

PREDICTING MILK TRAITS FROM SPECTRAL DATA USING BAYESIAN PROBABILISTIC PARTIAL LEAST SQUARES REGRESSION

BY SZYMON URBAS^{1,a}, PIERRE LOVERA^{2,c}, ROBERT DALY^{2,d}, ALAN O’RIORDAN^{2,e},
DONAGH BERRY^{3,f} AND ISOBEL CLAIRE GORMLEY^{1,b}

¹*School of Mathematics & Statistics, University College Dublin, ^aszymongurbas@gmail.com, ^bclaire.gormley@ucd.ie*

²*Tyndall National Institute, University College Cork, ^cpierre.lovera@tyndall.ie, ^drobert.daly@tyndall.ie,
^ealan.oriordan@tyndall.ie*

³*Animal & Grassland Research and Innovation Centre, Teagasc, ^fdonagh.berry@teagasc.ie*

High-dimensional spectral data—routinely generated in dairy production—are used to predict a range of traits in milk products. Partial least squares (PLS) regression is ubiquitously used for these prediction tasks. However, PLS regression is not typically viewed as arising from a probabilistic model, and parameter uncertainty is rarely quantified. Additionally, PLS regression does not easily lend itself to model-based modifications, coherent prediction intervals are not readily available, and the process of choosing the latent-space dimension, Q , can be subjective and sensitive to data size.

We introduce a Bayesian latent-variable model, emulating the desirable properties of PLS regression while accounting for parameter uncertainty in prediction. The need to choose Q is eschewed through a nonparametric shrinkage prior. The flexibility of the proposed Bayesian partial least squares (BPLS) regression framework is exemplified by considering sparsity modifications and allowing for multivariate response prediction.

The BPLS regression framework is used in two motivating settings: (1) multivariate trait prediction from mid-infrared spectral analyses of milk samples and (2) milk pH prediction from surface-enhanced Raman spectral data. The prediction performance of BPLS regression at least matches that of PLS regression. Additionally, the provision of correctly calibrated prediction intervals objectively provides richer, more informative inference for stakeholders in dairy production.

1. Introduction. Quality traits of animal-based products are important determinants of not only the nutritive value of the resulting product but also provide information on the optimal product processing strategy or product portfolio from the raw product. A rapid, low-cost, accurate and nondestructible method to quantify multiple quality parameters from a single sample, ideally in real time, is a requirement of the agri-food sector. Individual animal or bulk milk samples—taken globally from a range of different species—are subjected to spectroscopy analysis, usually in the mid-infrared region of the electromagnetic spectrum. Many studies have demonstrated how the generated spectrum can be used to predict milk quality (e.g., [De Marchi et al. \(2014\)](#)) and other animal level attributes including energy status ([Frizzarin et al. \(2023\)](#)) and methane emissions ([Coppa et al. \(2022\)](#)). The advantage of such a strategy is that the developed prediction model can be rapidly deployed to routinely collected samples to phenotype both the individual and bulk samples at little to no additional cost.

Partial least squares (PLS) regression is the industry-standard approach to developing prediction models using high-dimensional spectral data. This approach is also popular in the field of chemometrics (e.g., [Frank and Friedman \(1993\)](#), [Kourti \(2002\)](#)) and in other areas such as econometrics ([Naik and Tsai \(2000\)](#)) and genomics ([Boulesteix and Strimmer \(2007\)](#)).

Received July 2023; revised July 2024.

Key words and phrases. Milk trait prediction, spectral data, partial least squares, Bayesian factor analysis, high-dimensional statistics.

Partial least squares offers solutions to “*small N, large P*” prediction problems, where the number of observations in the training (calibration) data, N , is limited relative to the data dimensionality, P , and the predictor variables exhibit strong correlations.

The partial least squares (PLS) regression method quantifies the relationship between continuous predictor and response variables through a linear projection onto a lower-dimensional latent space. The particular projection is obtained by maximising the norm of the observed cross-covariance between the variables through algorithms such as NIPALS (Wold (1973)) or SIMPLS (de Jong (1993)).

Over the years, variants of PLS emerged to improve predictions under different settings; these include sparse PLS (sPLS, Chun and Keleş (2010)), robust PLS (Hubert and Branden (2003)), or kernel-based approaches (e.g., Lindgren, Geladi and Wold (1993)). An intrinsic limitation of PLS regression is that none of the optimisation algorithms arise as a solution to a particular inference problem; this is in contrast to principal component analysis (PCA, e.g., Chapter 12 of Bishop (2006)), a closely related dimension-reduction method, which was shown to be the maximum likelihood solution to a specific Gaussian factor model (Tipping and Bishop (1999)). To address this, approaches that emulate the PLS predictions through an appropriate probabilistic model have emerged; examples include Li et al. (2010), Vidaurre et al. (2013), Zheng, Song and Ge (2016), el Bouhaddani et al. (2018), and el Bouhaddani et al. (2022). Despite their coherent probabilistic frameworks, these prediction models focused on point predictions but did not consider their associated uncertainty, either through the assumed data-generating mechanism or as introduced by parameter uncertainty. Additionally, whilst in principle such methods could model a multivariate response, in practice they have limitations, for example, frequentist inference of probabilistic formulations of PLS requires limiting constraints on different parts of the models. Moreover, PLS regression requires the specification of the number of latent variables, which is usually informed through cross-validation based on the predictive error of a response variable—it is often more practical to fit multiple univariate models, leading to less coherent and interpretable regression.

Motivated by the need for coherent, richer inference when simultaneously predicting multiple milk traits from spectral data, we introduce Bayesian partial least squares (BPLS) regression, a latent-variable Gaussian model, which emulates PLS regression and addresses its limitations with regard to inference. The model inference is conducted in the Bayesian paradigm, thereby elegantly accounting for uncertainty. A nonparametric prior induces shrinkage, alleviating the difficult and often subjective choice of the latent dimension, particularly in settings involving multivariate traits. Finally, we exploit the model-based framework of BPLS by facilitating multivariate trait prediction and consider sparsity modifications providing increased model flexibility. Collectively these aspects of BPLS lead to a user-friendly and statistically-principled prediction method, applicable to a wide range of problems in the agri-food production sector. We demonstrate the performance of BPLS regression through synthetic examples prior to its applications to the real-life problem of predicting milk phenotypes from spectral data. Resulting point predictions are at least as accurate as those produced by the industry-standard PLS regression, with the provision of correctly calibrated prediction intervals objectively providing richer and more informative inference.

In what follows, Section 2 introduces the mid-infrared and surface-enhanced Raman spectral datasets from the Irish agri-food sector which motivate this study. Section 3 describes the BPLS regression model, its sparsity-inducing modifications, and its associated inference. Section 4 demonstrates the performance of BPLS regression on synthetic and benchmark data sets, while Section 5 explores the results of its motivating application to predict a range of milk traits from spectral data. Section 6 concludes with a discussion of the BPLS contribution and a delineation of some potential future research avenues.

The BPLS models and inference were implemented in R, and the code with which all results were generated is freely available at: <https://github.com/SzymonUrbas/bplsr>.

2. Milk spectral data from the Irish agri-food sector. While having general applicability, the BPLS methods introduced here are motivated by the problem of predicting a range of milk traits from spectral data as encountered in the Irish agri-food sector, whilst accounting for the inherent prediction uncertainty. Two motivating spectral datasets are considered: the first considers the prediction of various chemical and technological milk traits from mid-infrared (MIR) spectra, and the second focuses on predicting the pH of a milk sample using surface-enhanced Raman spectroscopy (SERS) data. In MIR data the predictors are based on spectra of absorbances, given as $\log(1/R_i)$, where R_i is the reflectance of the i th wavenumber in a given dataset; the reader is referred to [Beebe, Pell and Seasholtz \(1998\)](#) for further details. In SERS the predictors are logarithms of intensities of light scattering at different wavenumbers.

Throughout, we denote a predictor vector by $\mathbf{x} \in \mathbb{R}^P$, where P is the number of wavenumbers, and denote by $\mathbf{y} \in \mathbb{R}^R$ a vector of R responses or traits, noting that BPLS regression can consider multiple traits simultaneously, that is, $R \geq 1$. Each dataset has N samples, and \mathbf{X} denotes the $N \times P$ matrix whose n th row is \mathbf{x}_n^\top , with a similar definition for $\mathbf{Y} \in \mathbb{R}^{N \times R}$.

In what follows, all variables are standardised before inference. Additionally, cognisant of the practical utility of the BPLS predictions, traits \mathbf{Y} , which are strictly positive are transformed as outlined in Appendix A of the Supplementary Material ([Urbas et al. \(2024\)](#)).

2.1. Mid-infrared spectral data for milk trait prediction. Mid-infrared spectral data from 730 milk samples from 622 individual dairy cows of multiple breeds were available; data collection protocols and trait definitions are described in detail by [Visentin et al. \(2015\)](#) and [McDermott et al. \(2016\)](#). The spectral data were generated from the milk samples using the MilkoScan FT6000 (Foss Electronic A/S); the instrument reports the spectrum at 1060 different wavenumbers. As is standard, the high-noise regions of the spectral bands relating to water ($1710\text{--}1600\text{ cm}^{-1}$, $3690\text{--}2990\text{ cm}^{-1}$, and $> 3822\text{ cm}^{-1}$) were omitted. Here interest lies in simultaneously predicting from the spectra a set of $R = 3$ technical and chemical traits: heat stability (min), casein content (gL^{-1}), and rennet coagulation time (RCT, min). In the available data, casein content is supplied by an internal black-box prediction model of the instrument, while heat stability and RCT are physically measured in a laboratory; see [Visentin et al. \(2015\)](#) for details. Since the prediction methods considered here assume independent and identically distributed observations, we consider only one milk sample per cow and only those milk samples for which the complete set of $R = 3$ traits are available. The resulting dataset, therefore, comprises $N = 366$ milk samples, with $P = 531$ predictors and $R = 3$ responses. Figure 1 (left) shows the MIR spectra, while the traits are illustrated in




FIG. 1. The mid-infrared absorbance of milk samples (left) and the surface-enhanced Raman spectra from milk samples where intensity is scaled by the maximum observed value (right).




FIG. 2. Histograms of the heat stability (min), casein content (gL^{-1}), and rennet coagulation time (RCT, min) traits from the milk MIR spectral data.

Figure 2. This type of data is used for constructing prediction models for use by stakeholders across dairy production.

2.2. *Surface-enhanced Raman spectra from milk samples.* The intensity spectrum produced by surface-enhanced Raman spectroscopy (SERS) of a chemical on the surface of milk samples, along with the pH level of each sample, were collected by researchers at the Tyndall National Institute in Ireland; the resulting data are similar to those considered in Williams et al. (2016).

