



HAL
open science

Variational inference for probabilistic Poisson PCA

Julien Chiquet, Mahendra Mariadassou, Stephane Robin

► **To cite this version:**

Julien Chiquet, Mahendra Mariadassou, Stephane Robin. Variational inference for probabilistic Poisson PCA. *Annals of Applied Statistics*, 2018, 12 (4), pp.2674-2698. 10.1214/18-AOAS1177. hal-01941270

HAL Id: hal-01941270

<https://hal.science/hal-01941270>

Submitted on 12 May 2023

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.

VARIATIONAL INFERENCE FOR PROBABILISTIC POISSON PCA¹

BY JULIEN CHIQUET, MAHENDRA MARIADASSOU AND STÉPHANE ROBIN

AgroParisTech, INRA and Université Paris-Saclay

Many application domains, such as ecology or genomics, have to deal with multivariate non-Gaussian observations. A typical example is the joint observation of the respective abundances of a set of species in a series of sites aiming to understand the covariations between these species. The Gaussian setting provides a canonical way to model such dependencies but does not apply in general. We consider here the multivariate exponential family framework for which we introduce a generic model with multivariate Gaussian latent variables. We show that approximate maximum likelihood inference can be achieved via a variational algorithm for which gradient descent easily applies. We show that this setting enables us to account for covariates and offsets. We then focus on the case of the Poisson-lognormal model in the context of community ecology. We demonstrate the efficiency of our algorithm on microbial ecology datasets. We illustrate the importance of accounting for the effects of covariates to better understand interactions between species.

1. Introduction. Principal component analysis (PCA) is among the oldest and most popular tool for multivariate analysis. It basically aims at reducing the dimension of a large dataset made of continuous variables [Anderson (2003), Mardia, Kent and Bibby (1979)] in order to ease its interpretation and visualization. The methodology exploits the dependency structure between the variables to exhibit the few synthetic variables that best summarize the information content of the whole dataset, the principal components. In that sense PCA can be viewed as a way to better understand the dependency structure between the variables. From a purely algebraic point of view, PCA can be seen as a matrix-factorization problem where the data matrix is decomposed as the product of a loading matrix and a score matrix [Eckart and Young (1936)].

For statistical purposes PCA can also be cast in a probabilistic framework. Probabilistic PCA (pPCA) is a model-based version of PCA originally defined in a Gaussian setting in which the scores are treated as random hidden variables [Tipping and Bishop (1999), Minka (2000)]. It is closely related to factor analysis. As it involves hidden variables, maximum-likelihood estimates (MLE) can be obtained via an EM algorithm [Dempster, Laird and Rubin (1977)]. One major

Received March 2017; revised February 2018.

¹Supported in part by ANR Hydrogen (project ANR-14-CE23-0001) and Inra MEM Metaprogramme (Meta-omics and microbial ecosystems, projects LearnBioControl and Brassica-Dev).

Key words and phrases. Probabilistic PCA, Poisson-lognormal model, count data, variational inference.

interest of the probabilistic approach is that it allows to combine dimension reduction with other modeling tools, such as regression on some available covariates. Because observed variables can be affected by the variations of such covariates, the correction for their potential effects is desirable to avoid the presence of spurious correlations between the responses.

The Gaussian setting is obviously convenient as the dependency structure is entirely encoded in the covariance matrix, but pPCA has been extended to more general settings. Indeed, in many applications [Royle and Wikle (2005), Srivastava and Chen (2010)] Gaussian models need to be adapted to handle specific measurement types, such as binary or count data. For count data the multivariate Poisson distribution seems a natural counterpart of the multivariate normal. However, no canonical form exists for this distribution [Johnson, Kotz and Balakrishnan (1997)], and several alternatives have been proposed in the literature including Gamma-Poisson [Nelson (1985)] and lognormal-Poisson [Aitchison and Ho (1989), Izsák (2008)]. The latter takes advantage of the properties of the Gaussian distribution to display a larger panel of dependency structure than the former, but maximum likelihood-based inference raises some issues as the MLE of the covariance matrix is not always positive definite.

A series of works have contributed to extend PCA to a broader class of distributions, typically in the exponential family. The matrix factorization point of view has been adopted to satisfy a positivity constraint of the parameters [Lafond (2015)] and to minimize the Poisson loss function [Cao and Xie (2015)] or more general losses [Lee and Seung (2001)] consistent with exponential family noise. Sparse extensions have also been proposed [Witten, Tibshirani and Hastie (2009), Liu, Dobriban and Singer (2016)]. In a model-based context Collins, Dasgupta and Schapire (2001) suggest to minimize a Bregman divergence to get estimates of the scores. The divergence is chosen according to the distribution at hand, and a generic alternating minimization scheme is proposed. Salmon et al. (2014) consider a similar framework and use matrix factorization for the minimization of Bregman divergence. In both cases the scores are considered as fixed parameters. Mohamed, Ghahramani and Heller (2009) cast the same model in a Bayesian context and use Monte Carlo sampling for the inference. Acharya, Ghosh and Zhou (2015) consider Bayesian inference of the Gamma-Poisson distribution. A Bayesian version of PCA (where both loadings and scores are treated as random) is considered in Li and Tao (2010).

Landgraf (2015) reframes exponential family PCA as an optimization problem with some rank constraints and develops both a convex relaxation and a maximization-minimization algorithm for binomial and Poisson families. Finally, Zhou et al. (2012) and Zhou (2016) consider factor analysis in the more complex setting of negative-binomial families. Our approach differs from the previous ones as we only consider scores as random variables, whereas we consider the loadings as fixed parameters in the exact analog of Tipping and Bishop's pPCA.

As recalled above, in pPCA, the scores are treated as hidden variables. One of the main issues of non-Gaussian pPCA arises from the fact that their conditional distribution given the observed data is often intractable which hampers the use of an expectation–minimization (EM) strategy. Variational approximations [Jaakkola and Jordan (2000), Wainwright and Jordan (2008)] have become a standard tool to approximate such conditional distributions. Karlis (2005) uses such an approximation for the inference of the one-dimensional Poisson-lognormal model and derives a variational EM (VEM) algorithm. Hall, Ormerod and Wand (2011) provide a theoretical analysis of this approximation for the same model and prove the consistency of the estimators. Indeed, even the conditional distribution of one single hidden coordinate (given all others) is unknown which makes regular Gibbs sampling inaccessible. As a consequence Lee and Seung (2001) use moment estimates, whereas in a Bayesian context Li and Tao (2010) resort to a variational approximation of the conditional distribution.

Our contribution. We define a general framework for pPCA in the simple exponential family. The model we consider combines dimension reduction (via pPCA) and regression in order to account for known effects and focus on the remaining dependency structure. Scores are assumed to be Gaussian to allow a large panel of dependency structures. We put a special emphasis on the analysis of count data. We adopt a frequentist setting rather than a Bayesian approach to avoid nonscalable, computing-heavy Monte Carlo sampling. We use a variational approximation of the conditional distribution of the scores given the observed data to derive a variational lower bound of the likelihood. Since only the scores are assumed to be random, we can prove that this bound is biconcave, that is, concave in the model parameters and in the variational parameters but not jointly concave in general. Biconcavity allows us to design a gradient-based method rather than a (variational) EM algorithm traditionally used in this setting.

We illustrate the interest of our model on two examples of microbial ecology. We show that the proposed algorithm is efficient for large datasets such as these encountered in metagenomics. We also show the importance of accounting for covariates and offset in order to go beyond first-order effects. More specifically, we show how the proposed modeling allows us to distinguish between correlations that are caused by known covariates from those corresponding to an unknown structure and requiring further investigations.

The paper is organized as follows: in Section 2 we introduce pPCA for the exponential family and the variational framework that we consider. Section 3 generalizes the model in the manner of a generalized linear model in order to handle covariates and offsets. Section 4 is dedicated to the inference and optimization strategy. Section 5 details the special Poisson case, and Section 6 devises the visualization, an important issue for non-Gaussian PCA methods. Finally, Section 7 considers applications to two examples from metagenomics: the impact of a pathogenic fungi on microbial communities from tree leaves, and the impact of weaning on piglets gut microbiota.

2. A variational framework for probabilistic PCA in the exponential family. We start this section by stating the probabilistic framework associated to Gaussian probabilistic PCA. Then we show how it can be naturally extended to other exponential families. We finally derive variational lower bounds for the likelihood of pPCA and its gradient which are the building blocks of our inference strategy.

2.1. *Gaussian probabilistic PCA (pPCA).* The probabilistic version of principal component analysis or pPCA [Minka (2000), Mohamed, Ghahramani and Heller (2009), Tipping and Bishop (1999)] relates a sample of p -dimensional observation vectors \mathbf{Y}_i to a sample of q -dimensional vectors of latent variables \mathbf{W}_i in the following way:

$$(1) \quad \mathbf{Y}_i = \boldsymbol{\mu} + \mathbf{B}\mathbf{W}_i + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{\varepsilon}_i \sim \mathcal{N}(\mathbf{0}_p, \sigma^2 \mathbf{I}_p).$$

The parameter $\boldsymbol{\mu}$ allows the model to have *main effects*. The $p \times q$ matrix \mathbf{B} captures the dependence between latent and observed variables. Furthermore, the latent vectors are conventionally assumed to have independent Gaussian components with unit variance, that is, to say $\mathbf{W}_i \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)$. This ensures that there is no structure in the latent space. Model (1) can thus be restated as $\mathbf{Y}_i \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{B}\mathbf{B}^\top + \sigma^2 \mathbf{I}_p)$.

