



HAL
open science

Multivariate Bayesian Last Layer for Regression: Uncertainty Quantification and Disentanglement

Han Wang, Eiji Kawasaki, Guillaume Damblin, Geoffrey Daniel

► **To cite this version:**

Han Wang, Eiji Kawasaki, Guillaume Damblin, Geoffrey Daniel. Multivariate Bayesian Last Layer for Regression: Uncertainty Quantification and Disentanglement. 2024. cea-04567052

HAL Id: cea-04567052

<https://cea.hal.science/cea-04567052>

Preprint submitted on 2 May 2024

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Multivariate Bayesian Last Layer for Regression: Uncertainty Quantification and Disentanglement

Han WANG^{1,@}

Eiji KAWASAKI^{1,≈}

Guillaume DAMBLIN^{2,≈}

Geoffrey DANIEL^{2,≈}

¹Université Paris Saclay - CEA LIST, F-91120 Palaiseau Cedex, France

²Université Paris-Saclay, CEA, Service de Génie Logiciel pour la Simulation, 91191, Gif-sur-Yvette, France

@ Corresponding author: han.wang@cea.fr

≈ Equal contribution

Abstract

We present new Bayesian Last Layer models in the setting of multivariate regression under heteroscedastic noise, and propose an optimization algorithm for parameter learning. Bayesian Last Layer combines Bayesian modelling of the predictive distribution with neural networks for parameterization of the prior, and has the attractive property of uncertainty quantification with a single forward pass. The proposed framework is capable of disentangling the aleatoric and epistemic uncertainty, and can be used to transfer a canonically trained deep neural network to new data domains with uncertainty-aware capability.

1 INTRODUCTION

In the Deep Neural Network (DNN) literature, a deep interest has emerged for uncertainty quantification (UQ) of predictions [Gawlikowski et al., 2023]. Indeed, two types of uncertainty can affect the precision of DNN responses: aleatoric and epistemic uncertainty. The former is related to the possible randomness of the underlying physical process and the imperfection of measurement systems, while the latter concerns the modelling assumptions of DNN linked to limited data, in particular its architecture and the consequent approximation capacity. The safety of DNN deployment in industries is therefore conditioned upon providing trustworthy predictions, which can be embodied by confidence intervals representing the sum of aleatoric and epistemic uncertainties to be quantified.

To this end, several approaches have been put forward, including *e.g.* Bayesian DNN (BNN), Monte Carlo Dropout and Ensembles. More recently, another line of research that combines the power of Bayesian modelling and the approximation capacity of DNN, has been shown to drastically simplify its computation compared to the aforemen-

tioned computationally expensive methods. As remarkable developments in this direction, Deep Evidential Regression (DER) [Amini et al., 2020] employs DNNs uniquely for parameterizing the predictive distribution and enjoys closed form formulae for disentangling the epistemic and aleatoric uncertainties. Multivariate extension of DER has been proposed in [Meinert and Lavin, 2021], and the issue of non-identifiability of parameters and the reason of the effectiveness were analysed in [Meinert et al., 2023]. In this work we focus on Bayesian Last Layer (BLL), another hybrid method enjoying similar properties as DER, but with finer control in the Bayesian modelling part and results in different theoretical and algorithmic aspects.

Bayesian Last Layer BLL can be thought of as the technique of “randomizing” the last layer of a DNN, by treating its parameters as random variables. In the general case, BLL is simply a restricted type of BNN [Sharma et al., 2023]. However, when the last layer is linear, BLL becomes a straightforward generalization of Bayesian Linear Regression (BLR), where the covariate is replaced by the output of the penultimate layer, considered as a nonlinear feature transformation. Initially presented in Lázaro-Gredilla and R. [2010], BLL was, for example, used in Snoek et al. [2015] for Bayesian optimization of noisy black-box functions. In Ober and Rasmussen [2019], the performances of several techniques for joint estimation of the features (hidden layers) and parameters of the last layer were compared with one another on various datasets. Next, Watson et al. [2021] has developed a method using the model derivatives to avoid overconfident predictions of BLL outside the data distribution. Recently, Fiedler and Lucia [2023] and Harrison et al. [2023] extend BLL to multivariate outputs, but remain focused, like previous contributions, on homoscedastic data.

Contributions In this article we present new BLL models in the context of multivariate regression under heteroscedastic noise. To our knowledge, this is the first effort by BLL methods addressing both questions in a unified framework, thus making BLL capable of disentangling the aleatoric

and epistemic uncertainties in higher dimensions, as DER does. Our theoretical analysis is based on matrix-variate distributions which release the full degree of freedom in the marginal likelihood (or *evidence*), and open new paths toward efficient optimization algorithms. A possible application of the proposed framework is Transfer Learning, in which we transform a DNN trained canonically into a uncertainty-aware device up on new data domains.

Organization of the paper Starting in Section 2 by extending the classical homoscedastic BLR theory to obtain analytical formulae of UQ, we then establish in Section 3 BLL models and propose MAP as means of parameter estimation. Due to the loss of sample independence in the evidence, unbiased gradient estimate is no longer guaranteed on mini-batches, making the optimization procedure for BLL more challenging. As a solution, ELBO-EM, a numerically efficient variant of EM and well adapted for transfer learning, is proposed in Section 4. The proposed contributions are tested and validated in Section 5 through various numerical experiments.

2 BAYESIAN LINEAR REGRESSION REVISITED

For the development of BLL, we need to extend the classical Bayesian linear regression (BLR) model to the multivariate and heteroscedastic case. Many references exist on BLR (*e.g.* [Box and Tiao, 2011], [Bishop and Tipping, 2003], [Murphy, 2012], [Minka, 2010]) but most only concerns the homoscedastic case. Our presentation below is self-contained and concise, and we refer the reader to Appendix B for technical details. For quick reference, frequently used symbols are listed in Table 1.

Matrix variate Distributions A consequence of the multivariate regression setting we work in is the use of matrix variate distributions, including the Normal distribution \mathcal{N} , the Student’s T distribution \mathcal{T} , and the Inverse Wishart distribution \mathcal{W}^{-1} . For their definitions the conventions in [Gupta and Nagar, 1999] are followed and we collect the details in Appendix A.

2.1 MULTIVARIATE BASIS REGRESSION MODELS

Consider the multivariate linear regression model with basis between the input $x \in \mathbb{R}^n$ and the output $y \in \mathbb{R}^p$:

$$y = A\phi(x) + \varepsilon(x) \quad (1)$$

Here $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a nonlinear basis function, $A \sim p \times m$ is the regression matrix, and $\varepsilon(x)$ is a p -dimensional centered random noise depending on x . A bias term can be implicitly included by augmenting $\phi(x)$ with a 1. In the

Table 1: Symbols used in this paper.

Symbols	Definition
M, K, V	parameters of prior
Σ, ν	parameters of prior
ν_p	$\nu - p - 1$
X, Y	data matrices
D or $\sigma^2(X)$	$\text{diag}[-\sigma^2(x_i)-]$
Φ	$[-\phi(x_i)-]$
Ω	$D + \Phi^\top K \Phi$
S_{yx}	$YD^{-1}\Phi^\top + MK^{-1}$
S_{xx}	$K^{-1} + \Phi D^{-1}\Phi^\top$
S_{yy}	$YD^{-1}Y^\top + MK^{-1}M^\top$
$S_{y x}$	$S_{yy} - S_{yx}S_{xx}^{-1}S_{yx}^\top$
E	$Y - M\Phi$

following we study a restricted but more tractable version, in which $\varepsilon(x) := \sigma(x)\varepsilon$ is an isotropic heteroscedastic noise obtained from multiplying a *baseline* Gaussian white noise $\varepsilon \sim \mathcal{N}_{\text{vec}}(0, V)$ ¹ by a scalar $\sigma(x) > 0$, with the covariance matrix $V \sim p \times p$ possibly unknown. The model becomes homoscedastic when $\sigma(x)$ is constant.

Our modeling choice is not random. In fact, a scalar scaling $\sigma(x) > 0$ is possibly the only way to make the classical multivariate BLR a heteroscedastic theory. Actually, the compact matrix-based formulae of the multivariate BLR will break down if a matrix scaling $\sigma(x) \sim p \times p$ was chosen, which could also considerably increase the computation time and memory usage.

Matrix formulation Independent observations $\{(y_i, x_i)\}_{i=1:N}$ are realized with some fixed A and a underlying noise $\varepsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}_{\text{vec}}(0, V)$. Let the matrices $Y \sim p \times N$ and $X \sim n \times N$ be the column stack of data samples, and similarly for the feature matrix $\Phi := [-\phi(x_i)-]$ of dimension $m \times N$. The linear model in matrix form reads

$$Y = A\Phi + \mathcal{E}(X) \quad (2)$$

with the noise matrix $\mathcal{E}(X) := [-\sigma(x_i)\varepsilon_i-]$. For the ease of notation, hereafter we omit the conditional dependence on X in any probability distribution involving Y .

Following [Minka, 2010], in the simple case that V is known we specify the prior of A as

$$A \sim \mathcal{N}(M, V, K),$$

for a mean matrix $M \sim p \times m$ and a precision matrix $K \sim m \times m$. The covariance V here is the same as that of the baseline noise, and this is a well known technique in

¹The subscript indicates this as a vector-valued but not a matrix-valued distribution.

BLR to enable the conjugate prior. The likelihood of $Y|A$ is then

$$Y|A \sim \mathcal{N}(A\Phi, V, D).$$

Proposition 2.1. *Consider the regression model (2) with V known. The evidence (or marginal likelihood) is*

$$Y \sim \mathcal{N}(M\Phi, V, \Omega). \quad (3)$$

The posterior is

$$A|Y \sim \mathcal{N}(S_{y\mathbf{x}}S_{\mathbf{xx}}^{-1}, V, S_{\mathbf{xx}}^{-1}).$$

Given L new data points $X_* \sim n \times L$, the posterior predictive is

$$Y_*|Y \sim \mathcal{N}(S_{y\mathbf{x}}S_{\mathbf{xx}}^{-1}\Phi_*, V, D_* + \Phi_*^\top S_{\mathbf{xx}}^{-1}\Phi_*)$$

where $\Phi_* \sim m \times N$ and $D_* \sim L \times L$ are constructed in the same way as Φ, D .

For the case V is unknown and we specify an Inverse Wishart prior as

$$V \sim \mathcal{W}^{-1}(\Sigma, \nu) \quad (4)$$

with a degree of freedom $\nu \geq 2p + 1$ and a scale matrix Σ . The joint distribution of $[V \ A \ Y]$ can then be composed from $Y|(A, V)$, $A|V$ and V as

$$[V \ A \ Y] \sim \mathcal{N}(A\Phi, V, D) \times \mathcal{N}(M, V, K) \times \mathcal{W}^{-1}(\Sigma, \nu).$$

The following result is parallel to Proposition 2.1 and is generally considered more robust to outliers due to the heavy tail property of T distribution [Lange et al., 1989].

Proposition 2.2. *Consider the regression model (2) with V follows (4). The evidence is*

$$Y \sim \mathcal{T}(M\Phi, \Sigma, \Omega, \nu - 2p). \quad (5)$$

The posterior $[A \ V]|Y$ follows a compound Normal-Inverse Wishart distribution with the marginals

$$\begin{aligned} A|Y &\sim \mathcal{T}(S_{y\mathbf{x}}S_{\mathbf{xx}}^{-1}, \Sigma + S_{y|\mathbf{x}}, S_{\mathbf{xx}}^{-1}, \nu + N - 2p) \\ V|Y &\sim \mathcal{W}^{-1}(\Sigma + S_{y|\mathbf{x}}, \nu + N) \end{aligned}$$

Given L new data points $X_* \sim n \times L$, the posterior predictive is

$$Y_*|Y \sim \mathcal{T}(S_{y\mathbf{x}}S_{\mathbf{xx}}^{-1}\Phi_*, \Sigma + S_{y|\mathbf{x}}, D_* + \Phi_*^\top S_{\mathbf{xx}}^{-1}\Phi_*, \nu + N - 2p)$$

where $\Phi_* \sim m \times N$ and $D_* \sim L \times L$ are constructed in the same way as Φ, D .

Connections with existing models Up to some naming conventions, these two models cover the evidences and predictive distributions used in most existing BLL and DER methods (e.g. [Harrison et al., 2023] and [Fiedler and Lucia, 2023]), [Meinert and Lavin, 2021]). For example, the evidence term used in the multivariate DER can be recovered from the multivariate and non-regression (i.e. $x \equiv 1$) version of (5) as

$$y \sim \mathcal{T}_{\text{vec}}(m, (1+k)\Sigma, \nu - 2p) \quad (7)$$

where the parameters $m \in \mathbb{R}^p$, $k > 0$, $\nu \geq 1$ and Σ are then made x -dependent via a DNN. Notice that the full matrix-variate evidence (5) cannot be obtained from the vectorial evidence (7) by summing it on samples, due to the loss of sample independence at the evidential level.

2.2 UNCERTAINTY QUANTIFICATION

Proposition 2.1 and 2.2 can be used to disentangle and quantify the uncertainty in predictions. For this, let λ and θ represent respectively the parameters of the likelihood and the prior (i.e. hyperparameters). Applying the total variance decomposition on the posterior predictive gives a matrix equality:

$$\mathbb{V}[Y_*|Y, \theta] = \underbrace{\mathbb{E}[\mathbb{V}[Y_*|\lambda]|Y, \theta]}_{\text{Aleatoric}} + \underbrace{\mathbb{V}[\mathbb{E}[Y_*|\lambda]|Y, \theta]}_{\text{Epistemic}}$$

where $\mathbb{V}[Y] := \mathbb{E}[(Y - \mathbb{E}[Y])(Y - \mathbb{E}[Y])^\top]$ denotes the covariance matrix of a random matrix. We establish the following results as generalization of the disentanglement proposed in DER [Amini et al., 2020].

