

A Statistical Perspective of Inverse and Inverse Regression Problems

Debashis Chatterjee, Sourabh Bhattacharya

Interdisciplinary Statistical Research Unit, Indian Statistical Institute,
203 B. T. Road, Kolkata - 700108

ABSTRACT

Inverse problems, where in broad sense the task is to learn from the noisy response about some unknown function, usually represented as the argument of some known functional form, has received wide attention in the general scientific disciplines. However, in mainstream statistics such inverse problem paradigm does not seem to be as popular. In this article we provide a brief overview of such problems from a statistical, particularly Bayesian, perspective. We also compare and contrast the above class of problems with the perhaps more statistically familiar inverse regression problems, arguing that this class of problems contains the traditional class of inverse problems. In course of our review we point out that the statistical literature is very scarce with respect to both the inverse paradigms, and substantial research work is still necessary to develop the fields.

Keywords : Bayesian analysis, inverse problems, inverse regression problems, regularization, reproducing kernel Hilbert space (RKHS), palaeoclimate reconstruction

1. Introduction

The similarities and dissimilarities between inverse problems and the more traditional forward problems are usually not clearly explained in the literature, and often “ill-posed” is the term used to loosely characterize inverse problems. We point out that these two problems may have the same goal or different goal, while both consider the same model given the data. We first elucidate using the traditional case of deterministic differential equations, that the goals of the two problems may be the same. Consider a dynamical system

$$\frac{dx_t}{dt} = G(t, x_t, \theta), \quad (1.1)$$

where G is a known function and θ is a parameter. In the forward problem the goal is to obtain the solution $\chi_t \equiv \chi_t(\theta)$, given θ and the initial conditions, whereas, in the inverse problem, the aim is to obtain θ given the solution process χ_t . Realistically, the differential equation would be perturbed by noise, and so, one observes the data $y = \{y_t : t = 1, \dots, T\}$, where

$$y_t = \chi_t(\theta) + \varepsilon_t \quad (1.2)$$

for noise variables ε_t having some suitable independent and identical (*iid*) error distribution q , which we assume to be known for simplicity of illustration. A typical method of estimating θ , employed by the scientific community, is the method of calibration, where the solution of (1.1) would be obtained for each θ -value on a proposed grid of plausible values, and a set $\tilde{\mathbf{y}}(\theta) = \{\tilde{y}_t(\theta) : t = 1, \dots, T\}$ is generated from the model (1.2) for every such θ after simulating, for $i = 1, \dots, T$, $\tilde{\varepsilon}_t \stackrel{iid}{\sim} q$; then forming $\tilde{y}_t(\theta) = \chi_t(\theta) + \tilde{\varepsilon}_t$, and finally reporting that value θ in the grid as an estimate of the true values for which $\|\mathbf{y} - \tilde{\mathbf{y}}(\theta)\|$ is minimized, given some distance measure $\|\cdot\|$; maximization of the

correlation between y and $\tilde{\mathbf{y}}(\theta)$ is also considered. In other words, the calibration method makes use of the forward technique to estimate the desired quantities of the model. On the other hand, the inverse problem paradigm attempts to directly estimate θ from the observed data y usually by minimizing some discrepancy measure between y and $x(\theta)$, where $x(\theta) = \{\chi_t(\theta) : t = 1, \dots, T\}$. Hence, from this perspective the goals of both forward and inverse approaches are the same, that is, estimation of θ . However, the forward approach is well-posed, whereas, the inverse approach is often ill-posed. To clarify, note that within a grid, there always exists some $\hat{\theta}$ that minimizes the $\|\mathbf{y} - \tilde{\mathbf{y}}(\theta)\|$ among all the grid-values. In this sense the forward problem may be thought of as well-posed. However, direct minimization of the discrepancy between y and $x(\theta)$ with respect to Θ is usually difficult and for high-dimensional θ , the solution to the minimization problem is usually not unique, and small perturbations of the data causes large changes in the possible set of solutions, so that the inverse approach is usually ill-posed. Of course, if the minimization is sought over a set of grid values of θ only, then the inverse problem becomes well-posed.

From the statistical perspective, the unknown parameter θ of the model needs to be learned, in either classical or Bayesian way, and hence, in this sense there is no real distinction between forward and inverse problems. Indeed, statistically, since the data are modeled conditionally on the parameters, all problems where learning the model parameter given the data is the goal, are inverse problems. We remark that the literature usually considers learning unknown functions from the data in the realm of inverse problems, but a function is nothing but an infinite-dimensional parameter, which is a very common learning problem in statistics.

