

# Generalized Matrix Factorization: efficient algorithms for fitting generalized linear latent variable models to large data arrays

**Lukasz Kidziński**

LUKASZ.KIDZINSKI@STANFORD.EDU

*Department of Bioengineering  
Stanford University  
Stanford, CA 94305, USA*

**Francis K.C. Hui**

FRANCIS.HUI@ANU.EDU.AU

*Research School of Finance, Actuarial Studies and Statistics  
The Australian National University  
Canberra, ACT 2601, Australia*

**David I. Warton**

DAVID.WARTON@UNSW.EDU.AU

*School of Mathematics and Statistics  
and Evolution & Ecology Research Centre  
The University of New South Wales  
Sydney, NSW 2052, Australia*

**Trevor J. Hastie**

HASTIE@STANFORD.EDU

*Department of Statistics and Biomedical Data Science  
Stanford University  
Stanford, CA 94305, USA*

**Editor:** David Wipf

## Abstract

Unmeasured or latent variables are often the cause of correlations between multivariate measurements, which are studied in a variety of fields such as psychology, ecology, and medicine. For Gaussian measurements, there are classical tools such as factor analysis or principal component analysis with a well-established theory and fast algorithms. Generalized Linear Latent Variable models (GLLVMs) generalize such factor models to non-Gaussian responses. However, current algorithms for estimating model parameters in GLLVMs require intensive computation and do not scale to large data sets with thousands of observational units or responses. In this article, we propose a new approach for fitting GLLVMs to high-dimensional data sets, based on approximating the model using penalized quasi-likelihood and then using a Newton method and Fisher scoring to learn the model parameters. Computationally, our method is noticeably faster and more stable, enabling GLLVM fits to much larger matrices than previously possible. We apply our method on a data set of 48,000 observational units with over 2,000 observed species in each unit and find that most of the variability can be explained with a handful of factors. We publish an easy-to-use implementation of our proposed fitting algorithms.

**Keywords:** Generalized Linear Models, Generalized Linear Mixed-effect Models, Nuclear Norm, Penalized Quasi-Likelihood

## 1. Introduction

Latent factors underlying multivariate observations are of great interest in many applied disciplines. For example, in psychology or sociology, researchers measure multiple correlated test items to quantify certain constructs. They assume that responses can be described in terms of a small set of latent variables and that these latent variables can be interpreted as psychological or sociological traits (Bartholomew et al., 2011; Skrondal and Rabe-Hesketh, 2004; Spearman, 1961). In genetics, researchers measure gene expression in patients and expect that they correlate with intrinsic patient’s features, often not measurable directly (Stegle et al., 2012). In ecology, researchers observe sets of species in independent observational units (e.g., sites) and assume the existence of latent features associated with the abundance of species (e.g., representing a low-dimensional community composition space Warton et al., 2016; Ovaskainen et al., 2017).

Data from such experiments or observational studies can typically be expressed as a matrix of responses  $Y = [y_{ij}]$ , where rows  $i \in \{1, 2, \dots, n\}$  correspond to observational units (locations), subjects, etc. and columns  $j \in \{1, 2, \dots, m\}$  correspond to different responses such as species, genes, etc. Linear latent factors are then incorporated as a means of obtaining a low-rank approximation to the covariance structure of the rows in  $Y$ . Specifically, we assume that conditional on a set of latent variables  $u_i \in \mathbb{R}^p$ , with  $p \ll \min(m, n)$ , responses  $y_{ij}$  are independent observations. See Section 2 for more detailed notation and the model.

If all the responses are Gaussian in distribution, then we can find linear latent patterns with classical tools such as principal component analysis, singular value decomposition, or factor analytic models, and solutions can be computed very quickly even at scale. That is, many algorithms for fitting such large scale models have been suggested in the last two decades (Zou et al., 2006; Witten et al., 2009; Halko et al., 2011; Hirose and Yamamoto, 2015; Hirose and Imada, 2018). These algorithms are fast and scale well, since in this case there are known closed-form solutions which leverage from joint normality of the responses and latent variables. However, such theory and computation for Gaussian responses does not generalize easily to non-Gaussian cases.

Generalized Linear Latent Variable Models (GLLVMs, Moustaki and Knott, 2000; Niku et al., 2017) are a class of models which generalizes factor analysis to non-Gaussian responses. Specifically, they assume that responses follow distributions in the exponential family, where the mean for each response varies as a function of the observed covariates and the aforementioned set of latent features via a known link function. Model parameters are then usually estimated using Bayesian or maximum likelihood methods. Among a wide variety of Bayesian tools, practitioners use general purpose modeling software such as **Stan** (Carpenter et al., 2017), integrated nested Laplace approximations (INLA, Rue et al., 2009), as well as Bayesian Markov Chain Monte Carlo estimation with more specific software designed for GLLVMs (Tikhonov et al., 2020; Hui, 2020).

In this article, we focus on likelihood-based methods for GLLVMs. In particular, given the latent variable  $u_i$  for an observational unit  $i$ , all the responses  $y_{ij}$  are assumed to be conditionally independent of each other. Thus, the likelihood can be expressed as a product of  $m$  individual conditional likelihoods, after which we marginalize out the latent variables  $u_i$  to obtain the marginal likelihood function. The key problem in this approach is that the integral over  $u_i$  does not have a closed-form for non-Gaussian responses and must be

estimated or approximated by some means. To solve this problem, a number of methods have been proposed, including the Laplace method or some variation thereof (Huber et al., 2004; Bianconcini and Cagnone, 2012; Niku et al., 2017; Robin, 2019), numerical integration methods using adaptive quadrature (Rabe-Hesketh et al., 2002), and variational approximations (Hui et al., 2017b; Niku et al., 2019a).

While these approaches generally lead to very accurate solutions, they are computationally expensive for high-volume or high-dimensional problems and are difficult to parallelize, making them infeasible for solving large-scale problems. An alternative is to look for approximate solutions. One example of this comes from Pichler and Hartig (2021), who proposed estimating the multivariate joint likelihood using Monte Carlo sampling, along the lines of Hui et al. (2015), substantially improving computational performance thanks to the use of Graphical Processing Units (GPUs). To date their method has only been developed for binary responses. Huber et al. (2004) observed that if the latent scores are treated as fixed parameters, then estimates of them obtained through the Laplace method coincide with variables derived using Penalized Quasi-Likelihood estimation (PQL). This is a well-established method for estimating Generalized Linear Mixed Models (GLMMs; see Green, 1987; Breslow and Clayton, 1993; McGilchrist, 1994, for usages of PQL in the mixed model setting), and it is of interest to look at how effective PQL might be in the GLLVM context.

In this article, we propose a new approximate solution to the problem of fitting GLLVMs, based upon existing ideas for estimating parameters in GLMMs. In particular, we develop an alternating method which leverages the idea that a solution to a GLLVM can be approximated using Penalized Quasi-Likelihood estimation. This leads us to propose two algorithms: a direct Newton method with a simplified Hessian, and an Alternating Iteratively Reweighted Least Squares (AIRWLS) algorithm. As the name suggests, the AIRWLS algorithm uses classical Iteratively Reweighted Least Squares iterations by applying them alternately to rows and columns of  $Y$ .

Our contributions are threefold. First, we design an algorithm that is orders of magnitude faster than existing methods for fitting GLLVMs, but with similar accuracy if responses matrices are sufficiently large. Moreover, our approach can be further sped up by distributing it across many machines. Second, thanks to the much-improved computational performance of our method, we can find latent-variable decompositions of large matrices, enabling new directions of research. This is illustrated on an ecological data set on the coexistence of species across 48,000 observational units with over 2,000 responses in each unit. Third, we release easy-to-use software implementing the proposed approaches, which can be used as a drop-in replacement for other linear latent variable modeling approaches.

## 2. Model Formulation

We begin by providing a more precise mathematical formulation of the GLLVM. For observational unit  $1 \leq i \leq n$  and response  $1 \leq j \leq m$ , we assume

$$\begin{aligned} y_{ij} | \mu_{ij} &\sim \mathcal{F}(\mu_{ij}, \phi_j) \\ g(\mu_{ij}) &= \eta_{ij} = \beta_{0j} + x_i^\top \beta_j + u_i^\top \lambda_j, \end{aligned} \tag{1}$$

where  $x_i \in \mathbb{R}^d$  are a set of covariates observed for the  $i$ -th observational unit. The parameters  $\beta_{0j}$  are intercepts for each response, and  $\beta_j \in \mathbb{R}^d$  are response-specific regression coefficients corresponding to the covariates. Next, the vectors  $u_i \in \mathbb{R}^p$  denote the latent variables (also referred to as factor scores) for observational unit  $i$ , with  $\lambda_j \in \mathbb{R}^p$  the corresponding response-specific factor loadings. We assume that distribution  $\mathcal{F}(\mu_{ij}, \phi_j)$  is from the exponential family of distributions with mean  $\mu_{ij}$  and response-specific dispersion parameter  $\phi_j$ . The function  $g(\cdot)$  is a known link function, e.g., the logit link for binary responses and log link for Poisson responses. Finally, we use  $M^\top$  to denote the transposition of a matrix  $M$ .

In the GLLVM formulated above, we assume that: (A1)  $u_i \sim \mathcal{N}(0, I_p)$  where  $I_p$  is  $p \times p$  identity matrix, (A2)  $\Lambda = [\lambda_1, \lambda_2, \dots, \lambda_m]$  is a  $p \times m$  lower triangular with positive elements on the diagonal, and (A3) all observational units are independent and, conditioned on  $u_i$ , the responses  $y_{ij}$  are independent of each other. Critically, the second part of Assumption A3 implies that, conditional on  $u_i$ , the model as defined in Equation 1 specifies a Generalized Linear Model (GLM, McCullagh and Nelder, 1983) for each response  $j$  with mean  $\mu_{ij}$  and dispersion parameter  $\phi_j$ . Assumptions (A1) and (A2) are made in the context of GLLVMs to ensure parameter identifiability. In particular, without constraining  $\Lambda$  to be lower diagonal and its diagonal elements to be positive, we could rotate the vectors  $u_i$  and  $\lambda_j$  without changing the value of the linear predictor  $\eta_{ij}$  in Equation 1 (Huber et al., 2004). The independence of observational units as in Assumption (A3) is common in many formulations of GLLVMs (e.g., Hirose and Yamamoto, 2015; Hui et al., 2017b; Niku et al., 2017).

Throughout the remainder of the article, we use matrix notation whenever it is convenient and suitable. In particular, apart from  $\Lambda$  already defined, we use  $B = [\beta_1, \beta_2, \dots, \beta_m]$ ,  $X^\top = [x_1, x_2, \dots, x_n]$  and  $U^\top = [u_1, u_2, \dots, u_n]$  to define matrices corresponding to the regression coefficients, observed covariates, and factor loadings, respectively. We also denote the matrix of responses as  $Y = [y_{ij}]$  and the matrix of means  $M = [\mu_{ij}]$ . Finally, letting  $\Psi$  denote all the model parameters in the GLLVM, i.e. a vector formed from concatenating  $\beta_{0,i}$ ,  $\phi_j$ ,  $\beta_j$ ,  $\lambda_j$  for all  $i$  and  $j$ , then we let  $f(y_{ij}|u_i, \Psi)$  denote the probability density/mass function of  $y_{ij}$  given  $u_i$  and  $\Psi$ , corresponding to the distribution  $\mathcal{F}(\mu_{ij}, \phi_j)$  as defined in Equation 1.