In the following we consider an alternative formulation stated in a hierarchical framework. Despite its seemingly more complex nature, it lends itself nicely to generalizations. Formally,

$$(2) \quad \begin{aligned} &\text{latent space } (\mathbf{W}_i)_{i=1, \dots, n} \text{ i.i.d., } \mathbf{W}_i \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q) \\ &\text{parameter space } \mathbf{Z}_i = \boldsymbol{\mu} + \mathbf{B}\mathbf{W}_i, \\ &\text{observation space } (Y_{ij} | Z_{ij})_{i=1, \dots, n; j=1, \dots, p} \text{ indep., } Y_{ij} | Z_{ij} \sim \mathcal{N}(Z_{ij}, \sigma^2). \end{aligned}$$

In equation (2) \mathbf{Z}_i is a linear transform of \mathbf{W}_i , and the last layer $Y_{ij} | Z_{ij}$ simply corresponds to *observation noise*. Informally, the *latent* variables \mathbf{W}_i (in \mathbb{R}^q) are mapped to a linear subspace of the *parameter* space \mathbb{R}^p via the \mathbf{Z}_i which are then pushed into the *observation* space using Gaussian emission laws. The main idea of this paper is to replace Gaussian emission laws with univariate natural exponential families.

Note that the diagonal nature of the covariance matrix of $\boldsymbol{\varepsilon}_i$ specified in (1) now means that, *conditionally on* \mathbf{Z}_i , all components of \mathbf{Y}_i are independent. This is why we may consider univariate variables $Y_{ij} | Z_{ij}$ in Formulation (2). Although the observation noises are conditionally independent, the coordinates of a given \mathbf{Y}_i are not which makes the model genuinely multivariate. This is further emphasized in Section 5.1.

The loading matrix \mathbf{B} is a convenience parameter that is useful for both optimization and visualization of the model but not identifiable *per se*. Indeed any

orthogonal transformation of \mathbf{B} leads to the same model. Denoting $\mathbf{\Sigma} = \mathbf{B}\mathbf{B}^\top$, Model (2) can be rephrased as

$$(\mathbf{Z}_i)_{i=1,\dots,n} \text{ i.i.d., } \mathbf{Z}_i \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{\Sigma}),$$

$$(Y_{ij}|\mathbf{Z}_{ij})_{i=1,\dots,n; j=1,\dots,p} \text{ indep., } Y_{ij}|\mathbf{Z}_{ij} \sim \mathcal{N}(Z_{ij}, \sigma^2).$$

As a consequence, the identifiable parameters of the model are $\boldsymbol{\mu}$ and $\mathbf{\Sigma}$.

Hereafter and unless stated otherwise, index i refers to *observations* and ranges in $\{1, \dots, n\}$, index j refers to *variables* and ranges in $\{1, \dots, p\}$ and index k refers to *factors* and ranges in $\{1, \dots, q\}$.

2.2. *Natural exponential family (NEF).* The work in this study is based on essential properties of univariate *natural exponential families* (NEF) where the parameter is in canonical form. They include normal distribution with known variance, Poisson distribution, gamma distribution with known shape parameter (and therefore exponential distribution as a particular example) and binomial distribution with known number of trials. The probability density (or mass function) of a NEF can be written

$$(3) \quad f(x|\lambda) = \exp(x\lambda - b(\lambda) - a(x)),$$

where λ is the canonical parameter and b and a are known functions. The function b is well known to be convex (and analytic) over its domain and the mean and variance are easily deduced from b as

$$\mathbb{E}_\lambda[X] = b'(\lambda) \quad \text{and} \quad \mathbb{V}_\lambda[X] = b''(\lambda).$$

The canonical link function g is defined such that $g(b'(\lambda)) = \lambda$. The maximum likelihood estimate $\hat{\lambda}$ of λ from a single observation x is given by $\hat{\lambda} = \hat{\lambda}(x) = g(x)$ and satisfies

$$\mathbb{E}_{\hat{\lambda}(x)}[X] = b'(\hat{\lambda}(x)) = x.$$

2.3. *Probabilistic PCA for the exponential family.* We now extend pPCA from the Gaussian setting to more general NEF. The connection between the two versions is exactly the same as the connection between linear models and generalized linear models (GLM). Intuitively, we assume that (i) there exists a (low) q -dimensional (linear) subspace in the *natural canonical parameter space* where some latent variable \mathbf{Z}_i lie, and (ii) observations \mathbf{Y}_i are generated in the *observation space* according to some NEF distribution with parameter \mathbf{Z} . The latter is linked to $\mathbb{E}[\mathbf{Y}_i|\mathbf{Z}_i]$ through the canonical link function g . In the Gaussian case the link function is the identity and the parameter space can be identified with the observation space but this is not the case in general for other families. Formally, we

extend Model (2) to

$$\begin{aligned}
 & (\mathbf{W}_i)_{i=1,\dots,n} \text{ i.i.d., } \mathbf{W}_i \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q), \\
 & \mathbf{Z}_i = \boldsymbol{\mu} + \mathbf{B}\mathbf{W}_i, \\
 (4) \quad & (Y_{ij}|Z_{ij})_{i=1,\dots,n;j=1,\dots,p} \text{ indep.,} \\
 & p(Y_{ij}|Z_{ij}) = \exp(Y_{ij}Z_{ij} - b(Z_{ij}) - a(Y_{ij})).
 \end{aligned}$$

Note in particular that $g(\mathbb{E}[Y_{ij}|Z_{ij}]) = g(b'(Z_{ij})) = Z_{ij}$ and that an unconstrained estimate \tilde{Z}_{ij} of Z_{ij} is $\tilde{Z}_{ij} = g(Y_{ij})$. The vector $\boldsymbol{\mu}$ corresponds to main effects, \mathbf{B} to *rescaled* loadings in the parameter spaces and \mathbf{W}_i to scores of the i th observation in the low-dimensional latent subspace. The model specified in (4) is the same as the one specified in (2) but for the last data emission layer. Similar to Model (2), the first two lines of Model (4) can be combined into \mathbf{Z}_i i.i.d. such that $\mathbf{Z}_i \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = \mathbf{B}\mathbf{B}^\top$.

REMARK 1. As stated previously, \mathbf{B} is only identifiable through $\mathbf{B}\mathbf{B}^\top$ and therefore at best up to rotations in \mathbb{R}^q . Note that this limitation is shared with standard PCA. Intuitively, PCA finds a good q -dimensional affine approximation subspace $\boldsymbol{\mu} + \text{Span}(\mathbf{B})$ of \mathbf{Y} , but without additional constraints infinitely many bases \mathbf{B} can be used to parametrize this subspace. Orthogonality constraints and ordering of the principal components in decreasing order of variance are necessary to uniquely specify \mathbf{B} . Imposing them in standard PCA additionally allows one to leverage Eckart and Young’s theorem and reduce a q -dimensional approximation to a series of q unidimensional problems. It also entails nestedness: the best q -dimensional approximation is nested within the best $q + 1$ -dimensional one and so on. There is unfortunately no equivalent in exponential PCA. We therefore do not force \mathbf{B} to be orthogonal in our model. For visualization however, we perform orthogonalization to ensure consistency of the graphical outputs with standard PCA (see Section 6).

2.4. *Likelihood.* Note \mathbf{Y} (resp. \mathbf{W}) the $n \times p$ (resp. $n \times q$) matrix obtained by stacking the row vectors \mathbf{Y}_i^\top (resp. \mathbf{W}_i^\top). Conversely, for any matrix \mathbf{A} , A_i refers to the i th row of \mathbf{A} considered as a *column* vector. In matrix expression $\mathbf{Z} = \mathbf{1}_n \boldsymbol{\mu}^\top + \mathbf{W}\mathbf{B}^\top$. The observation matrix \mathbf{Y} only depends on \mathbf{Z} through $\boldsymbol{\mu}$, \mathbf{B} and \mathbf{W} , and the complete log-likelihood is therefore

$$\begin{aligned}
 \log p(\mathbf{Y}, \mathbf{W}; \boldsymbol{\mu}, \mathbf{B}) &= \sum_{i=1}^n \log p(\mathbf{Y}_i|\mathbf{W}_i; \boldsymbol{\mu}, \mathbf{B}) + \log p(\mathbf{W}_i) \\
 &= \sum_{i=1}^n \left[\sum_{j=1}^p Y_{ij}(\mu_j + \mathbf{B}_j^\top \mathbf{W}_i) - b(\mu_j + \mathbf{B}_j^\top \mathbf{W}_i) \right. \\
 &\quad \left. - a(Y_{ij}) - \sum_{k=1}^q \frac{W_{ik}^2 + \log(2\pi)}{2} \right]
 \end{aligned}$$

which can be stated in the following compact matrix form:

$$(5) \quad \log p(\mathbf{Y}, \mathbf{W}; \boldsymbol{\mu}, \mathbf{B}) = \mathbf{1}_n^\top [\mathbf{Y} \odot (\mathbf{1}_n \boldsymbol{\mu}^\top + \mathbf{W} \mathbf{B}^\top) - b(\mathbf{1}_n \boldsymbol{\mu}^\top + \mathbf{W} \mathbf{B}^\top)] \mathbf{1}_p - \frac{\|\mathbf{W}\|_F^2}{2} - \frac{nq}{2} \log(2\pi) - K(\mathbf{Y}),$$

where the function a and b are applied componentwise to vectors and matrices, \odot is the Hadamard product, and $K(\mathbf{Y}) = \mathbf{1}_n^\top a(\mathbf{Y}) \mathbf{1}_p$ is a constant depending only on \mathbf{Y} and not scaling with q .

We do not know how to integrate out \mathbf{W} and therefore cannot derive an analytic expression of $\log p(\mathbf{Y}; \boldsymbol{\mu}, \mathbf{B})$. Numerical approximation using Hermite-Gauss quadrature or MCMC techniques are possible but rely on computing np expectations of the form $\mathbb{E}[e^{a^\top \mathbf{u} - b(\alpha + c^\top \mathbf{u})}]$ for $\mathbf{u} \sim \mathcal{N}(0, \mathbf{I}_q)$, with b nonlinear, \mathbf{a} and \mathbf{c} arbitrary vectors and α a scalar depending on $\boldsymbol{\mu}$ and \mathbf{B} . This is likely to become computationally prohibitive as the dimension q of the latent integration space increases. A standard EM algorithm relying on $\mathbb{E}_{\mathbf{W}|\mathbf{Y}}[\log p(\mathbf{Y}, \mathbf{W}; \boldsymbol{\mu}, \mathbf{B})]$ is similarly not possible as it requires at least first and second order of $p(\mathbf{W}_i|\mathbf{Y}_i)$ which are unknown in general and as hard to compute as the previous expectations. We resort instead to a variational strategy and integrate out \mathbf{W} under a tractable approximation of $p(\mathbf{W}|\mathbf{Y})$.

