



Linear regression in the Bayesian framework

Thierry A. Mara

► **To cite this version:**

| Thierry A. Mara. Linear regression in the Bayesian framework. 2019. hal-02264944

HAL Id: hal-02264944

<https://hal.archives-ouvertes.fr/hal-02264944>

Submitted on 7 Aug 2019

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.

Linear regression in the Bayesian framework

Thierry Alex Mara*

Keywords: linear regression, Bayesian inference, generalized least-squares, ridge regression, LASSO, model selection criterion, KIC

Abstract

These notes aim at clarifying different strategies to perform linear regression from given dataset. Methods like the weighted and ordinary least squares, ridge regression or LASSO are proposed in the literature. The present article is my understanding of these methods which are, according to me, better unified in the Bayesian framework. The formulas to address linear regression with these methods are derived. The KIC for model selection is also derived in the end of the document.

*University of Reunion, [thierry.mara\(at\)univ-reunion.fr](mailto:thierry.mara@univ-reunion.fr)

1 Introduction

Let's consider an input-output relationship of the form:

$$y = \sum_{\alpha \in \mathcal{A}} a_{\alpha} \psi_{\alpha}(\mathbf{x}) + \epsilon \quad (1)$$

where $\alpha = \alpha_1 \dots \alpha_d$ is a multi-dimensional index, $\mathcal{A} \subset \mathbb{N}^d$ is a finite subset of multi-indexes and $\psi_{\alpha}(\mathbf{x})$ is some function independent of x_i if $\alpha_i = 0$, and ϵ an error term. Given the sets of observations $\mathbf{y} = \{y_n\}_{n=1}^N$ and $\mathbf{x} = \{x_{n1}, \dots, x_{nd}\}_{n=1}^N$, the matrix form of (1) is,

$$\mathbf{y} = \Psi \mathbf{a} + \boldsymbol{\epsilon}. \quad (2)$$

The error vector $\boldsymbol{\epsilon}$ accounts for i) possible *stochastic* error ϵ_s (e.g. measurement noise, stochastic model noise), ii) sampling error ϵ_d (as (\mathbf{y}, \mathbf{x}) is just one possible sample), and iii) truncation error ϵ_t (because $\text{Card}(\mathcal{A})$ is finite). We write $\epsilon = \epsilon_s + \epsilon_d + \epsilon_t$. We note that when one has to deal with y as the response of a *deterministic* mathematical model, stochastic errors $\epsilon_s = 0$. Sampling error can be assessed with several samples (\mathbf{y}, \mathbf{x}) . In practice, this is hardly possible, and one can refer to the *bootstrap* technique to assess ϵ_d [Efron and Tibshirani, 1993]. Both ϵ_d and ϵ_t are reducible by increasing the number of observations N (and the diversity of the observations of course).

The first issue addressed in this short note is: *how can we infer $\mathbf{a} = \{a_{\alpha} : \alpha \in \mathcal{A}\}$, given (\mathbf{y}, \mathbf{x}) and \mathcal{A} ?* The second one is *how to infer the best subset \mathcal{A} given (\mathbf{y}, \mathbf{x}) ?* These issues are addressed in a Bayesian framework which requires some assumptions about the error's probability density function (called the *likelihood function*) and some information about the unknown coefficients themselves (called *prior belief*). Having assumed the likelihood is $\epsilon \sim p_{\epsilon}(\epsilon)$, we can also write that $(y - \sum_{\alpha \in \mathcal{A}} a_{\alpha} \psi_{\alpha}(\mathbf{x})) \sim p_{\epsilon}(y - \sum_{\alpha \in \mathcal{A}} a_{\alpha} \psi_{\alpha}(\mathbf{x}))$ which is usually written: $p_{\mathbf{y}|\mathbf{a}, \mathcal{A}} = p(\mathbf{y}|\mathbf{a}, \mathcal{A})$ or simply $\mathcal{L}(\mathbf{y}|\mathbf{a}, \mathcal{A})$ (the \mathcal{L} standing for likelihood). The prior is denoted $p_{\mathbf{a}|\mathcal{A}} = p(\mathbf{a}|\mathcal{A})$.