To conclude this section, we note that the usual factor analytic model can be expressed in terms of Equation 1 if we set  $\mathcal{F}(\mu, \phi)$  to be a Gaussian distribution and have  $x_i$  represent only an intercept term. As the methods introduced in this paper for fitting in Equation 1 in the general form rely on analogous matrix factorization concepts, then in turn, we will refer to them as *generalized matrix factorization*.

### 3. Likelihood-based Estimation in GLLVM

Let  $f(y_i|u_i, \Psi)$  denote the conditional multivariate density function of the vector  $y_i = (y_{i1}, \dots, y_{im})^\top$  given the latent variables  $u_i$ . Then by the second part of Assumption (A3), we have  $f(y_i|u_i, \Psi) = \prod_j f(y_{ij}|u_i, \Psi)$ . We then integrate out the random latent variables

and write the marginal log-likelihood for the model defined in Equation 1 as:

$$\ell(\Psi) = \sum_{i=1}^n \ell_i(\Psi) = \sum_{i=1}^n \log \left( \int \prod_{j=1}^m f(y_{ij}|u_i, \Psi) \pi(u_i) du_i \right), \quad (2)$$

where  $\pi(u_i) \sim \mathcal{N}(0, I_p)$  and  $\ell_i$  is the marginal log-likelihood of the  $i$ -th observational unit alone.

Except for the special case where all the responses are Gaussian and the identity link function is used, the integral in Equation 2 can not be expressed in closed form, and has thus led to an extensive amount of research into overcoming this computational burden. For example, it can be directly computed using numerical integration methods, including Gauss-Hermite quadrature, adaptive quadrature, or Monte-Carlo integration. In brief, quadrature approaches aim at approximating the integral as a  $(2R - 1)$  polynomial by evaluating the function at  $R$  quadrature points. Adaptive quadrature shifts and scales locations of sampling points at each step to minimize the error (Rabe-Hesketh et al., 2002). Monte-Carlo methods sample the function within its domain and average the values. Importance sampling, a more sample-efficient method, was introduced for linear mixed models by Kuk (1999) and Skaug (2002) and can also be used here. While these methods can yield solutions close to exact, they are very expensive computationally and scale poorly with the number of latent variables.

Alternatively, we can approximate the value of the integral in Equation 2 using, say, a variational approximation or Taylor expansion (better known as the Laplace method), and maximize the approximated log-likelihood function instead. In the variational approximation, recently introduced in the context of GLLVMs by Hui et al. (2017b) and Niku et al. (2019a), Jensen’s inequality is applied to obtain a lower bound for the log-likelihood that is (closed to) fully closed form and thus computationally manageable. Maximizing this variational lower bound then gives an approximate solution to Equation 2. On the other hand, the Laplace method is based on taking a (typically second order) Taylor approximation of the log of the integrand in Equation 2 around its mode, and then integrating the approximated integrand (Huber et al., 2004; Bianconcini and Cagnone, 2012). Although estimates based on using the Laplace method may exhibit a non-negligible finite sample bias relative to aforementioned methods such as Monte-Carlo methods, they are consistent to order  $O(m^{-1})$  (Kass and Raftery, 1995; Vonesh, 1996).

While they are typically faster than their numerical integration counterparts, approximate likelihood methods such as the Laplace method still tend to be computationally prohibitive for problems where the response matrix consists of thousands of rows and/or columns, and may require days, weeks, or more time to converge (see Pichler and Hartig 2021 and the empirical study in Section 6.2 and Figure 5 for empirical demonstrations of this). Motivated by applications of GLLVMs and generalized matrix factorization to such high-dimensional settings, in this article, we leverage the structure of Equation 1 in two ways: (i) we show that by approximating the marginal log-likelihood using a Penalized Quasi-Likelihood approach, we can efficiently estimate all the required gradients and Hessians necessary for estimation of model parameters; (ii) thanks to this approximation, we demonstrate that the estimation problem can be decomposed to a set of smaller and relatively simple estimation problems based on individual rows and individual columns of

the response matrix, and thus treated in an alternating fashion. As a result, our proposed approach is extremely parallelizable, thus enabling further performance optimization.

#### 4. Penalized Quasi-Likelihood for GLLVMs

In the pursuit of an efficient algorithm for estimating the model parameters in Equation 1, we borrow ideas from Huber et al. (2004) who showed that, in the setting of GLLVMs, the maximum likelihood estimators of the latent variables derived from applying the Laplace method to the marginal likelihood in Equation 2 are equivalent to those based on maximizing a Penalized Quasi-Likelihood (PQL, Green, 1987; Breslow and Clayton, 1993). On the other hand, the estimates of  $\beta$  and  $u_i$  are not equivalent and are indeed more biased for the PQL approach relative to the Laplace method. However, it has nevertheless been proven that PQL produces asymptotically consistent estimates of the model parameters and random effects provided the size of each observational unit, which in this case corresponds to the number of responses, grows with the number of units (Nie, 2007; Hui et al., 2017a). This insight is particularly promising in the context of our work, since we focus on large-scale problems.

In what follows, we build on the derivation of PQL and Laplace methods proposed by Breslow and Clayton (1993) and Huber et al. (2004). Specifically, in order to speed up computation by orders of magnitude, we propose to omit the slowly-varying terms arising from applying the Laplace method to the GLLVM marginal log-likelihood function, and use a diagonal approximation in place of Hessians. Furthermore, by alternating the optimisation between columns and rows, our methods become amenable to parallel computing. We also investigate whether proposed approximations based on dropping slowly-varying terms lead to consistent estimation results. We show that the bias introduced by these modifications is small when the response matrices are sufficiently large. We introduce two algorithms for estimating model parameters using the PQL approach, and show substantial improvements in computational efficiency based on using either. It is important to point out that, while PQL has been developed and studied to a large degree for GLMMs, it has never been formally established and applied to GLLVMs. To our knowledge, this article is the first to do so.

##### 4.1 Derivation

We start by considering the log-likelihood  $\ell_i(\Psi)$  for the  $i$ -th observational unit. We follow standard GLM conventions popularized by McCullagh and Nelder (1983). Furthermore, for ease of notation, we assume distributions in the exponential family with some known cumulant function  $b(\cdot)$  and that the canonical link function is used. The developments below can be extended to the case of a non-canonical link function, at the expense of additional

tedious algebra. Let

$$\begin{aligned} \exp(\ell_i(\Psi)) &= \int \prod_j^m h(y_{ij}, \phi) \exp\left(\frac{y_{ij}\eta_{ij} - b(\eta_{ij})}{\phi_j}\right) \exp\left(-\frac{u_i^\top u_i}{2}\right) du_i \\ &\propto \int \exp\left(\sum_{j=1}^m \frac{y_{ij}\eta_{ij} - b(\eta_{ij})}{\phi_j} - \frac{u_i^\top u_i}{2}\right) du_i, \end{aligned} \quad (3)$$

where  $g(\mu_{ij}) = \eta_{ij}$  as defined in Equation 1. We can express the integral in Equation 3 in the form  $\int \exp(-L_i(u_i)) du_i$  in order to apply the Laplace method. Specifically, let  $\tilde{u}_i$  be the solution to  $\partial L_i(u)/\partial u_i = 0$ , i.e. the minimum of  $L_i(u)$ . Then the Laplace method yields

$$\ell_i(\Psi) \approx -\frac{1}{2} \log \left| \frac{\partial^2 L_i}{\partial u_i \partial u_i^\top}(\tilde{u}_i) \right| - L_i(\tilde{u}_i),$$

where  $|\cdot|$  denotes the determinant of a matrix.

In order to compute the derivatives of  $L_i$  with respect to the latent variables  $u_i$ , we first calculate the partial derivatives of the summands in  $L_i$ . Let  $v(\mu_{ij}) = 1/g'(\mu_{ij})$  be the variance function associated with the exponential family of distributions, e.g., for the Bernoulli distribution we have  $v(\mu) = \mu(1 - \mu)$ . Then, assuming the canonical link function is used, we have

$$\frac{\partial \eta_{ij}}{\partial u_i} = \frac{\mu_{ij}^\top(\eta_{ij})}{v(\mu_{ij})} = \lambda_j \quad \text{and} \quad \frac{\partial b(\eta_{ij})}{\partial u_i} = \frac{\mu_{ij}^\top(\eta_{ij})}{v(\mu_{ij})} = \lambda_j \mu_{ij}.$$

Therefore, we obtain

$$\frac{\partial L_i}{\partial u_i}(u_i) = -\sum_{j=1}^m \frac{\lambda_j (y_{ij} - \mu_{ij})}{\phi_j} + u_i \quad (4)$$

and

$$\frac{\partial^2 L_i}{\partial u_i \partial u_i^\top}(u_i) = \sum_{j=1}^m \frac{\lambda_j \lambda_j^\top v(\mu_{ij})}{\phi_j} + I_p. \quad (5)$$

Equation 5 can be rewritten as  $\Lambda W \Lambda^\top + I_p$ , where  $\Lambda = [\lambda_1, \dots, \lambda_m]$  is the  $p \times m$  matrix of factor loadings and  $W$  is a  $m \times m$  diagonal matrix with elements  $w_j = v(\mu_{ij})/\phi_j$  for  $j \in \{1, 2, \dots, m\}$ . Note that in the case of GLLVMs, the elements  $\{w_j\}_{j=1}^m$  correspond precisely to iterative weights coming from a GLM (McCullagh and Nelder, 1983, Chapter 2.5).

For the  $i$ -th observational unit then, we can write

$$\ell_i(\Psi) \approx -\frac{1}{2} \log |\Lambda W \Lambda^\top + I_p| + \sum_{j=1}^m \frac{1}{\phi_j} (y_{ij} \tilde{\eta}_{ij} - b(\tilde{\eta}_{ij})) - \frac{1}{2} \tilde{u}_i^\top \tilde{u}_i, \quad (6)$$

where  $\tilde{\eta}_{ij} = \beta_{0j} + x_i' \beta_j + \tilde{u}_i' \lambda_j$ . In the context of GLMMs, faced with a similar approximation to the log-likelihood of a cluster, Breslow and Clayton (1993) argued that since  $W$  varied

slowly as a function of the model parameters, then (for fixed  $\Lambda$ ) the first term in Equation 6 could be ignored. In the setting of GLLVMs however,  $\Lambda$  is not fixed. On the other hand, if we consider the normalized log-likelihood  $\sum_i \ell_i / (mn)$ , then we observe that for fixed  $m$  this first term is asymptotically negligible as  $n$  gets large. Moreover if both  $m$  and  $n$  are growing, then the first term in Equation 6 is asymptotically dominated by the second term; see also Demidenko (2013) and Hui et al. (2017a). This is also confirmed in our simulation study, reported in Section 7. Hence with applications of GLLVMs to high-dimensional data sets in mind, we chose also to (conveniently) ignore this term in our approximation so as to facilitate increases in computational efficiency. It is important to emphasize that our focus and thus methodological developments are driven by generalized matrix factorization for high-dimensional data sets where  $n$  and/or  $m$  are relatively large e.g., from the hundreds up to potentially tens of thousands. For small data sets, we recommend using more precise fitting methods such as numerical quadrature or (higher order) Laplace methods (Bianconcini and Cagnone, 2012).