2.5. *Variational bound of the likelihood.* Consider any product distribution $\tilde{p} = \otimes_{i=1}^n \tilde{p}_i$ on the \mathbf{Z}_i . The variational approximation relies on maximizing the following lower bound over a tractable set for \tilde{p} :

$$\log p(\mathbf{Y}; \boldsymbol{\mu}, \mathbf{B}) \geq J_q(\tilde{p}, \boldsymbol{\mu}, \mathbf{B}),$$

where

$$(6) \quad \begin{aligned} J_q(\tilde{p}, \boldsymbol{\mu}, \mathbf{B}) &\triangleq \log p(\mathbf{Y}; \boldsymbol{\mu}, \mathbf{B}) - KL(\tilde{p}(\mathbf{W}) \parallel p(\mathbf{W}|\mathbf{Y}; \boldsymbol{\mu}, \mathbf{B})) \\ &= \mathbb{E}_{\tilde{p}}[\log p(\mathbf{Y}, \mathbf{W}; \boldsymbol{\mu}, \mathbf{B}) - \log \tilde{p}(\mathbf{W})] \\ &= \sum_{i=1}^n \mathbb{E}_{\tilde{p}_i}[\log p(\mathbf{W}_i) + \log p(\mathbf{Y}_i|\mathbf{W}_i; \boldsymbol{\mu}, \mathbf{B}) - \log \tilde{p}_i(\mathbf{W}_i)], \end{aligned}$$

with term-by-term inequality

$$\begin{aligned} \log p(\mathbf{Y}_i; \boldsymbol{\mu}, \mathbf{B}) &\geq J_q(\tilde{p}_i, \boldsymbol{\mu}, \mathbf{B}) \\ &\triangleq \mathbb{E}_{\tilde{p}_i}[\log p(\mathbf{W}_i) + \log p(\mathbf{Y}_i|\mathbf{W}_i; \boldsymbol{\mu}, \mathbf{B}) - \log \tilde{p}_i(\mathbf{W}_i)]. \end{aligned}$$

In our variational approximation, we choose here the set \mathcal{Q} of product distribution of q -dimensional multivariate Gaussian with diagonal covariance matrices:

$$(7) \quad \mathcal{Q} = \left\{ \tilde{p} \triangleq \tilde{p}_{\mathbf{M}, \mathbf{S}}; \tilde{p}(\mathbf{w}) = \prod_{i=1}^n \tilde{p}_i(\mathbf{w}_i) \right\}$$

where $\tilde{p}_i = \mathcal{N}(\mathbf{m}_i, \text{diag}(\mathbf{s}_i \odot \mathbf{s}_i))$, $(\mathbf{m}_i, \mathbf{s}_i) \in \mathbb{R}^q \times \mathbb{R}_+^q$.

The $n \times q$ matrices \mathbf{M} and \mathbf{S} are obtained by respectively stacking \mathbf{m}_i^\top and \mathbf{s}_i^\top . Note that, by construction, $p(\mathbf{W}|\mathbf{Y})$ is a product distribution and that the approximation only stems from the functional form of each \tilde{p}_i , that is, multivariate normal with diagonal variance-covariance matrix. For such $\tilde{p} = \tilde{p}_{\mathbf{M},\mathbf{S}}$, results on first- and second-order moments of multivariate Gaussian show that

$$\begin{aligned} J_q(\boldsymbol{\mu}, \mathbf{B}, \mathbf{m}_i, \mathbf{s}_i) &\triangleq J_q(\tilde{p}_i, \boldsymbol{\mu}, \mathbf{B}) \\ &= \mathbf{Y}_i^\top (\boldsymbol{\mu} + \mathbf{B}\mathbf{m}_i) - \frac{1}{2} [\|\mathbf{m}_i\|_2^2 + \|\mathbf{s}_i\|_2^2] + \frac{1}{2} (\mathbf{2}_q^\top \log(\mathbf{s}_i) + q) \\ &\quad - \mathbf{1}_p^\top \mathbb{E}_{\tilde{p}_i} [b(\boldsymbol{\mu} + \mathbf{B}\mathbf{W}_i)] - K(\mathbf{Y}). \end{aligned}$$

Therefore,

$$\begin{aligned} J_q(\boldsymbol{\mu}, \mathbf{B}, \mathbf{M}, \mathbf{S}) &\triangleq J_q(\tilde{p}_{\mathbf{M},\mathbf{S}}, \boldsymbol{\mu}, \mathbf{B}) = \sum_{i=1}^n J_q(\boldsymbol{\mu}, \mathbf{B}, \mathbf{m}_i, \mathbf{s}_i) \\ (8) \quad &= \mathbf{1}_n^\top [\mathbf{Y} \odot (\mathbf{1}_n \boldsymbol{\mu}^\top + \mathbf{M}\mathbf{B}^\top) - \mathbb{E}_{\tilde{p}} [b(\mathbf{1}_n^\top \boldsymbol{\mu} + \mathbf{W}\mathbf{B}^\top)]] \mathbf{1}_p \\ &\quad - \frac{1}{2} \mathbf{1}_n^\top [\mathbf{M} \odot \mathbf{M} + \mathbf{S} \odot \mathbf{S} - 2 \log(\mathbf{S}) - \mathbf{1}_{n,q}] \mathbf{1}_q - K(\mathbf{Y}). \end{aligned}$$

Depending on the natural exponential family and thus the exact value of b in (8), we may have a fully explicit variational bound for the complete likelihood which paves the way for efficient optimization. In particular this is the case with the Poisson distribution that we investigate in further details in Section 5.

Before moving on to actual inference, we show how the framework introduced above can be extended to account for covariates and offsets.

3. Accounting for covariates and offsets. Multivariate analyses typically aim at deciphering dependencies between variables. Variations induced by the effect of covariates may be confounded with these dependencies. Therefore, it is desirable to account for such effects to focus on the residual dependencies. The rationale of our approach is to postulate the existence of a model similar to linear regression in the *parameter* space. We consider the general framework of linear regression with multivariate outputs, which encompasses multivariate analysis of variance.

3.1. *Model and likelihood.* Suppose that each observation i is associated to a known d -dimensional covariate vector \mathbf{X}_i . We assume that the covariates act linearly in the *parameter* space through a $p \times d$ regression matrix $\boldsymbol{\Theta}$, that is, \mathbf{X}_i is linearly related to \mathbf{Z}_i . It can be also useful to add an offset to model different sampling efforts and/or exposures. There is usually one known offset parameter O_{ij} per observation Y_{ij} , and this offset can be readily incorporated in our framework.

Thus, a natural generalization of (4) accounting for covariates and offsets is

$$\begin{aligned}
 & (\mathbf{W}_i)_{i=1,\dots,n} \text{ i.i.d. } \mathbf{W}_i \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q) \\
 & \mathbf{Z}_i = \mathbf{O}_i + \Theta \mathbf{X}_i + \mathbf{B} \mathbf{W}_i, \\
 & (Y_{ij} | Z_{ij})_{i=1,\dots,n; j=1,\dots,p} \text{ indep.}, \\
 & p(Y_{ij} | Z_{ij}) = \exp(Y_{ij} Z_{ij} - b(Z_{ij}) - a(Y_{ij})),
 \end{aligned}
 \tag{9}$$

where a column of ones can be added to the data matrix \mathbf{X} to get an intercept in the model. The log-likelihood can be computed from (9) like before to get

$$\begin{aligned}
 & \log p(\mathbf{Y}, \mathbf{W}; \mathbf{B}, \Theta, \mathbf{O}) \\
 & = \mathbf{1}_n^\top [\mathbf{Y} \odot (\mathbf{O} + \mathbf{X} \Theta^\top + \mathbf{W} \mathbf{B}^\top) - b(\mathbf{O} + \mathbf{X} \Theta^\top + \mathbf{W} \mathbf{B}^\top)] \mathbf{1}_p \\
 & - \frac{\|\mathbf{W}\|_F^2}{2} - \frac{nq}{2} \log(2\pi) - K(\mathbf{Y}),
 \end{aligned}
 \tag{10}$$

where the focus of inference is on \mathbf{B} and Θ while \mathbf{O} is known.

3.2. *Variational bound of the likelihood.* We can use the variational class \mathcal{Q} previously defined in (7) to lower bound the likelihood from equation (10). We first introduce the instrumental matrix \mathbf{A} which appears in many equations:

$$\begin{aligned}
 & \mathbf{A} = \mathbb{E}_{\tilde{p}} [b(\mathbf{O} + \mathbf{X} \Theta^\top + \mathbf{W} \mathbf{B}^\top)] \\
 & = \mathbb{E} [b(\mathbf{O} + \mathbf{X} \Theta^\top + (\mathbf{M} + \mathbf{S} \odot \mathbf{U}) \mathbf{B}^\top)] = \mathbb{E} [b(\mathbf{V})],
 \end{aligned}
 \tag{11}$$

where $\mathbf{V} = (\mathbf{O} + \mathbf{X} \Theta^\top + (\mathbf{M} + \mathbf{S} \odot \mathbf{U}) \mathbf{B}^\top)$, and \mathbf{U} is a $n \times q$ matrix with unit variance independent Gaussian components. \mathbf{V} can be interpreted as the variational counterpart of \mathbf{Z} .

Since \mathbf{O} is known, we drop it from the arguments of J_q and obtain the following lower bound which extends the bound from equation (8):

$$\begin{aligned}
 & J_q(\Theta, \mathbf{B}, \mathbf{M}, \mathbf{S}) = \mathbf{1}_n^\top (\mathbf{Y} \odot (\mathbf{O} + \mathbf{X} \Theta^\top + \mathbf{M} \mathbf{B}^\top) - \mathbf{A}) \mathbf{1}_p \\
 & - \frac{1}{2} \mathbf{1}_n^\top [\mathbf{M} \odot \mathbf{M} + \mathbf{S} \odot \mathbf{S} - 2 \log(\mathbf{S}) - \mathbf{1}_{n,q}] \mathbf{1}_q - K(\mathbf{Y}).
 \end{aligned}
 \tag{12}$$

4. Inference. As usual in the variational framework we aim to maximize the lower bound J_q , which we call the objective function in an optimization perspective. The optimization shall be performed on $\Theta, \mathbf{B}, \mathbf{M}, \mathbf{S}$. We only give results in the most general case (12) with covariates and offsets. All other cases are deduced by setting $\mathbf{O} = \mathbf{0}_{n \times p}$ and/or $\mathbf{X} = \mathbf{1}_n$ hereafter.

