_{jj}} &= \frac{\mathbb{E}[Y_j^2]}{\mathbb{E}[Y_j]} - 1 = \frac{\mathbb{E}[Y_j^2] - \mathbb{E}[Y_j]}{\mathbb{E}[Y_j]} \\
e^{\sigma_{jj}} &= \frac{\mathbb{E}[Y_j^3] - \mathbb{E}[Y_j^2] - 2\mathbb{E}[Y_j]A_j e^{\sigma_{jj}}}{\mathbb{E}[Y_j](A_j e^{\sigma_{jj}})^2} = \frac{\mathbb{E}[Y_j^3] - 3\mathbb{E}[Y_j^2] + 2\mathbb{E}[Y_j]}{(\mathbb{E}[Y_j^2] - \mathbb{E}[Y_j])^2 / \mathbb{E}[Y_j]} \\
e^{\mu_j} &= A_j e^{-\sigma_{jj}/2} = A_j e^{\sigma_{jj}} e^{-3\sigma_{jj}/2}
\end{aligned}$$

$$\begin{aligned}
&= (\mathbb{E}[Y_j^3] - 3\mathbb{E}[Y_j^2] + 2\mathbb{E}[Y_j])^{-3/2} (\mathbb{E}[Y_j^2] - \mathbb{E}[Y_j])^4 \mathbb{E}[Y_j]^{-1/2} \\
&= \frac{(\mathbb{E}[Y_j^2] - \mathbb{E}[Y_j])^4}{\sqrt{(\mathbb{E}[Y_j^3] - 3\mathbb{E}[Y_j^2] + 2\mathbb{E}[Y_j])^3 \mathbb{E}[Y_j]}} \\
e^{\sigma_{jk}} &= 1 + \frac{\text{Cov}[Y_j, Y_k]}{\mathbb{E}[Y_j]\mathbb{E}[Y_k]} \\
\pi_j &= 1 - \mathbb{E}[Y_j]/A_j = 1 - \frac{\mathbb{E}[Y_j]e^{\sigma_{jj}}}{A_j e^{\sigma_{jj}}} = 1 - \frac{\mathbb{E}[Y_j]^3 (\mathbb{E}[Y_j^3] - 3\mathbb{E}[Y_j^2] + 2\mathbb{E}[Y_j])}{(\mathbb{E}[Y_j^2] - \mathbb{E}[Y_j])^3}
\end{aligned}$$

Hence, each coordinate of θ can be expressed as a simple functions of the (first three) moments of the distribution p_θ of Y_j and thus $p_\theta = p_{\theta'} \Rightarrow \theta = \theta'$. \square

PROPOSITION 6 (Derivatives). $J^{(1)}$ has the following first order partial derivatives.

$$\begin{aligned}
\frac{\partial J^{(1)}}{\partial \Omega} &= \frac{n}{2} \Omega^{-1} - \frac{1}{2} \left[(M - \mathbf{X}B)^\top (M - \mathbf{X}B) + \bar{S}^2 \right], \\
\frac{\partial J^{(1)}}{\partial B} &= \mathbf{X}^\top \mathbf{X}B\Omega - \mathbf{X}^\top M\Omega, \\
\frac{\partial J^{(1)}}{\partial B^0} &= \mathbf{X}^{0\top} P - \mathbf{X}^{0\top} \left(\frac{\exp(\mathbf{X}^0 B^0)}{1 + e^{\mathbf{X}^0 B^0}} \right), \\
\frac{\partial J^{(1)}}{\partial M} &= (\mathbf{1}_{n,p} - P) \odot [Y - A] - (M - \mathbf{X}B)\Omega, \\
\frac{\partial J^{(1)}}{\partial S} &= S^\circledast - (\mathbf{1}_{n,p} - P) \odot S \odot A - S \text{Diag}(\Omega), \\
\frac{\partial J^{(1)}}{\partial P} &= P \odot (A + \mathbf{X}^0 B^0 - \text{logit}(P)) - \log(1 - P),
\end{aligned}$$

where S^\circledast denotes $\frac{1}{S}$ where the division is applied component-wise.

APPENDIX B: ADDITIONAL SIMULATIONS

Reconstruction error and computation time. We define the reconstruction error as the RMSE between \hat{Y} and Y where \hat{Y} is defined component-wise as

$$\hat{Y}_{ij} = \tilde{\mathbb{E}}[(1 - W_{ij}) \exp(Z_{ij})] = (1 - P_{ij}) \exp(O_{ij} + M_{ij} + S_{ij}^2/2).$$

We maintain the experimental protocol outlined in Section 5.1. Setting γ to $\gamma = 2$, following the procedure in Section 5.2, and fixing the zero-inflation parameter π^* at 0.3, as detailed in Section 5.3, we fix the sample size to $n = 500$. We then gradually increase the number of variables p by considering values in $\Lambda = 100, 200, \dots, 500$. For each dimension size in Λ , we simulate 30 distinct parameter sets θ , resulting in a total of 30×5 parameter combinations. Each algorithm is run for 1000 iterations, although convergence is typically achieved within 500 iterations. The results are presented in Figure 8.

Compared to the Poisson Log-Normal (PLN) model, both the Non-Analytic VA and Analytic VA exhibit approximately 2 times and 3 times longer computation times, respectively. This increased computational demand is attributed to the additional variational and model parameters required by the zero-inflated models. The notable disparity between Analytic and Non-Analytic VA stems from the computation of the Lambert function, which relies on a

computationally intensive fixed-point method. With 500 variables and 1000 iterations, the computation time averages around 30 seconds, indicating a reasonable computational burden even for several thousand variables.

Regarding the reconstruction error, all VA methods demonstrate similar performance regardless of the model choice, revealing a consistent pattern. Notably, the PLN model yields the best results, contrary to the findings in Figures 3, 2, and 1, where the root mean square error (RMSE) with respect to each model parameter significantly favored the zero-inflated VA methods. This discrepancy suggests that achieving good reconstruction error does not necessarily correlate with accurate parameter estimation.




FIG 8. Computation times (in seconds) and reconstruction error (in observation units) when the number of variables p increases. 1000 iterations were performed for each algorithm.