The SERS spectrum of the chemical varies with the pH of the solution in which it is immersed. The data set consists of intensity spectra of $N = 11$ milk samples at $P = 1733$ wavenumbers in the range $151.422\text{--}1799.820 \text{ cm}^{-1}$ and corresponding pH readings of each sample (hence, here $R = 1$). The milk samples are from two cartons of milk, five and six samples, respectively, with best-before dates one week apart. The cartons were kept refrigerated, and spectral and pH measurements were taken daily over the course of six days; to reduce measurement error, each spectrum is an average of approximately 20–30 measurements. Figure 1 (right) shows the SERS spectra; the pH values varied from 6.37–6.96 with a median value of 6.87. Interest lies in the potential of using the SERS data as a noninvasive and nondestructive means of inferring the pH of a milk sample. Such information is useful in dairy production, for example, to milk processors for spoilage detection or to farmers as milk pH may indicate mastitis infection (e.g., Kandeel et al. (2019)).

3. Model specification and inference.

3.1. *The Bayesian partial least squares regression model.* Partial least squares regression assumes a linear relationship between a set of covariates and a continuous multivariate response vector through a set of latent variables whose dimension Q is lower than P , the cardinality of the set of covariates. Here we consider PLS regression from a probabilistic modelling viewpoint and construct a principled statistical model, which assumes a similar linear relationship and imposes explicit distributional assumptions, that is, we formulate the following generative model for observations $\{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$:

$$(1) \quad \mathbf{x}_n = W\mathbf{z}_n + \boldsymbol{\varepsilon}_n, \quad \mathbf{y}_n = C\mathbf{z}_n + \boldsymbol{\eta}_n, \quad \mathbf{z}_n \sim N(\mathbf{0}, I_Q),$$

where $W \in \mathbb{R}^{P \times Q}$ and $C \in \mathbb{R}^{R \times Q}$ are loading matrices, and $\boldsymbol{\varepsilon}_n \sim N(\mathbf{0}, \Sigma)$ and $\boldsymbol{\eta}_n \sim N(\mathbf{0}, \Psi)$ are independent noise vectors. Here $N(\boldsymbol{\mu}, \Phi)$ denotes a multivariate Gaussian distribution with mean vector $\boldsymbol{\mu}$ and a covariance matrix Φ . Here the covariance matrices Σ and Ψ are assumed to be diagonal. Using terminology from the factor analysis literature, the individual diagonal elements of Σ and Ψ are referred to as *uniquenesses*, and the latent \mathbf{z}_n vectors are

known as *scores*. The specification of the latent dimension Q is not trivial and is discussed in detail in Section 3.2. The loading matrix C can be thought of as a regression matrix; that is, in the case of $R = 1$, it would be a vector of regression coefficients. As we conduct inference in the Bayesian paradigm, we refer to the model in (1) as Bayesian probabilistic partial least squares (BPLS) regression and let Θ denote the collection of all the model parameters where $\Theta = (W, C, \Sigma, \Psi)$. The distributional assumptions in (1) result in the following joint density of the model variables for a single observation:

$$(2) \quad \begin{aligned} p(\mathbf{x}_n, \mathbf{y}_n, \mathbf{z}_n | \Theta) &= p(\mathbf{x}_n | \mathbf{z}_n, W, \Sigma) p(\mathbf{y}_n | \mathbf{z}_n, C, \Psi) p(\mathbf{z}_n) \\ &\propto |\Sigma|^{-1/2} |\Psi|^{-1/2} \exp\left(-\frac{1}{2} \left((\mathbf{x}_n - W\mathbf{z}_n)^\top \Sigma^{-1} (\mathbf{x}_n - W\mathbf{z}_n) \right. \right. \\ &\quad \left. \left. + (\mathbf{y}_n - C\mathbf{z}_n)^\top \Psi^{-1} (\mathbf{y}_n - C\mathbf{z}_n) + \mathbf{z}_n^\top \mathbf{z}_n \right) \right). \end{aligned}$$

It follows that the posterior conditional distribution for the latent variable is $\mathbf{z}_n | \mathbf{x}_n, \mathbf{y}_n, \Theta \sim \mathbf{N}(\mathbf{a}, S)$, where $S = (I_Q + W^\top \Sigma^{-1} W + C^\top \Psi^{-1} C)^{-1}$ and $\mathbf{a} = S(W^\top \Sigma^{-1} \mathbf{x}_n + C^\top \Psi^{-1} \mathbf{y}_n)$. Letting $\mathcal{D} = ((\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N))$ denote a dataset of N independent and identically distributed observed samples, along with latent score variables $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_N)^\top$, the complete data likelihood is $p(\mathcal{D}, \mathbf{Z} | \Theta) = \prod_{n=1}^N p(\mathbf{x}_n, \mathbf{y}_n, \mathbf{z}_n | \Theta)$.

For Bayesian inference we assume a prior distribution π_0 on the model parameters and the posterior distribution follows from Bayes' Theorem: $\pi(\Theta, \mathbf{Z} | \mathcal{D}) \propto p(\mathcal{D}, \mathbf{Z} | \Theta) \pi_0(\Theta)$. The choice of π_0 used here results in tractable conjugate conditional posteriors, which allow exact sampling of $\Theta, \mathbf{Z} | \mathcal{D} \sim \pi$ via a Gibbs algorithm; inference details are given in Appendix B of the Supplementary Material. The elements of W and C are assigned conjugate normal priors with details given in Section 3.2. The uniquenesses are assigned conjugate inverse-gamma priors: $\sigma_p^{-2} \sim \text{Gamma}(A_\sigma, B_\sigma)$, $p = 1, \dots, P$, and $\psi_r^{-2} \sim \text{Gamma}(A_\psi, B_\psi)$, $r = 1, \dots, R$. We use $(A_\sigma, B_\sigma, A_\psi, B_\psi) = (2.5, 0.1, 2.5, 1.5)$; the shape parameters A_σ and A_ψ are taken to be > 1 to ensure the first two prior moments of the uniqueness variables are finite, whereas B_σ and B_ψ are chosen to reflect the typical scale of the noise observed in the corresponding variables. Prior sensitivity analyses in typical spectral data scenarios indicated no substantial difference in results when varying (A_σ, B_σ) , but the uniquenesses associated with \mathbf{Y} required a more informative prior to avoid Heywood problems;¹ for all data considered here (after standardisation), our choices for (A_ψ, B_ψ) were sufficiently general and stable. An alternative approach (Fruhworth-Schnatter and Lopes (2014)) is to use empirically-derived priors that vary by variable with, for example, $B_\sigma^p = (A_\sigma - 1) / [V^{-1}]_{pp}$, $p = 1, \dots, P$, where V is the sample covariance matrix of \mathbf{X} ; similar definitions follow for B_ψ^r , $r = 1, \dots, R$. However, experiments (not shown here) indicated that this approach generally resulted in poorer prediction accuracy. Either (or both) of the covariance matrices Σ and Φ could be assumed to be isotropic, as is the basis for probabilistic principal component analysis (Tipping and Bishop (1999)). This particular modification of BPLS regression (detailed in Appendix B of the Supplementary Material) gives a more parsimonious but less flexible model.

The BPLS regression model does not make assumptions about the orthonormality of the loading matrix columns or constrain Q to any specific range. This is in contrast to the probabilistic approach taken in el Bouhaddani et al. (2018), where $Q \leq \min\{P, R\}$ is required to ensure principled frequentist inference. In BPLS regression the need for model constraints is circumvented through the Bayesian priors which provide the required regularisation.

¹Without an appropriately regularising prior, the Bayesian posterior of an idiosyncratic variance term (here, σ_p^2 , $p = 1, \dots, P$; ψ_r^2 , $r = 1, \dots, R$) is multimodal with one mode located at 0. For a detailed discussion, see Section 3.2 of Fruhwirth-Schnatter and Lopes (2014).

The BPLS regression formulation is motivated by the specific variable structure assumed in PLS. By bringing it into a probabilistic modelling paradigm, without imposing hard constraints on latent variable dimensionality or loading matrix columns, we arrive at a formulation that can be seen as a special case of a Bayesian factor analysis (BFA) model (e.g., [Press and Shigemasa \(1989\)](#), [Song and Lee \(2001\)](#)) with a blocked covariance structure. Whilst the relationship is not as direct as that established in probabilistic PCA ([Tipping and Bishop \(1999\)](#)), it nonetheless admits the emulation of PLS using a probabilistic model. We exploit this connection by implementing a Bayesian nonparametric prior commonly seen in BFA models to avoid the need to specify the latent dimension (see Section 3.2). Additionally, by drawing on the regression aspect of (1), we can appropriately modify parts of the model using sparsity-inducing priors to improve overall predictive performance (see Sections 3.3 and 3.4). While [Vidaurre et al. \(2013\)](#) also proposes a Bayesian PLS-type model, the assumed data-generating mechanism does not emulate the relationship between the predictors and the response in the same way as standard PLS regression and so is fundamentally different to the BPLS regression model proposed here. Furthermore, [Vidaurre et al. \(2013\)](#) carry out approximate variational inference, with no quantification of the mean field approximation error in the predictions provided.

Regarding identifiability, an inherent feature of factor models is that the loading matrices, W and C , and scores \mathbf{Z} are not identifiable—rotating these in the latent space results in the same posterior density. However, while Procrustean rotations ([Hoff, Raftery and Handcock \(2002\)](#), [McParland et al. \(2017\)](#)) could be employed to resolve this, the quantities of interest here are the posterior predictive distributions of the response vectors, which are invariant to rotations (see Section 3.5), and so such an approach is possible but unnecessary.

3.2. *Inference on infinite latent dimensions.* Methods based on dimension-reduction, such as PLS, generally require the user to select the dimension of the latent variables (often termed the number of components) in the model; the choice can often be subjective or based on data-sensitive heuristics. A common approach for PLS is to use cross-validation (CV) in order to minimise the root mean squared error (RMSE) of the predictions; this is the default in standard statistical packages such as `ppls` ([Liland, Mevik and Wehrens \(2022\)](#)) and `sppls` ([Chung, Chun and Keles \(2019\)](#)) in R ([R Core Team \(2023\)](#)). These methods require many separate model fits, and one needs to specify the data folds which involves a trade-off between metric robustness and computational efficiency. Furthermore, there is no clear consensus as to how to choose Q when fitting a given model to multivariate responses, as is required in the motivating applications here. In practice, one either needs to fit each response variable separately or use subjective judgement when choosing a joint model.