4.1. *Inference strategy.* We first highlight the biconcavity of the objective function J_q . The major part of the proof is postponed to the [Appendix](#).

PROPOSITION 1. *The variational objective function $J_q(\Theta, \mathbf{B}, \mathbf{M}, \mathbf{S})$ is concave in (Θ, \mathbf{B}) for (\mathbf{M}, \mathbf{S}) fixed and vice versa.*

PROOF. Fix (\mathbf{M}, \mathbf{S}) in (12). The nonexplicit part of J_q , that is, to say $-\mathbf{1}_n^\top \mathbf{A} \mathbf{1}_p$, is concave in (Θ, \mathbf{B}) thanks to Lemma 2 (see the Appendix). By inspection the explicit part of J_q involves linear, quadratic and concave functions of (Θ, \mathbf{B}) and is also concave. The objective J_q is therefore concave in (Θ, \mathbf{B}) . The same is true for (\mathbf{M}, \mathbf{S}) when fixing (Θ, \mathbf{B}) . \square

A standard approach for maximizing biconcave functions is block coordinate descents, of which the Expectation–Maximization (EM) algorithm is a popular representative in the latent variable setting. It is especially powerful when we have access to closed formula for both the optimal (\mathbf{M}, \mathbf{S}) given (Θ, \mathbf{B}) (E-step) and the optimal (Θ, \mathbf{B}) given (\mathbf{M}, \mathbf{S}) (M-step). However, the nonlinear nature of $\mathbb{E}_{\tilde{p}}[b(\mathbf{O} + \mathbf{X}\Theta^\top + \mathbf{W}\mathbf{B}^\top)]$ combined with careful inspection of the objective function J_q shows that setting the derivatives of J_q to zero, even after fixing the variational or model parameters, does not lead to closed formula for (\mathbf{M}, \mathbf{S}) nor for (\mathbf{B}, Θ) . Nevertheless, since we may derive convenient expressions for the gradient ∇J_q (see next Section 4.2), we propose to rely on the globally convergent method of moving asymptotes (MMA) algorithm for gradient-based local optimization introduced by Svanberg (2002) and implemented in the NLOPT optimization library [Johnson (2011)]. In the general case (12), the total number of parameters to optimize $J_q(\Theta, \mathbf{B}, \mathbf{M}, \mathbf{S})$ is $p(d + q) + 2nq$. We use box constraints for the variational parameters \mathbf{S} [i.e., the standard deviations in (7) and thus only defined on \mathbb{R}_+^q]. The starting point is chosen according to the exact value of b .

4.2. *Blockwise gradients of J_q .* The blockwise gradient of $J_q(\Theta, \mathbf{B}, \mathbf{M}, \mathbf{S})$ can be expressed compactly in matrix notations. We skip the tedious but straightforward derivations and present only the resulting partial gradients. We introduce $\mathbf{A}' = \mathbb{E}[b'(\mathbf{V})]$, the natural counterpart to matrix \mathbf{A} given in (11). Intuitively, A'_{ij} is the conditional expectation of Y_{ij} under \tilde{p}_i . On top of that, we need two other matrices denoted \mathbf{A}'_1 and \mathbf{A}'_2 , defined as follows:

$$\mathbf{A}'_1 = \mathbb{E}[b'(\mathbf{V})^\top (\mathbf{S} \odot \mathbf{U})], \quad \mathbf{A}'_2 = \mathbb{E}[(b'(\mathbf{V})\mathbf{B}) \odot \mathbf{U}].$$

With those matrices the derivatives of J_q can be expressed compactly as

$$(13) \quad \begin{aligned} \frac{\partial J_q}{\partial \Theta} &= (\mathbf{Y} - \mathbf{A}')^\top \mathbf{X}, & \frac{\partial J_q}{\partial \mathbf{B}} &= (\mathbf{Y} - \mathbf{A}')^\top \mathbf{M} - \mathbf{A}'_1, \\ \frac{\partial J_q}{\partial \mathbf{M}} &= (\mathbf{Y} - \mathbf{A}')\mathbf{B} - \mathbf{M}, & \frac{\partial J_q}{\partial \mathbf{S}} &= [\mathbf{S}^\odot - \mathbf{A}'_2 - \mathbf{S}], \end{aligned}$$

where the $n \times q$ matrix \mathbf{S}^\odot is the elementwise inverse of \mathbf{S} , that is, $S^\odot_{ij} = S^{-1}_{ij}$ for all $i = 1, \dots, n, q = 1, \dots, Q$.

In the following the resulting parameter estimates will be denoted by $\widehat{\Theta}$ and $\widehat{\mathbf{B}}$, and the optimal variational parameters will be denoted by $\widetilde{\mathbf{M}}$ and $\widetilde{\mathbf{S}}$. We use different notation on purpose in order to distinguish model parameters from variational ones.

4.3. *About missing data.* When the data are missing at random (MAR), the sampling does not disturb the inference and it is sufficient to maximize the likelihood on the observed part of the data [Little and Rubin (2014)]. Our model can easily handle missing data under MAR conditions as follows: note $\Omega \subset \{1, \dots, n\} \times \{1, \dots, p\}$ the set of observed data and $\mathbf{\Omega}$ the matrix where $\Omega_{ij} = 1$ if $(i, j) \in \Omega$ and 0 otherwise. With this matrix $\mathbf{\Omega}$ the likelihood can be adapted from equation (10), and one has

$$\begin{aligned} \log p(\mathbf{Y}, \mathbf{W}; \mathbf{B}, \Theta, \mathbf{O}) &= \mathbf{1}_n^\top ((\mathbf{Y} \odot (\mathbf{O} + \mathbf{X}\Theta^\top + \mathbf{W}\mathbf{B}^\top) - b(\mathbf{O} + \mathbf{X}\Theta^\top + \mathbf{W}\mathbf{B}^\top)) \odot \mathbf{\Omega}) \mathbf{1}_p \\ &\quad - \frac{\|\mathbf{W}\|_F^2}{2} - \frac{nq}{2} \log(2\pi) - \text{tr}(\mathbf{\Omega}^\top a(\mathbf{Y})). \end{aligned}$$

The corresponding variational bound J_q and its partial derivatives are then simple adaptations from equations (12) and (13) where \mathbf{Y} (resp. \mathbf{A}, \mathbf{A}') is replaced with $\mathbf{Y} \odot \mathbf{\Omega}$ (resp. $\mathbf{A} \odot \mathbf{\Omega}, \mathbf{A}' \odot \mathbf{\Omega}$).

Note that it is strictly equivalent for the optimization method to use $(\mathbf{Y} - \mathbf{A}') \odot \mathbf{\Omega}$ or to impute missing Y_{ij} with A'_{ij} before using equation (13). Since \mathbf{A}' is computed as part of the gradient computation at each step, imputation of missing data is essentially a free by-product of the optimization method. Finally, note that $A'_{ij} = \mathbb{E}_{\tilde{p}_i}[Y_{ij}]$ so that the imputation makes intuitive sense. We are imputing Y_{ij} with its conditional expectation under the current variational parameters. Addressing not MAR conditions requires to take into account the sampling process leading to missing data in order to correctly unbiased the estimation. This is out of the scope of this paper.

4.4. *Variance estimation.* As mentioned above, only Θ and Σ are identifiable parameters, and an estimate of the later needs to be derived. Recall that Model (9) can be rephrased as $\mathbf{Z}_i \sim \mathcal{N}(\mathbf{O}_i + \Theta \mathbf{X}_i, \Sigma)$. It can be checked that the corresponding variational lower bound is maximal for

$$\widehat{\Sigma} = \frac{1}{n} \sum_i \mathbb{E}_{\tilde{p}}[(\mathbf{Z}_i - \mathbf{O}_i - \widehat{\Theta} \mathbf{X}_i)(\mathbf{Z}_i - \mathbf{O}_i - \widehat{\Theta} \mathbf{X}_i)^\top].$$

Since $\mathbb{E}_{\tilde{p}}(\mathbf{Z}_i) = \mathbf{O}_i + \widehat{\Theta} \mathbf{X}_i + \widetilde{\mathbf{M}}_i \widehat{\mathbf{B}}^\top$ and $\mathbb{V}_{\tilde{p}}(\mathbf{Z}_i) = \widehat{\mathbf{B}} \text{diag}(\tilde{s}_i \odot \tilde{s}_i) \widehat{\mathbf{B}}^\top$, we get

$$\widehat{\Sigma} = \widehat{\mathbf{B}} \left(\frac{1}{n} \widetilde{\mathbf{M}}^\top \widetilde{\mathbf{M}} + \widetilde{\mathbf{S}} \right) \widehat{\mathbf{B}}^\top,$$

where $\widetilde{\mathbf{S}} = n^{-1} \text{diag}[\mathbf{1}_n^\top (\widetilde{\mathbf{S}} \odot \widetilde{\mathbf{S}})]$. Observe that $\widehat{\Sigma}$ has rank q by construction.