Proposition 2.3. *Given data X, Y and new samples X_* , the aleatoric and epistemic uncertainty matrices of the predictive Y_* are given as follows.*

1. Under the setting of Proposition 2.1, with $\lambda = \{A\}$ and $\theta = \{M, K\}$, then

$$\begin{aligned} \mathbb{E}[\mathbb{V}[Y_*|\lambda]|Y, \theta] &= \text{trace}(D_*)V, \\ \mathbb{V}[\mathbb{E}[Y_*|\lambda]|Y, \theta] &= \text{trace}(\Phi_*^\top S_{\mathbf{xx}}^{-1}\Phi_*)V. \end{aligned}$$

2. Under the setting of Proposition 2.2, with $\lambda = \{A, V\}$ and $\theta = \{M, K, \Sigma\}$ with ν fixed, then

$$\begin{aligned} \mathbb{E}[\mathbb{V}[Y_*|\lambda]|Y, \theta] &= \frac{\text{trace}(D_*)}{\nu + N - 2p - 2}(\Sigma + S_{y|\mathbf{x}}), \\ \mathbb{V}[\mathbb{E}[Y_*|\lambda]|Y, \theta] &= \frac{\text{trace}(\Phi_*^\top S_{\mathbf{xx}}^{-1}\Phi_*)}{\nu + N - 2p - 2}(\Sigma + S_{y|\mathbf{x}}). \end{aligned}$$

The uncertainties for one single point x_* can be trivially obtained since the trace operator vanishes, and those of one specific dimension can be obtained by taking the corresponding diagonal element in these matrices.

3 BLL FOR REGRESSION

Bayesian Last Layer is a straightforward generalization of the regression models presented above, in which the basis ϕ and the heteroscedastic noise level σ are both parameterized by DNNs. Denote by $w := \{w^\phi, w^\sigma\}$ the weights of networks, and by θ the hyperparameters, the ensemble of parameters to be estimated is then $\xi := w \cup \theta$. This actually makes BLL a hybrid model, in that we use a frequentist estimate for w and an empirical Bayes approach for θ .

3.1 PARAMETER ESTIMATION

We specify here the log-evidence $\mathcal{L}_{\text{ev}}(\xi) := \ln \mathbb{P}_\xi(Y)$ for the model of unknown V (the case of known V can be handled similarly). With $\theta = \{M, K, \Sigma\}$ and the evidence expression given in Proposition 2.2, it holds

$$\mathcal{L}_{\text{ev}}(w, \theta) = \ln \mathcal{T}(Y; M\Phi, \Sigma, \Omega, \nu - 2p) \quad (10)$$

which equals to, up to some irrelevant constant

$$-\frac{N}{2} \ln |\Sigma| - \frac{p}{2} \ln |\Omega| - \frac{\nu_p + N}{2} \ln |I_p + \Sigma^{-1} E \Omega^{-1} E^\top|$$

with $\nu_p := \nu - p - 1$ and I_p the p -dimensional identity matrix. As in most existing BLL works, we propose solving the MAP type problem

$$\max_{\xi \in \mathcal{C}} \mathcal{L}_{\text{ev}}(\xi) + \mathcal{R}(\xi) \quad (11)$$

The regularization term $\mathcal{R}(\xi)$ can take form of a simple ℓ_2 -regularization on the weights *e.g.* $\mathcal{R}(\xi) \propto \|w\|^2$ or combined with a *hyperprior* $\mathcal{R}(\xi) \propto \ln \pi(\xi)$. Here \mathcal{C} denotes the constraint set we put on ξ , *e.g.* the covariance matrices K and Σ must live in the cone of positive semi-definite (PSD) matrix (or more specific cone if some structure is imposed on the covariance). For a gradient-based optimization method, the PSD cone constraint can be fulfilled *e.g.* via Cholesky decomposition, or using a full representation but projecting the gradient vector onto the PSD cone at each step.

Identifiability of parameters Despite the full matrix-variate formulation, we do not have rigorous proof on the unique identifiability of BLL parameters by maximizing $\mathcal{L}_{\text{ev}}(\xi)$ alone. In fact, with all other parameters fixed, the unique MLE solutions for M, Σ can be found to be (see Appendix D.2):

$$\hat{M} = Y \Omega^{-1} \Phi^\top (\Phi \Omega^{-1} \Phi^\top)^{-1}, \quad \hat{\Sigma} = \frac{\nu_p}{N} E \Omega^{-1} E^\top \quad (12)$$

which are both function of K (included in Ω). Unless K can be uniquely determined, these actually define paths on which the partial derivatives of \mathcal{L}_{ev} vanish, suggesting

that multiple MLE solutions exist for the full set of BLL parameters. This is similar to the situation in DER as pointed out in [Meinert et al., 2023], and introducing regularizations on θ can improve the numerical stability but cannot solve completely the issue. On the other hand, it should be noticed that the MLE $\hat{\Sigma}$ in (12) is actually a biased estimator, and the unbiased correction is

$$\hat{\Sigma} = \frac{\nu_p - p - 1}{N} E \Omega^{-1} E^\top$$

which can be checked by taking the expectation about Y . The MLE \hat{M} on the other hand is unbiased. See Appendix D.2 for details.

Mini-batches compatibility Notice from definition that the matrix Ω in (10) is non-diagonal in general, as consequence data samples are inseparable in \mathcal{L}_{ev} , which may result in biased gradient estimation on mini-batches. As a solution, we propose in Section 4 a mini-batches compatible variant of EM, called ELBO-EM, for the MAP problem (11). It combines Stochastic Gradient Descent (SGD) for w and analytical update for θ with guaranteed convergence within the framework of EM.

3.2 SCALAR BLL

The proposed multivariate BLL model can obviously be downgraded to a scalar version, by applying the multivariate model dimensionwisely. This amounts to assume the independence between dimensions or have a diagonal covariance structure for the baseline noise. At dimension j the scalar regression model is then $y = a_j^\top \phi_j(x) + \sigma_j(x) \varepsilon$. We place an inverse Gamma prior $v \sim \mathcal{G}^{-1}((\nu - 2)/2, \psi^2/2)$ for the variance of ε , which is just $\mathcal{W}^{-1}(\psi^2 I_1, \nu)$ in the scalar case, and a Normal prior $a_j \sim \mathcal{N}_{\text{vec}}(m_j, v K_j)$. Alternatively, we can assume $v = 1$ for ε and use the model of Proposition 2.1. In this case the log-evidence becomes

$$\mathcal{L}_{\text{ev}}(\xi_1, \dots, \xi_p) = \sum_j \ln \mathcal{N}_{\text{vec}}(Y_j; \Phi_j^\top m_j, D_j + \Phi_j^\top K_j \Phi_j) \quad (13)$$

with the learnable parameters $\xi_j := \{m_j, K_j, \phi_j, \sigma_j\}$. The dimensions can be linked to each other by letting $\phi_j = \phi$.

4 OPTIMIZATION BY ELBO-EM

The Expectation Maximization (EM) method originally employed in the context of MLE with missing observations, can also be applied in the Bayesian framework [Lange et al., 1989]. For BLL, the objective to be maximized is the log-evidence with parameters ξ , while the ‘‘missing observations’’ are the parameters of the likelihood denoted λ . EM is based on the substitution function

$$Q(\xi; \tilde{\xi}) := \mathbb{E}_{\lambda \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} [\ln \mathbb{P}_\xi(\lambda, Y)]$$

which can be shown fulfilling, using Jensen inequality,

$$Q(\xi; \tilde{\xi}) - Q(\tilde{\xi}; \tilde{\xi}) \leq \ln \mathbb{P}_\xi(Y) - \ln \mathbb{P}_{\tilde{\xi}}(Y)$$

Then at the iteration n , maximizing $Q(\cdot; \xi_{n-1})$ with ξ_{n-1} the parameters from the last iteration will define a sequence $\{\xi_n\}_n$ which monotonically increases the log-evidence. EM can also be used for the MAP problem (11): it suffices to add the term $\mathcal{R}(\xi)$ to the substitution function and modifies accordingly the M-step.

ELBO form objective Observe that the substitution function Q can be put in the equivalent form

$$\begin{aligned} \mathbb{E}_{\lambda \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} [\ln \mathbb{P}_\xi(Y|\lambda)] - \mathbb{D}_{\text{KL}} \left[\mathbb{P}_{\tilde{\xi}}(\lambda|Y) \parallel \mathbb{P}_\xi(\lambda) \right] \\ = Q_1(\xi; \tilde{\xi}) - Q_2(\xi; \tilde{\xi}) \end{aligned}$$

where a constant which is independent of ξ is omitted. One can easily recognize the ELBO objective with the posterior $\mathbb{P}_\xi(\lambda)$ as the variational distribution. We refer to the method of maximizing the ELBO objective, in place of the original Q function, as the ELBO-EM method. Notice that the idea of using ELBO as objective has already been proposed in [Harrison et al., 2023, Watson et al., 2021]. The difference in our approach is that ELBO is employed inside the EM framework, avoiding the need to specify the variational prior distribution. As we shall see later, the samples in Q_1 are in a separable form making it compatible with mini-batch SGD training, while explicit update formulae for θ can be obtained using the first order optimality condition on Q_2 .

Algorithm 1 ELBO-EM algorithm

Require: X, Y , the baseline covariance V or the degree of freedom ν .

Require: maximum epochs N_E

while $n < N_E$ **do**

(E-step) At the beginning of epoch n , compute Q_1 with $\tilde{\xi} = \xi_{n-1}$

(M-step) Apply SGD with mini-batches on Q_1 for optimization of w . At the end of epoch n , update θ using (15)

$\xi_n \leftarrow$ the updated ξ , and $n \leftarrow n + 1$

end while

4.1 ELBO-EM ALGORITHM

We establish hereafter the ELBO-EM algorithm for the case V is unknown with $\lambda = \{A, V\}$. The degree of freedom ν being fixed, the parameters to be estimated are $\xi = \{M, K, \Sigma, w^\phi, w^\sigma\}$, and their values from precedent step are denoted $\tilde{\xi} = \{\tilde{M}, \tilde{K}, \tilde{\Sigma}, \tilde{w}^\phi, \tilde{w}^\sigma\}$. The matrices $\tilde{S}_{y|x}, \tilde{S}_{xx}, \tilde{S}_{y|x}$ are constructed by definition as in Table 1 but with the parameters $\tilde{\xi}$. To develop the Q function of

the ELBO-EM form, we apply the posterior from Proposition 2.2:

$$\begin{aligned} \mathbb{P}_{\tilde{\xi}}(A, V|Y) &= \mathbb{P}_{\tilde{\xi}}(A|V, Y) \times \mathbb{P}_{\tilde{\xi}}(V|Y) \\ &= \mathcal{N}\left(A; \tilde{M}, V, \tilde{K}\right) \times \mathcal{W}^{-1}\left(V; \tilde{\Sigma} + \tilde{S}_{y|x}, \nu + N\right) \end{aligned}$$

Define the shorthand $\tilde{E} := Y - \tilde{S}_{y|x} \tilde{S}_{xx}^{-1} \Phi$ and $\tilde{F} := M - \tilde{S}_{y|x} \tilde{S}_{xx}^{-1}$, then it can be shown that²

$$Q_1(\xi; \tilde{\xi}) \cong -\frac{p}{2} \ln |D| - \frac{p}{2} \text{trace} \left(p \Phi D^{-1} \Phi^\top \tilde{S}_{xx}^{-1} \right) -$$

$$\frac{\nu_p + N}{2} \text{trace} \left(\left(\tilde{\Sigma} + \tilde{S}_{y|x} \right)^{-1} \tilde{E} D^{-1} \tilde{E}^\top \right)$$

$$Q_2(\xi; \tilde{\xi}) \cong \frac{p}{2} \left(\ln |K| + \text{trace} \left(K^{-1} \tilde{S}_{xx}^{-1} \right) \right) - \frac{\nu_p}{2} \ln |\Sigma| +$$

$$\frac{\nu_p + N}{2} \text{trace} \left(\left(\tilde{\Sigma} + \tilde{S}_{y|x} \right)^{-1} \left(\tilde{F} K^{-1} \tilde{F}^\top + \Sigma \right) \right)$$

Notice that the parameters w and θ are contained separated in Q_1 and Q_2 , that is $Q_1(\xi; \tilde{\xi}) = Q_1(w; \tilde{\xi})$, $Q_2(\xi; \tilde{\xi}) = Q_2(\theta; \tilde{\xi})$. Moreover, the samples in Q_1 are now in a separable form:

$$\begin{aligned} Q_1 &= -\frac{1}{2} \sum_i \left\{ p \ln \sigma^2(x_i) + \sigma^{-2}(x_i) \times \right. \\ &\quad \left. \left((\nu_p + N) \tilde{e}_i^\top \left(\tilde{\Sigma} + \tilde{S}_{y|x} \right)^{-1} \tilde{e}_i + p \phi(x_i)^\top \tilde{S}_{xx}^{-1} \phi(x_i) \right) \right\} \end{aligned}$$

with $\tilde{e}_n := y_n - \tilde{S}_{y|x} \tilde{S}_{xx}^{-1} \phi(x_n)$, which makes Q_1 compatible with mini-batches based SGD. For the Q_2 term, the gradient $\partial_\theta Q(\xi; \tilde{\xi})$ can be calculated explicitly by matrix calculus, and using the first order optimality condition $\partial_\theta Q(\xi; \tilde{\xi}) = 0$ we obtain the update formulae:

$$M = \tilde{S}_{y|x} \tilde{S}_{xx}^{-1}, \quad K = \tilde{S}_{xx}^{-1}, \quad \Sigma = \frac{\nu_p}{\nu_p + N} \left(\tilde{\Sigma} + \tilde{S}_{y|x} \right) \quad (15)$$

The case V that is known is handled similarly and leads to the same update formulae for M and K , see Appendix C.1. Notice that just as its MLE counterpart, the final estimate of Σ returned by ELBO-EM is also biased. For unbiased correction it should be multiplied by $(\nu_p - p - 1)/\nu_p$. The ELBO-EM algorithm is resumed in Algorithm 1.