After omitting the log-determinant term, we thus conclude that for large samples we can use the following approximation

$$\log \int \prod_{j=1}^m f(y_{ij}|u_i, \Psi) \pi(u_i) du_i \approx C + \sum_{j=1}^m \frac{1}{\phi_j} (y_{ij} \tilde{\eta}_{ij} - b(\tilde{\eta}_{ij})) - \frac{1}{2} \tilde{u}_i^\top \tilde{u}_i, \quad (7)$$

where  $\tilde{u}_i$  is a solution of Equation 4 and  $C$  is some constant as a function of the model parameters. Equation 7 has precisely the form of a PQL when applied to latent variable models. Moreover, this result implies that, provided  $m$  is sufficiently large, we can obtain a good approximation to the marginal log-likelihood function in Equation 2 as

$$L = - \sum_{i=1}^n L_i(\Psi) = - \sum_{i=1}^n \sum_{j=1}^m \frac{1}{\phi_j} (y_{ij} \tilde{\eta}_{ij} - b(\tilde{\eta}_{ij})) + \frac{1}{2} \sum_{i=1}^n u_i^\top u_i. \quad (8)$$

Importantly, we can solve the PQL as given by Equation 8 very efficiently using a Newton algorithm as we demonstrate in the following section. In particular, we discuss an approach inspired by iteratively reweighted least squares, where we alternate between the estimation of  $U$  and  $\Lambda$  (Section 4.2.1). Then, we introduce heuristics for estimating Hessians, which allows for a direct Newton algorithm on all the parameters, and substantially reduces the computations in each iteration of the Newton algorithm (Section 4.2.2).

## 4.2 Newton Algorithms

One approach to optimizing Equation 8 is via alternating minimization (Robin, 2019) with respect to  $U$  and  $\Lambda$ . For ease of notation, assume  $\phi_j = 1$  is known e.g., in the case of Poisson and Bernoulli distributed responses. In our proposed iterative algorithms, estimates of dispersion parameters can be updated after each iteration of the Newton algorithm, if required. Specifically for estimating  $\phi_j$  we can use a method of moments or maximum likelihood after each iteration of our proposed algorithms (Nelder and Wedderburn, 1972).



Assuming a canonical link and  $\phi_j = 1$ , Equations 4 and 8 imply that the derivative of  $L$  with respect to  $u_i$  is given by

$$\frac{\partial L}{\partial u_i} = \frac{\partial L_i}{\partial u_i} = - \sum_{j=1}^m (y_{ij} - \mu_{ij}) \lambda_j + u_i.$$

In the case of a non-canonical link a similar expression can be written involving additional weights. Likewise,

$$\frac{\partial^2 L}{\partial u_i \partial u_i^\top} = \sum_{j=1}^m v(\mu_{ij}) \lambda_j \lambda_j^\top + I_p, \quad (9)$$

and noting that

$$\frac{\partial^2 L}{\partial u_i \partial u_k^\top} = 0 \text{ for } k \neq i.$$

Hence the Hessian has the block diagonal form  $d^2L = \text{diag}(H_1, H_2, \dots, H_n)$  with

$$H_i = \partial^2 L / u_i \partial u_i^\top.$$

Similarly, we can straightforwardly calculate gradients and Hessians of the PQL in Equation 8 with respect to  $\lambda_j$  and  $\beta_j$  as follows

$$\begin{aligned} \frac{\partial L}{\partial \lambda_j} &= - \sum_{i=1}^n (y_{ij} - \mu_{ij}) u_i, \\ \frac{\partial L}{\partial \beta_j} &= - \sum_{i=1}^n (y_{ij} - \mu_{ij}) x_i \end{aligned}$$

and

$$\begin{aligned} \frac{\partial^2 L}{\partial \lambda_j \partial \lambda_j^\top} &= \sum_{i=1}^n v(\mu_{ij}) u_i u_i^\top \\ \frac{\partial^2 L}{\partial \beta_j \partial \beta_j^\top} &= \sum_{i=1}^n v(\mu_{ij}) x_i x_i^\top \\ \frac{\partial^2 L}{\partial \beta_j \partial \lambda_j^\top} &= \sum_{i=1}^n v(\mu_{ij}) x_i u_i^\top. \end{aligned}$$

Furthermore,

$$\frac{\partial^2 L}{\partial \lambda_i \partial \lambda_k^\top} = \frac{\partial^2 L}{\partial \beta_i \partial \beta_k^\top} = \frac{\partial^2 L}{\partial \beta_i \partial \lambda_k^\top} = 0 \text{ for } k \neq i.$$

Note that for fixed  $u_i$ , we can view  $(x_i^\top, u_i^\top)^\top$  as an enlarged covariate vector, with corresponding response-specific parameters  $(\beta_j^\top, \lambda_j^\top)^\top$ . In turn, the update step for each parameter  $\theta \in \{u_i, \lambda_j, \beta_j\}$  takes the form

$$\theta_{t+1} = \theta_t + s[-\nabla_\theta^2 L(\theta_t)]^{-1} \nabla_\theta L(\theta_t), \quad (10)$$

where  $\theta_t$  is the estimator of  $\theta$  in  $t$ -th iteration,  $\nabla_{\theta}L$  is the gradient of  $L$  with respect to  $\theta$ , and  $\nabla_{\theta}^2L(\theta)$  is the corresponding Hessian of  $L$  at  $\theta$ . The step size  $s > 0$  should be chosen carefully in order to guarantee convergence and speed up computation. In our implementations, we use Wolfe conditions (Wolfe, 1969) to address this issue, although we acknowledge that other methods for selecting the step size might further improve performance of our methods.

We now proceed to discuss two iterative algorithms for computing the update step in Equation 10 efficiently. The first approach uses Fisher scoring and produces an exact update, leveraging the fact that  $\nabla_{\theta}^2L(\theta)$  can be approximated (in fact, it is exact in the case where a canonical link is used) by the Fisher information matrix plus an identity matrix. The second approach uses only the diagonal of the Hessian  $\nabla_{\theta}^2L(\theta)$ , which is an approximate update that can be computed very quickly. In both approaches, after each update, we rotate the matrices  $U$  and  $\Lambda$  so as to satisfy the identifiability assumptions in Assumptions (A1) and (A2).

#### 4.2.1 FISHER SCORING AND ALTERNATING ITERATIVELY REWEIGHTED LEAST SQUARES

One approach for optimizing the PQL is to leverage the fact that, when we use the canonical link function, the negative Hessian matrices as defined in Section 4.2 are equivalent to the Fisher information (Nelder and Wedderburn, 1972) derived from the PQL in Equation 8. Specifically, conditional on  $(\beta_j, \lambda_j)$  for all  $j = 1, \dots, m$ , we can update each of the latent variables  $u_i$ ,  $i = 1, \dots, n$  in Equation 8 by solving  $n$  separate penalized GLMs (Green, 1987; Breslow and Clayton, 1993), where in each GLM the vector of  $m$  responses  $(y_{i1}, \dots, y_{im})^{\top}$  i.e., rows of  $Y$ , are treated as the ‘‘observations’’. Conversely, conditional on  $u_i$  for  $i = 1, \dots, n$ , we can update the response-specific coefficients and loadings  $(\beta_j, \lambda_j)$  for  $j = 1, \dots, m$  by solving  $m$  separate (unpenalized) GLMs, where the model matrix for each GLM is formed from the expanded predictor vectors  $(x_i^{\top}, u_i^{\top})^{\top}$  and the responses are given by the columns of  $Y$ .

Maximum likelihood for a single GLM is typically performed using the Newton algorithm implemented via iteratively reweighted least squares (IRWLS), or penalized IRWLS if a penalty is involved. Hence, minimizing the PQL in Equation 8 can be accomplished by applying alternating, parallel IRWLS algorithms applied to the rows and columns of the response matrix.

We now illustrate the derivation of the  $t$ -th IRWLS update step for  $u_i$  in more detail. Let  $u_i^{(t)}$  be the estimate of  $u_i$  in the  $t$ -th iteration,  $I(u_i^{(t)})$  be the (penalized) Fisher information matrix, and  $W_t$  be the iterative weight  $m$ -vector in the  $t$ -th iteration as defined above in Equation 6. Let  $\mu_i^{(t)} = (\mu_{i1}, \dots, \mu_{im})^{\top}$  be the  $i$ -th row of the matrix of means estimated in the  $t$ -th iteration, and similarly let  $Y_i = (y_{i1}, \dots, y_{im})^{\top}$  be the  $i$ -th row of the response matrix. We first derive the update step in Equation 10 with step size  $s = 1$ .

$$\begin{aligned}
 u_i^{(t+1)} &= u_i^{(t)} + [I(u_i^{(t)})]^{-1} \nabla L(u_i^{(t)}), \\
 &= u_i^{(t)} + [\Lambda W_t \Lambda^{\top} + I_p]^{-1} [\Lambda(Y_i - \mu_i^{(t)}) - u_i^{(t)}], \\
 &= [\Lambda W_t \Lambda^{\top} + I_p]^{-1} \Lambda W_t [\Lambda^{\top} u_i^{(t)} + W_t^{-1}(Y_i - \mu_i^{(t)})] \\
 &= [\Lambda W_t \Lambda^{\top} + I_p]^{-1} \Lambda W_t Z_t,
 \end{aligned} \tag{11}$$

where  $Z_t = \Lambda^\top u_i^{(t)} + W_t^{-1}(Y_i - \mu_i^{(t)})$  is a working response vector. Hence we obtain  $u_i^{(t+1)}$  by a ridge regression of  $Z_t$  on  $\Lambda^\top$  with weights  $W_t$ . When the step size  $s \neq 1$ , our update is instead  $u_i^{(t)} + s(u_i^{(t+1)} - u_i^{(t)})$ . Note that although the  $\beta_j$ 's (which are assumed to be fixed in the update) do not appear explicitly in these equations, the  $j$ th element of  $\mu_i^{(t)}$  includes  $x_i^\top \beta_j$  as part of its linear predictor i.e., an *offset* in GLM parlance.

We use a similar procedure for updating  $\lambda_j$  and  $\beta_j$ , conditional on the  $u_i$ 's. Specifically, the problem of estimating the response-specific coefficients and loadings in

$$g(\mu_{\cdot,j}) = X\beta_j + U\lambda_j,$$

where  $\mu_{\cdot,j} = (\mu_{1j}, \dots, \mu_{nj})^\top$  and  $U^\top = [u_1, \dots, u_n]$  can be rewritten as

$$g(\mu_{\cdot,j}) = (X, U)\gamma,$$

where  $(\cdot, \cdot)$  stands for horizontal concatenation of matrices and  $\gamma = (\beta_j^\top, \lambda_j^\top)^\top$ . Again we can solve it using IRWLS, although this time without the penalty term and with  $(X, U)$  in place of  $\Lambda$ . At the end of each iteration, we rotate  $U$  and  $\Lambda$  to fulfill the identifiability requirements given as part of Assumptions (A1)-(A2). For the full summary of our alternating two-step procedure, we refer to Algorithm 1. Importantly, note that for each  $i \in \{1, \dots, n\}$  in the PQL in Equation 8, the optimization problem can be decoupled and solved independently, allowing for the use of parallel computation to update the latent variables. Similarly, for each  $j \in \{1, \dots, m\}$  in Equation 8, the optimization problem can be decoupled and the updates of the response-specific coefficients and loadings can be performed in parallel.

---

**Algorithm 1:** ALTERNATING ITERATIVELY REWEIGHTED LEAST SQUARES

---

1. Initialize  $U, B, \Lambda$  randomly, where  $B = [\beta_1, \beta_2, \dots, \beta_m]$ .
  2. Repeat until the convergence condition:
    - (a) Perform one step of penalized IRWLS to regress rows of  $Y$  on  $\Lambda$ . Store regression coefficients as  $U$ .
    - (b) Perform one step of (unpenalized) IRWLS to regress columns of  $Y$  on  $(X, U)$ . Store regression parameters as  $(B^\top, \Lambda^\top)^\top$ .
    - (c) Transform data to comply with Assumptions (A1)-(A2):
      - i. Find a rotation  $\Theta$  such that  $\text{Cov}(U\Theta) = I_p$ , using, for example, principal component analysis,
      - ii. Compute  $U_0 = U\Theta$  and  $\Lambda_0 = \Theta^{-1}\Lambda$ ,
      - iii. Find a QR decomposition of  $\Lambda_0^\top = QR$ , such that  $R$  has positive elements on the diagonal,
      - iv. Return  $R^\top$  and  $U_0Q$  as new estimates of  $\Lambda$  and  $U$  respectively.
- 

As the convergence criterion in Algorithm 1, in our implementation we use the change in log-likelihood relative to the new log-likelihood, i.e., we stop when  $|L_{k-1} - L_k|/|L_k| < \varepsilon$ , where  $L_k$  is the log-likelihood in  $k$ -th iteration and  $\varepsilon$  is a sufficiently small value.