4.5. *Model selection.* The dimension q of the latent space itself needs to be estimated. To this aim, we adopt a penalized-likelihood approach, replacing the log-likelihood by its lower bound J_q . We consider two classical criteria: BIC [Schwarz (1978)] and ICL [Biernacki, Celeux and Govaert (2000)]. 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 $p(\mathbf{W}|\mathbf{Y})$ is intractable, we replace it with its variational approximation $\tilde{p}(\mathbf{W})$ to evaluate this entropy. The number of parameters in our model is $p(q + d)$ and the entropy of each W_i under \tilde{p}_i is $q \log(2\pi e)/2 + \sum_j \log(s_{ij})$. Based on this we define the following approximate BIC and ICL criteria:

$$(14) \quad \begin{aligned} \text{BIC}(q) &= J_q - \frac{1}{2}p(d + q) \log(n), \\ \text{ICL}(q) &= J_q - \frac{1}{2}p(d + q) \log(n) - \frac{nq}{2} \log(2\pi e) - \mathbf{1}_n^\top \log(\mathbf{S})\mathbf{1}_q. \end{aligned}$$

5. **Poisson family.** Each term of the expectation matrix A in (11) can be reduced to computing expectations of the form $\mathbb{E}[b(a + cU)]$ for a convex analytic function b , a standard Gaussian $U \sim \mathcal{N}(0, 1)$ and arbitrary scalars $(a, c) \in \mathbb{R}^2$. It can therefore be computed numerically efficiently using Gauss–Hermite quadrature [see, e.g., Press et al. (1989)]. However in the special case of Poisson-distributed observations, $b(x) = e^x$ and most of the expectations can be computed analytically leading to explicit formulas for equations (11), (12) and (13).

5.1. *Some features of Poisson pPCA.* The Poisson pPCA inherits some properties of the Poisson-lognormal distribution which states that the response vector \mathbf{Y}_i for sample i is generated such that $\mathbf{Z}_i \sim \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma})$ and the $(Y_{ij})_j$ are independent conditionally on \mathbf{Z}_i with $Y_{ij}|Z_{ij} \sim \mathcal{P}(\exp(Z_{ij}))$. The moments of the Y_{ij} 's are then

$$\begin{aligned} \mathbb{E}Y_{ij} &= e^{\mu_j + \sigma_j^2/2}, & \mathbb{V}Y_{ij} &= \mathbb{E}Y_{ij} + (e^{\sigma_j^2} - 1)(\mathbb{E}Y_{ij})^2, \\ \text{Cov}(Y_{ij}, Y_{ik}) &= (e^{\sigma_{jk}} - 1)\mathbb{E}Y_{ij}\mathbb{E}Y_{ik}. \end{aligned}$$

Consequently, the Poisson-lognormal model displays both overdispersion of each coordinate with respect to a Poisson distribution and pairwise correlations of arbitrary signs. In Poisson pPCA $\boldsymbol{\Sigma}$ is further assumed to have a low rank.

5.2. *Explicit form of A , J_q , and ∇J_q .* In the Poisson case the variational expectation of the nonlinear part involving b —the matrix of conditional expectations A —is equal to A' and can be expressed as

$$A = A' = \exp\left(\mathbf{O} + \mathbf{X}\boldsymbol{\Theta}^\top + \mathbf{M}\mathbf{B}^\top + \frac{1}{2}(\mathbf{S} \odot \mathbf{S})(\mathbf{B} \odot \mathbf{B})^\top\right).$$

The lower bound J_q and matrices A'_1, A'_2 appearing in (13) can be expressed simply from \mathbf{A} as

$$A'_1 = [A^\top (\mathbf{S} \odot \mathbf{S})] \odot \mathbf{B}, \quad A'_2 = 2[A(\mathbf{B} \odot \mathbf{B})] \odot \mathbf{S}.$$

5.3. Implementation details. We implemented our inference algorithm for the Poisson family in the R package *PLNmodels*, the last version of which is available on github (<https://github.com/jchiquet/PLNmodels>). Maximization of the variational bound J_q is done using the implementation found in the *nlopt* library [Johnson (2011)] of the globally convergent method of moving asymptotes algorithm for gradient-based local optimization [Svanberg (2002)]. We interface this algorithm to R [R Development Core Team (2008)] via the *nloptr* package [Ypma (2017)] and careful tuning of the parameters. All graphics are produced using the *ggplot2* package [Wickham (2009)].

The choice of a good starting value is crucial in iterative procedures as it helps the algorithm to start in the attractor field of a good local maximum and can substantially speed-up convergence. Here we initialize (Θ, \mathbf{B}) by fitting a linear model to $\log(1 + \mathbf{Y})$ then extracting the regression coefficients Θ_{LM} and the variance-covariance matrix Σ_{LM} of the Pearson residuals. We set $\Theta_0 = \Theta_{\text{LM}}$ and $\mathbf{B}_0 = (\Sigma_{\text{LM}}^{(q)})^{1/2}$ where $\Sigma_{\text{LM}}^{(q)}$ is the best rank q approximation of Σ_{LM} , as given by keeping the first q -dimensions of a SVD of Σ_{LM} . We set the other starting values as $\mathbf{M}_0 = \mathbf{S}_0 = \mathbf{0}_{n \times q}$.

6. Visualization.

6.1. Specific issues in non-Gaussian PCA. PCA is routinely used to visualize samples in a low dimensional space. Visualisation in exponential PCA shares many similarities with visualization in standard PCA, but important differences arise from the lack of validity of Eckart and Young's (1936) theorem in this setting:

- (i) In general, the parameter space \mathbb{R}^p defined in (4) is different from the observation space \mathbb{N}^p , as opposed to the special case of Gaussian PCA.
- (ii) Since principal components are not reconstructed incrementally, the corresponding subspaces need not be nested.
- (iii) The lack of constraints on \mathbf{B} means that *raw* scores may be correlated in the latent space, unlike their counterparts in standard PCA.

To address point (i), we provide representations in the parameter space as it has the Euclidean geometry with which practitioners are most familiar. Point (ii) is an inherent consequence of nonlinearity that has some consequences in terms of interpretation. Indeed, the ‘‘axis of maximum variance’’ of model with rank q is not the same as the first axis of model with rank $q + 1$. As for point (iii), we use an orthonormal coordinate system to represent samples in the Euclidean parameters rather than the ‘‘raw’’ results of the algorithms. The samples positions \mathbf{Z} can be

estimated with $\tilde{\mathbf{Z}} := \mathbf{O} + \mathbf{X}\hat{\Theta}^\top + \tilde{\mathbf{M}}\hat{\mathbf{B}}^\top$. $\tilde{\mathbf{Z}}$ is useful to assess goodness of fit and quality of the dimension reduction whereas $\tilde{\mathbf{P}} = \tilde{\mathbf{M}}\hat{\mathbf{B}}^\top$ is used to visualize and explore structure not already captured by the covariates.

6.2. *Quality of the dimension reduction.* A first important criterion in PCA is the amount of information that is preserved by the q -dimensional reduction. To this aim, we define a pseudo R^2 criterion, which compares the model at play to both a *null* model with no latent variables and a *saturated* model with one parameter per observation.

Formally, we define the matrix $\mathbf{\Lambda}^{(q)} = [\lambda_{ij}^{(q)}]$ where entry $\lambda_{ij}^{(q)} := \tilde{Z}_{ij}$ serves as an estimate of the canonical parameter of the distribution of Y_{ij} given in (3). We can thus define the log-likelihood ℓ_q of the observed data with

$$\ell_q = \sum_{i=1}^n \sum_{j=1}^p [Y_{ij}\lambda_{ij}^{(q)} - \exp(\lambda_{ij}^{(q)})] - K(\mathbf{Y}).$$

We can compare it to the log-likelihood of the saturated model ℓ_{\max} [replacing $\lambda_{ij}^{(q)}$ with $\lambda_{ij}^{\max} := \log(Y_{ij})$] and the log-likelihood ℓ_{\min} of the null model chosen here as a Poisson regression GLM with no latent structure (replacing λ_{ij} with $\lambda_{ij}^{\min} := o_{ij} + \hat{\Theta}\mathbf{X}_i$, where $\hat{\Theta}$ is estimated using a standard GLM). The resulting pseudo R^2 is defined as

$$(15) \quad R_q^2 = (\ell_q - \ell_{\min}) / (\ell_{\max} - \ell_{\min}).$$

This R^2 is a bit imperfect as it assumes Poisson counts, unlike the Poisson-lognormal in our model, but it is necessary to compute equivalents to *percentage of variance explained* to which practitioners have grown accustomed.

6.3. *Visualizing the latent structure.* The matrix $\tilde{\mathbf{P}} = \tilde{\mathbf{M}}\hat{\mathbf{B}}^\top$ encodes positions of the samples in the latent space using $\hat{\mathbf{B}}$ as basis and $\tilde{\mathbf{M}}$ as principal components. Since $\hat{\mathbf{B}}$ is not constrained whatsoever, the *raw* components are neither orthogonal nor sorted in decreasing order of variation. We therefore decompose $\tilde{\mathbf{P}}$ as $\tilde{\mathbf{P}} = \tilde{\mathbf{M}}_{\text{viz}}\hat{\mathbf{B}}_{\text{viz}}^\top$ with columns of $\hat{\mathbf{B}}_{\text{viz}}$ orthogonal and columns of $\tilde{\mathbf{M}}_{\text{viz}}$ sorted in decreasing order of variation and use $\tilde{\mathbf{M}}_{\text{viz}}$ as principal components for visualization purposes. Since $\tilde{\mathbf{P}}$ is already of low-rank q , this is achieved simply by doing a standard PCA of $\tilde{\mathbf{P}}$. Note also that using either $(\hat{\mathbf{B}}_{\text{viz}}, \tilde{\mathbf{M}}_{\text{viz}})$ or $(\hat{\mathbf{B}}, \tilde{\mathbf{M}})$ leaves $\tilde{\mathbf{\Sigma}}$ unchanged.

We then decompose the total variance along each component j as in standard PCA. The overall contribution of axis j is then $d_j \times R_q^2$, where d_j is the fraction of variance in the latent space explained by component j . Following the same line, we may plot the correlations between the columns of $\tilde{\mathbf{P}}$ and the components arising from its PCA to help with the interpretation of these components in terms of original variables.

7. Illustrations. In order to highlight the scalability of our variational approach for generalized pPCA and its flexibility for the inclusion of covariates, we analyze two microbiome sequence count datasets below. They consist in counts of microbial species (OTUs or operational taxonomic units) in a series of samples. Note that an intrinsic limitation of this sampling technology and of all marker-gene based metabarcoding analysis methods is that they do not give access to absolute cell counts in a sample. Indeed, these methods consist in sampling the DNA in the biological sample and only a fixed number of DNA fragments, referred to as the “sequencing depth,” is observed. Consequently, any multivariate analysis of such data aims to describe the dependencies between relative abundances [Tsilimigras and Fodor (2016), Gloor et al. (2017)], although some additional measures can be made to recover approximate absolute abundances [Smets et al. (2015), Vandeputte et al. (2017)].