We now explain when forward and inverse problems can differ in their aims, and are significantly different even from the statistical perspective. To give an example, consider the palaeoclimate reconstruction problem discussed in Haslett *et al.* [18] where the reconstruction of prehistoric climate at Glendalough in Ireland from fossil pollen is of interest. The model is built on the realistic assumption that pollen abundance depends upon climate, not the other way around. The compositional pollen data with the modern climates are available at many modern sites but the climate values associated with the fossil pollen data are missing. The inverse nature of the problem is associated with the fact that it is of interest to predict the fossil climate values, given the pollen assemblages. The forward problem would result, if given the fossil climate values (if known), the fossil pollen abundances (if unknown), were to be predicted.

Technically, given a data set y that depends upon covariates x , with a probability distribution $f(y|x, \theta)$ where θ is the model parameter, we call the problem ‘inverse’ if it is of interest to predict the corresponding unknown $\tilde{\chi}$ given a new observed \tilde{y} (see Bhattacharya and Haslett [9]), after eliminating θ . On the other hand, the more conventional forward problem considers the prediction of \tilde{y} for given $\tilde{\chi}$ with the same probability distribution, again, after eliminating the unknown parameter θ . This perspective clearly distinguishes the forward and inverse problems, as opposed to the other parameter-learning perspective discussed above, which is much more widely considered in the literature. In fact, with respect to predicting unknown covariates from the responses, mostly inverse linear regression, particularly in the classical set-up, has been considered in the literature. To distinguish the traditional inverse problems from the covariate-prediction perspective, we use the phrase ‘inverse regression’ to refer to the latter. Other examples of inverse regression are given in section 7.

Our discussion shows that statistically, there is nothing special about the existing literature on inverse problems that considers estimation of unknown (perhaps, infinite-dimensional) parameters, and the only class of problems that can be truly regarded as inverse problems as distinguished from forward problems are those which consider prediction of unknown covariates from the dependent response data. However, for the sake of completeness, the traditional inverse problems related to learning of unknown functions shall occupy a significant portion of our review.

The rest of the paper is structured as follows. In section 2 we discuss the general inverse model, providing several examples. In section 3 we focus on linear inverse problems, which constitute the most popular class of

inverse problems, and review the links between the Bayesian approach based on simple finite difference priors and the deterministic Tikhonov regularization. Connections between Gaussian process based Bayesian inverse problems and deterministic regularizations are reviewed in section 4. In section 5 we provide an overview of the connections between the Gaussian process based Bayesian approach and regularization using differential operators, which generalizes the discussion of section 3 on the connection between finite difference priors and the Tikhonov regularization. The Bayesian approach to inverse problems in Hilbert spaces is discussed in Section 6. We then turn attention to inverse regression problems, providing an overview of such problems and discussing the links with traditional inverse problems in section 7. Finally, we make concluding remarks in section 8.

2. Traditional inverse problem

Suppose that one is interested in learning about the function θ given the noisy observed responses y_n , where the relationship between y_n and θ is governed by following equation (2.1) :

$$y_i = G(\chi_i, \theta) + \varepsilon_i, \quad (2.1)$$

for $i = 1, \dots, n$, where χ_i are known covariates or design points, ε_i are errors associated with the i -th observation and G is a forward operator defined appropriately, which is usually allowed to be non-injective.

Note that since $\varepsilon_n = (\varepsilon_1, \dots, \varepsilon_n)^T$ is unknown, the noisy observation vector y_n itself may not be in the image set of G . If θ is a p -dimensional parameter, then there will often be situations when the number of equations is smaller than the number of unknowns, in the sense that $p > n$ (see, for example, Dashti and Stuart [14]). Modern statistical research is increasingly coming across such inverse problems termed as “ill-posed” which are not in the exact domain of statistical estimation procedures (O’Sullivan [27]) where the maximum likelihood solution or classical least squares may not be uniquely defined and with very bad perturbation sensitivity of the classical solution. However, although such problematic issues are said to characterize inverse problems, the problems in fact fall in the so-called “large p small n ” paradigm and has received wide attention in statistics; see, for example, Bühlmann and van de Geer [10], Giraud [17]. A key concept involved in handling such problems is inclusion of some appropriate penalty term in the discrepancy to be minimized with respect to θ . Such regularization methods are initiated by Tikhonov [34] and Tikhonov and Arsenin [35]. Under this method,

usually a criterion of the following form is chosen for the minimization purpose:

$$\frac{1}{n} \sum_{i=1}^n [y_i - G(x_i, \theta)]^2 + \lambda J(\theta), \lambda > 0. \quad (2.2)$$

The functional J is chosen such that highly implausible or irregular values of θ has large values (O’Sullivan [27]). Thus, depending on the problem at hand, $J(\theta)$ can be used to induce “sparsity” in an appropriate sense so that the minimization problem may be well-defined. We next present several examples of classical inverse problems based on Aster *et al.* [3].