The first issue is addressed by assessing the coefficient joint posterior density function (pdf inferred from Bayes rule), namely

$$p(\mathbf{a}|\mathbf{y}, \mathcal{A}) = \frac{p(\mathbf{y}|\mathbf{a}, \mathcal{A})p(\mathbf{a}|\mathcal{A})}{p(\mathbf{y}|\mathcal{A})}. \quad (3)$$

Because (1) is linear w.r.t. the vector of coefficients \mathbf{a} , straightforward calculations can be employed as opposed to more demanding approach like Markov chains Monte Carlo. These straightforward calculations are derived in the present document.

The second issue can be addressed by evaluating the Bayesian Model Evidence (BME) defined as:

$$p(\mathbf{y}|\mathcal{A}) = \int_{\mathbb{R}^P} p(\mathbf{y}, \mathbf{a}|\mathcal{A})d\mathbf{a} \quad (4)$$

where $P = \text{Card}(\mathcal{A})$ is the number of coefficients. Indeed, Eq.(4) represents how well the current model $(\mathbf{a}, \mathcal{A})$ explains (fits) the observed data \mathbf{y} . Finding the subset \mathcal{A} that maximizes (4) provides the answer to the second question.

2 Gaussian Likelihood

Let us assume that $p_\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_M)$, this leads to the following Gaussian likelihood:

$$p(\mathbf{y}|\mathbf{a}, \mathcal{A}) = (2\pi)^{-N/2}|\mathbf{C}_M|^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{y} - \Psi\mathbf{a})^t \mathbf{C}_M^{-1}(\mathbf{y} - \Psi\mathbf{a})\right) \quad (5)$$

where $|\cdot|$ stands for the determinant.

2.1 Uniform Prior & Maximum Likelihood Estimate

Let us assume further that the prior knowledge about the coefficient values is a rectangular domain $\Omega = [l_1, u_1] \times \dots \times [l_d, u_d]$. This means that $p_{\mathbf{a}|\mathcal{A}} = cste$ if $\mathbf{a} \in \Omega$, and zero elsewhere. As a consequence, the joint pdf of the coefficients is,

$$p(\mathbf{a}|\mathbf{y}, \mathcal{A}) \propto p(\mathbf{y}|\mathbf{a}, \mathcal{A}) \propto \exp\left(-\frac{1}{2}(\mathbf{y} - \Psi\mathbf{a})^t \mathbf{C}_M^{-1}(\mathbf{y} - \Psi\mathbf{a})\right), \mathbf{a} \in \Omega \quad (6)$$

Analytical Solution: The term in the exponential in (6),

$$(\mathbf{y} - \Psi\mathbf{a})^t \mathbf{C}_M^{-1}(\mathbf{y} - \Psi\mathbf{a}) = \underbrace{\mathbf{y}^t \mathbf{C}_M^{-1} \mathbf{y}}_{cste/\mathbf{a}} - \underbrace{\mathbf{y}^t \mathbf{C}_M^{-1} \Psi \mathbf{a}}_{=\mathbf{a}^t \Psi^t \mathbf{C}_M^{-1} \mathbf{y}} - \mathbf{a}^t \Psi^t \mathbf{C}_M^{-1} \mathbf{y} + \mathbf{a}^t \Psi^t \mathbf{C}_M^{-1} \Psi \mathbf{a}$$

the second underbrace true because \mathbf{C}_M is symmetric (and positive-definite) by definition. So, we get

$$\begin{aligned} (\mathbf{y} - \Psi\mathbf{a})^t \mathbf{C}_M^{-1}(\mathbf{y} - \Psi\mathbf{a}) &= \mathbf{a}^t \Psi^t \mathbf{C}_M^{-1} \Psi \mathbf{a} - 2\mathbf{a}^t \Psi^t \mathbf{C}_M^{-1} \mathbf{y} + c \\ &= \left([\Psi^t \mathbf{C}_M^{-1} \Psi]^{1/2} \mathbf{a} - [\Psi^t \mathbf{C}_M^{-1} \Psi]^{-1/2} \Psi^t \mathbf{C}_M^{-1} \mathbf{y}\right)^t \left([\Psi^t \mathbf{C}_M^{-1} \Psi]^{1/2} \mathbf{a} - [\Psi^t \mathbf{C}_M^{-1} \Psi]^{-1/2} \Psi^t \mathbf{C}_M^{-1} \mathbf{y}\right) + c \end{aligned}$$