Because of the different technical steps involved in library preparation, the sequencing depth is generally independent of the total cell count in the sample, and its variations across samples have no biological meaning. Therefore, the sequencing depth itself constitutes a nuisance parameter that we need to account for to avoid spurious correlations. To correct for varying depths across samples, we assume that average counts scale linearly with sequencing depth, although more sophisticated normalizations exist [Chen et al. (2018)]. In subsequent analysis the sequencing depth is just another covariate with a special status as we know its regression coefficient, and we therefore include it as an offset in the model. Offset is computed as the total sequencing depth before any filter is applied to OTUs. By doing so the observed counts within a sample are not linearly constrained to sum to sequencing depth.

7.1. *Impact of weaning on piglet microbiome.*

Description of the experiment. We considered the metagenomic dataset introduced in Mach et al. (2015). The dataset was obtained by sequencing the bacterial communities collected from the feces of 31 piglets at 5 points after birth ($n = 155$). The communities were sequenced using the hypervariable V3-V4 region of the 16S rRNA gene as metabarcoding marker gene and sequences were processed and clustered at the 97% identity level to form $p = 4031$ OTUs [see Mach et al. (2015) for details of bioinformatics preprocessing]. The dataset is thus a 155×4031 count table where entry (i, j) measures the relative abundance of OTU j in sample i as the number of sequences (originating from sample i) falling in sequence cluster j . One aim of this experiment is to understand the impact of weaning on gut microbiota. Weaning, and more generally diet changes, are well documented to strongly impact the gut microbiota and we therefore use weaning status as ground truth to check whether our method can recover known structure. We also use the example to test scalability and study how the method behaves when the number p of variables increases.




FIG. 1. Dataset from Mach et al. (2015). Running times averaged over 4 replicates of the PLNPCA function in R PLNmodels package. Single core Intel i7-4600U CPU 2.33 GHz, R 3.4.1, Linux Ubuntu 16.04.

Numerical experiments. To test the impact of the number of variables on the dimension of the latent subspace, we inferred q on nested subsets of the count table. We selected only the 3000, 2000, 1000, 500 and 100 most abundant OTUs and fitted a model with appropriate offset to each subset. The offsets were chosen as log-total read count of each sample, computed on the full OTU table. It reflects the fact that, *et ceteris paribus*, observed counts should be roughly twice as high in communities sequenced twice more. For context, the 2500 least abundant OTUs exhibit very high sparsity (less than 1% of nonnull counts): each has total abundance lower than 5 and more than half (1287) are seen only once. It is customary to remove such OTUs using abundance-based filters in microbiome studies. We expect them to behave like high-dimensional noise and strongly degrade structure recovery.

Figure 1 shows that running times increase sublinearly with q and linearly with p , as expected. Figure 2 additionally shows that low count OTUs act as high dimensional noise and hamper our ability to recover fine structure in the latent space, (the pseudo R^2 goes down from 95% and 68% and \hat{q} from 27 to 8) just like it would in high dimensional Gaussian PCA.

Impact of weaning. We focus on results obtained on the 500 most abundant OTUs, which account for 90.3% of the total counts. We emphasize that even doing so, the count table remains quite sparse, with 67% of null counts and 60% of positive counts lower than or equal to 5. The ICL criteria on this subset selects $\hat{q} = 25$ ($R^2 = 89.4\%$). The main structure present in the latent subspace is the strong and systematic impact of weaning (Figure 3, left), almost entirely captured




FIG. 2. Dataset from Mach et al. (2015). The minimum overall abundance of included OTUs (left panel), quality of approximation R_q^2 (central panel) and selected value \hat{q} (right panel) decreases when OTUs with low abundance are added to the dataset.

by Axis 1. The variable factor map highlights OTUs from two specific bacterial families: Lactobacillaceae (red) and Prevotellaceae (blue). The former are typically found in dairy products and thought to be transmitted to the piglets via breast milk. As expected, they are enriched in suckling piglets and negatively correlated with Axis 1. The latter produce enzymes that are essential to degrade cereals introduced in the diet after weaning. As reported in Mach et al. (2015), they are enriched after weaning and positively correlated with Axis 1. The method is thus




FIG. 3. Individual (left) and variable (right) maps corresponding to the first principal plane of the q -dimensional approximation. Weaning has a strong and systematic effect on gut microbiota composition, well captured by Axis 1. Bacterial families Prevotellaceae (red) and the Lactobacillaceae (blue) are two families well known to be affected by weaning and have a high correlation with Axis 1.

able to recover well known structures, cope with sparse count tables and account for varying sequencing depths.

7.2. Oak powdery mildew pathobiome.

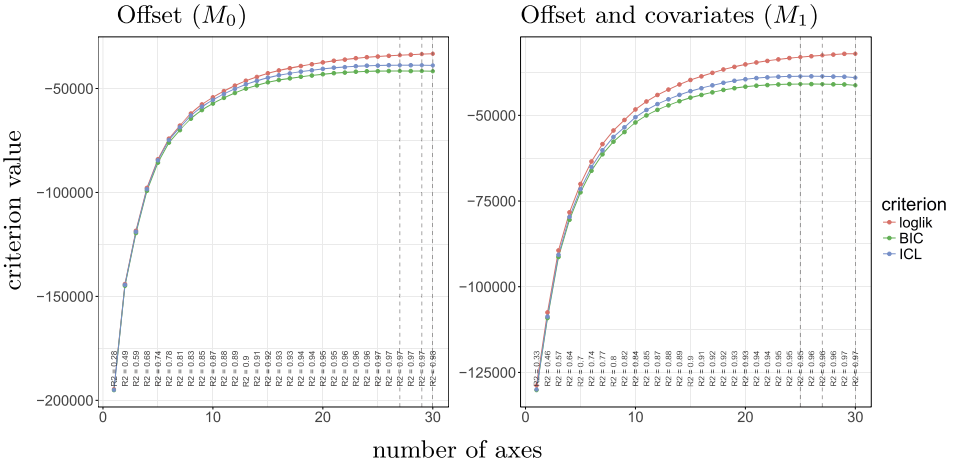
Description of the experiment. We considered the metagenomic dataset introduced in Jakuschkin et al. (2016). Similar to the Mach et al. (2015) dataset, it consists of microbial communities sampled on the surface of $n = 116$ oak leaves. Communities were sequenced with both the hypervariable V6 region of 16S rRNA as marker-gene for bacteria and the ITS1 as marker-gene for fungi. Sequences were cleaned, clustered at the 97% identity level to create OTUs and only the most abundant ones were kept [see Jakuschkin et al. (2016) for details of OTU picking and selection] resulting in a total of $p = 114$ OTUs (66 bacterial ones and 44 fungal ones). One aim of this experiment is to understand the association between the abundance of the fungal pathogenic species *E. alphitoides*, responsible for the oak powdery mildew, and the other species. Furthermore, the leaves were collected on three trees with different resistance levels to the pathogen. In addition to the sampling tree, several covariates, all thought to potentially structure the community, were measured for each leaf: orientation, distance to ground, distance to trunk, direction, etc.

We emphasize that our goal slightly differs from that of Jakuschkin et al. (2016) as these authors were interested in reconstructing the ecological network of the species, whereas our purpose is to summarize the species' dependency structure in low dimension. Our approach also differs from a methodological viewpoint as we jointly estimate the effect of the covariates Θ and the dependency structure Σ while they first corrected the observed counts for the effect of the covariates using a regression model before feeding the residuals from that regression to a network inference method. This two-steps procedure fails both to account for the fact that Θ is estimated and to propagate uncertainty from the first step to the second one.

Importance of the offset. The abundances Y_{ij} (where i denotes the leaf and j the OTUs) were measured separately for fungi and bacteria resulting in different sampling efforts for the two types of OTUs: the median total abundance were respectively 668 for bacteria and 2166 for fungi. To account for this, we define a different offset O_{ij} term for each OTU type. Offsets are still computed as the log-total sums of reads, including those of filtered out OTUs, for each OTU type.

Model selection. The three trees from which the leafs where collected were respectively susceptible, intermediately resistant (hereafter "intermediate") and resistant to mildew. We first fitted a null Poisson-lognormal model M_0 as defined in (9) with only an offset term. Alternatively, we considered model M_1 involving two covariates: the tree from which each leaf was collected from, and the orientation (0 = south-east, 1 = north-west) of its branch.

(a) Model selection



(b) Goodness of fit and Entropy

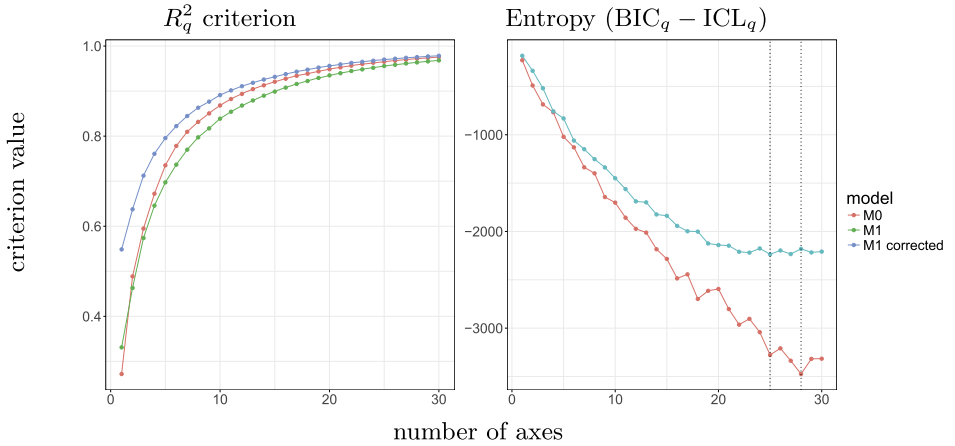
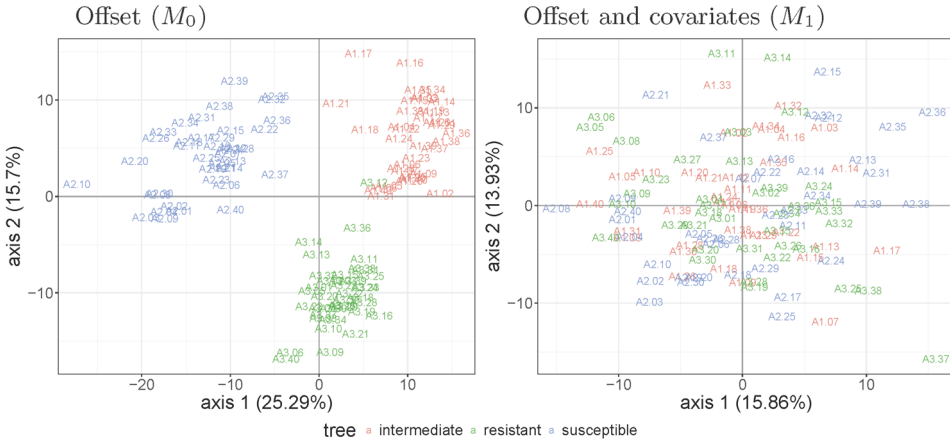


FIG. 4. Dataset from Jakuschkin *et al.* (2016). (a) Model selection criteria J_q , BIC_q and ICL_q for model M_0 (left) and M_1 (right); (b) R_q^2 criterion and entropy of $\tilde{p}(\mathbf{W})$.