2.1. *Examples of inverse problems*

2.1.1. *Vertical seismic profiling*

In this scientific field, one wishes to learn about the vertical seismic velocity of the material surrounding a borehole. A source generates downward-propagating seismic wavefront at the surface, and in the borehole, a string of seismometers sense these seismic waves. The arrival times of the seismic wavefront at each instrument are measured from the recorded seismograms. These times provide information on the seismic velocity for vertically traveling waves as a function of depth. The problem is nonlinear if it is expressed in terms of seismic velocities. However, we can linearize this problem *via* a simple change of variables, as follows. Letting z denote the depth, it is possible to parameterize the seismic structure in terms of slowness, $s(z)$, which is the reciprocal of the velocity $v(z)$. The observed travel time at depth z can then be expressed as:

$$t(z) = \int_0^z s(u)du = \int_0^\infty s(u)H(z - u)du, \quad (2.3)$$

where H is the Heaviside step function. The interest is

$$d(s) = \int_{-\pi/2}^{\pi/2} (\cos(s) + \cos(\theta))^2 \left(\frac{\sin(\pi (\sin(s) + \sin(\theta)))}{\pi (\sin(s) + \sin(\theta))} \right)^2 m(\theta)d\theta.$$

2.1.4. **Groundwater pollution source history reconstruction problem**

Consider the problem of recovering the history of groundwater pollution at a source site from later measurements of the contamination at downstream wells to which the contaminant plume has been transported by advection and diffusion. The mathematical model for contamination transport is given by the following advection-diffusion equation with respect to t and transported site χ :

$$\begin{aligned} \frac{\partial C}{\partial t} &= D \frac{\partial^2 C}{\partial x^2} - \nu \frac{\partial C}{\partial x} \\ C(0, t) &= C_{in}(t) \\ C(x, t) &\rightarrow 0 \text{ as } x \rightarrow \infty. \end{aligned}$$

to learn about $s(z)$ given observed $t(z)$. Theoretically,

$$s(z) = \frac{dt(z)}{dz},$$

but in practice, simply differentiating the observations need not lead to useful solutions because noise is generally present in the observed times $t(z)$, and naive differentiation may lead to unrealistic features of the solution.

2.1.2. **Estimation of buried line mass density from vertical gravity anomaly**

Here the problem is to estimate an unknown buried line mass density $m(\chi)$ from data on vertical gravity anomaly, $d(\chi)$, observed at some height, h . The mathematical relationship between $d(\chi)$ and $m(\chi)$ is given by

$$d(x) = \int_{-\infty}^{\infty} \frac{h}{[(u - x)^2 + h^2]^{\frac{3}{2}}} m(u)du.$$

As before, noise in the data renders the above linear inverse problem difficult. Variations of the above example has been considered in Aster *et al.* [3].

2.1.3. **Estimation of incident light intensity from diffracted light intensity**

Consider an experiment in which an angular distribution of illumination passes through a thin slit and produces a diffraction pattern, for which the intensity is observed. The data, $d(s)$, are measurements of diffracted light intensity as a function of the outgoing angle $-\pi/2 \leq s \leq \pi/2$. The goal here is to obtain the intensity of incident light on the slit, $m(\theta)$, as a function of the incoming angle $-\pi/2 \leq \theta \leq \pi/2$ using the following mathematical relationship:

In the equation, D is the diffusion coefficient, ν is the velocity of the groundwater flow, and $C_{in}(t)$ is the time history of contaminant injection at $\chi = 0$. The solution to the above advection-diffusion equation is given by

$$C(x, T) = \int_0^T C_{in}(t) f(x, T - t) dt,$$

where,

$$f(x, T - t) = \frac{x}{2\sqrt{\pi D(T - t)^3}} \exp \left[\frac{(x - \nu(T - t))^2}{4D(T - t)} \right]$$

It is of interest to learn about $C_{in}(t)$ from data observed on $C(\chi, T)$.

2.1.5. Transmission tomography

The most basic physical model for tomography assumes that wave energy traveling between a source and receiver can be considered to be propagating along infinitesimally narrow ray paths. In seismic tomography, if the slowness at a point χ is $s(\chi)$, and the ray path is known, then the travel time for seismic energy transiting along that ray path is given by the line integral along ℓ :

$$t = \int_{\ell} s(x(l))dl. \tag{2.4}$$

Learning of $s(\chi)$ from t is required. Note that (2.4) is a high-dimensional generalization of (2.3). In reality, seismic ray paths will be bent due to refraction and/or reflection, resulting in nonlinear inverse problem.

The above examples demonstrate the ubiquity of linear inverse problems. As a result, in the next section we take up the case of linear inverse problems and illustrate the Bayesian approach in details, also investigating connections with the deterministic approach employed by the general scientific community.