5 NUMERICAL EXPERIMENTS

5.1 TOY MODEL

In our first example we demonstrate the effectiveness of evidence maximization on a scalar function $f(x) = x^3/8$.

²The symbol \cong stands for the equality up to some irrelevant constant.

Samples are generated randomly and uniformly on the range $\{0.5 \leq |x| \leq 1.5\}$ following $y = f(x) + \sigma(x)\epsilon$ with a heteroscedastic noise level. Here the objective being maximized is the scalar BLL (13), and the covariance matrix K is represented internally via its Cholesky form. ϕ and σ are parameterized by two independent MLPs with 3 hidden layers and 32 units per layer using the Softplus activation. We generate 1000 samples and use a native SGD method, *e.g.* ADAM, with learning rate 10^{-3} and mini-batch size 128. Despite potentially biased gradient estimate, the convergence is generally observed after 1000 epochs.

The final prediction as well as the uncertainties, computed using the equations in Proposition B.4, are shown in Figure 1. Colors of the confidence bands from dark to light correspond to 1 to 3 times the predicted aleatoric uncertainty. From the predicted aleatoric uncertainty in the bottom right, it can be judged that the predicted aleatoric uncertainty is well calibrated while the epistemic uncertainty seems overconfident on the extrapolation zone. This can be explained by the large number of samples which is known to increase the precision of the aleatoric prediction while decrease the value of the epistemic one. More results of the same experiment using a sinus function can be found in Appendix E. For comparison, a fully stochastic Bayesian neural network posterior predictive distribution is shown in Figure 3 where we use a Gaussian density network [Bishop, 1994] as likelihood.

5.2 TRANSFER LEARNING

Rather than jointly training ϕ , σ and BLL from scratch, another interesting scenario is to augment a trained feature map with the uncertainty prediction capability, by transferring it onto some new data domain. We demonstrate this idea here using the cubic function above. Suppose a baseline neural network has first been trained in a standard fashion to fit the cubic function. By keeping all previous layers frozen, we replace its output layer with a small neural network followed by a BLL. The new DNN together with a network σ are then trained using the ELBO-EM algorithm.

In this experiment the baseline network was trained with 1000 random samples uniformly distributed on the range $\{|x| < 1.5\}$ with a small noise level $\simeq 0.1$. The new data domain is the same as in Section 5.1 but contains only 200 samples to make the experiment more challenging. The network σ is the same as in Section 5.1. We use the native SGD method as a warm start for the initialization of the matrix M and K for ELBO-EM. Results are depicted in Figure 2 where the predicted epistemic uncertainty is clearly more dominant on the extrapolation zone.

5.3 MULTIVARIATE TIME SERIES PREDICTION

In the last example we apply our proposed framework for the prediction of multivariate time series, using Beijing air quality data from [Chen, 2019]. The original time series contain 4 years of hourly measurement starting from Mars 2013. They are composed of 5 types of measurement as inputs: temperature, pressure, dew point temperature, precipitation and wind speed, as well as 6 types of concentration of air pollutants, including PM2.5, PM10, SO2, NO2, CO, and O3, as outputs. The original data was first averaged then down-sampled at the daily frequency. We extend the classical VARX³ model as for Bayesian basis regression to obtain the following model

$$y_t = A\phi(\tilde{x}_t) + \sigma(\tilde{x}_t)\epsilon_t$$

where $\tilde{x}_t := [x_t, \dots, x_{t-P+1}, y_{t-1}, \dots, y_{t-Q}]$ is the concatenation of historical inputs and outputs. BLL priors are placed on A and the covariance V of the baseline noise. The same architectures as in Section 5.1 are taken for ϕ and σ , and the input and auto-regression orders are fixed to $P = 2, Q = 1$. We maximize the log-evidence (10) and use the first 3 years data for training and the last 6 months for testing. Results of prediction are shown in Figure 4. The predicted uncertainty tends to be shared across all variables, for example for the peaks of pollution by the end of 2016 (mainly due to PM2.5, PM10 and CO). This clearly indicates a limit of our scalar-scaling based heteroscedastic model (1).

6 CONCLUSION

The BLL framework developed here is an alternative to existing solutions for multivariate regression problems under heteroscedastic noise. It provides closed-form formulas to disentangle aleatoric and epistemic uncertainties with a single forward pass. The ELBO-EM algorithm tailored for BLL benefits from unbiased gradient estimation on mini-batches and can be used for transfer learning.

In our experiments we noticed some critical issues, including the numerical instability occurred during the training, mainly due to the ill-conditioning of covariance matrices, and the overconfident predicted epistemic uncertainty on extrapolation zones. These issues could be improved by introducing regularizations or restrictions on the precision matrix K , and will be investigated in a follow-up work.

³Vectorial Auto-Regressive with eXogenous inputs

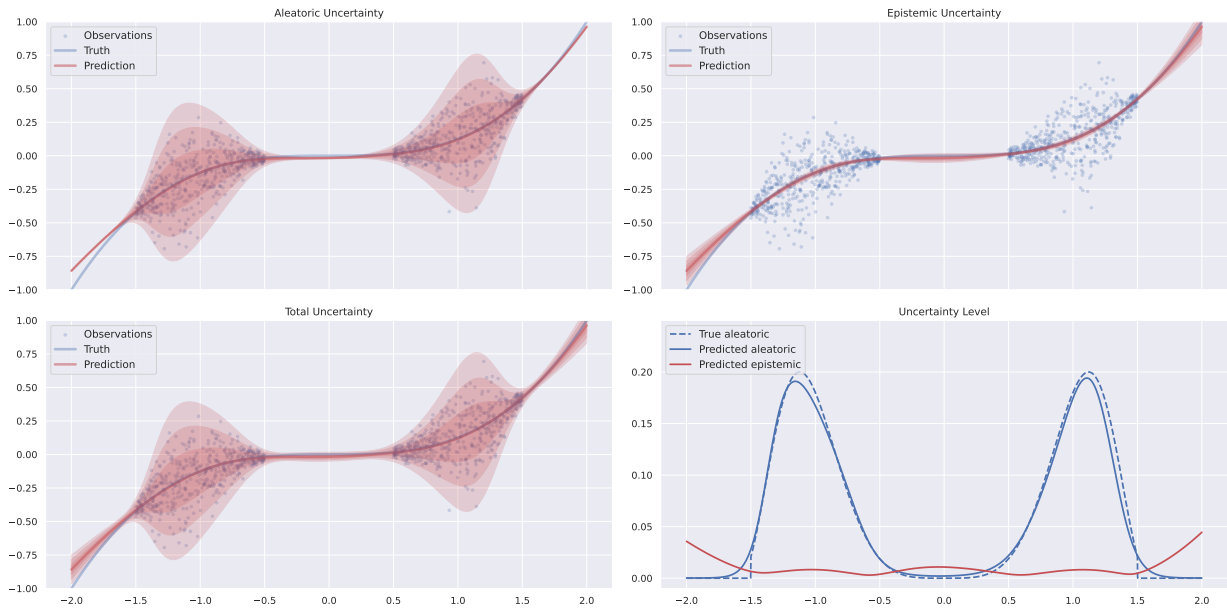


Figure 1: Joint learning of the cubic function $f(x) = x^3/8$ with 1000 samples using SGD.

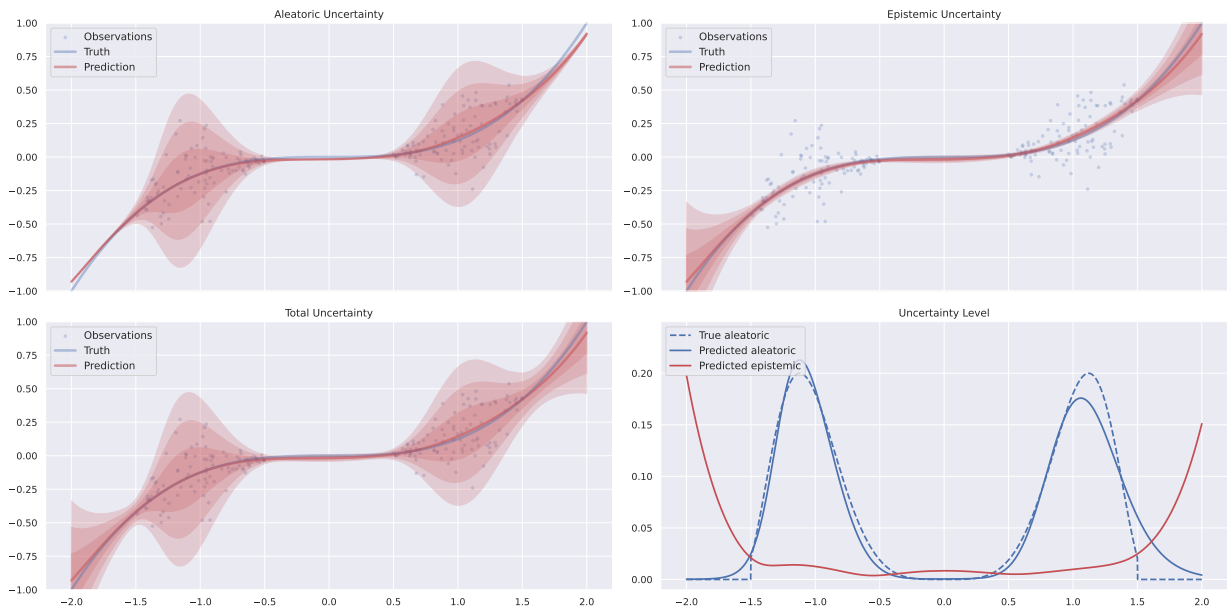


Figure 2: Transfer Learning with 200 samples using the ELBO-EM algorithm.

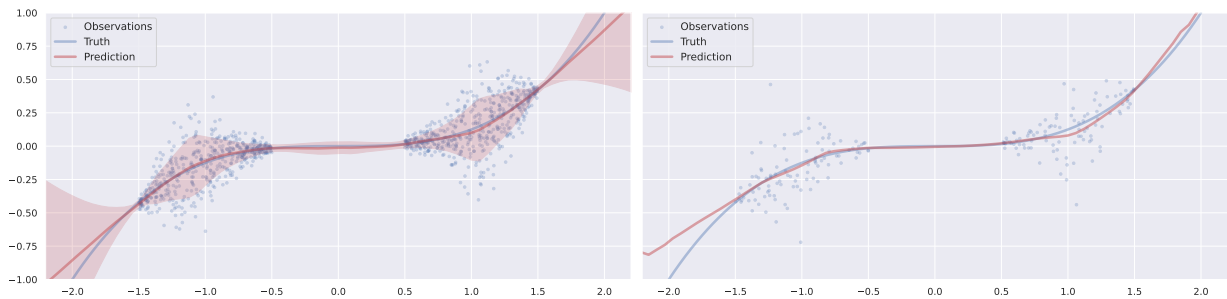


Figure 3: Uncertainty predicted by BNN with 1000 and 200 samples.