Figure 4(a) displays the lower bound J , the BIC and the ICL for model M_0 (left) and M_1 (right) as a function of the number of axes q considered. We observe that the J_q is always increasing and that both BIC and ICL criteria behave similarly. According to the ICL criterion, we selected $\hat{q}_0 = 28$ ($ICL = -38,619$) latent dimensions for model M_0 and $\hat{q} = 25$ ($ICL = -38,472$) for model M_1 . This suggest that the two models (with their respective optimum) provide a very similar fit.

We looked at the approximate posterior entropy in panel left of Figure 4(b): we observed that it is minimal near to the respective optimum in terms of model

(a) Individual Factor Maps and tree status



(b) Individual Factor Maps and distance to ground

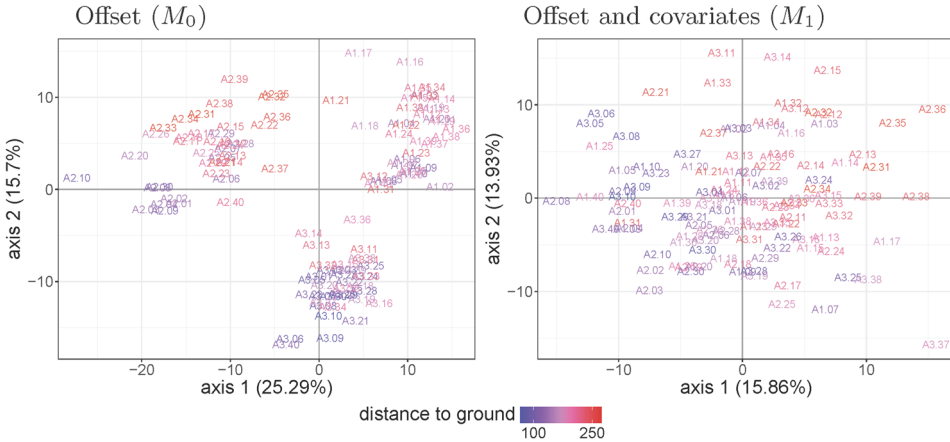


FIG. 5. Dataset from Jakuschkin et al. (2016). Scatter plot of the leaves on the first two principal components (left: M_0 , right: M_1) with colors corresponding to either tree status (a) or distance to ground (b). Accounting for tree status reveals an ecological gradient along distance to ground.

selection. This indicates that the selected dimensions are also optimal in terms of uncertainty on the latent variables.

Effect of the covariates. The choice between model M_0 and M_1 is mostly a matter of the type of dependency we analyze with each of them, as the former does not account for the covariates whereas the latter does. This is illustrated in Figure 5(a) (top) when plotting the first principal plane. In model M_0 (left), the leaves collected on each tree are clearly separated. As expected, taking the tree as a covariate (right) removes the tree effect from the principal plane.

Adding covariates in the model also allows us to explore second-order structuring effects that are masked by the strong first-order effect of the sampling tree. Figure 5(a) (bottom) thus shows that in addition to sampling tree, communities are structured by the distance of the leaf to the ground. The effect of covariates on the abundance of *E. alphitoides* were also consistent: the estimated parameters θ_{ij} associated with the intermediate and resistant trees were -3.94 and -7.05 , respectively, taking the susceptible tree as a reference.

We compared the respective estimates of Σ under M_0 (denoted $\widehat{\Sigma}_0$) and under M_1 ($\widehat{\Sigma}_1$) focusing on the correlations between *E. alphitoides* and the other OTUs. $\widehat{\Sigma}_0$ contains correlations between OTUs that are either due to marginal covariations between them or to the effect of the covariates, whereas the correlations in $\widehat{\Sigma}_1$ are corrected from the effect of covariates. We first observed a reduction of the variances (median = 0.175, mean = 0.303 in $\widehat{\Sigma}_0$; median = 0.087, mean = 0.176 in $\widehat{\Sigma}_1$), which proves the strong effect of covariates on the abundance of the different OTUs. We then ranked all species according to their correlation with the pathogene and found very different rankings M_0 and M_1 (Kendall's $\tau = 0.41$), showing that the covariates drastically change the apparent relationship between OTU abundances.

Percentage of variance. We now comment on use of the R_q^2 criterion defined in Section 6 to evaluate the proportion of variability captured by a model with q latent dimensions. R_q^2 compares the pseudo-likelihood ℓ_q^m obtained with q latent dimensions under model M_m ($m = 0, 1$) with the likelihoods ℓ_{\min}^m and ℓ_{\max}^m . We know that $\ell_{\max}^0 = \ell_{\max}^1$ whereas $\ell_{\min}^0 < \ell_{\min}^1$ because ℓ_{\min}^0 only relies on the offsets whereas ℓ_{\min}^1 accounts for both the offsets and the covariates. As a consequence, R_q^2 tends to be higher under M_0 than under M_1 for a given q . Right panel of Figure 4(b) compares the genuine R_q^2 under models M_0 and M_1 and the corrected version of R_q^2 under model M_1 using ℓ_{\min}^0 in place of ℓ_{\min}^1 . As expected, the corrected version of R_q^2 is always higher under M_1 than under M_0 . We also observe that, for both models, the proportion of variability captured by the latent space is quite high: $R_{28}^2 = 97.21\%$ for M_0 and $R_{25}^2 = 97.02\%$ for M_1 . We recall that $\hat{q}_0 = 28$ and $\hat{q}_1 = 25$ should both be compared with $p = 114$.

Variance of the variational conditional distribution. We recall that S_{ij} is the approximate conditional standard deviation of W_{ij} given the data. This parameter measures the precision of the location of individual i along the j th latent dimension. We can derive from them the approximate conditional variance of each Z_{ij} as $[\mathbf{B} \text{diag}(s_i \odot s_i) \mathbf{B}^T]_{jj}$. Figure 6 shows that this variance is much higher when the corresponding abundance Y_{ij} is low. Indeed, any large negative values of Z_{ij} yields a Poisson parameter close to zero and in turn a null Y_{ij} . As a consequence, large negative Z_{ij} cannot be predicted accurately. This is a natural consequence of the nonlinear nature of the exponential transform: large swaths of the parameter space are compressed to small regions of the observation space.

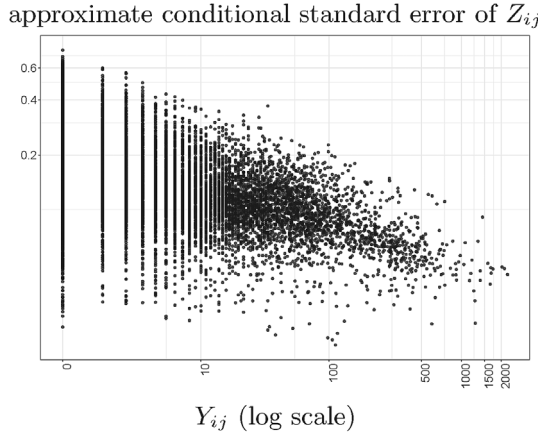


FIG. 6. Variational approximate conditional standard error of the Z_{ij} (y axis) as a function of the abundance Y_{ij} (x axis).

APPENDIX: CONVEXITY LEMMAS

LEMMA 1. For any vectors θ, x, m, s and b (with matching dimensions) and convex function f , if $u \sim \mathcal{N}(0, I)$ and $w = m + s \odot u \sim \mathcal{N}(m, \text{diag}(s \odot s))$, then the map $g : (\theta, m, s, b) \mapsto \mathbb{E}[f(\theta^T x + b^T w)]$ is convex in (θ, b) for (m, s) fixed and vice versa.

PROOF. Note $Z = \theta^T x + b^T w = (\theta^T x + b^T m) + b^T (s \odot u)$. The first-order derivative of g is

$$\nabla(\theta, b, m, s) = \mathbb{E}[f'(Z) [x \quad m + s \odot u \quad b \quad b \odot u]^T].$$

The second-order partial derivatives of g are:

$$\Psi_1(\theta, b) = \mathbb{E} \left[f''(Z) \begin{bmatrix} xx^T & x(m + u \odot s)^T \\ (m + s \odot u)x^T & (m + s \odot u)(m + s \odot u)^T \end{bmatrix} \right],$$

$$\Psi_2(m, s) = \mathbb{E} \left[f''(Z) \begin{bmatrix} bb^T & b(b \odot u)^T \\ (b \odot u)b^T & (b \odot u)(b \odot u)^T \end{bmatrix} \right].$$

And the associated quadratic form $\Phi_1(v, w) = (v, w)^T \Psi_1(\theta, b)(v, w)$ and $\Phi_2(v, w) = (v, w)^T \Psi_2(m, s)(v, w)$ can be simplified to

$$\Phi_1(v, w) = \mathbb{E}[f''(Z)(x^T v + (m + s \odot u)^T w)^2] \geq 0,$$

$$\Phi_2(v, w) = \mathbb{E}[f''(Z)(b^T v + (b \odot u)^T w)^2] \geq 0.$$

The Hessians Ψ_1 and Ψ_2 are thus semidefinite positive, which ends the proof. \square

LEMMA 2. For any matrices Θ , X , M , S and B (with matching dimensions) and convex function f , if $U = [U_1, \dots, U_n]^T$ where the U_i are i.i.d. and $U_i \sim \mathcal{N}(\mathbf{0}, I)$ and $W = M + S \odot U$. The map $g : (\Theta, M, S, B) \mapsto \mathbf{1}_n^T \mathbb{E}[f(X\Theta^T + WB^T)]\mathbf{1}_p$ is convex in (Θ, B) for (M, S) fixed and vice versa.