#### 4.2.2 QUASI-NEWTON WITH DIAGONAL HESSIANS

Instead of using the Fisher information, we can, of course, derive the updated step directly from Equation 10 by computing and inverting the Hessian explicitly. For each of the parameters  $u$ ,  $\lambda$ , and  $\beta$ , the corresponding Hessian matrices are block diagonal in structure. On the other hand, computing a Newton step requires inverting all blocks, which would be computationally expensive for large  $n$  since the Hessians are  $np \times np$  or  $mp \times mp$  matrices. To speed up computation then, we propose an alternative quasi-Newton method where we only use the diagonals of blocks in Hessians. These diagonals and their inverses in Equation 10 can be computed quickly, and indeed our empirical study later on shows that this approximation reduces computation time, despite increasing the number of steps required for convergence.

Note that the diagonal elements of the Hessian (Equation 9) can be computed by taking

$$\begin{aligned} \text{diag} \left( \frac{\partial^2 L}{\partial u_i \partial u_i^\top} \right) &= \text{diag} \left( \sum_{j=1}^m v(\mu_{ij}) \lambda_j \lambda_j^\top + I_p \right) \\ &= (\Lambda \circ \Lambda) v(\mu_{i,\cdot})^\top + \mathbf{1}_p, \end{aligned} \quad (12)$$

where  $\circ$  denotes the Hadamard product,  $\text{diag}(\cdot)$  denotes the diagonal of the given matrix,  $v(\mu_{i,\cdot}) = (v(\mu_{i1}), \dots, v(\mu_{im}))^\top$ , and  $\mathbf{1}_p = [1, 1, \dots, 1]^\top$ . The diagonal elements of Hessians of  $\lambda_j$  and  $\beta_j$  are derived analogously. Specifically, we obtain

$$\begin{aligned} \text{diag} \left( \frac{\partial^2 L}{\partial \beta_j \partial \beta_j^\top} \right) &= (X^\top \circ X^\top) v(\mu_{\cdot,j}), \\ \text{diag} \left( \frac{\partial^2 L}{\partial \lambda_j \partial \lambda_j^\top} \right) &= (U^\top \circ U^\top) v(\mu_{\cdot,j}), \end{aligned} \quad (13)$$

where  $v(\mu_{\cdot,j}) = (v(\mu_{1j}), \dots, v(\mu_{nj}))^\top$ . The full quasi-Newton approach thus follows the same steps as in Algorithm 1, except for steps 2(a) and 2(b) where we replace the AIRWLS update with an explicit implementation of Equation 10 with Hessians given by Equations 12 and 13.

### 4.3 Regularized Generalized Matrix Factorization

The above fitting algorithms for estimating GLLVMs and performing generalized matrix factorization has assumed that the number of latent variables,  $p$ , is known. In practice, however, we do not know the dimension of the latent space and we may need to estimate it from the data also. Methods for selecting the dimension include: cross-validation; information criteria (Bai and Ng, 2002; Hirose and Yamamoto, 2015); a somewhat arbitrary choice for the threshold of the variance explained (Smith et al., 2015); or a sparsity-inducing penalty added to the log-likelihood (Hui et al., 2018). We propose another smooth shrinkage parameter, motivated by regularized matrix factorization (Zou et al., 2006).

Instead of controlling the rank by explicitly choosing the number of latent variables, we can set a large upper bound on the number of latent variables (e.g.  $\sqrt{m}$ ) and then regularize the latent variables with an extra term  $\frac{1}{2} \|\Lambda\|_2^2$  added to the PQL criterion. If we control

penalties with a scaling parameter  $\gamma$ , then this leads to the regularized objective function

$$L_2(\Psi) = \sum_{i=1}^n \sum_{j=1}^m (y_{ij} \hat{\eta}_{ij} - b(\hat{\eta}_{ij})) + \frac{\gamma}{2} \|U\|_2^2 + \frac{\gamma}{2} \|\Lambda\|_2^2, \quad (14)$$

where  $\|\cdot\|_2$  is the Frobenius norm. Note that since the variance of  $u_i$  was arbitrarily set to  $I$  in (A1), the two penalties on the norms of  $U, V$  can be controlled with a single parameter  $\gamma$  without loss of generality. Srebro et al. (2005) show that solving Equation 14 with  $U$  and  $\Lambda$  of sufficiently high rank is equivalent to solving

$$L_*(\Psi) = \sum_{i=1}^n \sum_{j=1}^m (y_{ij} \hat{\eta}_{ij} - b(\hat{\eta}_{ij})) + \gamma \|M\|_*, \quad (15)$$

where  $\|\cdot\|_*$  denotes the nuclear norm and  $M = U\Lambda^\top$ . Equation 15 can be interpreted as a relaxed version of a rank constraint on  $M$  (Section 7.3.4, Hastie et al., 2019). This relaxation makes the problem convex. To tune dimensionality, we can control the penalty parameter  $\gamma$  in Equation 14. In particular, for sufficiently large  $\gamma$  some singular values of  $M$  vanish, effectively reducing dimensionality of the latent space (See Figure 7 for an illustrative example).

We illustrate empirical properties of this method of regularized generalized matrix factorization as part of our simulation study in Section 7.3.

#### 4.4 Missing Values

In certain practical situations, some elements of the response matrix  $Y$  may not be observed and have to be estimated. Such estimations are relatively straightforward using our proposed algorithms and involve little additional computational expense.

Based on the AIRWLS algorithm introduced in Section 4.2.1, in each step we can obtain estimates without imputating missing values as long as there are enough observations in a row or a column use for regression. Let  $\bar{u}_i^{t+1}$  denote an estimate of  $u_i^{t+1}$  using only the observed elements of the response vector  $Y_i$ , which we denote as  $\bar{Y}_i$ . Then we can rewrite the iteration step in Equation 11 as

$$\bar{u}_i^{t+1} = [\Lambda \bar{W}_t \bar{\Lambda}^\top + I_{\bar{p}}]^{-1} \bar{\Lambda} \bar{W}_t \bar{Z}_t,$$

where  $\bar{p}$  is the number of observed entries in  $Y_i$ ,  $\bar{Z}_t = \bar{\Lambda}^\top \bar{u}_i^{(t)} + \bar{W}_t^{-1} (\bar{Y}_i - \bar{\mu}_i^{(t)})$ , and  $\bar{W}_t, \bar{\Lambda}$  are matrices with truncated rows and columns corresponding to unobserved values in  $Y_i$ . This approach can be used for sparsely observed multi-response data.

On the other hand, the same approach can not be applied in the quasi-Newton method proposed in Section 4.2, since this method requires fully observed response matrix. However, following ideas from the SOFT-IMPUTE algorithm (Mazumder et al., 2010), in each iteration we can instead use predictions from the previous step to *impute* missing values and then compute gradients. Suppose a coordinate  $y_{i,j}$  is not observed. Then given estimates  $\beta_{0j}^{(t)}, \beta_j^{(t)}, \lambda_j^{(t)}, u_i^{(t)}$  in the  $t$ -th iteration of the algorithm, we set

$$\hat{y}_{i,j}^{(t)} = g^{-1} \left( \beta_{0j}^{(t)} + x_i^\top \beta_j^{(t)} + (u_i^{(t)})^\top \lambda_j^{(t)} \right).$$

Afterwards, we can run an iteration of the quasi-Newton algorithm on the estimated matrix  $\hat{Y}^{(t)}$  with values  $\hat{y}_{i,j}^{(t)}$  imputed in lieu of unobserved coordinates of  $Y$ .

#### 4.5 Cross-Validation

Imputation of missing values, as described in Section 4.4 enables us to straightforwardly employ cross-validation for assessing the fit and choosing model tuning parameters. While cross-validation on held-out coordinates is theoretically possible for most of the algorithms developed for estimating GLLVMs, its computational cost is prohibitive for the majority of them given fitting times even for a single GLLVM are already comparably quite long. However, thanks to the computational scalability of our approach compared to these other methods, cross-validation becomes more feasible in practice.

Suppose we want to use cross-validation for tuning the parameter  $\gamma$  in the regularized generalized matrix factorization in Section 4.3. To assess overall out-of-sample goodness-of-fit we can randomly sample elements of the observed response matrix  $Y$ , hold them out, and compute the out-of sample deviance of the predictions (See Section 5 for details). Depending on the objective, multiple techniques for sampling can be used, including uniform sampling of matrix entries, sampling based on response values, or sampling based on values of predictors. We use this method both for evaluating and comparing GLLVM fits on real data sets (Section 6.2), and in simulations for choosing the optimal shrinkage parameter (Section 7.3).

#### 4.6 Uncertainty Estimation

Resampling and simulation methods can be used both to assess uncertainty of the estimated parameters, as well to as quantify prediction uncertainty. For instance, we can repeatedly simulate data from a fitted model (i.e., a form of parametric bootstrap), fit models to the simulated data, and use the empirical covariance matrix based on the bootstrapped model parameter estimates as an estimate of uncertainty. Inference such as Wald-based confidence intervals or hypothesis tests for one or more model parameters then follows from this.

For measuring variability of predictive performance, we can randomly sample elements of the observed response matrix, hold them out, and compute deviance on the predictions. Alternatively, instead of randomly sampling elements, we can hold out entire rows of the response matrix and compare predictions based only on the fixed effects (Warton, 2008; Wang et al., 2012). Finally, we can sample rows or columns with replacement, to estimate uncertainty of estimates of the fixed effects. Importantly, regardless of the precise resampling or simulation technique employed, such an approach to quantifying uncertainty is computationally feasible thanks to the substantially reduced computation time of the proposed approach to fitting GLLVMs, relative to classical methods.

### 5. Evaluation Metrics

A baseline method for comparing the proposed algorithms for fitting GLLVMs is the R package `gllvm` (Niku et al., 2019b), applied with its default settings. This is a state-of-the-art package which uses variational approximations (Hui et al., 2017b) in conjunction with automatic differentiation (Niku et al., 2019a) to perform maximum approximate likelihood estimation for GLLVMs in a computationally efficiency manner. Since other approaches to the estimation of parameters in GLLVM are comparable or worse in terms of speed, then we only use the `gllvm` package for comparison in our numerical studies below.

We use a series of metrics to compare the quality of the GLLVM fits using the estimation methods proposed in this article, relative to fits obtained using the `gllvm` package. This is important because although the above algorithms minimize a version of penalized deviance, in certain applications, practitioners might instead be interested in estimates of the fixed effect parameters, the accuracy of the generalized matrix factorization in constructing the latent space, or out-of-sample predictive performance of the model.