In the proposed BPLS regression framework, we avoid the selection process and fit only one model by using an appropriate nonparametric prior. This obviates the need to choose Q and allows for posterior predictions, which reflect its inherent uncertainty, thus making the framework more objective, setting-agnostic and user-friendly. We employ the popular multiplicative gamma process (MGP) prior ([Bhattacharya and Dunson \(2011\)](#)), which assumes the effective dimensionality Q is infinite whilst enforcing shrinkage on the entries of the loading matrices (W and C), which increases with the column index. This is achieved by taking the prior precisions of the loadings to be a stochastically increasing sequence of gamma variables. Specifically, the MGP prior considered here is

$$w_{pq} | \phi_{pq}, \tau_q \sim N(0, \phi_{pq}^{-1} \tau_q^{-1}), \quad \phi_{pq} \sim \text{Gamma}(v_1^w, v_2^w), \quad p = 1, \dots, P, q \in \mathbb{N};$$

$$c_{rq} | \dot{\phi}_{rq}, \tau_q \sim N(0, \dot{\phi}_{rq}^{-1} \tau_q^{-1}), \quad \dot{\phi}_{rq} \sim \text{Gamma}(v_1^c, v_2^c), \quad r = 1, \dots, R, q \in \mathbb{N}, \quad \text{where}$$

$$\tau_q = \prod_{k=1}^q \delta_k, \quad \delta_1 = 1, \quad \text{and} \quad \delta_k \sim \text{Gamma}(\alpha, \beta), \quad k = 2, \dots, q.$$

The conditional posteriors of the hyperparameters $(\phi, \dot{\phi}, \delta)$ are tractable and can be inferred through a Gibbs sampling algorithm. The prior formulation considered here differs from that of [Bhattacharya and Dunson \(2011\)](#) in that, without loss of generality, we fix δ_1 as all prior precisions are multiplied by the term—empirically, no significant effect on the final model fit was observed. The MGP prior has been widely used in latent variable models (e.g., [Ovaskainen et al. \(2016\)](#), [Murphy, Viroli and Gormley \(2020\)](#)). The α parameter could be inferred; however, we observed that this can lead to unstable results. Whilst $\alpha > \beta$ leads to stochastically increasing τ , the reciprocals of its elements are not guaranteed to be stochastically decreasing. To overcome the issue, [Durante \(2017\)](#) advises fixing $\beta = 1$ and $\alpha > \beta + 1$ for reliable analyses. There is more flexibility when specifying the individual precision components ϕ_{pq} and $\dot{\phi}_{rq}$; adopting hyperparameters such that $\mathbb{E}[\phi_{pq}^{-1}]$, $\mathbb{E}[\dot{\phi}_{rq}^{-1}] \leq 1$ allows for sparser loading matrices, implying that $v_2^w \leq v_1^w - 1$ and $v_2^c \leq v_1^c - 1$. Here we set $(\alpha, \beta, v_1^w, v_2^w, v_1^c, v_2^c) = (2.2, 1, 2, 3, 2, 3)$; prior sensitivity analyses showed no substantial difference in results when varying these, provided the recommendations are followed.

For practical implementation of the MGP prior, we consider truncation at some appropriate dimension Q^* ; [Bhattacharya and Dunson \(2011\)](#) and [Gwee, Gormley and Fop \(2023\)](#) suggest using a reasonably large value. Here at the initialisation stage of inference, standard principal component analysis is performed on \mathbf{X} , and a conservative bound Q^* on the effective dimension is obtained; for example, one which explains 0.99 of the observed variation. For the spectral data considered here, this procedure gave Q^* between 8–27. This selection can be verified by examining the columns of the loading matrices during warm-up runs of the Markov chain Monte Carlo (MCMC) chains; the magnitudes of the elements of the latter columns should be, for example, less than 1% of the first few columns, indicating negligible contributions to the variation in the data. For the motivating datasets here, we found that 5000 iterations (after discarding the burn-in) were sufficient for reliably diagnosing the choice of Q^* . Additionally, the posterior distribution of τ can be used to determine if the shrinkage prior hyperparameters are appropriate; see Appendix C of the Supplementary Material. In the numerical experiments in Section 4, we consider synthetic data generated using $Q = 10$ latent variables, and for inference we truncate the infinite model at $Q^* = 15$. This is relatively large compared to what appears in, for example, [el Bouhaddani et al. \(2022\)](#) and is motivated by the intricacy of spectral data; dimension-reduction methods applied to spectral data in Sections 4.2 and 5 generally required between five and 20 latent variables.

An adaptive Gibbs algorithm such as that in [Bhattacharya and Dunson \(2011\)](#) could be used to automate the process of choosing Q^* ; however, here preliminary experiments showed no substantial computational gain from using the adaptive scheme as the adaptive steps increased the resulting MCMC burn-in.

3.3. Spike-and-slab BPLS regression. It is possible that some of the leading principal components of \mathbf{X} , accounting for much of its variation, play no role in the relevant regression in (1). Indeed, this is the basis for orthogonal partial least squares approaches (O2-PLS [Trygg and Wold \(2002, 2003\)](#)), where the latent variables are split into *joint* components, shared between both \mathbf{X} and \mathbf{Y} , and *specific* components, split into two further sets: ones only interacting with \mathbf{X} and ones only interacting with \mathbf{Y} . A probabilistic extension PO2-PLS ([el Bouhaddani et al. \(2022\)](#)) has a necessary constraint $R \leq Q$, and notably, this further requires the user to choose the dimensionality of each of the component sets.

The decision of whether a given component is shared between the two observed variable sets or only explains one equates to a variable selection problem on the space of the latent variables. We, therefore, propose an extension to BPLS which emulates O2-PLS by employing a modification based on spike-and-slab regression ([George and McCulloch \(1993\)](#)). Here as we strictly predict \mathbf{y} given \mathbf{x} , we only consider specific components relating to \mathbf{x} , which

corresponds to some columns of regression matrix C being set to zero. We augment the model with a binary, diagonal matrix $B = \text{diag}(b_1, \dots, b_Q)$, whose diagonal entries are independent Bernoulli random variables, each with a prior probability of success $p_0 \sim \text{Unif}(0, 1)$. The binary variates act as “switches” for their respective variables. The matrix is incorporated into the relevant regression part of (1), that is, $\mathbf{y} = CB\mathbf{z} + \boldsymbol{\eta}$, with the product CB providing the desired sparsity. The posterior sampling of $\{b_q\}_{q=1}^Q$ is based on exploiting the product form of the Gaussian likelihood (2) as detailed in Appendix B.1 of the Supplementary Material. We call this model variant ss-BPLS.

While the orthogonal PLS methods do not necessarily improve upon the predictions of standard PLS (Trygg and Wold (2002)), they can, however, offer more interpretable results by discarding parts of \mathbf{x} which are redundant in explaining the variation in \mathbf{y} . However, due to the rotational invariance of a Gaussian factor model, there is a lack of identifiability in the B matrix in the absence of hard constraints. Hence, whilst the posterior distribution of B here may not be useful in analysing the relationship between \mathbf{X} and \mathbf{Y} , the spike-and-slab mechanism can assist in avoiding overfitting where too many latent variables are required to explain the variation in the response vector. Sparse regression can aid the estimation of the Ψ matrix elements, producing more reliable prediction intervals.

3.4. *LASSO BPLS regression.* As we are interested in modelling multivariate \mathbf{y} vector, some latent variables may affect one subset of the responses and not another, implying sparsity on the individual elements of C , not on entire columns of C . This could be achieved by a LASSO-type penalty on the values of the entries (e.g., Choi, Zou and Oehlert (2010)). An explicit Bayesian prior of the form $\pi_0(C) \propto \exp(-\lambda \sum_{r,q} |c_{rq}|)$, where $\lambda > 0$ is a regularisation hyperparameter, leads to an intractable posterior density severely complicating the inference. However, as outlined in Park and Casella (2008), setting up a Bayesian prior structure $c_{rq} \sim N(0, \dot{\phi}_{rq}^{-1} \tau_q^{-1})$ and $\dot{\phi}_{rq}^{-1} \sim \text{Exp}(\frac{\lambda^2}{2})$ can give the desired ℓ_1 -penalty. It can be shown that

the marginal prior density of c_{rq} has the form $\pi_0(c_{rq}|\lambda, \boldsymbol{\tau}) \propto \frac{\lambda}{2\sqrt{\tau_q^{-1}}} e^{-\lambda|c_{rq}|/\sqrt{\tau_q^{-1}}}$, which leads

to tractable conditional posteriors; a posteriori, the individual $\dot{\phi}_{rq}$ parameters are inverse-Gaussian distributed. Furthermore, assigning a prior distribution $\lambda^2 \sim \text{Gamma}(A_\lambda, B_\lambda)$ leads to a conjugate conditional posterior, thus allowing inference of the regularisation hyperparameter; here we use $A_\lambda = B_\lambda = 1$, but other choices could be appropriate for specific settings. The details of the resulting MCMC scheme are given in Appendix B.2 of the Supplementary Material. We refer to this formulation of the model as L-BPLS.

Rockova (2023) argues that point predictions produced from the marginal posterior predictive of a Bayesian LASSO model have an inherent bias and are prone to overfitting as the number of variables increases. However, the same work shows that if the hierarchical LASSO-type prior is combined with another sparsity prior (e.g., spike-and-slab mixture), then intended predictive performance is recovered. Here we employ the MGP shrinkage prior which enforces shrinkage on the latter columns of the loading matrices and, as seen in the simulation studies in Sections 4 and 5, accurate predictive performance is indeed observed.

3.5. *Predictive distribution of the response.* The aim of BPLS regression is to accurately characterise the posterior predictive distribution of response traits \mathbf{y}_+ , given only new spectral samples \mathbf{x}_+ ; such distributions are readily accessible through the Bayesian framework. We first note the conditional distributions involving the latent variable \mathbf{z}_+ , that is, $\mathbf{z}_+|\mathbf{x}_+, \Theta \sim N(S_z W^\top \Sigma^{-1} \mathbf{x}_+, S_z)$, where $S_z = (I_Q + W^\top \Sigma^{-1} W)^{-1}$, which follows from the incomplete likelihood where only \mathbf{x}_+ is observed and $\mathbf{y}_+|\mathbf{z}_+, \Theta \sim N(C\mathbf{z}_+, \Psi)$, which follows from the assumed generative model (1). Integrating out the latent \mathbf{z}_+ , we find the marginal posterior predictive distribution $p(\mathbf{y}_+|\mathbf{x}_+, \Theta)$ is also Gaussian with $\mathbb{E}[\mathbf{y}_+|\mathbf{x}_+, \Theta] = CS_z W^\top \Sigma^{-1} \mathbf{x}_+$

and $\mathbb{V}[\mathbf{y}_+|\mathbf{x}_+, \Theta] = C S_z C^\top + \Psi$. Working in the Bayesian paradigm allows us to incorporate the uncertainty in Θ into the predictions by marginalising over the parameter posterior, $p(\mathbf{y}_+|\mathbf{x}_+, \mathcal{D}) = \int p(\mathbf{y}_+|\mathbf{x}_+, \Theta)\pi(\Theta|\mathcal{D}) d\Theta$. Expectations with respect to these marginal posterior predictive distributions can be approximated via Monte Carlo integration based on the MCMC output; Appendix D.1 of the Supplementary Material details this procedure.