Factorizing by $[\Psi^t \mathbf{C}_M^{-1} \Psi]^{1/2}$ in both parentheses yields,

$$(\mathbf{y} - \Psi \mathbf{a})^t \mathbf{C}_M^{-1} (\mathbf{y} - \Psi \mathbf{a}) = \left(\mathbf{a} - [\Psi^t \mathbf{C}_M^{-1} \Psi]^{-1} \Psi^t \mathbf{C}_M^{-1} \mathbf{y} \right)^t \Psi^t \mathbf{C}_M^{-1} \Psi \left(\mathbf{a} - [\Psi^t \mathbf{C}_M^{-1} \Psi]^{-1} \Psi^t \mathbf{C}_M^{-1} \mathbf{y} \right) + C$$

Replacing this result in (6) yields,

$$p(\mathbf{a} | \mathbf{y}, \mathcal{A}) = (2\pi)^{-P/2} |\tilde{\mathbf{C}}_{aa}|^{-1/2} \exp \left(-\frac{1}{2} (\mathbf{a} - \tilde{\mathbf{a}})^t \tilde{\mathbf{C}}_{aa}^{-1} (\mathbf{a} - \tilde{\mathbf{a}}) \right) \quad (7)$$

where

$\tilde{\mathbf{a}} = [\Psi^t \mathbf{C}_M^{-1} \Psi]^{-1} \Psi^t \mathbf{C}_M^{-1} \mathbf{y}$ is called the Maximum Likelihood Estimate (MLE)
 $\tilde{\mathbf{C}}_{aa} = [\Psi^t \mathbf{C}_M^{-1} \Psi]^{-1}$ is the covariance associated to \mathbf{a} and $P = \text{Card}(\mathcal{A})$ is the number of coefficients.

N.B.: Because the coefficients have been constrained within a finite rectangular domain Ω , and \mathbf{C}_M is assumed given, the posterior joint pdf should be written $p(\mathbf{a} | \mathbf{y}, \mathcal{A}, \mathbf{C}_M) = \mathcal{N}(\tilde{\mathbf{a}}, \tilde{\mathbf{C}}_{aa})$ if $\tilde{\mathbf{a}} \in \Omega$ and $p(\mathbf{a} | \mathbf{y}, \mathcal{A}, \mathbf{C}_M) = 0$ elsewhere.

2.2 Homoscedastic Gaussian Error & Ordinary Least Squares

Assuming homoscedastic Gaussian error, that is, setting $\mathbf{C}_M = \sigma_\epsilon^2 \mathbf{I}_N$ where \mathbf{I}_N stands for $N \times N$ identity matrix, yields

$$p(\mathbf{a} | \mathbf{y}, \mathcal{A}, \sigma_\epsilon^2) = \mathcal{N}(\tilde{\mathbf{a}}, \tilde{\mathbf{C}}_{aa})$$

where

$\tilde{\mathbf{a}} = [\sigma_\epsilon^{-2} \Psi^t \mathbf{I}_N \Psi]^{-1} \sigma_\epsilon^{-2} \Psi^t \mathbf{y} = [\Psi^t \Psi]^{-1} \Psi^t \mathbf{y}$ which is the Ordinary Least-Square estimator,

and $\tilde{\mathbf{C}}_{aa} = \sigma_\epsilon^2 [\Psi^t \Psi]^{-1}$.

In practice σ_ϵ^2 is unknown, and because (5) becomes in this case,

$$p(\mathbf{y} | \mathbf{a}, \mathcal{A}, \sigma_\epsilon^2) = (2\pi\sigma_\epsilon^2)^{-N/2} \exp \left(-\frac{1}{2\sigma_\epsilon^2} (\mathbf{y} - \Psi \mathbf{a})^t (\mathbf{y} - \Psi \mathbf{a}) \right) \quad (8)$$

$$\Leftrightarrow -2 \ln(p(\mathbf{y} | \tilde{\mathbf{a}}, \mathcal{A}, \sigma_\epsilon^2)) = c + N \ln(\sigma_\epsilon^2) + \frac{(\mathbf{y} - \Psi \tilde{\mathbf{a}})^t (\mathbf{y} - \Psi \tilde{\mathbf{a}})}{\sigma_\epsilon^2}$$

we can infer that its MLE $\tilde{\sigma}_\epsilon^2$ is given by,

$$-2 \left. \frac{d \ln(p(\mathbf{y} | \tilde{\mathbf{a}}, \mathcal{A}, \sigma_\epsilon^2))}{d\sigma_\epsilon^2} \right|_{\sigma_\epsilon^2 = \tilde{\sigma}_\epsilon^2} = 0 = \frac{N}{\tilde{\sigma}_\epsilon^2} - \frac{(\mathbf{y} - \Psi \tilde{\mathbf{a}})^t (\mathbf{y} - \Psi \tilde{\mathbf{a}})}{\tilde{\sigma}_\epsilon^4}$$

$$\Leftrightarrow \tilde{\sigma}_\epsilon^2 = \frac{(\mathbf{y} - \Psi \tilde{\mathbf{a}})^t (\mathbf{y} - \Psi \tilde{\mathbf{a}})}{N}$$