**Deviance for evaluating fit to responses.** Following Nelder and Wedderburn (1972), we define deviance of the GLLVM as

$$D(Y, \hat{M}) = 2 \sum_{i=1}^n \sum_{j=1}^m \left( \log f(y_{ij} | \hat{\mu}_{ij}, \hat{\phi}_j) - \log f(y_{ij} | \hat{\mu}_{0,ij}, \hat{\phi}_{0,j}) \right),$$

where the  $\hat{\mu}_{ij}$ 's and  $\hat{\phi}_j$ 's are the predicted values and estimated dispersion parameters from the fitted GLLVM, the  $\hat{\mu}_{0,ij}$ 's and  $\hat{\phi}_{0,j}$ 's are the predicted values and dispersion parameters based on the saturated GLLVM, and  $f(y_{ij} | \mu_{ij}, \phi_j) = f(y_{ij} | u_i, \Psi)$  is the probability density/mass function of the responses assumed in Equation 1.

The absolute value of deviance is usually hard to interpret. Therefore, we choose to calculate the ratio of deviances between either two fitted models or between the fitted model and the null model. For a null model denoted as  $M_{null}$ , we refer to  $1 - D(Y, \hat{M})/D(Y, M_{null})$  as a fraction of null deviance explained by  $\hat{M}$ . The closer the fraction is to one, the more deviance is explained by the fitted GLLVM.

**Procrustes error for evaluating fit of the latent space.** In certain applications we are interested in how accurate the generalized matrix factorization method recovers the latent space. Since vectors spanning latent spaces are not unique, then we use a metric that rotates them before comparison. In particular, we follow Niku et al. (2019a) and (for the purpose of this article) define the Procrustes error as

$$P(\Lambda_0, \hat{\Lambda}) = \min_{\Omega} \|\Lambda_0 - \Omega \hat{\Lambda}\|_F, \text{ subject to } \Omega^T \Omega = I,$$

where  $\|\cdot\|_F$  is the Frobenius norm,  $\Lambda_0$  denotes for the true value of the loading matrix,  $\hat{\Lambda}$  is the estimated loadings, and  $\Omega$  is a rotation matrix. Note that unlike the deviance metric defined above, the Procrustes error is only computable in simulation studies where we have access to the ground-truth factor loadings.

**Mean squared errors for evaluating fixed effect coefficients.** As a basic measure of the accuracy in estimating the fixed effects in the GLLVM, i.e., the  $\beta_j$ 's in Equation 1, we compute the mean squared error of estimated values with respect to true values as a normalized Frobenius norm,

$$MSE(B_0, \hat{B}) = \|B_0 - \hat{B}\|_2^2 / md,$$

where  $d \times m$  matrix  $B_0$  denotes the true value of the regression coefficient matrix, and  $\hat{B}$  denotes the estimated value. As with Procrustes error, this metric requires access to the true  $B_0$  and, as such, it is also only applicable in simulation studies.

## 6. Data Studies

We report two applications from Ecology, with two distributions of responses: Poisson (Section 6.1) and Bernoulli (Section 6.2). In the first study we analyze a  $41 \times 30$  matrix of ant species, while in the second one we analyze a much larger  $48,331 \times 2,211$  matrix of plant species. We compare three methods for estimating GLLVMs: AIRWLS as described in Section 4.2.1, the quasi-Newton method as discussed in Section 4.2.2 and which we refer to it as *Newton*, and the variational approximations approach using the `gllvm` package (Niku et al., 2019b).

### 6.1 Study 1: Abundance of Ants

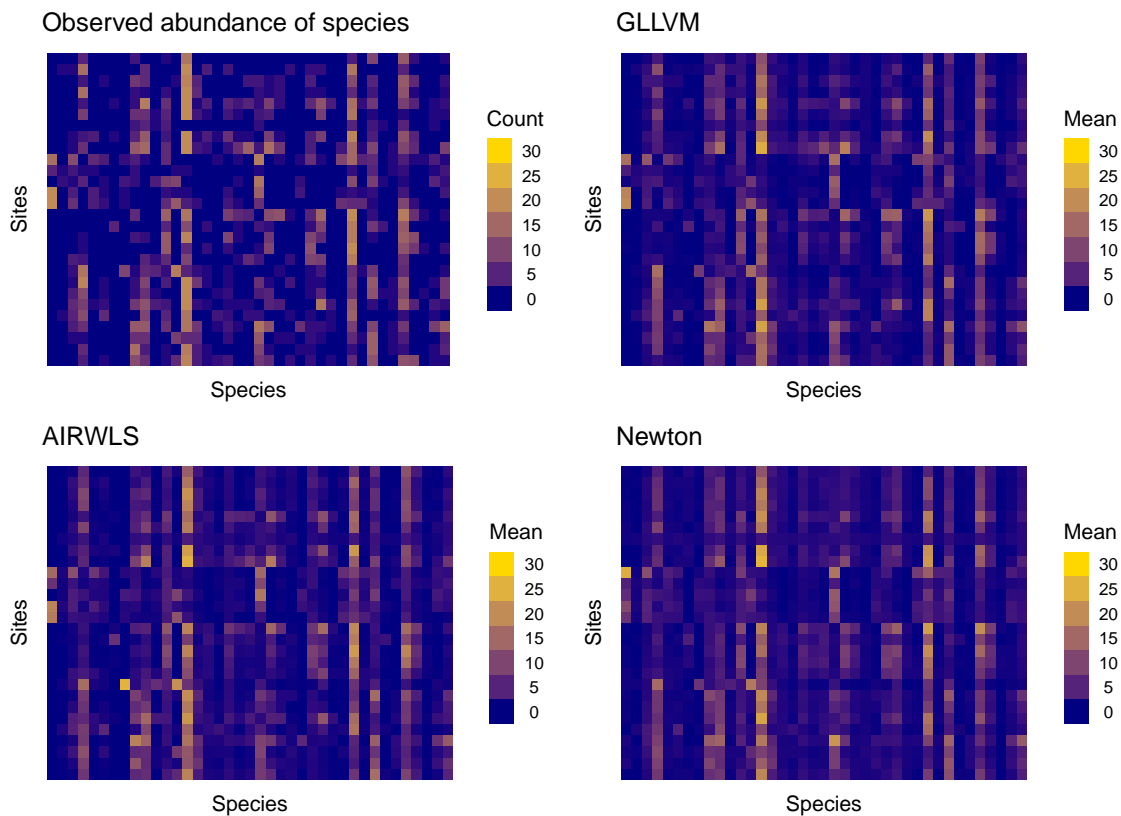


Figure 1: Qualitative validation of our algorithms on a data set of abundance of 41 ant species measured at 30 observational sites. In our experiment, we assumed responses are Poisson-distributed and we fitted GLLVMs using the `gllvm` package (top right), along with our proposed methods AIRWLS (bottom left) and Newton (bottom right). We found that all three methods capture qualitatively similar features of the observed response matrix (top left). This observation is confirmed in our quantitative analysis later on.



In order to validate our algorithms, we start by analyzing a small data set of 41 ant species measured at 30 study sites in March–April 2008 (Gibb and Cunningham, 2011). Together with the abundance of ants at these sites, researchers measured a number of environmental variables at each site (% cover of shrubs, bare ground, coarse woody debris, etc.). For a full report on the data acquisition methodology and the study design, we refer the reader to Gibb and Cunningham (2011). Here, we focus on illustrating how GLLVMs can be used for identifying intrinsic environmental factors that are not expressed in measured habitat structure variables. For ordination purposes, we set the number of latent variables in the GLLVM to  $p = 2$ . In Figure 1, we present the observed response matrix along with the predicted mean abundance for all species using the three methods of estimation.

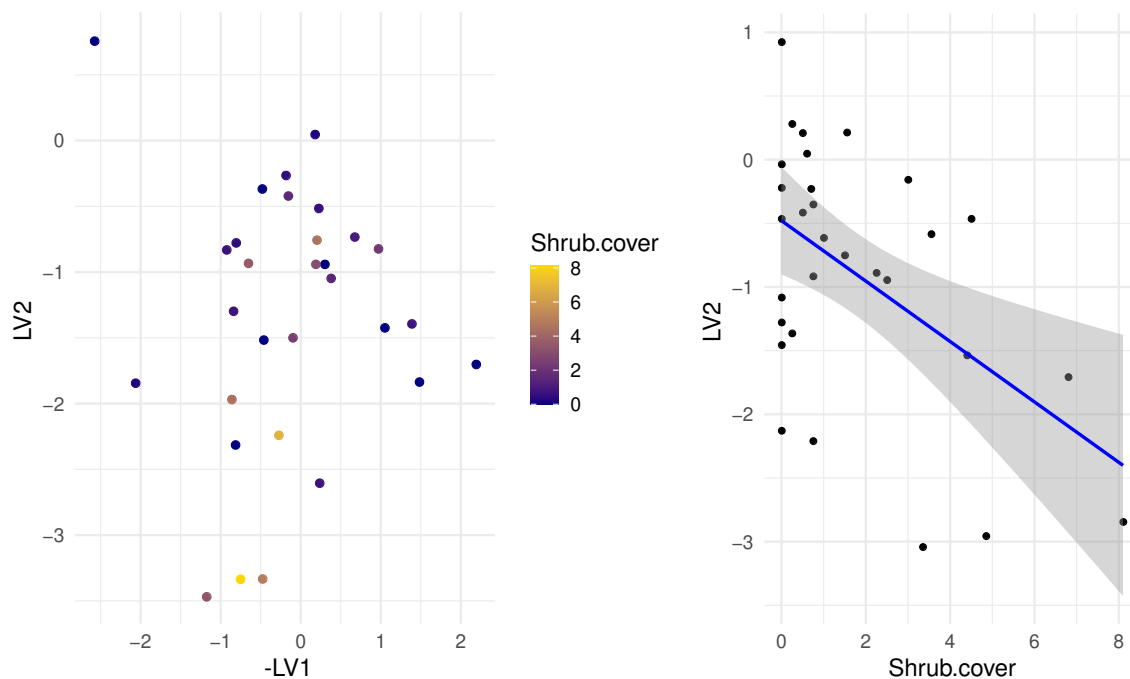


Figure 2: Validation of the generalized matrix factorization. In the data set of abundance of ants, we held out a known feature of observational sites (`Shrub.cover`) and compared it to predicted latent variables/factor scores. We used the AIRWLS algorithm for model fitting. An ordination plot showing the latent space spanned by the estimated factor scores, colored by their corresponding value of `Shrub.cover`, is on the left panel. In the right panel, we found that the second predicted latent variable correlates strongly with the held-out predictor.

Next, to further validate the method, we removed one of the environmental variables and examined if some of its variability was captured by a latent variable (which could be interpreted as a missing covariate). That is, we hypothesized that some of variability explained by the variable can be explained by the inclusion of a latent variable. Here, we arbitrarily chose to remove `Shrub.cover`. In Figure 2, we present an ordination plot of the

predicted latent variable scores, as well as a scatterplot of the second latent variable versus `Shrub.cover`. In the latter, the Pearson correlation coefficient between the two was  $-0.49$ , and it was statistically significant with  $p$ -value equal  $0.005$  as based on a two-sided  $t$ -test.

All three estimation approaches produced similar results in terms of deviance explained: `gllvm` 79%, AIRWLS 79%, and Newton 75%. Turning to computation time the variational approximation approach using the `gllvm` package fit converged in 2.3 seconds, while the quasi-Newton method took 0.2 seconds, and the AIRWLS algorithm converged in 0.5 seconds on 1 CPU and in 0.04 seconds when run in parallel across 8 CPUs.

## 6.2 Study 2: Large Scale Coexistence of Species

In this study, we analyzed data from systematic flora surveys along the east coast of New South Wales, using data obtained from the New South Wales Government (NSW Department of Planning Industry and Environment). Transects of the fixed area were exhaustively searched, and all plant species found in them were identified to species, where possible. We were interested in understanding co-occurrence patterns of different plants. A total of 48,737 transects were surveyed over the last two decades, and 4841 species have been recorded from each transect. We first filtered out columns and rows with less than 0.1% positive responses, leaving us with 48,331 observational units and 2,211 species. Each element in the response matrix is a binary response, reflecting the recorded presence-absence of a species at a specific location. Along with the species records, we also have 9 covariates describing the local environment.