In what follows, the point predictions from BPLS regression models are given as $\mathbb{E}[\mathbf{y}_+|\mathbf{x}_+, \mathcal{D}]$, but other options are available such as predicting from the posterior mode; a comparison in Appendix D.2 of the Supplementary Material shows there is little substantive difference between the approaches. The prediction intervals for each observation are constructed from equal-tail quantile intervals of the posterior predictive distribution. These are very similar to the highest posterior density intervals; the predictive distributions are unimodal and symmetric due to the distributional assumptions of the model.

4. Numerical experiments. In this section we evaluate the empirical performance of the BPLS regression framework on synthetic and benchmark data and compare its performance to other well-utilised methods: partial least squares where Q is chosen through cross-validation and the “one-sigma” heuristic approach (van der Voet (1994)), denoted PLS and PLS-1s, respectively, sparse partial least squares (sPLS), principal component regression (PCR), and LASSO (Tibshirani (1996)) and ridge regressions (Hoerl and Kennard (2020)). The number of components for PLS and PCR methods was chosen using the default 10-fold cross-validation using the `pls` package in R; sPLS used the default cross-validation in the `spls` package. We assess performance through the root mean squared error of the predictions (RMSEP) on a given test set. The probabilistic partial least squares (PPLS) regression model (el Bouhaddani et al. (2018)) is not applied to the synthetic data here as a setting where $R < Q$ is of interest, which PPLS does not permit. The PPLS model is, however, used in comparisons on real data where the latent dimension is unknown (see Sections 4.2 and 5).

The Gibbs sampling algorithm is used for BPLS regression inference, with initial values and the iterations of the Markov chains detailed in Appendix B.3 of the Supplementary Material. For synthetic data examples, a conservative choice of $Q^* = 15$ was used when inferring BPLS models; for analyses on real spectral datasets, we employed the method discussed in Section 3.2. To assess convergence, we consider the mixing of the entries of $\mathbf{Y}^{\text{test}} = C S_z W^\top \Sigma^{-1} \mathbf{X}^{\text{test}}$; this product is invariant to the different rotations of the latent space. The chains are run long enough such that after appropriately discarding the burn-in, we are left with at least 1000 effective samples of each element. This condition indicates the chains are mixing sufficiently well; see, for example, Vats and Knudson (2021). Here this was safely achieved in all scenarios with a burn-in of 9000 iterations and a further 21,000 iterations. The BPLS models are fitted to multivariate responses, whereas the competitors are fitted to each univariate response element; preliminary experiments indicated no benefit in using multiple univariate-response BPLS regression models.

4.1. Synthetic datasets. The performance of BPLS regression is assessed by comparison to that of competitors across different data scenarios. To emulate the motivating practical applications, we consider a multivariate response, where $R = 4$, and set $Q = 10$ with $N \in \{50, 500\}$ and $P \in \{100, 1000\}$. The loading matrices’ entries are simulated using via $w_{pq} \stackrel{\text{iid}}{\sim} N(0, 1/Q)$ and $c_{pq} \stackrel{\text{iid}}{\sim} N(0, 1/Q)$; this controls the marginal prior variance contribution to the components of \mathbf{x} and \mathbf{y} . We consider two isotropic noise scenarios: low-noise and high-noise, corresponding to $\sigma^2 = \psi^2 = 0.1$ and $\sigma^2 = \psi^2 = 0.5$, respectively. Appendix E.1 of the Supplementary Material provides additional results where the C matrix is sparse; all BPLS regression variants were, in fact, robust in such settings.




FIG. 3. The RMSEP across 10 synthetic datasets for different settings of N_{train} and P under the BPLS and existing methods. Low-noise and high-noise scenarios correspond to $\sigma^2 = \psi^2 = 0.1$ and $\sigma^2 = \psi^2 = 0.5$, respectively.

Under the hyperparameter settings outlined in Section 3.2, each model is fitted to a training dataset of size N_{train} , simulated with the above specifications, and then standardised. The fitted models are then used to make predictions for $N_{\text{test}} = 1000$ unseen test samples, and the RMSEP is given as the average error over the response elements; by construction they are on the same scale. This setup is repeated for a total of 10 synthetic data sets.

Figure 3 illustrates how the BPLS performance compares to other methods. In scenarios where the data have small $N_{\text{train}} = 50$ and $P = 100$, the BPLS, ss-BPLS, and L-BPLS models provide a clear advantage in terms of prediction accuracy over their competitors. As the number of training samples increases to $N_{\text{train}} = 500$, all PLS-type methods perform similarly and tend to be better than the standard LASSO and ridge regression competitors. Figure 3 shows that the impact of the nonparametric prior resulting in bias in an infinite-dimensional model does not necessarily disappear as the sample size grows (Diaconis and Freedman (1986), Rousseau (2016)). The presence of high levels of noise decreases the accuracy of the competing methods more substantially, as compared to the BPLS regression approaches. Appendix C of the Supplementary Material provides details on the shrinkage parameter posteriors showing the BPLS regression model correctly identifies the latent dimensionality of the observed data. An additional comparison of the BPLS methods to standard Bayesian factor analysis on the most challenging synthetic data setting (i.e., $P = 1000$, $N = 50$, high noise) is given in Appendix F.1 of the Supplementary Material.

In addition to strong predictive performance, the BPLS models provide uncertainty quantification through the posterior predictive distributions. Table 1 gives the estimated coverage probabilities for test-set observations under various simulation settings. All models achieve the desired coverage when larger samples are available ($N_{\text{train}} = 500$), but when only a small number of samples is available, the prediction interval (PI) width is underestimated for BPLS

TABLE 1

Estimated prediction coverage of 95% prediction intervals produced by different BPLS regression models under different data settings: mean (sd). At each iteration the median coverage value is taken across the elements of the multivariate response vector. Low-noise and high-noise scenarios correspond to $\sigma^2 = \psi^2 = 0.1$ and $\sigma^2 = \psi^2 = 0.5$, respectively

		Low noise		High noise	
		P = 100	P = 1000	P = 100	P = 1000
BPLS	N = 50	0.929 (0.014)	0.880 (0.026)	0.911 (0.029)	0.844 (0.056)
	N = 500	0.948 (0.007)	0.945 (0.005)	0.945 (0.007)	0.945 (0.008)
ss-BPLS	N = 50	0.927 (0.014)	0.891 (0.029)	0.906 (0.028)	0.845 (0.041)
	N = 500	0.948 (0.009)	0.945 (0.004)	0.945 (0.006)	0.944 (0.007)
L-BPLS	N = 50	0.966 (0.008)	0.962 (0.009)	0.938 (0.021)	0.937 (0.019)
	N = 500	0.952 (0.005)	0.952 (0.002)	0.950 (0.005)	0.951 (0.002)

and ss-BPLS. Upon further investigation, the too narrow PIs correspond to underestimation of the Ψ matrix. The PIs produced by L-BPLS are a little too conservative when the noise is low but too narrow when it is high; indeed, this suggests that a sparse C matrix leads to a Ψ posterior which relies heavily on the prior. There did not appear any meaningful difference in the prediction interval widths achieving the same coverage.

4.2. *Benchmark spectral data.* Here we compare existing approaches and the proposed BPLS regression methods on near-infrared (NIR) spectral data of grain mash samples considered in Liebmann, Friedl and Varmuza (2009); data are available from the R package *chemometrics* (Filzmoser and Varmuza (2017)). Each of the $N = 166$ samples in the data contains $P = 235$ predictor variables, given as the approximate first derivative of the curve made up of absorbance values, and $R = 2$ traits: ethanol concentration and glucose concentration (both measured in gL^{-1}). This data set is well utilised as a benchmark for dimension-reduction prediction methods (e.g., Filzmoser, Gschwandtner and Todorov (2012)).

We consider a similar setup to that in Frizzarin et al. (2021): the dataset is divided into four approximately equal folds; three folds are used as the training set $\mathcal{D}^{\text{train}} = (\mathbf{X}^{\text{train}}, \mathbf{Y}^{\text{train}})$ with N_{train} observations, and the remaining fold is treated as the test set, $\mathcal{D}^{\text{test}} = (\mathbf{X}^{\text{test}}, \mathbf{Y}^{\text{test}})$ with N_{test} observations; all four fold combinations are considered. As in Section 4.1, the BPLS regression models are fitted to multivariate response vectors, whereas competitor methods each involve multiple univariate response fits. In addition to the competitor methods used in Section 4.1, we also consider the frequentist probabilistic partial least squares model (PPLS, el Bouhaddani et al. (2018)) fitted to the multivariate response. In PCR we identify the components with significant regression coefficients, and in BPLS regression models, we let \widehat{Q}_{y_r} be the effective dimension of the regression of response element y_r ($r = 1, \dots, R$), that is, $\widehat{Q}_{y_r} = \sum_{q=1}^{\infty} \mathbb{I}_{\{|\hat{c}_{r,q}| \geq 0.05c^*\}}$, where $\hat{c}_{r,q}$ are the posterior mean estimates of the entries of C and $c^* = \max_{r,q} |\hat{c}_{r,q}|$. Performance metrics are given as empirical averages of the RMSEP across the four fold combinations, and we also report the estimates of the effective dimensionality of the data (i.e., number of latent variables or components) for each method as an average over the four folds.