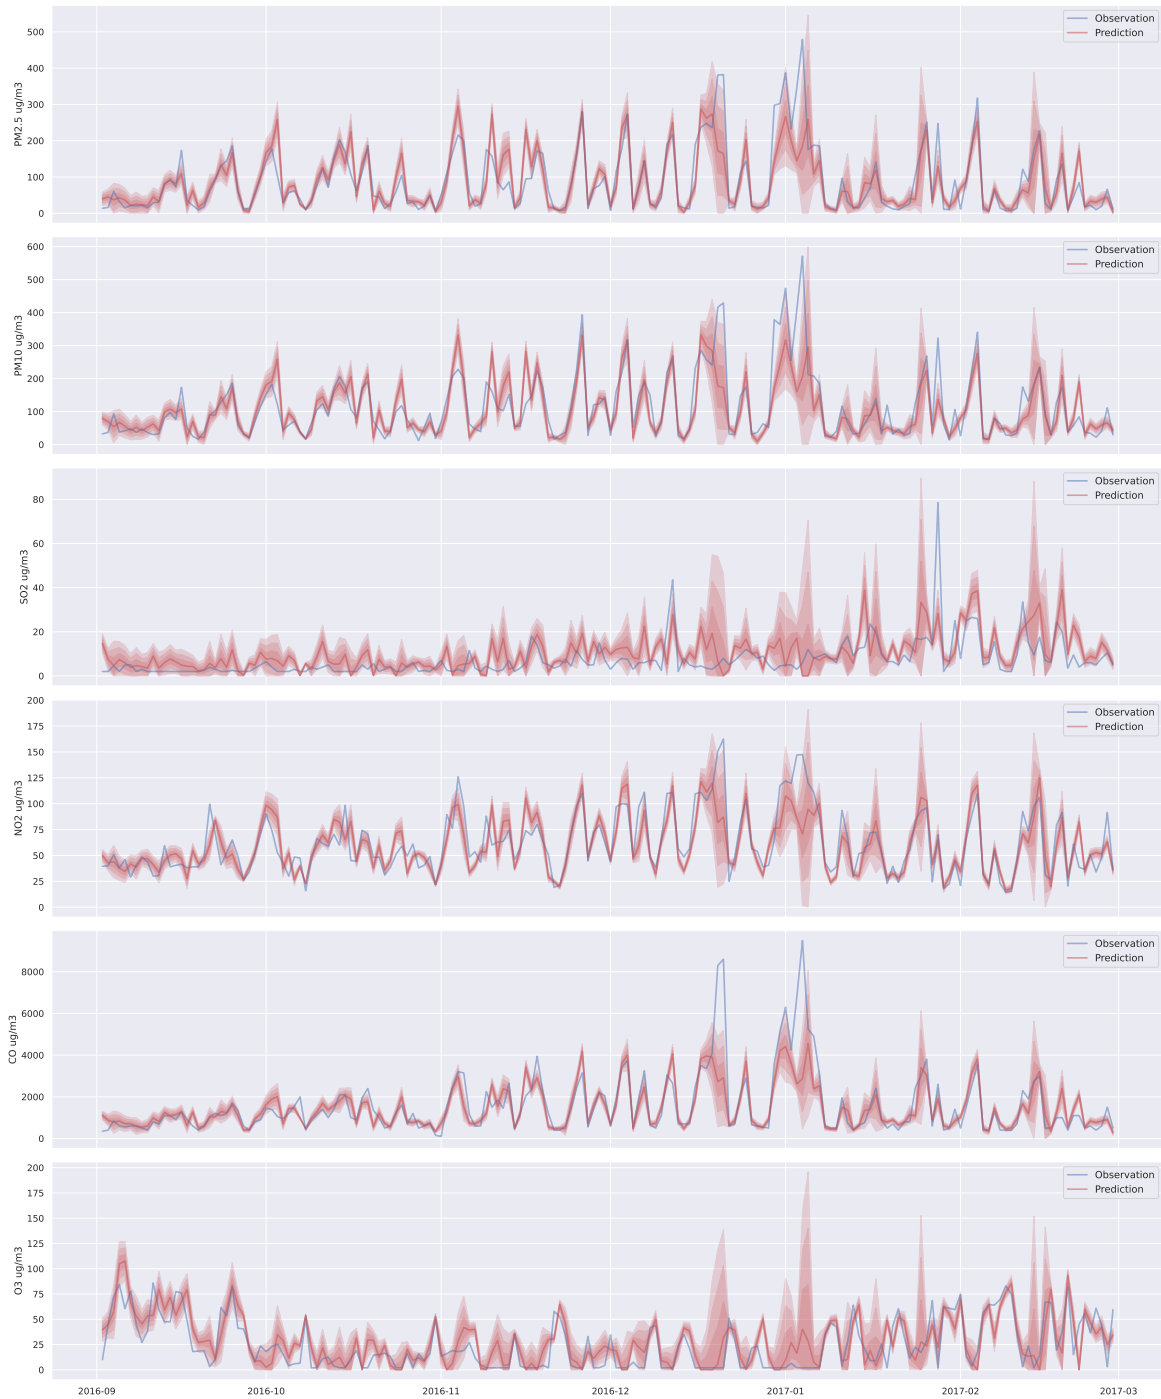


Figure 4: Predictions of Beijing air quality using a nonlinear (1,2)-VARX BLL model.

References

- Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. *Advances in Neural Information Processing Systems*, 33:14927–14937, 2020.
- Sudipto Banerjee. Bayesian linear model: Gory details. 2008.
- Christopher M. Bishop. Mixture density networks. Workingpaper, Aston University, 1994.
- Christopher M. Bishop and Michael E. Tipping. Bayesian regression and classification. *Nato Science Series sub Series III Computer And Systems Sciences*, 190:267–288, 2003. Publisher: IOS PRESS.
- George E. P. Box and George C. Tiao. *Bayesian Inference in Statistical Analysis*. John Wiley & Sons, January 2011. ISBN 978-1-118-03144-5. Google-Books-ID: T8Askeyk1k4C.
- Song Chen. Beijing Multi-Site Air-Quality Data. UCI Machine Learning Repository, 2019. DOI: <https://doi.org/10.24432/C5RK5G>.
- Felix Fiedler and Sergio Lucia. Improved uncertainty quantification for neural networks with Bayesian last layer. *arXiv preprint arXiv:2302.10975*, 2023. URL <https://arxiv.org/abs/2302.10975>.
- Jakob Gawlikowski, Cedrique Rovile Njiteucheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56:1513–1589, 2023.
- A. K. Gupta and D. K. Nagar. *Matrix Variate Distributions*. CRC Press, October 1999. ISBN 978-1-58488-046-2. Google-Books-ID: PQOYN7P1loC.
- James Harrison, John Willes, and Jasper Snoek. Variational Bayesian Last Layers. In *Fifth Symposium on Advances in Approximate Bayesian Inference*, 2023. URL <https://openreview.net/forum?id=lAlyfxQ2pS>.
- Kenneth L. Lange, Roderick JA Little, and Jeremy MG Taylor. Robust statistical modeling using the t distribution. *Journal of the American Statistical Association*, 84(408): 881–896, 1989. Publisher: Taylor & Francis.
- M. Lázaro-Gredilla and Figueiras-Vidal Aníbal R. Marginalized neural network mixtures for large-scale regression. *IEEE Transactions on Neural Networks*, 21:1345–1351, 2010.
- Nis Meinert and Alexander Lavin. Multivariate deep evidential regression. *arXiv preprint arXiv:2104.06135*, 2021.
- Nis Meinert, Jakob Gawlikowski, and Alexander Lavin. The Unreasonable Effectiveness of Deep Evidential Regression. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(8):9134–9142, June 2023. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v37i8.26096. URL <http://arxiv.org/abs/2205.10060>. arXiv:2205.10060 [cs, stat].
- Thomas Minka. Bayesian linear regression. Technical report, Citeseer, 2010.
- Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT Press, August 2012. ISBN 978-0-262-01802-9. Google-Books-ID: NZP6AQAAQBAJ.
- Sebastian W. Ober and Carl Edward Rasmussen. Benchmarking the Neural Linear Model for Regression, December 2019. URL <http://arxiv.org/abs/1912.08416>. arXiv:1912.08416 [cs, stat].
- Mrinank Sharma, Sebastian Farquhar, Eric Nalisnick, and Tom Rainforth. Do bayesian neural networks need to be fully stochastic? In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 7694–7722. PMLR, 25–27 Apr 2023. URL <https://proceedings.mlr.press/v206/sharma23a.html>.
- J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, Md. M. A. Patwary, M. Prabhat, and R. Adams. Scalable bayesian optimization using deep neural networks. *Statistics*, 02 2015.
- Joe Watson, Jihao Andreas Lin, Pascal Klink, Joni Pajarinen, and Jan Peters. Latent derivative Bayesian last layer networks. In *International Conference on Artificial Intelligence and Statistics*, pages 1198–1206. PMLR, 2021. URL <http://proceedings.mlr.press/v130/watson21a.html>.

Multivariate Bayesian Last Layer for Regression

Appendix

Han WANG^{1,@}

Eiji KAWASAKI^{1,≈}

Guillaume DAMBLIN^{2,≈}

Geoffrey DANIEL^{2,≈}

¹Université Paris Saclay - CEA LIST, F-91120 Palaiseau Cedex, France

²Université Paris-Saclay, CEA, Service de Génie Logiciel pour la Simulation, 91191, Gif-sur-Yvette, France

[@]Corresponding author: han.wang@cea.fr

[≈]Equal contribution

A SOME MATRIX VARIATE DISTRIBUTIONS

We collect here some important matrix variate distributions. Different but equivalent definitions can be found in the literature and the conventions we followed here are from the standard reference on this matter [Gupta and Nagar, 1999], which are slightly different from those of [Minka, 2010].

Notations Let $\text{vec}(X)$ be the concatenation of column vectors of a matrix X , and $X \otimes Y$ be the Kronecker product. For a square matrix A , we define the operator

$$\text{etr}(A) := \exp(\text{trace}(A))$$

and denote by $|A|$ its determinant.

A.1 MATRIX-NORMAL DISTRIBUTION

A random matrix $A \sim p \times n$ follows a matrix-Normal distribution $\mathcal{N}(M, \Sigma, \Psi)$ iff $\text{vec}(A)$ follows a multivariate Normal distribution $\mathcal{N}(M, \Psi \otimes \Sigma)$. Here $M \sim p \times n$ is the mean, and the matrices $\Sigma \sim p \times p$ and $\Psi \sim n \times n$ are both symmetric and positive definite. It holds $A \sim \mathcal{N}(M, \Sigma, \Psi)$ iff $A^\top \sim \mathcal{N}(M^\top, \Psi, \Sigma)$. The p.d.f. is

$$\mathcal{N}(A; M, \Sigma, \Psi) = (2\pi)^{-\frac{pn}{2}} |\Sigma|^{-\frac{n}{2}} |\Psi|^{-\frac{p}{2}} \text{etr} \left(-\frac{1}{2} \Sigma^{-1} (A - M) \Psi^{-1} (A - M)^\top \right)$$

Moreover, for any square matrix Q of appropriate dimensions it holds

$$\mathbb{E} \left[(A - M) Q (A - M)^\top \right] = \text{trace} (Q^\top \Psi) \Sigma \tag{A.16}$$

as well as

$$\mathbb{E} \left[(A - M)^\top Q (A - M) \right] = \text{trace} (Q^\top \Sigma) \Psi \tag{A.17}$$

A.2 MATRIX-T DISTRIBUTION

Given the mean $M \sim p \times n$, the symmetric and positive definite matrices $\Sigma \sim p \times p$ and $\Omega \sim n \times n$, as well as the constant $\nu > 0$, a random matrix $T \sim p \times n$ follows a matrix-T distribution $\mathcal{T}(M, \Sigma, \Omega, \nu)$ if it has the p.d.f.

$$\mathcal{T}(T; M, \Sigma, \Omega, \nu) = \frac{\Gamma_p \left(\frac{\nu+n+p-1}{2} \right)}{\Gamma_p \left(\frac{\nu+p-1}{2} \right) \pi^{\frac{np}{2}}} |\Sigma|^{-\frac{n}{2}} |\Omega|^{-\frac{p}{2}} \left| I_p + \Sigma^{-1} (T - M) \Omega^{-1} (T - M)^\top \right|^{-\left(\frac{\nu+n+p-1}{2} \right)}$$

where Γ_p is the Multivariate Gamma function defined as

$$\Gamma_p(a) := \pi^{\frac{p(p-1)}{4}} \prod_{j=1}^p \Gamma\left(a + \frac{1-j}{2}\right)$$

It holds $T \sim \mathcal{T}(M, \Sigma, \Omega, \nu)$ iff $T^\top \sim \mathcal{T}(M^\top, \Omega, \Sigma, \nu)$. Moreover, T distribution is closely related to Normal distribution: $T \sim \mathcal{T}(M, \nu\Sigma, \Omega, \nu)$ converges in distribution to $\mathcal{N}(M, \Sigma, \Omega)$ as $\nu \rightarrow \infty$. Parallel to (A.16) and (A.17), it holds for any square matrix Q of appropriate dimensions and for $\nu > 2$

$$\mathbb{E} \left[(T - M) Q (T - M)^\top \right] = \frac{1}{\nu - 2} \text{trace} (Q^\top \Omega) \Sigma \quad (\text{A.18})$$

as well as

$$\mathbb{E} \left[(T - M)^\top Q (T - M) \right] = \frac{1}{\nu - 2} \text{trace} (Q^\top \Sigma) \Omega \quad (\text{A.19})$$

Multivariate-T distribution

In the case $n = 1$, i.e. T is a p -dimension column vector, the matrix Ω becomes a scalar ω^2 which can be absorbed in Σ . This gives the multivariate-T distribution with the p.d.f.

$$\mathcal{T}_{\text{vec}}(T; M, \Sigma, \nu) = \frac{\Gamma\left(\frac{\nu+p}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right) \pi^{\frac{p}{2}}} |\Sigma|^{-\frac{1}{2}} \left(1 + (T - M)^\top \Sigma^{-1} (T - M)\right)^{-\left(\frac{\nu+p}{2}\right)}$$

which can be obtained from the matrix-T p.d.f. as $\mathcal{T}(T; M, \Sigma/\omega^2, \omega^2, \nu)$.

Remark A.1. Basic linear algebra identities show that

$$\begin{aligned} \left| I_p + \Sigma^{-1} (T - M) \Omega^{-1} (T - M)^\top \right| &= |\Sigma|^{-1} \times \left| \Sigma + (T - M) \Omega^{-1} (T - M)^\top \right| \\ &= \left| I_n + \Omega^{-1} (T - M)^\top \Sigma^{-1} (T - M) \right| \\ &= |\Omega|^{-1} \times \left| \Omega + (T - M)^\top \Sigma^{-1} (T - M) \right| \end{aligned}$$

Depending on the value of p and n , in practice some of these identities can be more efficient for evaluating numerically the p.d.f. than others.