3. Linear inverse problem

The motivating examples and discussions in this section are based on Bui-Thanh [11].

Let us consider the following one-dimensional integral equation on a finite interval as in equation (3.1):

$$G(x, \theta) = \int K(x, t) \theta(t) dt, \tag{3.1}$$

where $K(\chi, \cdot)$ is some appropriate, known, real-valued function given χ . Now, let the dataset be $y_n = (y_1, y_2, \dots, y_n)^T$. Then for a known system response $K(\chi_i, t)$ for the dataset, the equation can be written as follows :

$$y_i = \int G(x_i, \theta) + \epsilon_i ; \quad i \in \{1, 2, \dots, n\} \tag{3.2}$$

As a particular example, let $G(x, \theta) = \int_0^1 K(x, t) \theta(t) dt$, where $K(\chi, t) = \frac{1}{\sqrt{2\pi\psi^2}} \exp\{-\frac{1}{2\psi^2}(x-t)^2\}$ is the Gaussian kernel and $\theta : [0; 1] \mapsto \mathbb{R}$ is to be learned given the data y_n and $x_n = (\chi_1, \dots, \chi_n)^T$. We first illustrate the Bayesian approach and draw connections with the traditional approach of Tikhonov's regularization when the integral in G is discretized. In this regard, let $\chi_i = (i-1)/n$, for $i = 1, \dots, n$. Letting $\theta = (\theta(\chi_1), \dots, \theta(\chi_n))^T$ and K be the $n \times n$ matrix with the (i, j) -th element $K(\chi_i, \chi_j)/n$, and $\epsilon_n = (\epsilon_1, \dots, \epsilon_n)^T$, the discretized version of (3.2) can be represented as

$$y_n = K \theta + \epsilon_n \tag{3.3}$$

We assume that $\epsilon_n \sim N_n(0_n, \sigma^2 I_n)$, that is, an n -

variate normal with mean 0_n , an n -dimensional vector with all components zero, and covariance $\sigma^2 I_n$, where I_n is the n -th order identity matrix.

3.1. Smooth prior on θ

To reflect the belief that the function θ is smooth, one may presume that

$$\theta(x_i) = \frac{\theta(x_{i-1}) + \theta(x_{i+1})}{2} + \tilde{\epsilon}_i, \tag{3.4}$$

where, for $i = 1, \dots, n$, $\tilde{\epsilon}_i \stackrel{iid}{\sim} N(0, \tilde{\sigma}^2)$. Thus, *a priori*, $\theta(\chi_i)$ is assumed to be an average of its nearest neighbors to quantify smoothness, with an additive random perturbation term. Letting

$$L = \frac{1}{2} \begin{pmatrix} -1 & 2 & -1 & 0 & \dots & \dots \\ 0 & -1 & 2 & -1 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & -1 & 2 & -1 \end{pmatrix} \tag{3.5}$$

and $\tilde{\epsilon} = (\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n)^T$, it follows from (3.4) that $L\theta = \tilde{\epsilon}$, (3.6)

Now, noting that the Laplacian of a twice-differentiable real-valued function f with independent arguments z_1, \dots, z_k is given by $\Delta f = \sum_{i=1}^k \frac{\partial^2 f}{\partial z_i^2}$, we have

$$\Delta \theta(x_j) \approx n^2 (L\theta)_j, \tag{3.7}$$

where $(L\theta)_j$ is the j -th element of $L\theta$.

However, the rank of L is $n-1$, and boundary conditions on the Laplacian operator is necessary to ensure positive definiteness of the operator. In our case, we assume that $\theta \equiv 0$ outside $[0, 1]$, so that we now assume $\theta(0) = \frac{\theta(x_1)}{2} + \bar{\epsilon}_0$ and $\theta(x_n) = \frac{\theta(x_{n-1})}{2} + \bar{\epsilon}_n$, where $\bar{\epsilon}_0$ and $\bar{\epsilon}_n$ are *iid* $N(0, \bar{\sigma}^2)$. With this modification, the prior on θ is given by

$$\pi(\theta) \propto \exp\left(-\frac{1}{2\bar{\sigma}^2} \|\tilde{L}\theta\|^2\right), \tag{3.8}$$

where $\|\cdot\|$ is the Euclidean norm and

$$\tilde{L} = \frac{1}{2} \begin{pmatrix} 2 & -1 & 0 & 0 & \dots & \dots \\ -1 & 2 & -1 & 0 & \dots & \dots \\ 0 & -1 & 2 & -1 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & -1 & 2 & -1 \\ 0 & 0 & \dots & 0 & -1 & 2 \end{pmatrix} \tag{3.9}$$