As shown in Table 2, in the benchmark grain mash NIR spectral data, all three of the BPLS-based methods outperform the existing standard approaches in terms of prediction accuracy, with the L-BPLS approach performing optimally. The BPLS methods had lower variation in the \widehat{Q}_{y_r} estimates, which suggests that the inaccuracy for other methods here may be due to the incorrect choice of the latent dimension. A comparison of BPLS methods to

TABLE 2

Accuracy of predictions of the standardised glucose and ethanol traits under the proposed BPLS and existing methods on the benchmark grain mash NIR spectral data. Estimates are given as “mean (standard deviation)” over the four test folds, with bold font indicating optimal performance. *Due to numerical instability, \hat{Q}_y was picked manually

	Glucose		Ethanol	
	RMSEP	\hat{Q}_y	RMSEP	\hat{Q}_y
BPLS	0.358 (0.045)	8.5 (0.6)	0.061 (0.009)	7.5 (1.3)
ss-BPLS	0.395 (0.041)	9.0 (1.4)	0.061 (0.005)	7.3 (0.5)
L-BPLS	0.353 (0.053)	14.0 (0.8)	0.060 (0.008)	8.3 (0.5)
PLS	0.422 (0.048)	8.3 (1.3)	0.072 (0.012)	10.3 (1.3)
PLS-1s	0.447 (0.043)	10.8 (3.6)	0.066 (0.009)	10.5 (1.7)
sPLS	0.442 (0.065)	16.0 (3.7)	0.070 (0.008)	13.0 (2.3)
PCR	0.440 (0.038)	18.0 (1.2)	0.069 (0.006)	22.3 (2.5)
LASSO	0.428 (0.018)	–	0.080 (0.012)	–
RIDGE	0.727 (0.033)	–	0.264 (0.031)	–
PPLS	0.923 (0.064)	*1.0 (0.0)	0.491 (0.038)	*1.0 (0.0)

standard Bayesian factor analysis approaches on the NIR dataset can be found in Appendix F.1 of the Supplementary Material. Additionally, Figure 4 illustrates point predictions and 95% prediction intervals from the L-BPLS model for 40 random test samples in one test fold of the NIR grain spectral data. The natural provision of coherent prediction intervals under the BPLS approach, particularly in the case of glucose (Figure 4 (left)), provides richer inference, enabling more informed decision-making. For example, as noted in Liebmann, Friedl and Varmuza (2009), it is desirable to dilute the grain mash to achieve a glucose concentration of $\leq 1\text{g/L}^{-1}$ for fermentation. If NIR spectral measurements are used instead of more expensive gold standard tests, there will be uncertainty in the final glucose predictions. Treating these predictions as infallible could lead to grain mash samples which have too high a concentration of glucose for processing purposes. To mitigate this, the processor could in-




FIG. 4. Point predictions and 95% prediction intervals of glucose and ethanol content from the L-BPLS model in one test fold of the NIR grain spectra dataset; 40 random test samples are given in increasing order of respective trait quantities for clarity. In addition, true values and predictions from standard PLS are provided.

TABLE 3

Accuracy of predictions of the standardised heat stability, casein content, and RCT traits from the milk MIR spectral data using each of the proposed BPLS and existing methods. Estimates are given as “mean (standard deviation)” of the RMSEP over the four test folds, with bold font indicating optimal performance. *Due to numerical instability, \hat{Q}_y was picked manually

	Heat stability		Casein content		RCT	
	RMSEP	\hat{Q}_y	RMSEP	\hat{Q}_y	RMSEP	\hat{Q}_y
BBPLS	0.875 (0.172)	6.5 (2.5)	0.376 (0.027)	6.8 (1.5)	0.815 (0.139)	5.5 (1.9)
ss-BPLS	0.874 (0.164)	4.3 (0.5)	0.376 (0.031)	6.3 (1.3)	0.813 (0.141)	5.3 (1.0)
L-BPLS	0.823 (0.156)	15.8 (1.0)	0.360 (0.041)	15.8 (2.1)	0.782 (0.126)	18.3 (3.0)
PLS	0.896 (0.117)	5.5 (0.6)	0.395 (0.028)	4.5 (0.6)	0.863 (0.141)	5.5 (1.7)
PLS-1s	0.849 (0.125)	7.8 (1.5)	0.397 (0.022)	4.5 (0.6)	0.810 (0.105)	7.3 (1.3)
sPLS	0.879 (0.173)	14.5 (4.0)	0.376 (0.044)	8.5 (4.4)	0.801 (0.128)	14.3 (3.5)
PCR	0.877 (0.112)	11.3 (1.9)	0.402 (0.024)	4.8 (1.5)	0.824 (0.089)	14.3 (7.2)
LASSO	0.846 (0.121)	–	0.406 (0.024)	–	0.845 (0.125)	–
RIDGE	0.992 (0.145)	–	0.446 (0.027)	–	0.925 (0.108)	–
PPLS	0.987 (0.127)	*2.0 (0.0)	0.682 (0.082)	*2.0 (0.0)	0.970 (0.095)	*2.0 (0.0)

stead employ a more conservative dilution protocol using an upper quantile of the posterior predictive distribution as the glucose content prediction, thereby ensuring robustness.

5. Predicting milk traits from MIR and SERS spectral data. We use the developed BPLS regression methods to predict milk traits from the motivating spectral datasets described in Section 2. We use the same hyperparameter settings and MCMC diagnostics as in Section 4, and we use the same four-fold training-test approach and compare the different methods in the same way as detailed in Section 4.2.

5.1. Predicting technological and protein traits from MIR milk spectral data. The technological traits of both heat stability and RCT as well as a protein trait related to the casein content of the milk samples were jointly predicted using the proposed BPLS methods from the milk MIR spectral data (see Section 2.1). As shown in Table 3, the traits are predicted with similar accuracy under the proposed BPLS models and existing methods, with L-BPLS offering the most accurate point predictions for all three traits; further prediction accuracy results comparing BPLS regression to Bayesian factor analysis approaches appear in Appendix F.1 of the Supplementary Material. While accurate milk trait prediction is important, BPLS models offer additional utility in the form of prediction intervals, as seen in Figure 5. The benefit of using the proposed BPLS methodology is exemplified in the heat stability predictions (Figure 5, left): where PLS produces very inaccurate point predictions for milk samples 26 and 40, the predictions from our model are substantially closer to the true trait value. Further, when the BPLS predictions are farther from the true trait values, the uncertainty is conveyed through the provision of coherent prediction intervals (e.g., milk sample 30).

The quantification of the uncertainty associated with the predictions can be particularly useful for milk processors segregating incoming milk streams for different value markets, for example, cheese, butter, or infant formula. Such segregation may prevent the undesirable introduction of inferior milk streams into a pool intended for a specific value product. The provision of reliable prediction intervals can give processors confidence that a particular standard has been met: if a given point prediction is above an acceptable threshold but there is substantial posterior predictive mass below it, the product is likely to fail to meet the standard. In such cases the processor may commission secondary analyses using a gold

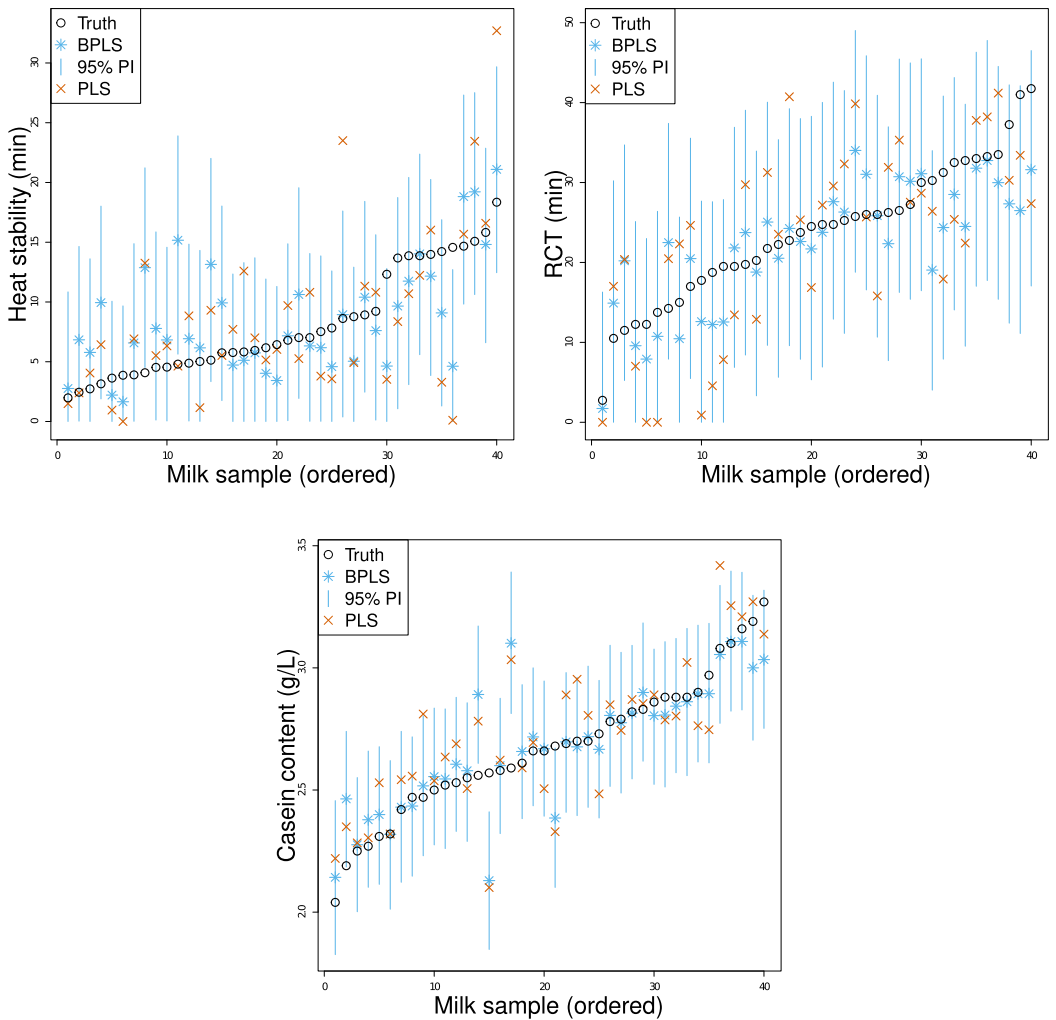


FIG. 5. Point predictions and 95% prediction intervals for the heat stability, RCT, and casein content traits under the L-BPLS model for one test fold of the MIR milk spectral data; 40 random test-set milk samples are given in increasing order of respective trait quantities. True trait values and predictions from standard PLS are also provided.

standard method on a sample of the milk pool, for example, measuring heat stability using an oil bath [Dumpler, Hupertz and Kulozik \(2020\)](#). Additionally, stringent food labelling policies require that the properties (e.g., nutritional values) of a given product reside within some acceptable margin of error. The provision of prediction intervals here would give processors confidence that a given product’s stated properties are sufficiently likely to correspond to what is being delivered. The BPLS approach is, therefore, beneficial to stakeholders in dairy production, aiding more informed management decisions by taking into account the uncertainty in the predictions, particularly when incorrect decisions may have health and/or financial repercussions.