By noticing that (8) yields

$$p(\sigma_\epsilon^2 | \mathbf{y}, \tilde{\mathbf{a}}, \mathcal{A}) = (2\pi)^{-N/2} (\sigma_\epsilon^{-2})^{N/2} \exp\left(-\frac{(\mathbf{y} - \Psi \tilde{\mathbf{a}})^t (\mathbf{y} - \Psi \tilde{\mathbf{a}})}{2} \sigma_\epsilon^{-2}\right) \quad (9)$$

we can infer that,

$$\boxed{p(\sigma_\epsilon^{-2} | \mathbf{y}, \tilde{\mathbf{a}}, \mathcal{A}) \propto (\sigma_\epsilon^{-2})^{k-1} \exp\left(-\frac{\sigma_\epsilon^{-2}}{\theta}\right) = \Gamma\left(\frac{N+2}{2}, \frac{2}{(\mathbf{y} - \Psi \tilde{\mathbf{a}})^t (\mathbf{y} - \Psi \tilde{\mathbf{a}})}\right)} \quad (10)$$

known as Gamma distribution whose mode is $\tilde{\sigma}_\epsilon^2$. Therefore, under homoscedastic Gaussian likelihood assumption the MLE of the posterior covariance is, $\tilde{\mathbf{C}}_{aa} = \tilde{\sigma}_\epsilon^2 [\Psi^t \Psi]^{-1}$ with still $\tilde{\mathbf{a}} = [\Psi^t \Psi]^{-1} \Psi^t \mathbf{y}$.

2.3 Gaussian Prior & Ridge Regression

Let us assume now that $p_{\mathbf{a}|\mathcal{A}} = \mathcal{N}(\mathbf{a}_0, \mathbf{C}_{aa})$. We note that Ω is no longer a rectangular domain but mostly a hyper-ellipsoid. Then, the joint posterior distribution of the coefficients is written,

$$p(\mathbf{a} | \mathbf{y}, \mathcal{A}) \propto p(\mathbf{y} | \mathbf{a}, \mathcal{A}) p(\mathbf{a} | \mathcal{A}) \quad (11)$$

$$\Leftrightarrow p(\mathbf{a} | \mathbf{y}, \mathcal{A}) \propto \exp\left(-\frac{1}{2} [(\mathbf{y} - \Psi \mathbf{a})^t \mathbf{C}_M^{-1} (\mathbf{y} - \Psi \mathbf{a}) + (\mathbf{a} - \mathbf{a}_0)^t \mathbf{C}_{aa}^{-1} (\mathbf{a} - \mathbf{a}_0)]\right) \quad (12)$$

Analytical Solution: Let us once again develop the term in the exponential, we get

$$\begin{aligned} (\mathbf{y} - \Psi \mathbf{a})^t \mathbf{C}_M^{-1} (\mathbf{y} - \Psi \mathbf{a}) + (\mathbf{a} - \mathbf{a}_0)^t \mathbf{C}_{aa}^{-1} (\mathbf{a} - \mathbf{a}_0) &= \mathbf{a}^t \Psi^t \mathbf{C}_M^{-1} \Psi \mathbf{a} - 2\mathbf{a}^t \Psi^t \mathbf{C}_M^{-1} \mathbf{y} \\ &\quad + \mathbf{a}^t \mathbf{C}_{aa}^{-1} \mathbf{a} - 2\mathbf{a}^t \mathbf{C}_{aa}^{-1} \mathbf{a}_0 + c \end{aligned}$$