Rather than assuming zero boundary conditions, more generally one may assume that $\theta(0)$ and $\theta(\chi_n)$ are

distributed as $N\left(0, \frac{\tilde{\sigma}^2}{\delta_0^2}\right)$ and $N\left(0, \frac{\tilde{\sigma}^2}{\delta_n^2}\right)$, respectively. The resulting modified matrix is then given by

$$\hat{\mathbf{L}} = \frac{1}{2} \begin{pmatrix} 2\delta_0 & 0 & 0 & 0 & \cdots & \cdots \\ -1 & 2 & -1 & 0 & \cdots & \cdots \\ 0 & -1 & 2 & -1 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & -1 & 2 & -1 \\ 0 & 0 & \cdots & 0 & 0 & 2\delta_n \end{pmatrix} \quad (3.10)$$

To choose δ_0 and δ_n one may assume that

$$Var[\theta(0)] = \frac{\tilde{\sigma}^2}{\delta_0^2} = Var[\theta(x_n)] = \frac{\tilde{\sigma}^2}{\delta_n^2} =$$

$$Var[\theta(x_{[n/2]})] = \tilde{\sigma}^2 \boldsymbol{\epsilon}_{[n/2]}^T \left(\hat{\mathbf{L}}^T \hat{\mathbf{L}}\right)^{-1} \boldsymbol{\epsilon}_{[n/2]},$$

where $[n/2]$ is the largest integer not exceeding $n/2$, and $\boldsymbol{\epsilon}_{[n/2]}$ is the $[n/2]$ -th canonical basis vector in \mathbb{R}^{n+1} . It follows that

$$\delta_0^2 = \delta_n^2 = \frac{1}{\boldsymbol{\epsilon}_{[n/2]}^T \left(\hat{\mathbf{L}}^T \hat{\mathbf{L}}\right)^{-1} \boldsymbol{\epsilon}_{[n/2]}}.$$

Since this requires solving a non-linear equation (since $\tilde{\mathbf{L}}$ contains δ_0 and δ_n), for avoiding computational complexity one may simply employ the approximation

$$\delta_0^2 = \delta_n^2 = \frac{1}{\tilde{\boldsymbol{\epsilon}}_{[n/2]}^T \left(\tilde{\mathbf{L}}^T \tilde{\mathbf{L}}\right)^{-1} \tilde{\boldsymbol{\epsilon}}_{[n/2]}},$$

where $\tilde{\mathbf{L}}$ is given by (3.9).

3.2. Non-smooth prior on θ

To begin with, let us assume that θ has several points of discontinuities on the grid of points $\{\chi_0, \dots, \chi_n\}$. To reflect this information in the prior, one may assume that $\theta(0) = 0$ and for $i = 1, \dots, n$, $\theta(\chi_i) = \theta(\chi_{i-1}) + \bar{\epsilon}_i$, where, as before, $\bar{\epsilon}_i$ are iid $N\left(0, \bar{\sigma}^2\right)$. Then, with

$$\mathbf{L}^* = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & \cdots \\ -1 & 1 & 0 & 0 & \cdots & \cdots \\ 0 & -1 & 1 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & -1 & 1 & 0 \\ 0 & 0 & \cdots & 0 & - & 1 \end{pmatrix} \quad (3.11)$$

the prior is given by

$$\pi(\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2\tilde{\sigma}^2} \|\mathbf{L}^* \boldsymbol{\theta}\|^2\right) \quad (3.12)$$

One may also flexibly account for any particular big jump. For instance, if for some $l \in \{0, \dots, n\}$, the jump $\theta(\chi_l) - \theta(\chi_{l-1})$ is particularly large compared to the other jumps, then it can be assumed that $\theta(\chi_l) = \theta(\chi_{l-1}) + \epsilon_l^*$,

with $\epsilon_l^* \sim N\left(0, \frac{\bar{\sigma}^2}{\xi^2}\right)$, where $\xi < 1$. Letting D_l be the

diagonal matrix with ξ^2 being the l -th diagonal element and 1 being the other diagonal elements, the prior is then given by

$$\pi(\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2\tilde{\sigma}^2} \|\mathbf{D}_l \mathbf{L}^* \boldsymbol{\theta}\|^2\right). \quad (3.13)$$

A more general prior can be envisaged where the number and location of the jump discontinuities are unknown. Then we may consider a diagonal matrix $D = \text{diag}\{\xi_1, \dots, \xi_n\}$, so that conditionally on the hyperparameters ξ_1, \dots, ξ_n , the prior on θ is given by

$$\pi(\boldsymbol{\theta} | \xi_1, \dots, \xi_n) \propto \exp\left(-\frac{1}{2\tilde{\sigma}^2} \|\mathbf{D} \mathbf{L}^* \boldsymbol{\theta}\|^2\right) \quad (3.14)$$

Prior on ξ_1, \dots, ξ_n may be considered to complete the specification. These may also be estimated by maximizing the marginal likelihood obtained by integrating out θ , which is known as the ML-II method; see Berger [6]. Calvetti and Somersalo [12] also advocate likelihood based methods.