PROOF. The function g is a sum of functions of the form $g_{ij} : (\Theta, M, S, B) \mapsto \mathbb{E}[f(X_i^T \Theta_j + B_j^T(M_i + S_i \odot U))]$. The result follows from Lemma 1. \square

Acknowledgements. We thank Corinne Vacher and Nuria Mach for providing the data and discussing the results. Two anonymous reviewers gave insightful comments that helped improve the quality of this manuscript. We also thank Charlie Pauvert for his feedback on the code and Elsa Teulière for finding extra typos in the paper.

REFERENCES

- ACHARYA, A., GHOSH, J. and ZHOU, M. (2015). Nonparametric Bayesian factor analysis for dynamic count matrices. In *AISTATS*.
- AITCHISON, J. and HO, C.-H. (1989). The multivariate Poisson-log normal distribution. *Biometrika* **76** 643–653. [MR1041409](#)
- ANDERSON, T. W. (2003). *An Introduction to Multivariate Statistical Analysis*, 3rd ed. Wiley-Interscience, Hoboken, NJ. [MR1990662](#)
- BIERNACKI, C., CELEUX, G. and GOVAERT, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Trans. Pattern Anal. Mach. Intell.* **22** 719–25.
- CAO, Y. and XIE, Y. (2015). Poisson matrix completion. In *2015 IEEE International Symposium on Information Theory (ISIT)* 1841–1845. IEEE, New York.
- CHEN, J., KING, E., DEEK, R., WEI, Z., YU, Y., GRILL, D. and BALLMAN, K. (2018). An omnibus test for differential distribution analysis of microbiome sequencing data. *Bioinformatics* **34** 643–651.
- COLLINS, M., DASGUPTA, S. and SCHAPIRE, R. E. (2001). A generalization of principal components analysis to the exponential family. In *Advances in Neural Information Processing Systems* 617–624.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B* **39** 1–38. [MR0501537](#)
- ECKART, C. and YOUNG, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika* **1** 211–218.
- GLOOR, G. B., MACKLAIM, J. M., PAWLOWSKY-GLAHN, V. and EGOZCUE, J. J. (2017). Microbiome datasets are compositional: And this is not optional. *Front. Microbiol.* **8** 2224.
- HALL, P., ORMEROD, J. T. and WAND, M. P. (2011). Theory of Gaussian variational approximation for a Poisson mixed model. *Statist. Sinica* **21** 369–389. [MR2796867](#)
- IZSÁK, R. (2008). Maximum likelihood fitting of the Poisson lognormal distribution. *Environ. Ecol. Stat.* **15** 143–156. [MR2399077](#)
- JAAKKOLA, T. S. and JORDAN, M. I. (2000). Bayesian parameter estimation via variational methods. *Stat. Comput.* **10** 25–37.
- JAKUSCHKIN, B., FIEVET, V., SCHWALLER, L., FORT, T., ROBIN, C. and VACHER, C. (2016). Deciphering the pathobiome: Intra- and interkingdom interactions involving the pathogen *Erysiphe alphitoides*. *Microb. Ecol.* **72** 870–880.

- JOHNSON, S. G. (2011). The NLOpt nonlinear-optimization package. Available at <http://ab-initio.mit.edu/nlopt>.
- JOHNSON, N. L., KOTZ, S. and BALAKRISHNAN, N. (1997). *Discrete Multivariate Distributions*. Wiley, New York. [MR1429617](#)
- KARLIS, D. (2005). EM algorithm for mixed Poisson and other discrete distributions. *Astin Bull.* **35** 3–24. [MR2142681](#)
- LAFOND, J. (2015). Low rank matrix completion with exponential family noise. Preprint. Available at [arXiv:1502.06919](https://arxiv.org/abs/1502.06919).
- LANDGRAF, A. J. (2015). Generalized principal component analysis: Dimensionality reduction through the projection of natural parameters. Ph.D. thesis, Ohio State Univ., Columbus, OH. [MR3450322](#)
- LEE, D. D. and SEUNG, H. S. (2001). Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems* 556–562.
- LI, J. and TAO, D. (2010). Simple exponential family PCA. In *AISTATS* 453–460.
- LITTLE, R. J. A. and RUBIN, D. B. (2014). *Statistical Analysis with Missing Data*, Wiley-Interscience, Hoboken, NJ. [MR1925014](#)
- LIU, L. T., DOBRIBAN, E. and SINGER, A. (2016). ePCA: High dimensional exponential family PCA. Preprint. Available at [arXiv:1611.05550](https://arxiv.org/abs/1611.05550).
- MACH, N., BERRI, M., ESTELLÉ, J., LEVENEZ, F., LEMONNIER, G., DENIS, C., LEPLAT, J.-J., CHEVALEYRE, C., BILLON, Y., DORÉ, J., ROGEL-GAILLARD, C. and LEPAGE, P. (2015). Early-life establishment of the swine gut microbiome and impact on host phenotypes. *Environ. Microbiol. Rep.* **7** 554–569.
- MARDIA, K. V., KENT, J. T. and BIBBY, J. M. (1979). *Multivariate Analysis*. Academic Press, New York. [MR0560319](#)
- MINKA, T. P. (2000). Automatic choice of dimensionality for PCA. In *NIPS* **13** 598–604.
- MOHAMED, S., GHAHRAMANI, Z. and HELLER, K. A. (2009). Bayesian exponential family PCA. In *Advances in Neural Information Processing Systems* 1089–1096.
- NELSON, J. F. (1985). Multivariate gamma-Poisson models. *J. Amer. Statist. Assoc.* **80** 828–834.
- PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T. and FLANNERY, B. P. (1989). *Numerical Recipes: The Art of Scientific Computing. Code CD-ROM v 2.06 with UNIX Single-Screen License*, 3rd ed. Cambridge Univ. Press, Cambridge. [MR1414682](#)
- R DEVELOPMENT CORE TEAM (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. Available at <http://www.R-project.org>.
- ROYLE, J. A. and WIKLE, C. K. (2005). Efficient statistical mapping of avian count data. *Environ. Ecol. Stat.* **12** 225–243. [MR2144403](#)
- SALMON, J., HARMANY, Z., DELEDALLE, C.-A. and WILLETT, R. (2014). Poisson noise reduction with non-local PCA. *J. Math. Imaging Vision* **48** 279–294. [MR3152105](#)
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Ann. Statist.* **6** 461–464. [MR0468014](#)
- SMETS, W., LEFF, J. W., BRADFORD, M. A., MCCULLEY, R. L., LEBEER, S. and FIERER, N. (2015). A method for simultaneous measurement of soil bacterial abundances and community composition via 16s rRNA gene sequencing. *PeerJ PrePrints* **3** e1318v1.
- SRIVASTAVA, S. and CHEN, L. (2010). A two-parameter generalized Poisson model to improve the analysis of RNA-seq data. *Nucleic Acids Res.* **38** e170–e170.
- SVANBERG, K. (2002). A class of globally convergent optimization methods based on conservative convex separable approximations. *SIAM J. Optim.* **12** 555–573. [MR1885575](#)
- TIPPING, M. E. and BISHOP, C. M. (1999). Probabilistic principal component analysis. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **61** 611–622. [MR1707864](#)
- TSILIMIGRAS, M. C. B. and FODOR, A. A. (2016). Compositional data analysis of the microbiome: Fundamentals, tools, and challenges. *Ann. Epidemiol.* **26** 330–335.

- VANDEPUTTE, D., KATHAGEN, G., D'HOE, K., VIEIRA-SILVA, S., VALLES-COLOMER, M., SABINO, J., WANG, J., TITO, R. Y., DE COMMER, L., DARZI, Y. et al. (2017). Quantitative microbiome profiling links gut community variation to microbial load. *Nature* **551** 507–511.
- WAINWRIGHT, M. J. and JORDAN, M. I. (2008). Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.* **1** 1–305.
- WICKHAM, H. (2009). *Ggplot2: Elegant Graphics for Data Analysis*. Springer, New York. Available at <http://ggplot2.org>.
- WITTEN, D. M., TIBSHIRANI, R. and HASTIE, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics* **10** 515–534.
- YPMA, J. (2017). R interface to NLOpt, v. 1.0.4. Available at <https://github.com/jyypma/nloptr>.
- ZHOU, M. (2016). Nonparametric Bayesian negative binomial factor analysis. Preprint. Available at [arXiv:1604.07464](https://arxiv.org/abs/1604.07464).
- ZHOU, M., HANNAH, L., DUNSON, D. B. and CARIN, L. (2012). Beta-negative binomial process and Poisson factor analysis. In *AISTATS* **22** 1462–1471.

J. CHIQUET
S. ROBIN
UMR MIA-PARIS
AGROPARISTECH
INRA
UNIVERSITÉ PARIS-SACLAY
16, RUE CLAUDE BERNARD
75231 PARIS CEDEX 05
FRANCE
E-MAIL: julien.chiquet@inra.fr
robin@agroparistech.fr

M. MARIADASSOU
UNITÉ MAIAGE
INRA
UNIVERSITÉ PARIS-SACLAY
BÂT. 233 ET 210
DOMAINE DE VILVERT
78352 JOUY-EN-JOSAS CEDEX
FRANCE
E-MAIL: mahendra.mariadassou@inra.fr