First, to estimate the number of latent variables  $p$ , we fitted a model with  $p = 20$  on the full data set using the proposed quasi-Newton method in Section 4.2.2, and looked for a drop-off in size of the singular values from this fit. Based on the scree plot (Figure 3, left panel) we set  $p = 3$ .

Next we evaluated the predictive performance of estimated model (in particular, to assess the importance of including latent variables) by leaving out 1,000 elements in the response matrix and computing out-of-sample measures of predictive performance, as described in Section 5. Specifically, we held out 500 elements for which the model predicted 1 with at least 0.5 probability and 500 elements for which the model predicted 0 with at least 0.5 probability. We evaluated the fit by calculating the out-sample deviance and area under the receiver operator curve (AUC), and compared performance from only the fixed effect model i.e., with  $p = 0$ , and the full model with  $p = 3$  latent variables (Figure 3, right panel).

Due to the large scale of the problem, we were unable to use the `gllvm` package. Therefore, we tested only the Newton algorithm and it converged in three hours on commodity hardware. Our model with only fixed effects had  $AUC = 0.72$  and explained 39% of the out-of-sample deviance in the held-out data set. By contrast, the full model containing  $p = 3$  latent variables models had  $AUC = 0.87$  and explained 58% of the out-of-sample deviance (Figure 3, right panel).

Next, we examined how the size of the data set influenced the performance of the various estimation algorithms. We sub-sampled rows and columns from the response matrix and fitted models to the subsetted data. Specifically, we set  $\rho \in \{0.005, 0.01, 0.015, \dots, 0.065\}$  and sampled  $\lfloor \rho \cdot n \rfloor$  rows and  $\lfloor \rho \cdot m \rfloor$  columns, i.e. from 0.5% to 6.5% of the total number of rows and columns in  $Y$ . We did not use more than 6.5% of rows and columns because at

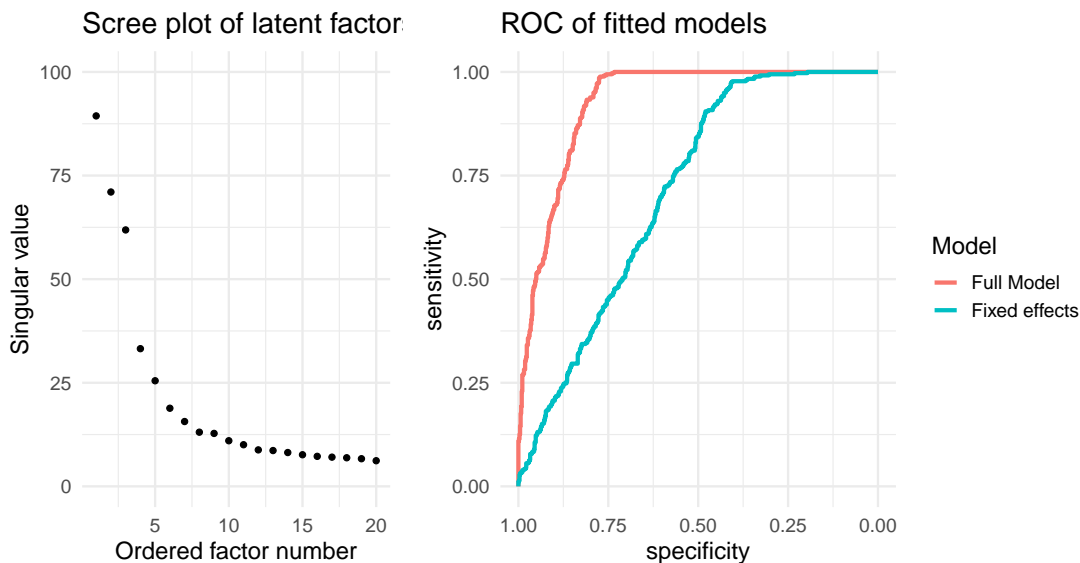


Figure 3: Model selection in the large-scale coexistence of species data set. We used the quasi-Newton method to fit the model. To choose the number of factors, we used a scree plot, i.e., we plotted the singular values of the latent decomposition (left), defined as the squared roots of the diagonal elements of  $\Lambda\Lambda^\top$ . The drop between the third and fourth values motivated the choice of the dimension of the latent space  $p = 3$ . To validate if the latent space contained meaningful information, we compared ROC curves computed on a held-out data set of 1,000 matrix entries and found a substantial increase in predictive power for the full model compared to the model only using fixed effects (right).

this size the `gllvm` package was already having difficulty converging to a solution. Based on the scree plot of singular values of the full model (Figure 3), we fixed the number of latent variables at  $p = 3$  across this simulation. Also, in order to use a metric comparable across different samplings, we used the mean deviance i.e.,  $D(Y, \hat{M})$  standardized by the number of rows and columns in the corresponding response matrix.

The proposed quasi-Newton method performed best both in terms of deviance and computation time, although for larger fractions of the full data set AIRWLS performed very similarly in terms of mean deviance (Figure 4). Elementary linear extrapolation suggests that computing a solution for the full data set via `gllvm` would take at least 2 weeks (if it converged at all). In light of this, we chose not to attempt to fit `gllvm` to the full data set due to anticipated memory constraints.

## 7. Simulations

We conducted a numerical study to empirically investigate how the number of responses, the number of latent variables, and the distribution of responses influence the performance of the proposed algorithms for estimating GLLVMs.

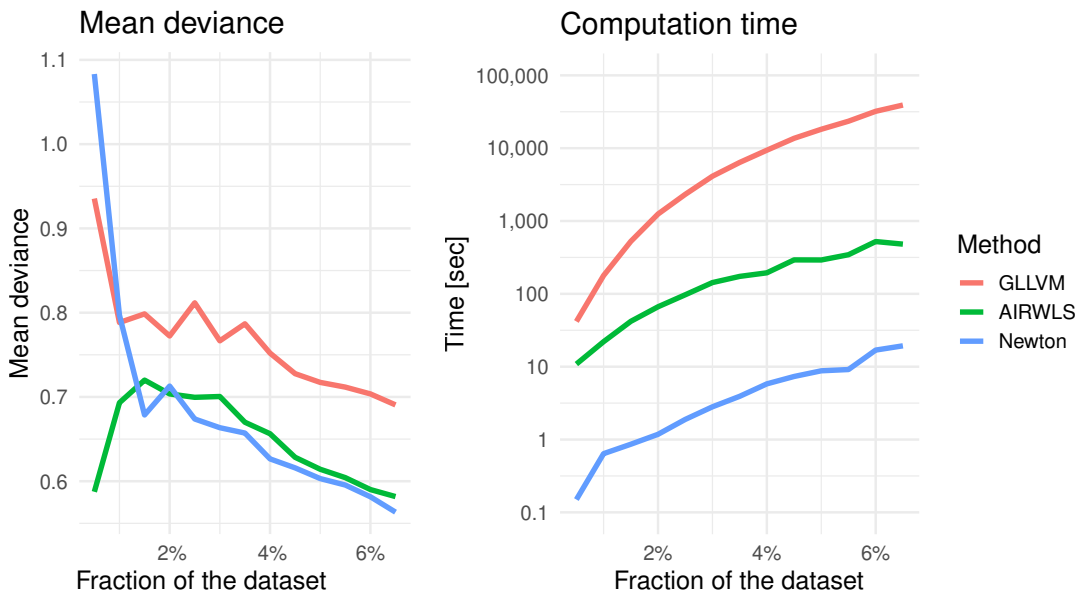


Figure 4: Mean deviance and computation time as a function of the data set size. We sampled fractions of columns and rows from the large data set of coexistence of species, and compared three methods for fitting GLLVMs to the sub-sampled subsets. A fraction  $\rho \in \{0.01, \dots, 0.065\}$  corresponded to a response matrix with  $\lfloor \rho \cdot 48737 \rfloor$  rows and  $\lfloor \rho \cdot 2211 \rfloor$  columns. We fixed the number of latent variables at  $p = 3$ . Both AIRWLS and Newton methods outperformed the baseline `gllvm` implementation both in terms of explained deviance (left) and computation time (right). We only report results up to 6.5% of the data set, since applying the baseline `gllvm` to larger data sets was not computationally feasible. We note however that the quasi-Newton method converged in three hours on the full data set.

## 7.1 Simulation Design

We designed our simulations to mimic the setting from the data study described in Section 6.1. To that end, we started with a GLLVM fitted to the ant abundance data. We used sample estimates of the covariance matrix of  $X$  and  $\Lambda$  as the basis for sampling new multivariate Gaussian variables for  $X$  and  $\Lambda$ . For  $U$  and  $B$ , we constructed matrices by randomly simulating each element independently from a standard Gaussian distribution. Using these quantities, we were able to construct a matrix of fitted values  $\mu_{ij}$ , from which we can then simulate a response matrix following Equation 1. We considered different combinations of the number of observational units  $n$ , the number of responses in each unit  $m$ , the number of latent variables  $p$ , and the distribution of responses (assuming a canonical link all response distributions). In particular, we ran simulations with  $n, m \in \{100, 200, 300, 500\}$ ,  $p \in \{2, 3, 5, 10\}$ , and considered Poisson and binomial distributions. Additionally, we tested

a small sample setting with  $n \in \{5, 10, 15\}$ ,  $m = 20$ , and  $p = 2$ . For each set of simulation parameters, we repeated the experiment with 100 generated data sets.

As in Section 6, for each simulated data set, we fitted the GLLVM using the `gllvm` package, along with the proposed quasi-Newton and AIRWLS methods. For all three methods, we used the same stopping criterion, with a relevant error tolerance equal to  $10^{-3}$ . We evaluated performance using mean deviance, Procrustes error for the fit of the latent space, and mean squared error of the estimated response-specific coefficients, all of which were detailed in Section 5.

## 7.2 Results

Aggregating across the combinations of  $n$ ,  $m$  and  $p$ , considered, both `gllvm` and the quasi-Newton method produced similar performance in terms of deviance explained, Procrustes error, and the MSE of fixed effects, while AIRWLS performed slightly better across these three measures (Figure 5). The main gain of the proposed methods, however, comes in computation time: using the same stopping criteria, computation times differed by several orders of magnitude, and on average the quasi-Newton algorithm took 51 seconds to compute, AIRWLS method took 132 seconds, while `gllvm` 51.6 minutes (Figure 5, bottom right).

Next, we investigated how the size of the response matrix, number of latent variables, and response distribution affected model performance and computation time. Overall, we found that in the small data experiment ( $m = 20, n \in \{5, 10, 15\}$ ), variational approximations using the `gllvm` packaged outperformed the proposed AIRWLS and quasi-Newton methods for  $n = 5$ , with deviance explained equal to 0.922, 0.733 and 0.65 respectively. However, all methods performed similarly in terms of mean deviance explained once  $n > 10$ . In terms of the number of latent variables, we found that as  $p$  increased both AIRWLS and quasi-Newton methods started to outperform the `gllvm` package in terms of deviance explained, and the gap in computation time increased in their favor (see Figure 6). An extreme example of this is at  $p = 10$ ,  $n = 500$ , and  $m = 300$ , where we found that the `gllvm` package took approximately 2.5 days to converge, while AIRWLS converged in 40 seconds and resulted in a more accurate fit in terms of deviance explained.