Remark A.2. Unlike matrix-Normal case, matrix-T distribution cannot be obtained by rearranging vectors of multivariate-T distribution. In fact, the columns (resp. rows) of T are never mutually independent (see Theorem 4.3.9 of Gupta and Nagar [1999]), and are only uncorrelated iff Ω (resp. Σ) is diagonal, as can be observed easily from (A.19) (resp. (A.18)). This puts constraints on the form of loss function in fitting a T distribution, depending on whether the sample independence is assumed or not.

A.3 WISHART DISTRIBUTION

Given a matrix $\Sigma \sim p \times p$ being symmetric positive definite, and a constant $\nu \geq p$, a symmetric positive definite random matrix S follows the Wishart distribution $\mathcal{W}(\Sigma, \nu)$ if it has the p.d.f.

$$\mathcal{W}(S; \Sigma, \nu) = \frac{1}{2^{\frac{\nu p}{2}} \Gamma_p\left(\frac{\nu}{2}\right)} \cdot \frac{|S|^{\frac{\nu-p-1}{2}}}{|\Sigma|^{\frac{\nu}{2}}} \text{etr} \left(-\frac{1}{2} \Sigma^{-1} S \right)$$

A.4 INVERSE WISHART DISTRIBUTION

A random matrix $S \sim p \times p$ follows an Inverse Wishart distribution $\mathcal{W}^{-1}(\Psi, \nu)$ with $\nu \geq 2p + 1$ iff S^{-1} follows a Wishart distribution $\mathcal{W}(\Psi^{-1}, \nu - p - 1)$. It has the p.d.f.

$$\mathcal{W}^{-1}(S; \Psi, \nu) = \frac{1}{2^{\frac{(\nu-p-1)p}{2}} \Gamma_p\left(\frac{\nu-p-1}{2}\right)} \cdot \frac{|\Psi|^{\frac{\nu-p-1}{2}}}{|S|^{\frac{\nu}{2}}} \text{etr} \left(-\frac{1}{2} S^{-1} \Psi \right)$$

and has the mean $\Psi / (\nu - 2p - 2)$ and the mode Ψ / ν .

We state a result on the compound Normal-Inverse Wishart distribution:

Proposition A.1. *Let $V \sim p \times p$ follow $\mathcal{W}^{-1}(\Psi, \nu)$, and $A \sim p \times n$ conditioned on V follow $\mathcal{N}(M, V, K)$ for some fixed K . Then the joint distribution of $[V \ A]$ is*

$$[V \ A] \sim \mathcal{N}(M, V, K) \times \mathcal{W}^{-1}(\Psi, \nu).$$

The marginal distribution of A is

$$A \sim \mathcal{T}(M, \Psi, K, \nu - 2p).$$

and the conditional of $V|A$ is

$$V|A \sim \mathcal{W}^{-1}\left(\Psi + (A - M)K^{-1}(A - M)^\top, \nu + n\right)$$

Proof. Skipped. □

Inverse Gamma distribution

In the scalar case Ψ becomes a scalar ψ^2 , and $\mathcal{W}^{-1}(s; \psi^2 I_1, \nu)$ is reduced to the inverse Gamma distribution $\mathcal{G}^{-1}(s; (\nu - 2)/2, \psi^2/2)$ with the p.d.f.

$$\mathcal{G}^{-1}(x; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)x^{\alpha+1}} \exp(-\beta/x)$$

B BAYESIAN LINEAR REGRESSION MODEL

We review and extend classical results about Bayesian linear regression (BLR) model. Many references exist on BLR, see *e.g.* Bishop and Tipping [2003], Box and Tiao [2011], chapter 7.6 of Murphy [2012], the notes Banerjee [2008] and Minka [2010]. We follow the multivariate case exposed in Minka [2010] by extending it to the heteroscedastic case, and present the computational details in a more self-contained and polished way, by exploiting the conjugacy and conditionals of some matrix distributions. The results presented here apply also on the model of basis regression.

We consider the multivariate linear regression model between the input $x \in \mathbb{R}^n$ and the output $y \in \mathbb{R}^p$:

$$y = Ax + \varepsilon(x) \quad (\text{B.20})$$

Here $\varepsilon(x) := \sigma(x)\varepsilon$ is a heteroscedastic noise obtained from multiplying a Gaussian white noise $\varepsilon \sim \mathcal{N}_{\text{vec}}(0, V)$ by a scalar $\sigma(x) > 0$. The covariance matrix $V \sim p \times p$ is considered a nuisance parameter, and $A \sim p \times n$ is the unknown regression matrix to be estimated. The model becomes homoscedastic when $\sigma(x)$ is constant, in that case we take $\sigma(x) = 1$. Suppose that for a **fixed** (but unknown) A , independent observations $\{(y_i, x_i)\}_{i=1:N}$ have been realized with the underlying noise $\varepsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}_{\text{vec}}(0, V)$, and let the matrices $Y = [-y_i-]$, $X = [-x_i-]$ be horizontal concatenation of column vectors together with $\mathcal{E} = [-\varepsilon_i-]$. Then the linear model in a matrix form is

$$Y = AX + \mathcal{E}(X) \quad (\text{B.21})$$

with the noise matrix given by

$$\mathcal{E}(X) := [-\sigma(x_i)\varepsilon_i-] = \mathcal{E}\sigma(X) \quad (\text{B.22})$$

where $\sigma(X)$ is the diagonal matrix $\text{diag}[-\sigma(x_i)-]$. Hereafter X will be treated as fixed. Any distribution involving Y depends also on X , although in the following we omit the conditioning on X for the ease of notation.

Non-regression case The non-regression case (*i.e.* without the covariate) refers to the data fitting model

$$y = A + \varepsilon$$

with $A \in \mathbb{R}^p$ a column vector representing the mean of y , and ε a homoscedastic noise. This is covered by the model (B.20) by setting $x = 1$ and $\sigma(x) = 1$. Then in (B.21) we have $X = [-1-] = \mathbb{1}_N^\top$ and $AX = [-A-] \sim p \times N$. All results below of Section B.1 and B.2 can be readily reformulated for the non-regression case.

B.1 CASE 1: V KNOWN

First suppose V is known. We specify the prior of A as

$$A \sim \mathcal{N}(M, V, K), \quad (\text{B.23})$$

and the likelihood of $Y|A$ as

$$Y|A \sim \mathcal{N}(AX, V, \sigma^2(X)),$$

with $\sigma^2(X)$ the square of $\sigma(X)$ introduced in (B.22). In the homoscedastic case $\sigma^2(X)$ reduces to the identity matrix I_N . Besides V , in the following $M, K, \sigma^2(\cdot)$ are also considered fixed.

Shorthand We assign some shorthands to the following matrices which will be used frequently:

$$\begin{aligned} S_{\mathbf{y}\mathbf{x}} &:= Y\sigma^2(X)^{-1}X^\top + MK^{-1} \\ S_{\mathbf{x}\mathbf{x}} &:= K^{-1} + X\sigma^2(X)^{-1}X^\top \\ \Omega &:= \sigma^2(X) + X^\top KX \end{aligned}$$

and it can be checked that

$$S_{\mathbf{y}\mathbf{x}}S_{\mathbf{x}\mathbf{x}}^{-1} = M + (Y - MX)\Omega^{-1}X^\top K$$

moreover, by the Woodbury identity,

$$\begin{aligned} S_{\mathbf{xx}}^{-1} &= K - KX\Omega^{-1}X^\top K \\ \Omega^{-1} &= \sigma^2(X)^{-1} - \sigma^2(X)^{-1}X^\top S_{\mathbf{xx}}^{-1}X\sigma^2(X)^{-1} \end{aligned}$$

Proposition B.1. *It holds the following:*

1. *The joint distribution of $[A Y] \sim p \times (n + N)$ is*

$$[A Y] \sim \mathcal{N} \left([M \ MX], V, \begin{bmatrix} K & KX \\ X^\top K & \Omega \end{bmatrix} \right) \quad (\text{B.24})$$

2. *The evidence is*

$$Y \sim \mathcal{N}(MX, V, \Omega) \quad (\text{B.25})$$

3. *The posterior is*

$$A|Y \sim \mathcal{N}(S_{\mathbf{yx}}S_{\mathbf{xx}}^{-1}, V, S_{\mathbf{xx}}^{-1}) \quad (\text{B.26})$$

4. *Given L new data points $X_* \sim n \times L$, the posterior predictive is*

$$Y_*|Y \sim \mathcal{N}(S_{\mathbf{yx}}S_{\mathbf{xx}}^{-1}X_*, V, \sigma^2(X_*) + X_*^\top S_{\mathbf{xx}}^{-1}X_*)$$

where $\sigma^2(X_*)$ is the $L \times L$ diagonal matrix formed similarly as $\sigma^2(X)$. Moreover, the posterior predictive knowing additionally A is

$$Y_*|Y, A \sim \mathcal{N}(AX_*, V, \sigma^2(X_*))$$

which is the same as $Y_*|A$. In other words, Y_* , Y and independent conditioned on A .

Proof. Given the prior and the likelihood we know that $[A Y]$ follows a matrix-Normal distribution of form

$$\mathcal{N} \left([M \ MX], \Sigma, \begin{bmatrix} \Psi_{11} & \Psi_{12} \\ \Psi_{12}^\top & \Psi_{22} \end{bmatrix} \right)$$

The expressions of conditional matrix-Normal distributions allow to identify these parameters as $\Sigma = V$, $\Psi_{11} = K$, $\Psi_{12} = KX$ and $\Psi_{22} = \sigma^2(X) + X^\top KX$. This yields the joint distribution (B.24), from which we establish easily the other results, using again the conditionals of matrix-Normal distribution, see Theorem 2.3.12 of Gupta and Nagar [1999].

□

B.2 CASE 2: V UNKNOWN

We suppose now V is unknown and specify the prior as:

$$V \sim \mathcal{W}^{-1}(\Sigma, \nu)$$

with $\nu \geq 2p + 1$ and a symmetric positive definite Σ as fixed parameters. Conditioned on V , the prior of A and the likelihood are the same as in the case where V is fixed. The prior's parameters $\{M, K, \nu, \Sigma\}$ together with $\sigma^2(\cdot)$ are considered fixed. The joint distribution of $[V \ A \ Y]$ is then composed by $Y|A, V$, $A|V$ and V :

$$[V \ A \ Y] \sim \mathcal{N}(AX, V, \sigma^2(X)) \times \mathcal{N}(M, V, K) \times \mathcal{W}^{-1}(\Sigma, \nu)$$

Shorthand In addition to previous short-hands, we define

$$\begin{aligned} S_{yy} &:= Y \sigma^2(X)^{-1} Y^\top + M K^{-1} M^\top \\ S_{y|x} &:= S_{yy} - S_{yx} S_{xx}^{-1} S_{yx}^\top \\ S_{y|x}^A &:= [(A - M) (Y - MX)] \begin{bmatrix} K & KX \\ X^\top K & \Omega \end{bmatrix}^{-1} \begin{bmatrix} (A - M)^\top \\ (Y - MX)^\top \end{bmatrix} \end{aligned}$$

and it can be checked that

$$\begin{aligned} S_{y|x} &= (Y - MX) \Omega^{-1} (Y - MX)^\top \\ S_{y|x}^A &= (Y - AX) \sigma^2(X)^{-1} (Y - AX)^\top + (A - M) K^{-1} (A - M)^\top \\ &= S_{yy} + A S_{xx} A^\top - (S_{yx} A^\top + A S_{yx}^\top) \end{aligned}$$

In particular, the formula of block matrix inverse gives

$$\begin{bmatrix} K & KX \\ X^\top K & \Omega \end{bmatrix}^{-1} = \begin{bmatrix} S_{xx} & -X \sigma^2(X)^{-1} \\ -\sigma^2(X)^{-1} X^\top & \sigma^2(X)^{-1} \end{bmatrix}$$

Proposition B.2. *It holds the following*

1. *The joint distribution of $[A Y]$ is*

$$[A Y] \sim \mathcal{T} \left([M \ MX], \Sigma, \begin{bmatrix} K & KX \\ X^\top K & \Omega \end{bmatrix}, \nu - 2p \right) \quad (\text{B.28})$$

2. *The evidence is*

$$Y \sim \mathcal{T}(MX, \Sigma, \Omega, \nu - 2p) \quad (\text{B.29})$$

the marginal of A is

$$A \sim \mathcal{T}(M, \Sigma, K, \nu - 2p) \quad (\text{B.30})$$

3. *The posterior $[A V]|Y$ follows a compound Normal-Inverse Wishart distribution with the marginals*

$$A|Y \sim \mathcal{T}(S_{yx} S_{xx}^{-1}, \Sigma + S_{y|x}, S_{xx}^{-1}, \nu + N - 2p) \quad (\text{B.31})$$

$$V|Y \sim \mathcal{W}^{-1}(\Sigma + S_{y|x}, \nu + N) \quad (\text{B.32})$$

Moreover, conditioned additionally on V , the posterior of A becomes

$$A|Y, V \sim \mathcal{N}(S_{yx} S_{xx}^{-1}, V, S_{xx}^{-1})$$

which is the same as (B.26), and conditioned additionally on A , the posterior of V becomes

$$V|Y, A \sim \mathcal{W}^{-1}(\Sigma + S_{y|x}^A, \nu + N + n) \quad (\text{B.33})$$

4. *Given L new data points $X_* \sim n \times L$, the posterior predictive is*

$$Y_*|Y \sim \mathcal{T}(S_{yx} S_{xx}^{-1} X_*, \Sigma + S_{y|x}, \sigma^2(X_*) + X_*^\top S_{xx}^{-1} X_*, \nu + N - 2p) \quad (\text{B.34})$$

Moreover, the posterior predictive knowing additionally A is

$$Y_*|Y, A \sim \mathcal{T}(AX_*, \Sigma + S_{y|x}^A, \sigma^2(X_*), \nu + N + n - 2p) \quad (\text{B.35})$$

which is different from $Y_|A$ since*

$$Y_*|A \sim \mathcal{T}(AX_*, \Sigma + (A - M) K^{-1} (A - M)^\top, \sigma^2(X_*), \nu + n - 2p).$$

Proof. Various paths can lead to these same results and below is just one of them. The followings have the same distribution as their counterpart in case 1:

- (a) $[A Y]|V$ is distributed as (B.24);
- (b) $A|V$ is distributed as (B.23);
- (c) $Y|V$ is distributed as (B.25).