that we can rearrange as follows,

$$\begin{aligned} (\mathbf{y} - \Psi \mathbf{a})^t \mathbf{C}_M^{-1} (\mathbf{y} - \Psi \mathbf{a}) + (\mathbf{a} - \mathbf{a}_0)^t \mathbf{C}_{aa}^{-1} (\mathbf{a} - \mathbf{a}_0) &= \\ &= \mathbf{a}^t (\Psi^t \mathbf{C}_M^{-1} \Psi + \mathbf{C}_{aa}^{-1}) \mathbf{a} - 2\mathbf{a}^t (\Psi^t \mathbf{C}_M^{-1} \mathbf{y} + \mathbf{C}_{aa}^{-1} \mathbf{a}_0) + c \\ &= \left[[\Psi^t \mathbf{C}_M^{-1} \Psi + \mathbf{C}_{aa}^{-1}]^{1/2} \mathbf{a} - [\Psi^t \mathbf{C}_M^{-1} \Psi + \mathbf{C}_{aa}^{-1}]^{-1/2} (\Psi^t \mathbf{C}_M^{-1} \mathbf{y} + \mathbf{C}_{aa}^{-1} \mathbf{a}_0) \right]^2 + c \\ &= \left[\left(\mathbf{a} - [\Psi^t \mathbf{C}_M^{-1} \Psi + \mathbf{C}_{aa}^{-1}]^{-1} (\Psi^t \mathbf{C}_M^{-1} \mathbf{y} + \mathbf{C}_{aa}^{-1} \mathbf{a}_0) \right)^t [\Psi^t \mathbf{C}_M^{-1} \Psi + \mathbf{C}_{aa}^{-1}]^{1/2} \right]^2 + c \end{aligned}$$

with $[\cdot]^2 = [\cdot]^t[\cdot]$. We conclude that

$$\boxed{p(\mathbf{a}|\mathbf{y}, \mathcal{A}) = (2\pi)^{-P/2} |\hat{\mathbf{C}}_{aa}|^{-1/2} \exp\left(-\frac{1}{2} (\mathbf{a} - \hat{\mathbf{a}})^t \hat{\mathbf{C}}_{aa}^{-1} (\mathbf{a} - \hat{\mathbf{a}})\right)} \quad (13)$$

where

$\hat{\mathbf{a}} = [\Psi^t \mathbf{C}_M^{-1} \Psi + \mathfrak{m} \mathbf{C}_{aa}^{-1}]^{-1} (\Psi^t \mathbf{C}_M^{-1} \mathbf{y} + \mathbf{C}_{aa}^{-1} \mathbf{a}_0)$ is called the Maximum A Posteriori estimate (MAP). When $\mathbf{C}_{aa} = \lambda \mathbf{I}_P$ this solution is known as the ridge regression estimator (Hoerl and Kennard [1970]).

$\hat{\mathbf{C}}_{aa} = [\Psi^t \mathbf{C}_M^{-1} \Psi + \mathbf{C}_{aa}^{-1}]^{-1}$ is the covariance associated to $\hat{\mathbf{a}}$.

N.B.: For the same reason as previously, the posterior joint pdf should be denoted $p(\mathbf{a}|\mathbf{a}_0, \mathbf{C}_{aa}, \mathbf{C}_M, \mathbf{y}, \mathcal{A}) = \mathcal{N}(\hat{\mathbf{a}}, \hat{\mathbf{C}}_{aa})$. When $\mathbf{C}_M = \sigma_\epsilon^2 \mathbf{I}_N$, it is not necessary to postulate σ_ϵ^2 as the latter can be determined simultaneously with $(\hat{\mathbf{a}}, \hat{\mathbf{C}}_{aa})$ as shown in the previous subsection (§ 2.2).

2.4 Laplace Prior & LASSO

Let us consider now that the prior distribution is the Laplace one, namely $p_{\mathbf{a}|\mathcal{A}} = \mathcal{Laplace}(\mathbf{a}_0, \Lambda_{aa})$ with $\Lambda_{aa} = \text{diag}(\lambda_1, \dots, \lambda_P)$, $\lambda_i > 0$. Then, the joint posterior distribution of the coefficients becomes,

$$p(\mathbf{a}|\mathbf{y}, \mathcal{A}) \propto \exp\left(-\frac{1}{2} (\mathbf{y} - \Psi \mathbf{a})^t \mathbf{C}_M^{-1} (\mathbf{y} - \Psi \mathbf{a}) - \text{sgn}(\mathbf{a} - \mathbf{a}_0)^t \Lambda_{aa} (\mathbf{a} - \mathbf{a}_0)\right) \quad (14)$$

performing the following transformation $\mathbf{b} = \mathbf{a} - \mathbf{a}_0$, yields,

$$p(\mathbf{b}|\mathbf{y}, \mathcal{A}) \propto \exp -\frac{1}{2} \left[\left(\mathbf{C}_M^{-1/2} (\mathbf{y} - \Psi \mathbf{b} - \Psi \mathbf{a}_0) \right)^2 + 2 \text{sgn}(\mathbf{b})^t \Lambda_{aa} \mathbf{b} \right] \quad (15)$$