Finally, GLLVMs with binary responses (and using the logistic link) tended to be more difficult to fit than the Poisson distributions, and this occurred for all three estimation methods we tested. For example, keeping all other model parameters the same, the average mean deviance explained for GLLVMs fitted using `gllvm`, AIRWLS, and Newton was 0.383, 0.434, and 0.412 respectively for the Poisson distribution, but this decreased to 0.092, 0.127, and 0.125 respectively for the binomial distribution. The generally poorer fit observed for binary responses is consistent with what is commonly observed for GLMs, and reflects the general lack of information in binary response as well as potential issues such as quasi or complete-separation in the fitting process.

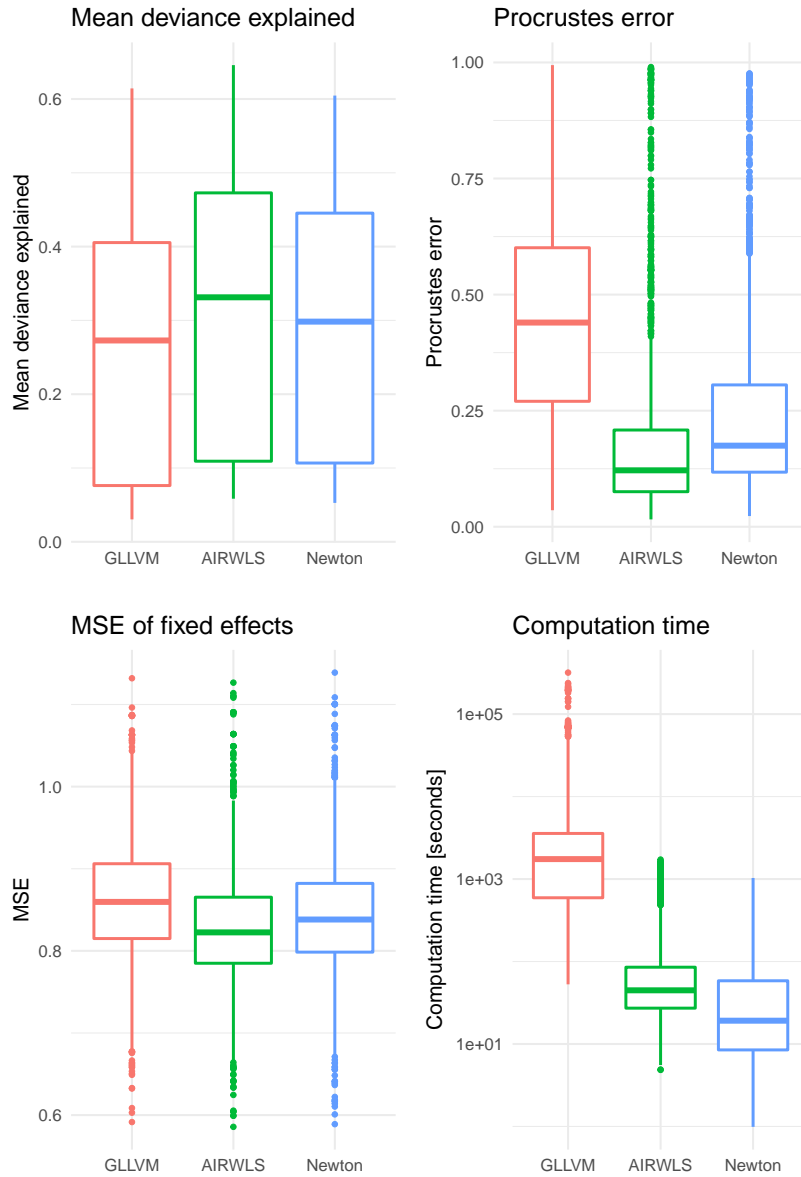


Figure 5: Comparison of the proposed AIRWLS and quasi-Newton methods with the `gllvm` package for fitting GLLVMs. In this simulation study, we sampled data such that covariances of  $X$  and  $V$  matched those of the sample estimates from the ant abundance data. We varied  $n, m \in \{100, 200, 300, 500\}$ ,  $p \in \{2, 3, 5, 10\}$  and the distribution of responses (Poisson or Binomial). Based on generating 100 data sets for each combination of simulation parameters, we found that the AIRWLS slightly outperformed other methods on all metrics (mean deviance explained, Procrustes error, and MSE of fixed effects), while the Newton algorithm was superior in terms of computation time.

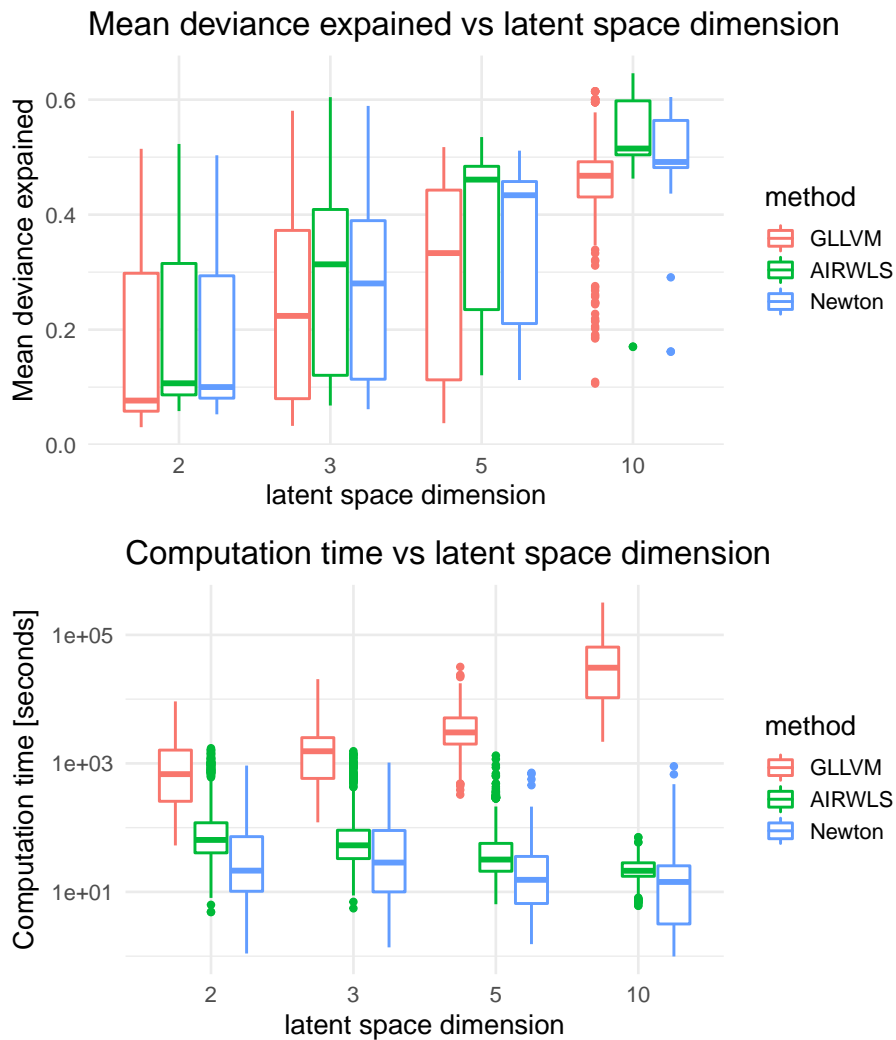


Figure 6: Assessment of how the number of latent variables  $p$  affects computation time and accuracy in terms of deviance explained. We found that as  $p$  grows, the proposed methods became more accurate than `gllvm`, while the difference in computation time also grew.

### 7.3 Model Selection: Unknown $p$ and Regularization

Our cross-validation framework introduced in Section 5 allows us to choose the tuning parameters, particularly for predictive applications. As an illustration of this, we simulated responses using the process described in Section 7.1, and compared two approaches for constraining and choosing the dimension of the latent space: smooth regularization based on sliding the parameter  $\gamma$  as described in Section 4.3, and an explicit rank constraint based on selecting  $p$  from a set of candidate integer values. We used 20-fold cross-validation to select both  $\gamma$  and  $p$ , and used out-of-sample deviance as the performance measure. Finally, we compared the predictive performance based on these selected parameter values by fitting the corresponding GLLVMs to a full training set containing 95% of the observations and then computing the deviance on the remaining hold-out set containing 5% of the observations.

In more detail, we set  $n = m = 100$ , the true dimension of the latent space as  $p = 2$ , and simulated Poisson responses. In both model selection approaches, we used the quasi-Newton method for fitting GLLVMs, and kept all parameters other than  $p$  and  $\gamma$  equal. We implemented smooth regularization as defined in Equation 14 with  $\gamma \in \{0, 1, \dots, 60\}$ , while for the rank-constrained models we considered  $p \in \{1, 2, \dots, 50\}$ .

Based on repeating the entire simulation above twenty times, we found that both methods of model selection achieved very similar results, with mean deviance 1.272 ( $sd = 0.130$ ) for the smooth regularization approach, and 1.279 ( $sd = 0.187$ ) for the rank-constrained approach. We conclude (in particular) that the cross-validation and smooth regularization approach for selecting  $p$  is a promising approach for tuning GLLVMs, made possible by the computational gains from our PQL-based estimation methods. We leave the thorough analysis of the theoretical properties of cross-validation and regularization of  $\gamma$  as an avenue of future research.

## 8. Discussion

For estimating GLLVMs, our proposed PQL-based methods are orders of magnitude faster than current state-of-the-art algorithms for estimating model parameters for GLLVMs. They can be decomposed and parallelized across multiple machines, leveraging modern open parallel computing platforms such as Apache Spark (Zaharia et al., 2010). It should be emphasized though that PQL-based methods are expected to work well in the settings that motivated this article, i.e., fitting GLLVMs to high-dimensional data sets. While estimates from PQL can exhibit more finite sample bias relative to say, the Laplace method, it has nevertheless been proven that they generally perform well when the number of observations per random effect,  $m$  in our setting, becomes large (Nie, 2007; Hui et al., 2017a; Hui, 2021). This is precisely the scenario of interest here, since one of the main reasons behind fitting a latent variable model, in general, is to perform matrix factorization and approximate a covariance matrix between  $m$  responses when  $m$  is too large to use an unstructured estimate.