3.3. Posterior distribution

For convenience, let us generically denote the matrices $\mathbf{L}, \tilde{\mathbf{L}}, \hat{\mathbf{L}}, \mathbf{L}^*, \mathbf{D}_l \mathbf{L}^*, \mathbf{D} \mathbf{L}^*$, by $\Gamma^{1/2}$. Then it can be easily verified that the posterior of θ admits the following generic form:

$$\pi(\boldsymbol{\theta} | \mathbf{y}_n, \mathbf{x}_n) \propto \exp\left\{-\left[\frac{1}{2\sigma^2} \|\mathbf{y}_n - \mathbf{K}\boldsymbol{\theta}\|^2 + \frac{1}{2\tilde{\sigma}^2} \|\Gamma^{-1/2} \boldsymbol{\theta}\|^2\right]\right\} \quad (3.15)$$

Note that the exponent of the posterior is of the form of the Tikhonov functional, which we denote by $T(\boldsymbol{\theta})$. The maximizer of the posterior, commonly known as the *maximum a posteriori* (MAP) estimator, is given by $\boldsymbol{\theta}_{MAP} = \arg \max_{\boldsymbol{\theta}} \pi(\boldsymbol{\theta} | \mathbf{y}_n, \mathbf{x}_n) = \arg \min_{\boldsymbol{\theta}} T(\boldsymbol{\theta})$.

$$(3.16)$$

In other words, the deterministic solution to the inverse problem obtained by Tikhonov's regularization is nothing but the Bayesian MAP estimator in our context.

Writing $\mathbf{H} = \frac{1}{\sigma^2} \mathbf{K}^T \mathbf{K} + \frac{1}{\sigma^2} \mathbf{\Gamma}^{-1}$, which is the Hessian of the Tikhonov functional (regularized misfit), and writing $\|\cdot\|_{\mathbf{H}} = \|\mathbf{H}^{\frac{1}{2}} \cdot\|$, it is clear that (3.15) can be simplified to the Gaussian form, given by

$$\pi(\boldsymbol{\theta} | \mathbf{y}_n, \mathbf{x}_n) \propto \exp \left\{ - \left\| \boldsymbol{\theta} - \frac{1}{\sigma^2} \mathbf{H}^{-1} \mathbf{K}^{-1} \mathbf{y}_n \right\|_{\mathbf{H}}^2 \right\} \quad (3.17)$$

It follows from (3.17) that the inverse of the Hessian of the regularized misfit is the posterior covariance itself. From the above posterior it also trivially follows that

$$\hat{\boldsymbol{\theta}}_{MAP} = \frac{1}{\sigma^2} \mathbf{H}^{-1} \mathbf{K}^{-1} \mathbf{y}_n = \frac{1}{\sigma^2} \left(\frac{1}{\sigma^2} \mathbf{K}^T \mathbf{K} + \frac{1}{\sigma^2} \mathbf{\Gamma}^{-1} \right)^{-1} \mathbf{K}^T \mathbf{Y}_n, \quad (3.18)$$

which coincides with the Tikhonov solution for linear inverse problems. The connection between the traditional deterministic Tikhonov regularization approach with Bayesian analysis continues to hold even if the likelihood is non-Gaussian.

3.4. Exploration of the smoothness conditions

For deeper investigation of the smoothness conditions, let us write

$$\hat{\boldsymbol{\theta}}_{MAP} = \arg \min_{\boldsymbol{\theta}} T(\boldsymbol{\theta}) = \sigma^2 \left(\frac{1}{2} \|\mathbf{y}_n - \tilde{\mathbf{y}}_n\|^2 + \frac{1}{2} \varrho \|\tilde{\mathbf{\Gamma}}^{\frac{1}{2}} \boldsymbol{\theta}\|^2 \right), \quad (3.19)$$

where $\tilde{\mathbf{y}}_n = \mathbf{K}\boldsymbol{\theta}$, $\varrho = \sigma^2/\tilde{\sigma}^2$ and $\tilde{\mathbf{\Gamma}}^{\frac{1}{2}} = \mathbf{\Gamma}^{-\frac{1}{2}}$. Now, from (3.7) it follows that for the smooth priors with the zero boundary conditions, our Tikhonov functional discretizes

$$T_{\infty}(\boldsymbol{\theta}) = \frac{1}{2} \|\mathbf{y}_n - \tilde{\mathbf{y}}_n\|^2 + \frac{1}{2} \varrho \|\Delta\boldsymbol{\theta}\|_{L^2(0,1)}^2, \quad (3.20)$$

where $\|\cdot\|_{L^2(0,1)}^2 = \int_0^1 (\cdot)^2 dt$.