By developing the term in bracket, and **by assuming that $\text{sgn}(\mathbf{b})$ remains unchanged within the overall posterior solutions**, we get:

$$\mathbf{b}^t \Psi^t \mathbf{C}_M^{-1} \Psi \mathbf{b} - 2 \mathbf{b}^t (\Psi^t \mathbf{C}_M^{-1} (\mathbf{y} - \mathbf{a}_0) - \Lambda_{aa} \text{sgn}(\mathbf{b})) + c$$

which can be factorized as follows,

$$\begin{aligned} & \left[(\Psi^t \mathbf{C}_M^{-1} \Psi)^{1/2} \mathbf{b} - (\Psi^t \mathbf{C}_M^{-1} \Psi)^{-1/2} (\Psi^t \mathbf{C}_M^{-1} (\mathbf{y} - \mathbf{a}_0) - \Lambda_{aa} \text{sgn}(\mathbf{b})) \right]^2 + c \\ \Leftrightarrow & \left[\left(\mathbf{b} - (\Psi^t \mathbf{C}_M^{-1} \Psi)^{-1} (\Psi^t \mathbf{C}_M^{-1} (\mathbf{y} - \mathbf{a}_0) - \Lambda_{aa} \text{sgn}(\mathbf{b})) \right) (\Psi^t \mathbf{C}_M^{-1} \Psi)^{1/2} \right]^2 + c \end{aligned}$$

From this result, it can be concluded that,

$$\boxed{p(\mathbf{a}|\mathbf{y}, \mathcal{A}) = (2\pi)^{-P/2} |\hat{\mathbf{C}}_{aa}|^{-1/2} \exp\left(-\frac{1}{2} (\mathbf{a} - \hat{\mathbf{a}})^t \hat{\mathbf{C}}_{aa}^{-1} (\mathbf{a} - \hat{\mathbf{a}})\right)} \quad (16)$$

with the assumption that $\text{sgn}(\hat{\mathbf{a}} - \mathbf{a}_0)$ is constant (to be checked a posteriori),

$\hat{\mathbf{a}} = \mathbf{a}_0 + (\Psi^t \mathbf{C}_M^{-1} \Psi)^{-1} (\Psi^t \mathbf{C}_M^{-1} (\mathbf{y} - \mathbf{a}_0) - \Lambda_{aa} \text{sgn}(\mathbf{a} - \mathbf{a}_0))$ is the Maximum A Posteriori estimate (MAP) also called in this case *Least Absolute Shrinkage and Selection Operator* (LASSO, Tibshirani [1996]) when $\Lambda_{aa} = \lambda \mathbf{I}_P$.

$\hat{\mathbf{C}}_{aa} = (\Psi^t \mathbf{C}_M^{-1} \Psi)^{-1}$ is the covariance of \mathbf{a} .

N.B.: For the same reason as previously, the posterior joint pdf should be $p(\mathbf{a}|\mathbf{a}_0, \Lambda_{aa}, \mathbf{C}_M, \mathbf{y}, \mathcal{A}) = \mathcal{N}(\hat{\mathbf{a}}, \hat{\mathbf{C}}_{aa})$. Solving the problem is not trivial as one cannot guess $\text{sgn}(\mathbf{a} - \mathbf{a}_0)$ (approximated by $\text{sgn}(\hat{\mathbf{a}} - \mathbf{a}_0)$) before computing $\hat{\mathbf{a}}$. From the computational standpoint, this requires a loop over the criterion that $\text{sgn}(\hat{\mathbf{a}} - \mathbf{a}_0)$ remains unchanged.

3 Model selection

In this section, we discuss the model selection issue which boils down, in the present note, to finding the subset \mathcal{A} such that the Bayesian model evidence, that is,

$$p(\mathbf{y}|\mathcal{A}) = \int_{\mathbb{R}^P} p(\mathbf{y}, \mathbf{a}|\mathcal{A}) d\mathbf{a} = \int_{\mathbb{R}^P} p(\mathbf{y}|\mathbf{a}, \mathcal{A}) p(\mathbf{a}|\mathcal{A}) d\mathbf{a} \quad (17)$$

is maximal.