All being matrix-Normal distributed, therefore applying Proposition A.1 on each gives

- (a) the joint marginal (B.28), as well as the posterior update (B.33);
- (b) the posterior marginal (B.30);
- (c) the evidence (B.29) as well as the posterior marginal (B.32).

Since $[A V]|Y$ is the product of $A|V, Y$ and $V|Y$, which are distributed respectively as (B.26) and (B.32), its distribution is a compound Normal-Inverse Wishart. Therefore calling again Proposition A.1 gives the marginal (B.31). Applying (B.29) on $[Y_* Y]$ gives

$$[Y_* Y] \sim \mathcal{T} \left([MX_* MX], \Sigma, \begin{bmatrix} \sigma^2(X_*) + X_*^\top K X_* & X_*^\top K X \\ X^\top K X_* & \sigma^2(X) + X^\top K X \end{bmatrix}, \nu - 2p \right)$$

the posterior predictive (B.34) then follows from the conditionals of matrix-T distribution, see Theorem 4.3.9 of Gupta and Nagar [1999]. Finally (B.35) can be established in the same way by working on the joint $[A Y Y_*]$ using (B.28). \square

B.3 QUANTIFICATION OF UNCERTAINTIES

The total variance decomposition can be easily extended to the posterior predictive as stated in the following result.

Proposition B.3 (Decomposition of uncertainties). *Let λ be the parameters of the likelihood, and θ be those of the prior. It holds*

$$\mathbb{V} [Y_* | Y, \theta] = \underbrace{\mathbb{E} [\mathbb{V} [Y_* | \lambda] | Y, \theta]}_{\text{Aleatoric}} + \underbrace{\mathbb{V} [\mathbb{E} [Y_* | \lambda] | Y, \theta]}_{\text{Epistemic}} \quad (\text{B.36})$$

Proof. Under the assumption of sample independence, the distribution of Y_* conditioned on Y, λ, θ is the same as conditioned on λ alone. Therefore

$$p(Y_* | Y, \theta) = \int p(Y_* | \lambda) p(\lambda | Y, \theta) d\lambda$$

which leads to

$$\mathbb{V} [Y_* | Y, \theta] = \mathbb{E} \left[\mathbb{E} \left[(Y_* - \mathbb{E} [Y_* | Y, \theta]) (Y_* - \mathbb{E} [Y_* | Y, \theta])^\top | \lambda \right] | Y, \theta \right]$$

Writing $Y_* - \mathbb{E} [Y_* | Y, \theta] = (Y_* - \mathbb{E} [Y_* | \lambda]) + (\mathbb{E} [Y_* | \lambda] - \mathbb{E} [Y_* | Y, \theta])$, the cross terms vanish since

$$\mathbb{E} \left[(Y_* - \mathbb{E} [Y_* | \lambda]) (\mathbb{E} [Y_* | \lambda] - \mathbb{E} [Y_* | Y, \theta])^\top | \lambda \right] = 0,$$

by consequence

$$\begin{aligned} \mathbb{V} [Y_* | Y, \theta] &= \mathbb{E} \left[\mathbb{E} \left[(Y_* - \mathbb{E} [Y_* | \lambda]) (Y_* - \mathbb{E} [Y_* | \lambda])^\top | \lambda \right] | Y, \theta \right] + \\ &\quad \mathbb{E} \left[(\mathbb{E} [Y_* | \lambda] - \mathbb{E} [Y_* | Y, \theta]) (\mathbb{E} [Y_* | \lambda] - \mathbb{E} [Y_* | Y, \theta])^\top | Y, \theta \right] \end{aligned}$$

where the two terms of RHS correspond respectively to the aleatoric and the epistemic uncertainty in (B.36), using the fact $\mathbb{E} [\mathbb{E} [Y_* | \lambda] | Y, \theta] = \mathbb{E} [Y_* | Y, \theta]$. \square

Now we apply this result on the linear regression model discussed above, to compute explicitly the aleatoric and epistemic uncertainty of posterior predictive.

Proposition B.4 (Uncertainties case 1: V known). *Let $\lambda = \{A\}$ and $\theta = \{M, K\}$ be the likelihood and prior parameters as in Section B.1. Given the data Y , the aleatoric and epistemic uncertainty of the predictive Y_* are respectively*

$$\mathbb{E}[\mathbb{V}[Y_*|\lambda] | Y, \theta] = \text{trace}(\sigma^2(X_*)) V, \quad (\text{B.37a})$$

$$\mathbb{V}[\mathbb{E}[Y_*|\lambda] | Y, \theta] = \text{trace}(X_*^\top S_{\mathbf{x}\mathbf{x}}^{-1} X_*) V \quad (\text{B.37b})$$

and add up to $\mathbb{V}[Y_*|Y, \theta]$.

Proof. The identity in (A.16) allows to compute the covariance for matrix Normal distribution. By Proposition B.1 we know $Y_*|\lambda \sim \mathcal{N}(AX_*, V, \sigma^2(X_*))$, therefore $\mathbb{E}[Y_*|\lambda] = AX_*$ and $\mathbb{V}[Y_*|\lambda] = \text{trace}(\sigma^2(X_*)) V$. Then the aleatoric uncertainty (B.37a) follows since V is constant. Since $A|Y, \theta \sim \mathcal{N}(S_{\mathbf{y}\mathbf{x}} S_{\mathbf{x}\mathbf{x}}^{-1}, V, S_{\mathbf{x}\mathbf{x}}^{-1})$, therefore

$$\mathbb{V}[AX_*|Y, \theta] = \mathbb{E}\left[\left(A - S_{\mathbf{y}\mathbf{x}} S_{\mathbf{x}\mathbf{x}}^{-1}\right) X_* X_*^\top \left(A - S_{\mathbf{y}\mathbf{x}} S_{\mathbf{x}\mathbf{x}}^{-1}\right)^\top\right]$$

and the epistemic uncertainty (B.37b) follows by applying (A.16) again. \square

Proposition B.5 (Uncertainties case 2: V unknown). *Let $\lambda = \{A, V\}$ and $\theta = \{M, K, \nu, \Sigma\}$ be the likelihood and prior parameters as in Section B.2. Given the data Y , the aleatoric and epistemic uncertainty of the predictive Y_* are respectively*

$$\mathbb{E}[\mathbb{V}[Y_*|\lambda] | Y, \theta] = \frac{\text{trace}(\sigma^2(X_*))}{\nu + N - 2p - 2} (\Sigma + S_{\mathbf{y}|\mathbf{x}})$$

$$\mathbb{V}[\mathbb{E}[Y_*|\lambda] | Y, \theta] = \frac{\text{trace}(X_*^\top S_{\mathbf{x}\mathbf{x}}^{-1} X_*)}{\nu + N - 2p - 2} (\Sigma + S_{\mathbf{y}|\mathbf{x}}),$$

and add up to $\mathbb{V}[Y_*|Y, \theta]$.

Proof. Same as the proof of Proposition B.4, it holds $Y_*|\lambda \sim \mathcal{N}(AX_*, V, \sigma^2(X_*))$. Moreover, $V|Y, \theta$ and $A|Y, \theta$ are given respectively by (B.32) and (B.31). The aleatoric uncertainty then follows from the mean of Inverse Wishart distribution (see Appendix A.4), and the epistemic uncertainty follows by applying (A.18). \square

Remark B.1. A rough estimation shows $S_{\mathbf{x}\mathbf{x}}^{-1} = O(1/N)$ and $S_{\mathbf{y}|\mathbf{x}} = O(N)$. Therefore as the number of data samples increases, the effect of the prior decreases in the uncertainties. Moreover, the level of aleatoric uncertainty remains stable but the epistemic uncertainty shrinks to 0. These observations are consistent with those in the literature.

C ELBO-EM ALGORITHM

In this section we develop the ELBO-EM algorithm for the maximization of evidence, also called the type-2 MLE problem. The Expectation Maximization (EM) method is often employed in the context of MLE with missing observations. It is based on the substitution function

$$Q(\xi; \tilde{\xi}) := \mathbb{E}_{\lambda \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} [\ln \mathbb{P}_{\xi}(\lambda, Y)]$$

which can be shown fulfilling, using Jensen inequality,

$$Q(\xi; \tilde{\xi}) - Q(\tilde{\xi}; \tilde{\xi}) \leq \ln \mathbb{P}_{\xi}(Y) - \ln \mathbb{P}_{\tilde{\xi}}(Y)$$

Then at the iteration n , maximizing $Q(\cdot; \xi_{n-1})$ with ξ_{n-1} the parameters from the last iteration will define a sequence $\{\xi_n\}_n$ which monotonically increases the log-evidence. In the context of BLL, we will show that the substitution function can be put into an ELBO form in which the hyperparameters θ and NN's parameters w are separated. The M-step is then carried out by alternating the mini-batches based SGD for w and closed-form update formulae for θ .

EM for MAP EM can also be used for the MAP problem (11). One only needs to modify the substitution function to

$$Q_{\text{MAP}}(\xi; \tilde{\xi}) := \mathbb{E}_{\lambda \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} [\ln \mathbb{P}(\xi, \lambda, Y)] = Q(\xi; \tilde{\xi}) + \ln \pi(\xi)$$

which fulfills consequently

$$Q_{\text{MAP}}(\xi; \tilde{\xi}) - Q_{\text{MAP}}(\tilde{\xi}; \tilde{\xi}) \leq \ln \mathbb{P}(\xi|Y) - \ln \mathbb{P}(\tilde{\xi}|Y),$$

and adapt accordingly the M-step.

C.1 CASE 1: V KNOWN

In this case the ‘‘unobserved’’ quantities is just the system matrix A . Let $\xi = \theta \cup w = \{M, K, w^\phi, w^\sigma\}$, and $\tilde{\xi} = \tilde{\theta} \cup \tilde{w} = \{\tilde{M}, \tilde{K}, \tilde{w}^\phi, \tilde{w}^\sigma\}$. The substitution function is then

$$Q(\xi; \tilde{\xi}) = \mathbb{E}_{A \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} [\ln \mathbb{P}_{\xi}([A, Y])]$$

where the posterior $\mathbb{P}_{\tilde{\xi}}(A|Y) = \mathcal{N}(A; \tilde{S}_{\text{yx}} \tilde{S}_{\text{xx}}^{-1}, V, \tilde{S}_{\text{xx}}^{-1})$ comes from (B.26). Rewriting Q in the ELBO form

$$\begin{aligned} Q(\xi; \tilde{\xi}) &= \underbrace{\mathbb{E}_{A \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} [\ln \mathbb{P}_{\xi}(Y|A)]}_{=: Q_1(\xi; \tilde{\xi})} - \underbrace{\mathbb{D}_{\text{KL}}[\mathbb{P}_{\tilde{\xi}}(A|Y) || \mathbb{P}_{\xi}(A)]}_{=: Q_2(\xi; \tilde{\xi})} + \mathbb{E}_{A \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} [\ln \mathbb{P}_{\tilde{\xi}}(A|Y)] \\ &\simeq Q_1(\xi; \tilde{\xi}) - Q_2(\xi; \tilde{\xi}), \end{aligned}$$

we notice that the hyperparameters θ and the NN's parameters w are separated. Actually, the Q_1 term is

$$\begin{aligned} Q_1(\xi; \tilde{\xi}) &\simeq -\frac{p}{2} \ln |D| - \frac{1}{2} \mathbb{E}_{A \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} \left[\text{trace} \left(V^{-1} (Y - A\Phi) D^{-1} (Y - A\Phi)^\top \right) \right] \\ &= -\frac{p}{2} \ln |D| - \frac{1}{2} \text{trace} \left(V^{-1} \left(Y - \tilde{S}_{\text{yx}} \tilde{S}_{\text{xx}}^{-1} \Phi \right) D^{-1} \left(Y - \tilde{S}_{\text{yx}} \tilde{S}_{\text{xx}}^{-1} \Phi \right)^\top \right) \\ &\quad - \frac{1}{2} \text{trace} \left(V^{-1} \mathbb{E}_{A \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} \left[\left(A - \tilde{S}_{\text{yx}} \tilde{S}_{\text{xx}}^{-1} \right) \Phi D^{-1} \Phi^\top \left(A - \tilde{S}_{\text{yx}} \tilde{S}_{\text{xx}}^{-1} \right)^\top \right] \right) \end{aligned}$$

Using the identity (A.16), the expectation in the last term evaluates to $\text{trace}(\Phi D^{-1} \Phi^\top \tilde{S}_{\text{xx}}^{-1}) V$. Finally we get

$$Q_1(\xi; \tilde{\xi}) \simeq -\frac{p}{2} \ln |D| - \frac{1}{2} \left(\text{trace} \left(V^{-1} \tilde{E} D^{-1} \tilde{E}^\top \right) + p \text{trace} \left(\Phi D^{-1} \Phi^\top \tilde{S}_{\text{xx}}^{-1} \right) \right)$$

with $\tilde{E} := Y - \tilde{S}_{\mathbf{y}\mathbf{x}}\tilde{S}_{\mathbf{x}\mathbf{x}}^{-1}\Phi$, or equivalently in a separable form:

$$Q_1 \simeq -\frac{1}{2} \sum_i \left(p \ln \sigma^2(x_i) + \sigma^{-2}(x_i) \left(\tilde{e}_i^\top V^{-1} \tilde{e}_i + p \phi(x_i)^\top \tilde{S}_{\mathbf{x}\mathbf{x}}^{-1} \phi(x_i) \right) \right).$$

with $\tilde{e}_i := y_i - \tilde{S}_{\mathbf{y}\mathbf{x}}\tilde{S}_{\mathbf{x}\mathbf{x}}^{-1}\phi(x_i)$, which makes Q_1 compatible with mini-batches based SGD.