On the other hand, for the non-smooth prior (3.12), rather than discretizing $\nabla\boldsymbol{\theta}$, that is, the gradient of $\boldsymbol{\theta}$, is discretized. In other words, for non-smooth priors, our Tikhonov functional discretizes

$$T_{\infty}(\boldsymbol{\theta}) = \frac{1}{2} \|\mathbf{y}_n - \tilde{\mathbf{y}}_n\|^2 + \frac{1}{2} \varrho \|\nabla\boldsymbol{\theta}\|_{L^2(0,1)}^2. \quad (3.21)$$

Hence, realizations of prior (3.12) is less smooth compared to those of our smooth priors. However, the realizations (3.12) must be continuous. The priors given by (3.13) and (3.14) also support continuous functions as long as the hyperparameters are bounded away from zero. These facts, although clear, can be rigorously justified by functional analysis arguments, in particular, using the Sobolev imbedding theorem (see, for example, Arbogast and Bona [1]).

4. Links between Bayesian inverse problems based on Gaussian process prior and deterministic regularizations

In this section, based on Rasmussen and Williams [31], we illustrate the connections between deterministic regularizations such as those obtained from differential operators as above, and Bayesian inverse problems based

on the very popular Gaussian process prior on the unknown function. A key tool for investigating such relationship is the reproducing kernel Hilbert space (RKHS).

4.1. RKHS

We adopt the following definition of RKHS provided in Rasmussen and Williams [31]:

Definition 4.1 (RKHS). Let \mathcal{H} be a Hilbert space of real functions θ defined on an index set \mathfrak{X} . Then \mathcal{H} is called an RKHS endowed with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ (and norm $\|\theta\|_{\mathcal{H}} = \langle \theta, \theta \rangle_{\mathcal{H}}$) if there exists a function $\mathcal{K} : \mathfrak{X} \times \mathfrak{X} \mapsto \mathbb{R}$ with the following properties:

- (a) for every \mathfrak{X} , $\mathcal{K}(\cdot, \cdot) \in \mathcal{H}$, and
- (b) \mathcal{K} has the reproducing property $\langle \theta(\cdot), \mathcal{K}(\cdot, x) \rangle_{\mathcal{H}} = \theta(x)$.

Observe that since $\mathcal{K}(\cdot, x), \mathcal{K}(\cdot, x') \in \mathcal{H}$, it follows that $\langle \mathcal{K}(\cdot, x), \mathcal{K}(\cdot, x') \rangle_{\mathcal{H}} = \mathcal{K}(x, x')$. The Moore-Aronszajn theorem asserts that the RKHS uniquely determines \mathcal{K} , and vice versa. Formally,

Theorem 1 (Aronszajn [2]). *Let \mathfrak{X} be an index set. Then for every positive definite function $\mathcal{K}(\cdot, \cdot)$ on $\mathfrak{X} \times \mathfrak{X}$ there exists a unique RKHS, and vice versa.*

Here, by positive definite function $\mathcal{K}(\cdot, \cdot)$ on $\mathfrak{X} \times \mathfrak{X}$, we mean $\int \mathcal{K}(x, x') g(x) g(x') d\nu(x) d\nu(x') > 0$ for all non-zero functions $g \in L_2(\mathfrak{X}, \nu)$, where $L_2(\mathfrak{X}, \nu)$ denotes the space of functions square-integrable on \mathfrak{X} with respect to the measure ν .

Indeed, the subspace \mathcal{H}_0 of \mathcal{H} spanned by the functions $\{\mathcal{K}(\cdot, x_i); i = 1, 2, \dots\}$ is dense in \mathcal{H} in the sense that every function in \mathcal{H} is a pointwise limit of a

Cauchy sequence from \mathcal{H}_0 .

To proceed, we require the concepts of eigenvalues and eigenfunctions associated with kernels. In the following section we provide a briefing on these.

4.2. Eigenvalues and eigenfunctions of kernels

We borrow the statements of the following definition of eigenvalue and eigenfunction, and the subsequent statement of Mercer’s theorem from Rasmussen and Williams [31].

Definition 4.2. A function $\psi(\cdot)$ that obeys the integral equation

$$\int_{\mathcal{X}} C(x, x')\psi(x)d\nu(x) = \lambda\psi(x'), \tag{4.1}$$

is called an eigenfunction of the kernel C with eigenvalue λ with respect to the measure ν .

We assume that the ordering is chosen such that $\lambda_1 \geq \lambda_2 \geq \dots$. The eigen functions are orthogonal with respect to ν and can be chosen to be normalized so that $\int_{\mathcal{X}} \psi_i(\mathbf{x})\psi_j(\mathbf{x})d\nu(x) = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise.