There are many model selection criterion proposed in the literature, one can cite the Bayesian information criterion (BIC, Schwartz [1978]), the Akaike information criterion (AIC, Akaike [1973]), the Deviance information criterion (DIC, Spiegelhalter et al. [2002]). In this note, we only consider the Kashyap information criterion (KIC, Kashyap [1982]). This criterion was derived in a Bayesian framework and is particularly suited when the input-output relationship is linear as considered in the present work (see Eq.(1)). In Schöniger et al. [2014], it is demonstrated that KIC usually outperforms the BIC and the AIC especially when the model is linear and the error is Gaussian. For the sake of completeness, the KIC at the MAP is derived in the next section (there is also a KIC at the MLE not discussed here).

3.1 Kashyap information criterion

We recall that the solution of (3) is $\mathbf{a}|\mathbf{y}, \mathcal{A} = \mathcal{N}(\hat{\mathbf{a}}, \hat{\mathbf{C}}_{aa})$ under the assumption of Gaussian likelihood and prior. Let us derive the second-order Taylor series of $\ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A}))$ around the MAP, we get

$$\begin{aligned} \ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A})) &\simeq \ln(p(\mathbf{y}, \hat{\mathbf{a}}|\mathcal{A})) + (\mathbf{a} - \hat{\mathbf{a}})^t \underbrace{\left[\frac{\partial \ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A}))}{\partial \mathbf{a}} \right]_{\mathbf{a}=\hat{\mathbf{a}}}}_{=0} \\ &+ \frac{1}{2}(\mathbf{a} - \hat{\mathbf{a}})^t \left[\frac{\partial^2 \ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A}))}{\partial^2 \mathbf{a}} \right]_{\mathbf{a}=\hat{\mathbf{a}}} (\mathbf{a} - \hat{\mathbf{a}}) \end{aligned} \quad (18)$$

Proof that the 2nd term is zero: By noting that $\ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A})) = \ln\left(p(\mathbf{a}|\mathbf{y}, \mathcal{A}) \underbrace{p(\mathbf{y}|\mathcal{A})}_{\text{indep. of } \mathbf{a}}\right)$,

we get $\left. \frac{\partial \ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A}))}{\partial \mathbf{a}} \right|_{\mathbf{a}=\hat{\mathbf{a}}} = \left. \frac{\partial \ln(p(\mathbf{a}|\mathbf{y}, \mathcal{A}))}{\partial \mathbf{a}} \right|_{\mathbf{a}=\hat{\mathbf{a}}} = 0$ by definition of the MAP.

So we come up with,

$$\ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A})) \simeq \ln(p(\mathbf{y}, \hat{\mathbf{a}}|\mathcal{A})) - \frac{1}{2}(\mathbf{a} - \hat{\mathbf{a}})^t \boldsymbol{\Sigma}^{-1}(\mathbf{a} - \hat{\mathbf{a}}) \quad (19)$$

with $\boldsymbol{\Sigma}^{-1} = - \left[\frac{\partial^2 \ln(p(\mathbf{y}, \mathbf{a}|\mathcal{A}))}{\partial^2 \mathbf{a}} \right]_{\mathbf{a}=\hat{\mathbf{a}}}$. Replacing (19) in (17) under the Laplace approximation yields,

$$p(\mathbf{y}|\mathcal{A}) = p(\mathbf{y}, \hat{\mathbf{a}}|\mathcal{A}) \int_{\mathbb{R}^P} \exp\left(-\frac{1}{2}(\mathbf{a} - \hat{\mathbf{a}})^t \boldsymbol{\Sigma}^{-1}(\mathbf{a} - \hat{\mathbf{a}})\right) d\mathbf{a} \quad (20)$$

N.B.: The strict equality comes from the Laplace approximation that assumes that the pdf is Gaussian around the MAP.