As is good statistical practice, to validate the reported prediction intervals the prediction interval coverage is reported in [Table 4](#) for each of the three BPLS methods. Here L-BPLS interval widths are narrower for the heat stability trait compared to the other two methods—this could be due to the additional bias introduced by the LASSO-type penalty. All three methods have slightly overconfident intervals for the RCT trait, with BPLS exhibiting the strongest bias, suggesting potentially inadequate model fit; the coverage could be further improved by

TABLE 4

Medians, across four data folds, of the 95% prediction interval coverage estimates from the milk MIR spectral data. Estimates are based on the proportion of true test sample traits contained within the predicted intervals for the three BPLS models

	Heat stability	Casein content	RCT
BPLS	0.940	0.951	0.924
ss-BPLS	0.946	0.957	0.940
L-BPLS	0.912	0.951	0.940

adjusting the variable transformation. It is noteworthy that the estimated prediction interval coverage for casein was very close to the target 95%. Casein constitutes approximately 80% of total milk protein (Farrell et al. (2004)), and the total protein content of milk can be predicted very accurately with standard methods (Bonfatti, Di Martino and Carnier (2011)). For all three traits, the Gaussianity assumption appears appropriate, particularly for casein. The slightly too narrow prediction intervals for heat stability and RCT suggest a possibly more nuanced relationship between these technological traits of milk and the MIR spectra. Furthermore, there can be ambiguity about the ground-truth measurements of the traits themselves (e.g., measuring the heat stability of milk using an oil bath is arguably subjective), naturally limiting the accuracy of any statistical model.

Further analyses of the MIR data where two quarters of data are used for training and the remaining two are used for testing can be found in Appendix F.2 of the Supplementary Material.

5.2. *Predicting pH from SERS milk spectral data.* As described in Section 2.2, interest lies in predicting the pH of milk samples from their SERS intensity data. The suite of BPLS models and existing methods were used to predict the pH, with the BPLS models offering a clear improvement over the competitors (see Table 5); additional performance comparisons between BPLS models and Bayesian factor analysis approaches are given in Appendix F.1 of the Supplementary Material. The best-performing method in terms of RMSEP is again L-BPLS, which results in a 34.7% decrease in the error compared to the popular sPLS method.

TABLE 5

Accuracy of pH predictions from the milk SERS data under the BPLS and existing methods. Estimates are given as “mean (standard deviation)” over the four test folds, with bold font indicating optimal performance. *Due to numerical instability, \hat{Q}_y was picked manually

	pH	
	RMSEP	\hat{Q}_y
BPLS	0.627 (0.248)	6.0 (1.4)
ss-BPLS	0.644 (0.272)	4.3 (1.0)
L-BPLS	0.614 (0.232)	5.5 (0.6)
PLS	0.959 (0.330)	1.3 (0.5)
PLS-1s	0.959 (0.330)	*0.3 (1.5)
sPLS	0.941 (0.293)	4.0 (1.6)
PCR	1.024 (0.374)	0.3 (0.5)
LASSO	0.985 (0.471)	–
RIDGE	0.946 (0.453)	–
PPLS	0.783 (0.573)	*1.0 (0.0)

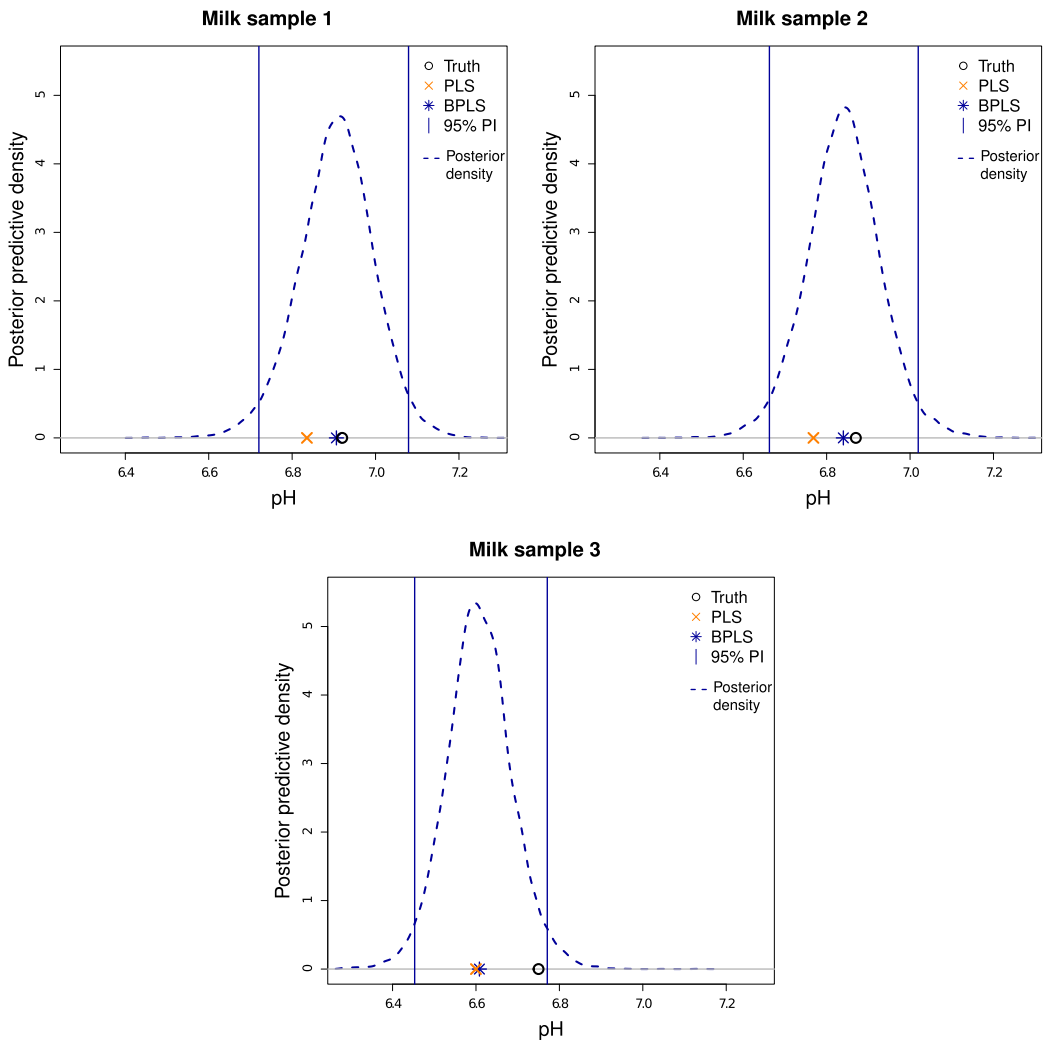


FIG. 6. Posterior predictive densities (blue, dashed) of pH for three test milk samples (with posterior predictive means and 95% prediction intervals marked) obtained from the L-BPLS model from one test fold of the milk pH SERS dataset, along with standard PLS predictions.

This is considered a particularly challenging data setting as $N = 11 \ll P = 1733$; the BPLS methods demonstrated particularly strong performance in such settings in Section 4.

Figure 6 illustrates the posterior predictive densities of pH, which are naturally available under the proposed L-BPLS model, for three milk samples from one test fold of the SERS data; posterior predictive means and 95% prediction intervals are highlighted, as is the prediction from PLS regression. The L-BPLS predictions were more accurate than the PLS regression approach and, even when the predictions from L-BPLS were not close to the gold standard value, the latter resided within the 95% prediction interval (e.g., milk sample 3, Figure 6 (right)).

Access to posterior predictive distributions for milk sample traits enables the end-users (here milk producers or processors) to make probabilistic statements about the milk pH, especially in relation to milk spoilage and animal health. A point prediction indicating a low pH could be an indication of mastitis (Kandeel et al. (2019)). However, the presence of high uncertainty in the form of a wide prediction interval would dissuade immediate remedial action, minimising unnecessary medical interventions through the use of antibiotics with

downstream repercussions on antimicrobial resistance. Instead, the end user may pay more attention to the animal or may perform additional tests (e.g., California mastitis test; Schalm and Noorlander (1957)).

6. Discussion. We address the problem statement in the global agri-food industry in relation to the simultaneous prediction of multivariate phenotype vectors using high-dimensional spectral data, specifically in the context of dairy production. There is a need for accurate and informative milk trait prediction methodologies which can utilise the spectral data that are routinely generated, as in the Irish agri-food sector. To achieve this, we introduce a probabilistic Gaussian latent-variable model, BPLS, which emulates the popular partial least squares regression algorithm, along with two variants of the BPLS model inspired by sparse regression methods. Posterior inference for the suite of BPLS methods was carried out through a MCMC algorithm. While numerical experiments explored the performance of the BPLS models, improved prediction accuracy of milk traits from the motivating MIR and SERS spectral data sets was achieved under the proposed BPLS models compared to the industry-standard PLS approaches. Importantly, the BPLS methods bestow additional benefits, such as the coherent provision of prediction intervals for the traits, which can assist in informed decision-making by milk producers and processors. They obviate the need for the user to select a number of latent components and facilitate the prediction of traits in a multivariate fashion. In preliminary analyses, when separate BPLS models were fitted to individual milk traits, the overall predictions worsened indicating a clear benefit from the multivariate modelling approach, which is facilitated by the BPLS's basis in a statistical model. This is in contrast to multivariate PLS-type models where picking Q can be particularly challenging—one may be forced to overfit for one trait to compensate for another; this is illustrated in Appendix F.3 of the Supplementary Material. The BPLS methods also performed well in the SERS data setting which is a “small N , large P ” scenario; the regularisation provided through the use of Bayesian priors allowed inference, even in extreme settings where the number of parameters was over a hundred times larger than the number of milk samples. In summary, treating PLS regression from a probabilistic viewpoint, drawing on Bayesian nonparametric priors to obviate model selection challenges, and incorporating sparse regression strategies results in a single, coherent BPLS regression modelling framework for prediction of dairy traits from spectral data, with uncertainty inherently provided.

Based on the thorough empirical comparisons, we found that the LASSO variant of BPLS had the most reliable prediction accuracy for the datasets considered here. The sparsity structure arising from the combination of Bayesian LASSO and shrinkage priors resulted in a model with very accurate posterior predictive means. Whilst L-BPLS was also robust in estimating the predictive variance in the synthetic datasets, as indicated by the coverage probability results, it underperformed in the MIR spectral data setting where ss-BPLS achieved more desirable coverage. As is often the case, no single model works best in all situations. When working with spectral data, L-BPLS may be more favourable if point prediction accuracy is more important than robust uncertainty quantification, but if the opposite is the case, ss-BPLS may be the more robust choice. Based on these results, the recommendation would be to fit the two models to the given spectral data problem and validate for a given objective (e.g., point prediction or coverage probability) prior to deployment.