The following result will be used for the simplification of the Q_2 term:

Lemma C.1. *Let $A \sim p \times n$ follow a matrix Normal distribution $\mathbb{P}_\xi = \mathcal{N}(M, V, K)$ with the parameters $\xi = \{M, V, K\}$. For the parameters $\tilde{\xi} = \{\tilde{M}, \tilde{V}, \tilde{K}\}$, the Kullback-Leibler divergence between $\mathbb{P}_{\tilde{\xi}}(A)$ and $\mathbb{P}_\xi(A)$ equals to, up to some constant which is independent of ξ :*

$$\begin{aligned} \mathbb{D}_{\text{KL}} \left[\mathbb{P}_{\tilde{\xi}}(A) \parallel \mathbb{P}_\xi(A) \right] \simeq & \frac{n}{2} \ln |V| + \frac{p}{2} \ln |K| + \frac{1}{2} \text{trace} \left(V^{-1} \tilde{V} \right) \text{trace} \left(K^{-1} \tilde{K} \right) + \\ & \frac{1}{2} \text{trace} \left(V^{-1} \left(M - \tilde{M} \right) K^{-1} \left(M - \tilde{M} \right)^\top \right) \end{aligned}$$

Proof. Skipped. □

In our case $V = \tilde{V}$ is fixed, so we get

$$Q_2(\xi; \tilde{\xi}) \simeq \frac{p}{2} \ln |K| + \frac{p}{2} \text{trace} \left(K^{-1} \tilde{S}_{\mathbf{x}\mathbf{x}}^{-1} \right) + \frac{1}{2} \text{trace} \left(V^{-1} \tilde{F} K^{-1} \tilde{F}^\top \right)$$

with the shorthand $\tilde{F} := M - \tilde{S}_{\mathbf{y}\mathbf{x}}\tilde{S}_{\mathbf{x}\mathbf{x}}^{-1}$. Matrix calculus gives

$$\begin{aligned} \frac{\partial Q_2}{\partial M} &= V^{-1} \left(M - \tilde{S}_{\mathbf{y}\mathbf{x}}\tilde{S}_{\mathbf{x}\mathbf{x}}^{-1} \right) K^{-1} \\ \frac{\partial Q_2}{\partial K} &= \frac{p}{2} \left(K^{-1} - K^{-1} \tilde{S}_{\mathbf{x}\mathbf{x}}^{-1} K^{-1} \right) - \frac{1}{2} K^{-1} \tilde{F}^\top V^{-1} \tilde{F} K^{-1} \end{aligned}$$

Now setting $\partial_M Q_2 = 0$ and $\partial_K Q_2 = 0$ jointly gives

$$M = \tilde{S}_{\mathbf{y}\mathbf{x}}\tilde{S}_{\mathbf{x}\mathbf{x}}^{-1}, \quad K = \tilde{S}_{\mathbf{x}\mathbf{x}}^{-1}.$$

C.2 CASE 2: V UNKNOWN

In this case the ‘‘unobserved’’ quantities is the system matrix A and the baseline noise covariance V . Let $\xi = \theta \cup w = \{M, \Sigma, K, w^\phi, w^\sigma\}$, and $\tilde{\xi} = \tilde{\theta} \cup \tilde{w} = \{\tilde{M}, \tilde{\Sigma}, \tilde{K}, \tilde{w}^\phi, \tilde{w}^\sigma\}$. The substitution function is then

$$\begin{aligned} Q(\xi; \tilde{\xi}) &= \mathbb{E}_{[A, V] \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} \left[\ln \mathbb{P}_\xi([A, V, Y]) \right] \\ &= \underbrace{\mathbb{E}_{V \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} \left[\mathbb{E}_{A \sim \mathbb{P}_{\tilde{\xi}}(\cdot|V, Y)} \left[\ln \mathbb{P}_\xi(Y|A, V) \right] \right]}_{=: Q_1(\xi; \tilde{\xi})} - \underbrace{\mathbb{D}_{\text{KL}} \left[\mathbb{P}_{\tilde{\xi}}([A, V]|Y) \parallel \mathbb{P}_\xi([A, V]) \right]}_{=: Q_2(\xi; \tilde{\xi})} \\ &\quad + \mathbb{E}_{[A, V] \sim \mathbb{P}_{\tilde{\xi}}(\cdot|Y)} \left[\ln \mathbb{P}_{\tilde{\xi}}([A, V]|Y) \right] \\ &\simeq Q_1(\xi; \tilde{\xi}) - Q_2(\xi; \tilde{\xi}) \end{aligned}$$

where the posterior $\mathbb{P}_{\tilde{\xi}}([A, V]|Y) = \mathcal{W}^{-1} \left(V; \tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}}, \nu + N \right) \times \mathcal{N} \left(A; \tilde{S}_{\mathbf{y}\mathbf{x}}\tilde{S}_{\mathbf{x}\mathbf{x}}^{-1}, V, \tilde{S}_{\mathbf{x}\mathbf{x}}^{-1} \right)$ comes from (B.32).

For a p -dimensional covariance matrix $W \sim \mathcal{W}^{-1}(\Sigma, \nu)$, it holds $\mathbb{E}[W^{-1}] = (\nu - p - 1)\Sigma^{-1}$. Using this fact allows us to obtain the Q_1 term from the case 1 above. We get

$$Q_1 \simeq -\frac{p}{2} \ln |D| - \frac{\nu_p + N}{2} \text{trace} \left(\left(\tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}} \right)^{-1} \tilde{E} D^{-1} \tilde{E}^\top \right) - \frac{p}{2} \text{trace} \left(\Phi D^{-1} \Phi^\top \tilde{S}_{\mathbf{x}\mathbf{x}}^{-1} \right)$$

or equivalently in a separable form:

$$Q_1 \simeq -\frac{1}{2} \sum_i \left(p \ln \sigma^2(x_i) + \sigma^{-2}(x_i) \left((\nu_p + N) \tilde{e}_i^\top \left(\tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}} \right)^{-1} \tilde{e}_i + p \phi(x_i)^\top \tilde{S}_{\mathbf{xx}}^{-1} \phi(x_i) \right) \right).$$

where $\nu_p = \nu - p - 1$.

The following result will be used for the simplification of the Q_2 term:

Lemma C.2. *Let $V \sim p \times p$ follow an Inverse Wishart distribution $\mathcal{W}^{-1}(\Sigma, \nu)$ and $A \sim p \times n$ conditioned on V follow a matrix Normal distribution $\mathcal{N}(M, V, K)$. Denote by $\xi = \{M, K, \Sigma, \nu\}$ and $\tilde{\xi} = \{\tilde{M}, \tilde{K}, \tilde{\Sigma}, \tilde{\nu}\}$. Then the KL-divergence between the joint distribution $\mathbb{P}_{\tilde{\xi}}(A, V)$ and $\mathbb{P}_{\xi}(A, V)$ equals to, up to some constant which is independent of ξ :*

$$\mathbb{D}_{\text{KL}} \left[\mathbb{P}_{\tilde{\xi}}([A, V]) \parallel \mathbb{P}_{\xi}([A, V]) \right] \simeq \frac{p}{2} \left(\ln |K| + \text{trace} \left(K^{-1} \tilde{K} \right) \right) + \frac{\tilde{\nu}_p}{2} \text{trace} \left\{ \tilde{\Sigma}^{-1} \left((M - \tilde{M}) K^{-1} (M - \tilde{M})^\top + \Sigma \right) \right\} - \frac{\nu_p}{2} \ln |\Sigma|$$

with $\tilde{\nu}_p := \tilde{\nu} - p - 1, \nu_p := \nu - p - 1$.

Proof. Skipped. □

Applying this on Q_2 with the prior $\mathbb{P}_{\xi}([A, V]) = \mathcal{W}^{-1}(V; \Sigma, \nu) \times \mathcal{N}(A; M, V, K)$ gives

$$Q_2(\xi; \tilde{\xi}) \cong \frac{p}{2} \left(\ln |K| + \text{trace} \left(K^{-1} \tilde{S}_{\mathbf{xx}}^{-1} \right) \right) - \frac{\nu_p}{2} \ln |\Sigma| + \frac{\nu_p + N}{2} \text{trace} \left(\left(\tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}} \right)^{-1} \left(\tilde{F} K^{-1} \tilde{F}^\top + \Sigma \right) \right)$$

Matrix calculus gives

$$\begin{aligned} \frac{\partial Q_2}{\partial M} &= (\nu_p + N) \left(\tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}} \right)^{-1} \left(M - \tilde{S}_{\mathbf{yx}} \tilde{S}_{\mathbf{xx}}^{-1} \right) K^{-1} \\ \frac{\partial Q_2}{\partial K} &= \frac{p}{2} \left(K^{-1} - K^{-1} \tilde{S}_{\mathbf{xx}}^{-1} K^{-1} \right) - \frac{\nu_p + N}{2} K^{-1} \tilde{F}^\top \left(\tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}} \right)^{-1} \tilde{F} K^{-1} \\ \frac{\partial Q_2}{\partial \Sigma} &= -\frac{\nu_p}{2} \Sigma^{-1} + \frac{\nu_p + N}{2} \left(\tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}} \right)^{-1} \end{aligned}$$

Now setting these derivatives to 0 jointly gives

$$M = \tilde{S}_{\mathbf{yx}} \tilde{S}_{\mathbf{xx}}^{-1}, \quad K = \tilde{S}_{\mathbf{xx}}^{-1}, \quad \Sigma = \frac{\nu_p}{\nu_p + N} \left(\tilde{\Sigma} + \tilde{S}_{\mathbf{y}|\mathbf{x}} \right)$$

D MLE FOR BLL

In this section we study the maximum likelihood estimators for the BLL models.

D.1 CASE 1: V KNOWN

Recall that the log-evidence (3) in this case reads

$$\mathcal{L} = \text{cst} - \frac{N}{2} \ln |V| - \frac{p}{2} \ln |\Omega| - \frac{1}{2} \text{trace} (V^{-1} E \Omega^{-1} E^\top)$$

with the shorthands

$$\Omega := D + \Phi^\top K \Phi, \quad E := Y - M \Phi$$

Using matrix calculus to get the following derivatives

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial M} &= V^{-1} E \Omega^{-1} \Phi^\top \\ \frac{\partial \mathcal{L}}{\partial V} &= -\frac{1}{2} (N V^{-1} - V^{-1} E \Omega^{-1} E^\top V^{-1}) \\ \frac{\partial \mathcal{L}}{\partial \Omega} &= -\frac{1}{2} (p \Omega^{-1} - \Omega^{-1} E^\top V^{-1} E \Omega^{-1}) \\ \frac{\partial \mathcal{L}}{\partial K} &= \Phi \frac{\partial \mathcal{L}}{\partial \Omega} \Phi^\top \end{aligned}$$

Setting $\frac{\partial \mathcal{L}}{\partial M}$, $\frac{\partial \mathcal{L}}{\partial V}$, $\frac{\partial \mathcal{L}}{\partial \Omega}$ to 0 yields

$$\begin{aligned} \hat{M} &= Y \Omega^{-1} \Phi^\top (\Phi \Omega^{-1} \Phi^\top)^{-1} \\ \hat{V} &= \frac{1}{N} E \Omega^{-1} E^\top \\ \hat{\Omega} &= \frac{1}{p} E^\top V^{-1} E \end{aligned}$$

which are unique unbiased estimator if K , ϕ , σ are fixed.

Assuming $\Phi \Phi^\top$ invertible, then a sufficient condition for $\frac{\partial \mathcal{L}}{\partial K} = 0$ is $\frac{\partial \mathcal{L}}{\partial \Omega} = 0$, which yields

$$\hat{K} = \Phi^\dagger{}^\top \left(\frac{1}{p} E^\top V^{-1} E - D \right) \Phi^\dagger$$

where we defined the generalized inverse $\Phi^\dagger := \Phi^\top (\Phi \Phi^\top)^{-1}$. However this estimate for K is numerically unreliable: in fact, the middle term may fail to be positive definite. The ELBO-EM algorithm, on the other hand, has a update formula for K with guaranteed positive definiteness.