Recalling that $p(\mathbf{y}|\mathcal{A}) = \frac{p(\mathbf{y}, \mathbf{a}|\mathcal{A})}{p(\mathbf{a}|\mathbf{y}, \mathcal{A})}$, it is straightforward to guess that, at the MAP ($\boldsymbol{\Sigma} = \hat{\mathbf{C}}_{aa}$), the integral in (20) is equal to,

$$\int_{\mathbb{R}^P} \exp\left(-\frac{1}{2}(\mathbf{a} - \hat{\mathbf{a}})^t \boldsymbol{\Sigma}^{-1}(\mathbf{a} - \hat{\mathbf{a}})\right) d\mathbf{a} = \frac{1}{p(\hat{\mathbf{a}}|\mathbf{y}, \mathcal{A})} = (2\pi)^{P/2} |\hat{\mathbf{C}}_{aa}|^{1/2}.$$

We conclude that,

$$\boxed{p(\mathbf{y}|\mathcal{A}) = p(\mathbf{y}, \hat{\mathbf{a}}|\mathcal{A})(2\pi)^{P/2} |\hat{\mathbf{C}}_{aa}|^{1/2}} \quad (21)$$

The Kashyap information criterion is defined as the deviance, namely,

$$KIC_{\mathcal{A}} = -2 \ln(p(\mathbf{y}|\mathcal{A})) = -2 \ln(p(\mathbf{y}, \hat{\mathbf{a}}|\mathcal{A})) - P \ln(2\pi) - \ln(|\hat{\mathbf{C}}_{aa}|) \quad (22)$$

$$\Leftrightarrow \boxed{KIC_{\mathcal{A}} = -2 \ln(p(\mathbf{y}|\hat{\mathbf{a}}, \mathcal{A})) - 2 \ln(p(\hat{\mathbf{a}}|\mathcal{A})) - P \ln(2\pi) - \ln(|\hat{\mathbf{C}}_{aa}|)} \quad (23)$$

Conclusion: Under the assumption that the model is linear and that the error is Gaussian, the best subset of multi-indexes \mathcal{A} (or best model) is the one with the lowest KIC.

4 Concluding remarks

OLS (Gaussian homoscedastic likelihood+uniform prior) provides unbiased estimates of the coefficients. But it sometimes faces convergence issues, especially when the y -data have outliers. In that case, imposing informative prior (ridge regression or LASSO) might help overcoming this issue although providing biased estimates. Moreover, the choice of the prior's hyperparameter can be non trivial. LASSO (Laplace prior) is known to provide very sparse linear models meaning that many coefficients are set to zero and can be withdrawn from the model though. But in my experience with the KIC model selection criterion, ridge regression (Gaussian prior) performs as well as LASSO in most cases.

The formulas derived in the present document have served to derive the Bayesian sparse polynomial chaos expansion (BSPCE) algorithm in Shao et al. [2017]. BSPCE takes the form of Eq.(1) with ψ_{α} as tensor-product of univariate orthogonal polynomials. The latter has proven to be very efficient for performing global sensitivity analysis of computer model responses under the assumption that the input dataset (namely, \mathbf{x}) is sampled from independent distributions. The algorithm of BSPCE relies on a stepwise linear regression strategy. At each step, a new candidate is added to the current subset \mathcal{A} and the associated KIC is assessed. If the latter is worse than the previous one, the new element is withdrawn from the subset. Notably, the KIC_{MAP} is implemented with Gaussian prior (ridge regression) assigned to \mathbf{a}_{α} that favours low-dimensional and low-degree polynomial elements (specific choice of \mathbf{C}_{aa}).

References

B. Efron and R. Tibshirani. *An introduction to bootstrap*. Chapman & Hall, 1 edition, 1993. ISBN 9780412042317.

- A. E. Hoerl and R. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67, 1970.
- R. Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society, Serie B*, 58(1):267–288, 1996.
- G. Schwartz. Estimating the dimensions of a model. *Annal of Statistics*, 6: 461–464, 1978.
- H. Akaike. Information theory as an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory*, pages 267–281. Akademiai Kiado, 1973.
- D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. van der Linde. Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society*, 64(4):583–639, October 2002. doi: 0.1111/1467-9868.00353.
- R. L. Kashyap. Optimal choice of AR and MA parts in autoregressive moving average models. *IEEE Trans. Pattern Anal. Machine Intell.*, 4(2):99–104, 1982.
- A. Schöniger, T. Wöhling, L. Samaniego, and W. Novak. Model selection on solid ground: Rigorous comparison of nine ways to evaluate Bayesian evidence. *Water Resources Research*, 50:9484–9513, 2014. doi: 10.1002/2014WR016062.
- Q. Shao, A. Younes, M. Fahs, and T. A. Mara. Bayesian sparse polynomial chaos expansion for global sensitivity analysis. *Computer Methods in Applied Mechanics & Engineering*, 318:474–496, 2017.