While technological and protein milk traits were considered here, the BPLS framework does not rely on a particular type of milk trait. Indeed, BPLS methods could be used for predictions of quantities not directly related to milk such as cattle methane output (Coppa et al. (2022)), body energy status of cattle (Frizzarin et al. (2023)), cattle reproductive performance (Toledo-Alvarado et al. (2021)), or presence of disease (Denholm et al. (2020)). Equally, the BPLS models are not confined to mid-infrared or surface-enhanced Raman spectra from milk

samples and could be applied to problems involving ultraviolet, visible, or near-infrared spectral data (e.g., [Aernouts et al. \(2011\)](#), [Behkami et al. \(2019\)](#)). The underpinning of BPLS by a coherent probabilistic model would aid researchers in predicting a wide range of traits in the agri-food domain.

The BPLS models, as introduced, and many of their existing competitors assume independent and identically distributed observations, conditional on the model parameters. This can be an incorrect modelling assumption, particularly for problems encountered in the agri-food context. It is not uncommon for agri-food datasets to contain multiple samples relating to a single source; for example, in milk MIR spectral datasets, several milk samples typically come from one cow. The BPLS's probabilistic formulation elegantly lends itself to facilitate modifications to appropriately model such dependency structures; for example, it could allow for a hierarchical structure with random effects, appropriately accounting for first-order interactions. Such extensions are less feasible in existing, nonprobabilistic methods such as PLS.

Given the motivating application, in this work we focused on predicting real-valued traits; however, a great deal of interest lies in problems involving the prediction of binary responses from spectral data, for example, predicting a cow's pregnancy status ([Brand et al. \(2021\)](#)) or diet type ([Coppa et al. \(2021\)](#)). Generalised linear model-oriented extensions to the BPLS methods would naturally and appropriately allow for such categorical responses, with random effects again permitting faithful modelling of repeated measures data. Indeed, in [Frizzarin et al. \(2021\)](#), separate PLS methods were used to predict the RCT and heat stability traits from milk samples; samples were subsequently classified into groups based on high and low levels of the traits with PLS discriminant analysis ([Barker and Rayens \(2003\)](#)) relied on for the latter task. Both tasks could be simultaneously achieved using the same posterior predictive distributions from a single BPLS model fit.

Given the practical utility of the developed methods in dairy production, one potential concern is the scalability of the BPLS framework as the number of variables, P , grows. A naive R implementation of the Gibbs algorithm was sufficiently computationally efficient for the MIR and SERS data, with run times of 18 and 30 minutes respectively on a Dell Latitude 5430 laptop, equipped with a 4.40 GHz Intel core i5-1235U processor and 32 GB of RAM. In practice, regression models would not require frequent refitting as the collection of traits can be costly and takes time. More sophisticated future models with complex hierarchical structures would likely require approximate inference schemes to alleviate the computational burden that comes with the growing parameter space (e.g., [Hansen et al. \(2023\)](#)); however, caution should be taken to ensure a controlled approximation error. Relatedly, the BPLS framework does not strictly depend on the infinitely-dimensional nonparametric shrinkage prior. A finite-dimensional variant could be employed by performing model selection based on information criteria such as BIC-MCMC ([Frühwirth-Schnatter \(2011\)](#)); albeit, fitting multiple models would come at a steep computational cost.

Acknowledgments. The authors would like to thank the members of the VistaMilk Science Foundation Ireland (SFI) Research Center and, in particular, the members of the VistaMilk Spectroscopy Working Group for discussions that contributed to this work.

The authors are grateful to the Editors and the two anonymous reviewers for comments and suggestions that have materially improved the article.

Funding. This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) and the Department of Agriculture, Food and Marine on behalf of the Government of Ireland under grant number (16/RC/3835) and the SFI Insight Research Centre under grant number (SFI/12/RC/2289_P2).

SUPPLEMENTARY MATERIAL

Technical appendix (DOI: [10.1214/24-AOAS1947SUPPA](https://doi.org/10.1214/24-AOAS1947SUPPA); .pdf). Document containing details on: (i) transformation of positive response variables; (ii) inference using MCMC algorithms; (iii) illustration of how the model fit achieves desired shrinkage; (iv) additional details and results relating to the posterior predictive distributions of the responses; (v) additional simulation results; and (vi) additional benchmark and milk trait data results.

Code (DOI: [10.1214/24-AOAS1947SUPPB](https://doi.org/10.1214/24-AOAS1947SUPPB); .zip). BPLS models and inference were implemented in R.

REFERENCES

- AERNOUTS, B., POLSHIN, E., LAMMERTYN, J. and SAEYS, W. (2011). Visible and near-infrared spectroscopic analysis of raw milk for cow health monitoring: Reflectance or transmittance? *J. Dairy Sci.* **94** 5315–5329. <https://doi.org/10.3168/jds.2011-4354>
- BARKER, M. and RAYENS, W. (2003). Partial least squares for discrimination. *Journal of Chemometrics: A Journal of the Chemometrics Society* **17** 166–173.
- BEEBE, K. R., PELL, R. J. and SEASHOLTZ, M. B. (1998). *Chemometrics: A Practical Guide*. Wiley, New York; Chichester.
- BEHKAMI, S., ZAIN, S. M., GHOLAMI, M. and KHIR, M. F. A. (2019). Classification of cow milk using artificial neural network developed from the spectral data of single- and three-detector spectrophotometers. *Food Chem.* **294** 309–315. <https://doi.org/10.1016/j.foodchem.2019.05.060>
- BHATTACHARYA, A. and DUNSON, D. B. (2011). Sparse Bayesian infinite factor models. *Biometrika* **98** 291–306. [MR2806429 https://doi.org/10.1093/biomet/asr013](https://doi.org/10.1093/biomet/asr013)
- BISHOP, C. M. (2006). *Pattern Recognition and Machine Learning. Information Science and Statistics*. Springer, New York. [MR2247587 https://doi.org/10.1007/978-0-387-45528-0](https://doi.org/10.1007/978-0-387-45528-0)
- BONFATTI, V., DI MARTINO, G. and CARNIER, P. (2011). Effectiveness of mid-infrared spectroscopy for the prediction of detailed protein composition and contents of protein genetic variants of individual milk of Simmental cows. *J. Dairy Sci.* **94** 5776–5785. <https://doi.org/10.3168/jds.2011-4401>
- BOULESTEIX, A.-L. and STRIMMER, K. (2007). Partial least squares: A versatile tool for the analysis of high-dimensional genomic data. *Brief. Bioinform.* **8** 32–44. <https://doi.org/10.1093/bib/bbl016>
- BRAND, W., WELLS, A. T., SMITH, S. L., DENHOLM, S. J., WALL, E. and COFFEY, M. P. (2021). Predicting pregnancy status from mid-infrared spectroscopy in dairy cow milk using deep learning. *J. Dairy Sci.* **104** 4980–4990. <https://doi.org/10.3168/jds.2020-18367>
- CHOI, J., ZOU, H. and OEHLERT, G. (2010). A penalized maximum likelihood approach to sparse factor analysis. *Stat. Interface* **3** 429–436. [MR2754740 https://doi.org/10.4310/SII.2010.v3.n4.a1](https://doi.org/10.4310/SII.2010.v3.n4.a1)
- CHUN, H. and KELES, S. (2010). Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **72** 3–25. [MR2751241 https://doi.org/10.1111/j.1467-9868.2009.00723.x](https://doi.org/10.1111/j.1467-9868.2009.00723.x)
- CHUNG, D., CHUN, H. and KELES, S. (2019). spls: Sparse Partial Least Squares (SPLS) regression and classification. R package version 2.2-3.
- COPPA, M., MARTIN, B., HULIN, S., GUILLEMIN, J., GAUZENTES, J. V., PECOU, A. and ANDUEZA, D. (2021). Prediction of indicators of cow diet composition and authentication of feeding specifications of protected designation of origin cheese using mid-infrared spectroscopy on milk. *J. Dairy Sci.* **104** 112–125. <https://doi.org/10.3168/jds.2020-18468>
- COPPA, M., VANLIERDE, A., BOUCHON, M., JURQUET, J., MUSATI, M., DEHARENG, F. and MARTIN, C. (2022). Methodological guidelines: Cow milk mid-infrared spectra to predict reference enteric methane data collected by an automated head-chamber system. *J. Dairy Sci.* **105** 9271–9285. <https://doi.org/10.3168/jds.2022-21890>
- DE JONG, S. (1993). SIMPLS: An alternative approach to partial least squares regression. *Chemom. Intell. Lab. Syst.* **18** 251–263.
- DE MARCHI, M., TOFFANIN, V., CASSANDRO, M. and PENASA, M. (2014). Invited review: Mid-infrared spectroscopy as phenotyping tool for milk traits1. *J. Dairy Sci.* **97** 1171–1186. <https://doi.org/10.3168/jds.2013-6799>
- DENHOLM, S., BRAND, W., MITCHELL, A., WELLS, A., KRZYZELEWSKI, T., SMITH, S., WALL, E. and COFFEY, M. (2020). Predicting bovine tuberculosis status of dairy cows from mid-infrared spectral data of milk using deep learning. *J. Dairy Sci.* **103** 9355–9367.
- DIACONIS, P. and FREEDMAN, D. (1986). On the consistency of Bayes estimates. *Ann. Statist.* **14** 1–67. [MR0829555 https://doi.org/10.1214/aos/1176349830](https://doi.org/10.1214/aos/1176349830)