D.2 CASE 2: V UNKNOWN

First, introduce the following shorthands:

$$E := Y - M \Phi, \quad H := \Sigma + E \Omega^{-1} E^\top, \quad C_\nu := \frac{\Gamma_p \left(\frac{\nu_p + N}{2} \right)}{\Gamma_p \left(\frac{\nu_p}{2} \right) \pi^{\frac{Np}{2}}}, \quad \nu_p := \nu - p - 1$$

The log-evidence (B.29) can be rewritten as

$$\mathcal{L} = \ln C_\nu + \frac{\nu_p}{2} \ln |\Sigma| - \frac{p}{2} \ln |\Omega| - \frac{\nu_p + N}{2} \ln |H|.$$

and matrix calculus gives

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial M} &= (\nu_p + N)H^{-1}E\Omega^{-1}\Phi^\top \\ \frac{\partial \mathcal{L}}{\partial \Sigma} &= \frac{1}{2}(\nu_p \Sigma^{-1} - (\nu_p + N)H^{-1}) \\ \frac{\partial \mathcal{L}}{\partial K} &= \frac{1}{2}\Phi(-p\Omega^{-1} + (\nu_p + N)\Omega^{-1}E^\top H^{-1}E\Omega^{-1})\Phi^\top\end{aligned}$$

Setting the derivatives of M, Σ to zero yields the unique estimates

$$\begin{aligned}\hat{M} &= Y\Omega^{-1}\Phi^\top (\Phi\Omega^{-1}\Phi^\top)^{-1} \\ \hat{\Sigma} &= \frac{\nu_p}{N}E\Omega^{-1}E^\top\end{aligned}$$

Notice that \hat{M} is the same as in case 1, while the estimate $\hat{\Sigma}$ is biased. In fact, taking the expectation about Y and using (A.18) gives

$$\mathbb{E}[\hat{\Sigma}] = \left(\frac{\nu - p - 1}{\nu - 2p - 2}\right)\Sigma.$$

so the unbiased correction is

$$\hat{\Sigma} = \left(\frac{\nu - 2p - 2}{N}\right)E\Omega^{-1}E^\top$$

Using Woodbury identity we find a sufficient condition for $\frac{\partial \mathcal{L}}{\partial K} = 0$ to be

$$D + \Phi^\top K\Phi = \left(\frac{\nu_p + N - p}{p}\right)E^\top \Sigma^{-1}E$$

Assuming $\Phi\Phi^\top$ invertible, we obtain also an estimate for K :

$$\hat{K} = \Phi^\dagger \left(\left(\frac{\nu_p + N - p}{p} \right) E^\top \Sigma^{-1} E - D \right) \Phi^\dagger$$

which is biased and numerically unreliable as in case 1.

E NUMERICAL EXPERIMENTS

E.1 ELBO-EM VS. NATIVE SGD

As explained in Section 3.1, with the parameters w and K being fixed the MLE has the unique solution (12). This setting allows us to compare the convergence of ELBO-EM with that of a native SGD, on the estimation of the hyperparameters M, Σ . The native SGD here minimizes the log-evidence on mini-batches with respect to M, Σ using automatic differentiation (AD), despite the inseparable form of samples which can incur biased gradient estimate as already mentioned.. In our experiment (see Appendix E.1) we observed via Monte-Carlo tests that the convergence of ELBO-EM to the MLE solution is much faster than that of the native SGD method.

Data samples for this experiment are generated in an hierarchical way following the model (1). We fix the dimensions $p = 3, n = 4$, and the degree of freedom $\nu = 2p + 1$. The matrices M, V, K are sampled from the standard Normal and Wishart distributions, and are used to draw the system matrix $A \sim \mathcal{N}(M, V, K)$. For simplicity $\phi(x) = x$ is taken as basis in this experiment. We then generate 100 input samples as $x_i \stackrel{\text{iid}}{\sim} \mathcal{N}_{\text{vec}}(0, 10I_n)$ and y_i from the model (1) with a homoscedastic noise $\sigma(x) \equiv 1$. Both ELBO-EM and SGD are initialized with $M = 0, \Sigma = I_p$ and with a K fixed during iteration, and run until convergence. We observed the estimated M, Σ remain close to the MLE solution whatever how K is fixed. Their results are then compared to the MLE solution computed using (12) by evaluating the relative error of Frobenius norm $\|A - B\|_F / \|B\|_F$.

The experiment is repeated 1000 times using always $K = I_n$ as fixed value and ELBO-EM converges in general within 7 steps (tolerance $< 10^{-6}$), and the average relative error of estimate is of order 10^{-4} for M and 10^{-5} for Σ . On the other hand, the native SGD method converges much slower.

E.2 ELBO-EM

We provide more results of the ELBO-EM algorithm using the cubic and the sinus function. Two independent MLPs of 3 layers with 32 hidden units are trained jointly with the BLL matrices M, K . The results are show in Figure 5.

E.3 TRANSFER LEARNING

In this example a NN is first trained on 500 samples of the sinus function, then transferred to the data of the cubic function with 200 samples, by freezing all precedent layers and replacing the last layer by a BLL. The result is shown in Figure 6.

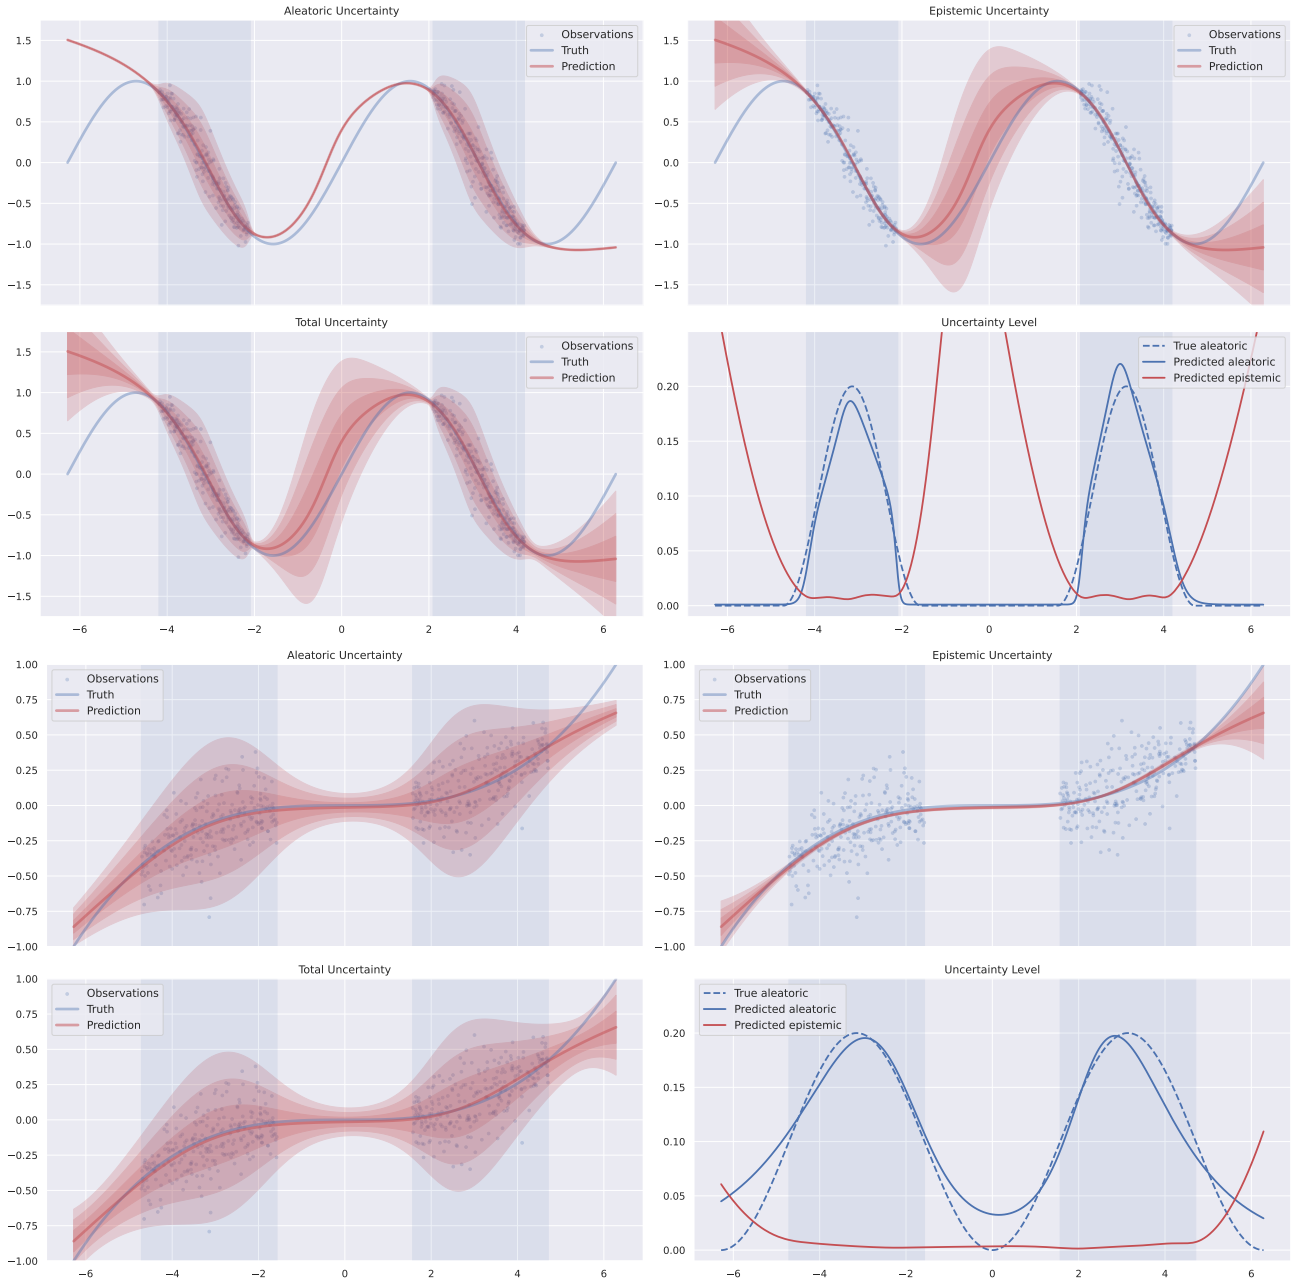


Figure 5: ELBO-EM training on the sinus and the cubic function, using 500 samples.

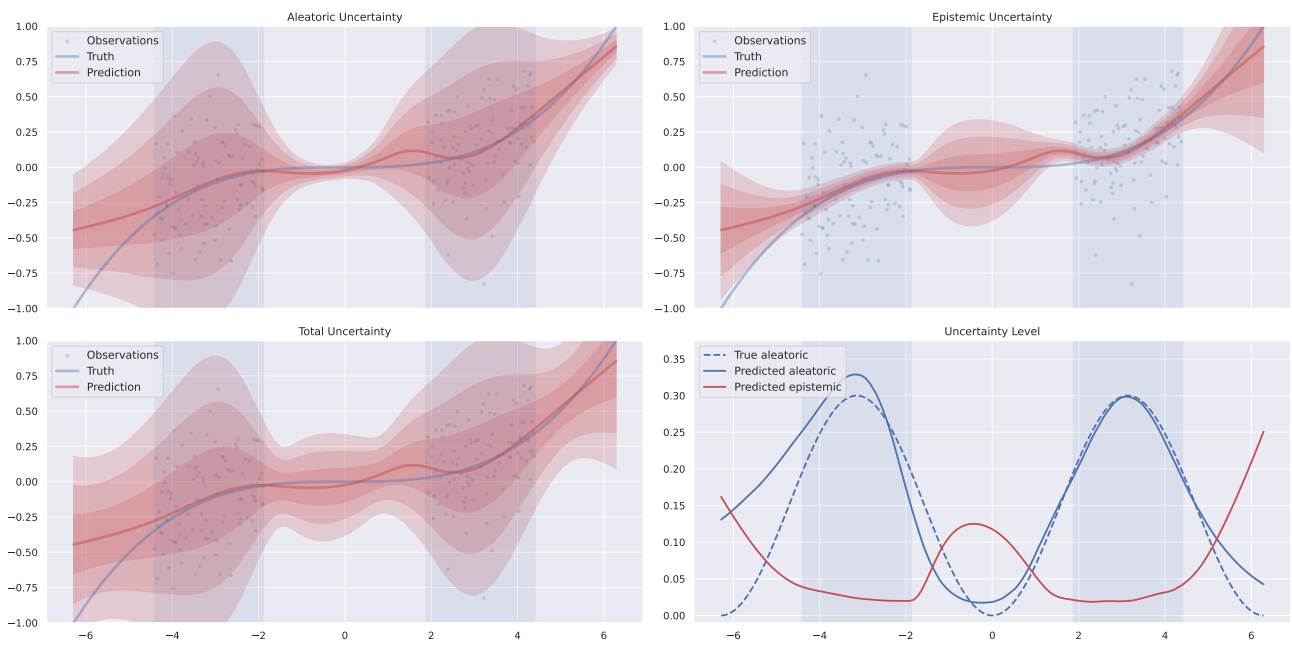


Figure 6: Transfer learning.