Our algorithms are elementary to implement using existing GLM routines. Specifically, we provide an R implementation via the open-source package `gmf`<sup>1</sup> enabling integration with existing workflows and further extension of our algorithms. Our package can be used as a

---

1. Visit <http://github.com/kidzik/gmf/>



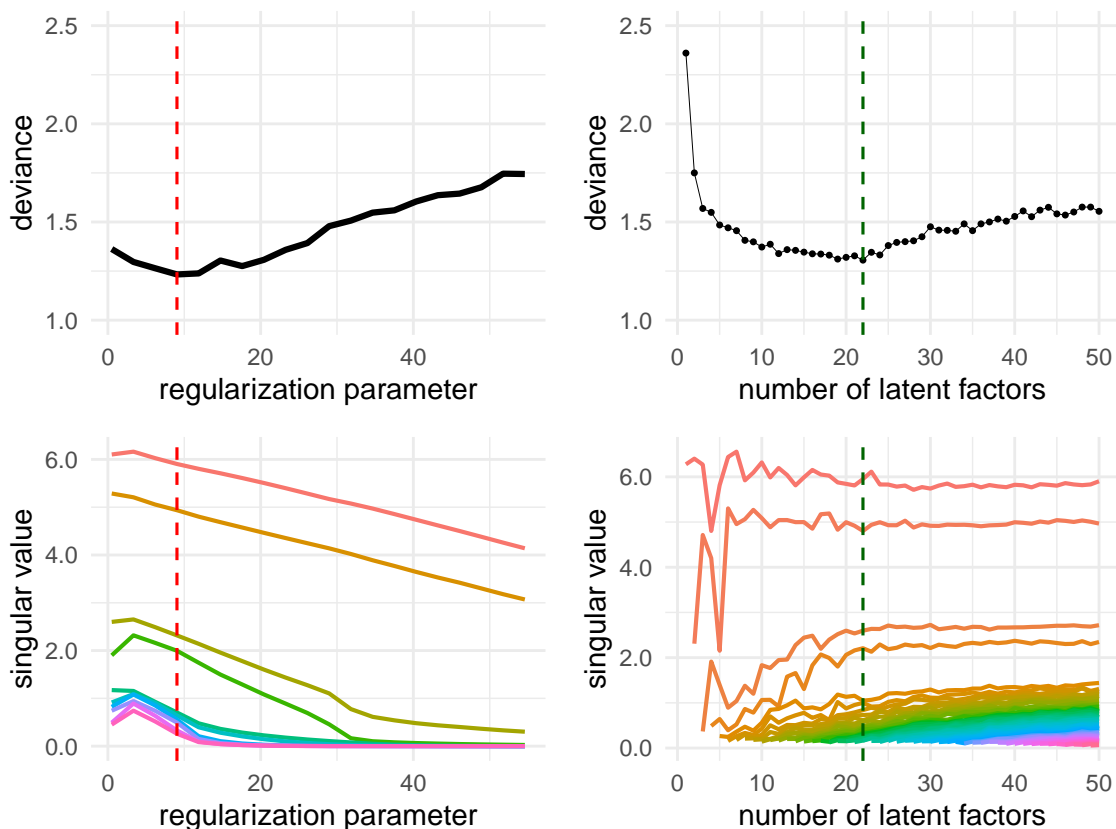


Figure 7: Model selection using the cross-validation framework. We present two methods for choosing the complexity of the model: smooth regularization (left panels) and rank constraint (right panels). The 20-fold cross-validation allowed us to select hyper-parameters minimizing the deviance (dashed lines). The number of nonzero singular values for a given  $\gamma$  is the selected rank. As expected, in the smooth regularization model (bottom left), the singular values all shrank when the regularizing parameter increased, while they remained stable when we added more factors in the rank-constrained model (bottom right).

drop-in replacement for widely-adopted heuristic methods such as, for example, fitting principal component analysis to log-transformed count data. Since our package is several orders of magnitude faster than current publicly available alternatives, then it has the potential to open new opportunities for cutting-edge applications of and research into GLLVMs to very high-dimensional multivariate abundance data.

Throughout this work, we have illustrated the applicability of our methodology in the context of ecology. However, similar problems can be found in other disciplines whenever we are interested in extracting latent factors underlying certain responses. Our methods are particularly useful when matrices are large, for example, in studies of behavior of subjects

online with thousands of individuals and items or web pages they view, or count data in single-cell RNA sequencing.

## Acknowledgments

Lukasz Kidziński was supported by the Mobilize Center grant U54 EB020405 from the National Institute of Health. Francis K.C. Hui was supported by an Australian Research Council Fellowship (DE200100435). David I. Warton was supported by the Australian Research Council’s Discovery Project Scheme (project DP210101923). Trevor J. Hastie was partially supported by grants DMS-2013736 And IIS, 1837931 from the National Science Foundation, and grant 5R01 EB, 001988-21 from the National Institutes of Health.

## References

- Jushan Bai and Serena Ng. Determining the number of factors in approximate factor models. *Econometrica*, 70:191–221, 2002.
- David J Bartholomew, Martin Knott, and Irimi Moustaki. *Latent Variable Models and Factor Analysis: A Unified Approach*, volume 904. John Wiley & Sons, 2011.
- Silvia Bianconcini and Silvia Cagnone. Estimation of generalized linear latent variable models via fully exponential laplace approximation. *Journal of Multivariate Analysis*, 112:183–193, 2012.
- Norman E Breslow and David G Clayton. Approximate inference in generalized linear mixed models. *Journal of the American statistical Association*, 88:9–25, 1993.
- Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76, 2017.
- Eugene Demidenko. *Mixed models: Theory and Applications with R*. John Wiley & Sons, 2013.
- Heloise Gibb and Saul A Cunningham. Habitat contrasts reveal a shift in the trophic position of ant assemblages. *Journal of Animal Ecology*, 80:119–127, 2011.
- Peter J Green. Penalized likelihood for general semi-parametric regression models. *International Statistical Review*, 55:245–259, 1987.
- Nathan Halko, Per-Gunnar Martinsson, Yoel Shkolnisky, and Mark Tygert. An algorithm for the principal component analysis of large data sets. *SIAM Journal on Scientific computing*, 33:2580–2594, 2011.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: the Lasso and Generalizations*. Chapman and Hall/CRC, 2019.
- Kei Hirose and Miyuki Imada. Sparse factor regression via penalized maximum likelihood estimation. *Statistical Papers*, 59:633–662, 2018.

- Kei Hirose and Michio Yamamoto. Sparse estimation via nonconcave penalized likelihood in factor analysis model. *Statistics and Computing*, 25:863–875, 2015.
- Philippe Huber, Elvezio Ronchetti, and Maria-Pia Victoria-Feser. Estimation of generalized linear latent variable models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66:893–908, 2004.
- Francis K. C. Hui. *boral: Bayesian Ordination and Regression AnaLysis*, 2020. R package version 1.9.
- Francis K C Hui. On the use of a penalized quaslikelihood information criterion for generalized linear mixed models. *Biometrika*, 108:353–365, 2021.
- Francis K C Hui, Samuel Mueller, and A H Welsh. Joint selection in mixed models using regularized PQL. *Journal of the American Statistical Association*, 112:1323–1333, 2017a.
- Francis K C Hui, David I Warton, John T Ormerod, Viivi Haapaniemi, and Sara Taskinen. Variational approximations for generalized linear latent variable models. *Journal of Computational and Graphical Statistics*, 26:35–43, 2017b.
- Francis K C Hui, Emi Tanaka, and David I Warton. Order selection and sparsity in latent variable models via the ordered factor lasso. *Biometrics*, 74:1311–1319, 2018.
- Francis KC Hui, Sara Taskinen, Shirley Pledger, Scott D Foster, and David I Warton. Model-based approaches to unconstrained ordination. *Methods in Ecology and Evolution*, 6(4):399–411, 2015.
- Robert E Kass and Adrian E Raftery. Bayes factors. *Journal of the American Statistical Association*, 90:773–795, 1995.
- Anthony YC Kuk. Laplace importance sampling for generalized linear mixed models. *Journal of Statistical Computation and Simulation*, 63:143–158, 1999.
- Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11:2287–2322, 2010.
- P McCullagh and J A Nelder. *Generalized Linear Models*. Chapman and Hall, 1983.
- CA McGilchrist. Estimation in generalized mixed models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56:61–69, 1994.
- Irini Moustaki and Martin Knott. Generalized Latent Trait Models. *Psychometrika*, 65:391–411, 2000.
- John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135:370–384, 1972.
- Lei Nie. Convergence rate of MLE in generalized linear and nonlinear mixed-effects models: Theory and applications. *Journal of Statistical Planning and Inference*, 137:1787–1804, 2007.

- Jenni Niku, David I Warton, Francis K C Hui, and Sara Taskinen. Generalized linear latent variable models for multivariate count and biomass data in ecology. *Journal of Agricultural, Biological and Environmental Statistics*, 22:498–522, 2017.
- Jenni Niku, Wesley Brooks, Riki Herliansyah, Francis K C Hui, Sara Taskinen, and David I Warton. Efficient estimation of generalized linear latent variable models. *PloS one*, 14: e0216129, 2019a.
- Jenni Niku, Francis K C Hui, Sara Taskinen, and David I Warton. gllvm: Fast analysis of multivariate abundance data with generalized linear latent variable models in R. *Methods in Ecology and Evolution*, 10:2173–2182, 2019b.
- Otso Ovaskainen, Gleb Tikhonov, Anna Norberg, F Guillaume Blanchet, Leo Duan, David Dunson, Tomas Roslin, and Nerea Abrego. How to make more out of community data? a conceptual framework and its implementation as models and software. *Ecology Letters*, 20:561–576, 2017.
- Maximilian Pichler and Florian Hartig. A new joint species distribution model for faster and more accurate inference of species associations from big community data. *Methods in Ecology and Evolution*, 12:2159–2173, 2021.
- Sophia Rabe-Hesketh, Anders Skrondal, and Andrew Pickles. Reliable estimation of generalized linear mixed models using adaptive quadrature. *The Stata Journal*, 2:1–21, 2002.
- Geneviève Robin. *Low-rank methods for heterogeneous and multi-source data*. PhD thesis, École Polytechnique, Paris, 2019.
- Håvard Rue, Sara Martino, and Nicolas Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71:319–392, 2009.
- Hans J Skaug. Automatic differentiation to facilitate maximum likelihood estimation in nonlinear random effects models. *Journal of Computational and Graphical Statistics*, 11: 458–470, 2002.
- Anders Skrondal and Sophia Rabe-Hesketh. *Generalized Latent Variable Modeling: Multi-level, Longitudinal, and Structural Equation Models*. CRC Press, 2004.
- Alison B Smith, Aanandini Ganesalingam, Haydn Kuchel, and Brian R Cullis. Factor analytic mixed models for the provision of grower information from national crop variety testing programs. *Theoretical and Applied Genetics*, 128:55–72, 2015.
- Charles Spearman. "general intelligence" objectively determined and measured. pages 59–73, 1961.
- Nathan Srebro, Noga Alon, and Tommi S Jaakkola. Generalization error bounds for collaborative prediction with low-rank matrices. In *Advances In Neural Information Processing Systems*, volume 17, pages 1321–1328, 2005.

- Oliver Stegle, Leopold Parts, Matias Piipari, John Winn, and Richard Durbin. Using probabilistic estimation of expression residuals (peer) to obtain increased power and interpretability of gene expression analyses. *Nature protocols*, 7:500, 2012.
- Gleb Tikhonov, Øystein H Opedal, Nerea Abrego, Alekski Lehtikoinen, Melinda MJ de Jonge, Jari Oksanen, and Otso Ovaskainen. Joint species distribution modelling with the R-package Hmsc. *Methods in ecology and evolution*, 11:442–447, 2020.
- Edward F Vonesh. A note on the use of Laplace’s approximation for nonlinear mixed-effects models. *Biometrika*, 83:447–452, 1996.
- YI Wang, Ulrike Naumann, Stephen T Wright, and David I Warton. mvabund—an R package for model-based analysis of multivariate abundance data. *Methods in Ecology and Evolution*, 3:471–474, 2012.
- David I Warton. Penalized normal likelihood and ridge regularization of correlation and covariance matrices. *Journal of the American Statistical Association*, 103:340–349, 2008.
- David I Warton, F Guillaume Blanchet, Robert B O’Hara, Otso Ovaskainen, Sara Taskinen, Steven C Walker, and Francis K C Hui. So many variables: joint modeling in community ecology. *Trends in Ecology & Evolution*, pages 766–779.
- David I Warton, F Guillaume Blanchet, Robert O’Hara, Otso Ovaskainen, Sara Taskinen, Steven C Walker, and Francis K C Hui. Extending joint models in community ecology: A response to Beissinger et al. *Trends in Ecology & Evolution*, 31:737–738, 2016.
- Daniela M Witten, Robert Tibshirani, and Trevor Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10:515–534, 2009.
- Philip Wolfe. Convergence conditions for ascent methods. *SIAM review*, 11:226–235, 1969.
- Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, Ion Stoica, et al. Spark: Cluster computing with working sets. *HotCloud*, 10:95, 2010.
- Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15:265–286, 2006.