- DUMPLER, J., HUPPERTZ, T. and KULOZIK, U. (2020). Invited review: Heat stability of milk and concentrated milk: Past, present, and future research objectives. *J. Dairy Sci.* **103** 10986–11007. <https://doi.org/10.3168/jds.2020-18605>
- DURANTE, D. (2017). A note on the multiplicative gamma process. *Statist. Probab. Lett.* **122** 198–204. <https://doi.org/10.1016/j.spl.2016.11.014>
- EL BOUHADDANI, S., UH, H.-W., HAYWARD, C., JONGBLOED, G. and HOUWING-DUISTERMAAT, J. (2018). Probabilistic partial least squares model: Identifiability, estimation and application. *J. Multivariate Anal.* **167** 331–346. <https://doi.org/10.1016/j.jmva.2018.05.009>
- EL BOUHADDANI, S., UH, H.-W., JONGBLOED, G. and HOUWING-DUISTERMAAT, J. (2022). Statistical integration of heterogeneous omics data: Probabilistic two-way partial least squares (PO2PLS). *J. R. Stat. Soc. Ser. C. Appl. Stat.* **71** 1451–1470. <https://doi.org/10.1111/rssc.12583>
- FARRELL, H. M. J., JIMENEZ-FLORES, R., BLECK, G. T., BROWN, E. M., BUTLER, J. E., CREAMER, L. K., HICKS, C. L., HOLLAR, C. M., NG-KWAI-HANG, K. F. et al. (2004). Nomenclature of the proteins of cows' milk—sixth revision. *J. Dairy Sci.* **87** 1641–1674.
- FILZMOSER, P., GSCHWANDTNER, M. and TODOROV, V. (2012). Review of sparse methods in regression and classification with application to chemometrics. *J. Chemom.* **26** 42–51.
- FILZMOSER, P. and VARMUZA, K. (2017). Chemometrics: Multivariate statistical analysis in chemometrics. R package version 1.4.2.
- FRANK, I. E. and FRIEDMAN, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35** 109–135. <https://doi.org/10.1080/00401706.1993.10485033>
- FRIZZARIN, M., GORMLEY, I., BERRY, D., MURPHY, T., CASA, A., LYNCH, A. and MCPARLAND, S. (2021). Predicting cow milk quality traits from routinely available milk spectra using statistical machine learning methods. *J. Dairy Sci.* **104** 7438–7447.
- FRIZZARIN, M., GORMLEY, I. C., BERRY, D. P. and MCPARLAND, S. (2023). Estimation of body condition score change in dairy cows in a seasonal calving pasture-based system using routinely available milk mid-infrared spectra and machine learning techniques. *J. Dairy Sci.* **106** 4232–4244. <https://doi.org/10.3168/jds.2022-22394>
- FRÜHWIRTH-SCHNATTER, S. (2011). Dealing with label switching under model uncertainty. In *Mixtures: Estimation and Applications*. Wiley Ser. Probab. Stat. 213–239. Wiley, Chichester. <https://doi.org/10.1002/9781119995678.ch10>
- FRÜHWIRTH-SCHNATTER, S. and LOPES, H. F. (2014). Parsimonious Bayesian factor analysis when the number of factors is unknown. Technical Report No. 345. Insper Working Paper.
- GEORGE, E. I. and MCCULLOCH, R. E. (1993). Variable selection via Gibbs sampling. *J. Amer. Statist. Assoc.* **88** 881–889.
- GWEE, X. Y., GORMLEY, I. C. and FOP, M. (2023). A latent shrinkage position model for binary and count network data. *Bayesian Anal.* 1–29. <https://doi.org/10.1214/23-BA1403>
- HANSEN, B., AVALOS-PACHECO, A., RUSSO, M. and VITO, R. D. (2023). Fast variational inference for Bayesian factor analysis in single and multi-study settings.
- HOERL, A. E. and KENNARD, R. W. (1970). Ridge regression: Applications to nonorthogonal problems. *Technometrics* **12** 69–82.
- HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.* **97** 1090–1098. <https://doi.org/10.1198/016214502388618906>
- HUBERT, M. and BRANDEN, K. V. (2003). Robust methods for partial least squares regression. *Journal of Chemometrics: A Journal of the Chemometrics Society* **17** 537–549.
- KANDEEL, S. A., MEGAHED, A. A., EBEID, M. H. and CONSTABLE, P. D. (2019). Ability of milk pH to predict subclinical mastitis and intramammary infection in quarters from lactating dairy cattle. *J. Dairy Sci.* **102** 1417–1427. <https://doi.org/10.3168/jds.2018-14993>
- KOURTI, T. (2002). Process analysis and abnormal situation detection: From theory to practice. *IEEE Control Syst. Mag.* **22** 10–25.
- LI, S., GAO, J., NYAGILO, J. O. and DAVE, D. P. (2010). Eigenspectra, a robust regression method for multiplexed Raman spectra analysis. In *2010 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* 525–530. IEEE.
- LIEBMANN, B., FRIEDL, A. and VARMUZA, K. (2009). Determination of glucose and ethanol in bioethanol production by near infrared spectroscopy and chemometrics. *Anal. Chim. Acta* **642** 171–178. <https://doi.org/10.1016/j.aca.2008.10.069>
- LILAND, K. H., MEVIK, B.-H. and WEHRENS, R. (2022). pls: Partial least squares and principal component regression. R package version 2.8-1.
- LINDGREN, F., GELADI, P. and WOLD, S. (1993). The kernel algorithm for PLS. *J. Chemom.* **7** 45–59.
- MCDERMOTT, A., VISENTIN, G., MARCHI, M. D., BERRY, D. P., FENELON, M. A., O'CONNOR, P. M., KENNY, O. A. and MCPARLAND, S. (2016). Prediction of individual milk proteins including free amino acids

- in bovine milk using mid-infrared spectroscopy and their correlations with milk processing characteristics. *J. Dairy Sci.* **99** 3171–3182. <https://doi.org/10.3168/jds.2015-9747>
- MCPARLAND, D., PHILLIPS, C. M., BRENNAN, L., ROCHE, H. M. and GORMLEY, I. C. (2017). Clustering high-dimensional mixed data to uncover sub-phenotypes: Joint analysis of phenotypic and genotypic data. *Stat. Med.* **36** 4548–4569. [MR3731239 https://doi.org/10.1002/sim.7371](https://doi.org/10.1002/sim.7371)
- MURPHY, K., VIROLI, C. and GORMLEY, I. C. (2020). Infinite mixtures of infinite factor analysers. *Bayesian Anal.* **15** 937–963. [MR4132655 https://doi.org/10.1214/19-BA1179](https://doi.org/10.1214/19-BA1179)
- NAIK, P. and TSAI, C.-L. (2000). Partial least squares estimator for single-index models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **62** 763–771. [MR1796290 https://doi.org/10.1111/1467-9868.00262](https://doi.org/10.1111/1467-9868.00262)
- OVASKAINEN, O., ABREGO, N., HALME, P. and DUNSON, D. (2016). Using latent variable models to identify large networks of species-to-species associations at different spatial scales. *Methods Ecol. Evol.* **7** 549–555.
- PARK, T. and CASELLA, G. (2008). The Bayesian lasso. *J. Amer. Statist. Assoc.* **103** 681–686. [MR2524001 https://doi.org/10.1198/016214508000000337](https://doi.org/10.1198/016214508000000337)
- PRESS, S. J. and SHIGEMASU, K. (1989). Bayesian inference in factor analysis. In *Contributions to Probability and Statistics* 271–287. Springer, New York. [MR1024336](https://doi.org/10.1007/978-1-4613-0263-9_15)
- ROCKOVA, V. (2023). Adaptive Bayesian predictive inference. Preprint. Available at [arXiv:2309.02369](https://arxiv.org/abs/2309.02369).
- ROUSSEAU, J. (2016). On the frequentist properties of Bayesian nonparametric methods. *Annu. Rev. Stat. Appl.* **3** 211–231.
- SCHALM, O. W. and NOORLANDER, D. O. (1957). Experiments and observations leading to development of the California mastitis test. *J. Am. Vet. Med. Assoc.* **130** 199–204.
- SONG, X.-Y. and LEE, S.-Y. (2001). Bayesian estimation and test for factor analysis model with continuous and polytomous data in several populations. *Br. J. Math. Stat. Psychol.* **54** 237–263.
- R CORE TEAM (2023). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](https://doi.org/10.2307/2346178)
- TIPPING, M. E. and BISHOP, C. M. (1999). Probabilistic principal component analysis. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **61** 611–622. [MR1707864 https://doi.org/10.1111/1467-9868.00196](https://doi.org/10.1111/1467-9868.00196)
- TOLEDO-ALVARADO, H., PÉREZ-CABAL, M. A., TEMPELMAN, R. J., CECCHINATO, A., BITTANTE, G., DE LOS CAMPOS, G. and VAZQUEZ, A. I. (2021). Association between days open and milk spectral data in dairy cows. *J. Dairy Sci.* **104** 3665–3675. <https://doi.org/10.3168/jds.2020-19031>
- TRYGG, J. and WOLD, S. (2002). Orthogonal projections to latent structures (O-PLS). *Journal of Chemometrics: A Journal of the Chemometrics Society* **16** 119–128.
- TRYGG, J. and WOLD, S. (2003). O2-PLS, a two-block (X–Y) latent variable regression (LVR) method with an integral OSC filter. *J. Chemom.* **17** 53–64.
- URBAS, S., LOVERA, P., DALY, R., O’RIORDAN, A., BERRY, D. and GORMLEY, I. C. (2024). Supplement to “Predicting milk traits from spectral data using Bayesian probabilistic partial least squares regression.” <https://doi.org/10.1214/24-AOAS1947SUPPA>, <https://doi.org/10.1214/24-AOAS1947SUPPB>
- VAN DER VOET, H. (1994). Comparing the predictive accuracy of models using a simple randomization test. *Chemom. Intell. Lab. Syst.* **25** 313–323.
- VATS, D. and KNUDSON, C. (2021). Revisiting the Gelman-Rubin diagnostic. *Statist. Sci.* **36** 518–529. [MR4323050 https://doi.org/10.1214/20-sts812](https://doi.org/10.1214/20-sts812)
- VIDAURRE, D., VAN GERVEN, M. A. J., BIELZA, C., LARRAÑAGA, P. and HESKES, T. (2013). Bayesian sparse partial least squares. *Neural Comput.* **25** 3318–3339. [MR3154315 https://doi.org/10.1162/NECO_a_00524](https://doi.org/10.1162/NECO_a_00524)
- VISENTIN, G., MCDERMOTT, A., MCPARLAND, S., BERRY, D. P., KENNY, O. A., BRODKORB, A., FENELON, M. A. and MARCHI, M. D. (2015). Prediction of bovine milk technological traits from mid-infrared spectroscopy analysis in dairy cows. *J. Dairy Sci.* **98** 6620–6629. <https://doi.org/10.3168/jds.2015-9323>
- WILLIAMS, A., FLYNN, K. J., XIA, Z. and DUNSTAN, P. R. (2016). Multivariate spectral analysis of pH SERS probes for improved sensing capabilities. *J. Raman Spectrosc.* **47** 819–827.
- WOLD, H. (1973). Nonlinear iterative partial least squares (NIPALS) modelling: Some current developments. In *Multivariate Analysis, III (Proc. Third Internat. Sympos., Wright State Univ., Dayton, Ohio, 1972)* 383–407. Academic Press, New York. [MR0343487](https://doi.org/10.1007/978-1-4613-0263-9_15)
- ZHENG, J., SONG, Z. and GE, Z. (2016). Probabilistic learning of partial least squares regression model: Theory and industrial applications. *Chemom. Intell. Lab. Syst.* **158** 80–90.