The following well-known theorem (see, for example, König [21]) expresses the positive definite

$$\lambda_j(C)\psi_j(x') = \int_{\mathcal{X}} C(x, x')\psi_j(x)d\nu(x) \approx \frac{1}{n} \sum_{i=1}^n C(x_i, x')\psi_j(x_i, x'), \tag{4.2}$$

where, for $i = 1, \dots, n, x_i \sim \nu$, assuming that ν is a probability measure. Now substituting $x' = x_i; i = 1, \dots, n$ in (4.2) yields the following approximate eigen system for the matrix Σ_n :

$$\Sigma_n \mathbf{u}_j \approx n\lambda_j(C)\mathbf{u}_j, \tag{4.3}$$

where the i -th component of \mathbf{u}_j is given by

$$u_{ij} = \frac{\psi_j(x_i)}{\sqrt{n}}. \tag{4.4}$$

Since Ψ_j are normalized to have unit norm, it holds that

$$\mathbf{u}_j^T \mathbf{u}_j = \frac{1}{n} \sum_{i=1}^n \psi_j^2(x_i) \approx \int_{\mathcal{X}} \psi_j^2(x)d\nu(x) = 1. \tag{4.5}$$

From (4.5) it follows that

$$\lambda_j(\Sigma_n) \approx n\lambda_j(C). \tag{4.6}$$

Indeed, Theorem 3.4 of Baker [5] shows that $n^{-1} \lambda_j(\Sigma_n) \rightarrow \lambda_j(C)$, as $n \rightarrow \infty$.

For our purposes the main usefulness of the RKHS framework is that $\|\theta\|_{\mathcal{H}}^2$ can be perceived as a generalization of $\theta^T K^{-1} \theta$, where $\theta = (\theta(\chi_1), \dots, \theta(\bar{\chi}_n))^T$ and $K = (K(\chi_i, \chi_j))_{i,j=1, \dots, n}$ is the $n \times n$ matrix with (i, j) -th element $K(\chi_i, \chi_j)$.

4.3. Inner product

Consider a real positive semidefinite kernel $\mathcal{K}(\mathfrak{X}, \mathfrak{X}')$ with an eigenfunction expansion $\mathcal{K}(\mathfrak{X}, \mathfrak{X}')$

kernel C in terms of its eigenvalues and eigenfunctions.

Theorem 2 (Mercer’s theorem). *Let (\mathfrak{X}, ν) be a finite measure space and $C \in L_{\infty}(\mathfrak{X}^2, \nu^2)$ be a positive definite kernel. By $L_{\infty}(\mathfrak{X}^2, \nu^2)$ we mean the set of all measurable functions $C : \mathfrak{X}^2 \mapsto \mathbb{R}$ which are essentially bounded, that is, bounded up to a set of ν^2 -measure zero. For any function C in this set, its essential supremum, given by in $\{C \geq 0 : |C(\mathfrak{X}_1, \mathfrak{X}_2)| < C; \text{ for almost all } (\mathfrak{X}_1, \mathfrak{X}_2) \in \mathfrak{X} \times \mathfrak{X}\}$ serves as the norm $\|C\|$.*

Let $\psi \in L_2(\mathfrak{X}, \nu)$ be the normalized eigenfunctions of C associated with the eigenvalues $\lambda_j(C) > 0$. Then

(a) the eigenvalues $\{\lambda_j(C)\}_{j=1}^{\infty}$ are absolutely summable.

(b) $\mathcal{C}(x, x') = \sum_{j=1}^{\infty} \lambda_j(C)\psi_j(\mathbf{x})\bar{\psi}_j(x')$ holds ν^2 -almost everywhere, where the series converges absolutely and uniformly ν^2 -almost everywhere. In the above, $\bar{\psi}_j$ denotes the complex conjugate of ψ_j .

It is important to note the difference between the eigenvalue $\lambda_j(C)$ associated with the kernel C and $\lambda_j(\Sigma_n)$ where Σ_n denotes the $n \times n$ Gram matrix with (i, j) -th element $C(\mathfrak{X}_i, \mathfrak{X}_j)$. Observe that (see Rasmussen and Williams [31]):

$\sum_{i=1}^N \lambda_i \phi_i(\chi)\phi_i(\chi')$ relative to a measure μ . Mercer’s theorem ensures that the eigenfunctions are orthonormal with respect to μ , that is, we have $\int \phi_i(x)\phi_j(x)d\mu(x) = \delta_{ij}$. Consider a Hilbert space of linear combinations of the eigenfunctions, that is, $\theta(x) = \sum_{i=1}^N \theta_i \phi_i(x)$ with $\sum_{i=1}^N \frac{\theta_i^2}{\lambda_i} < \infty$. Then the inner product $(\theta_1, \theta_2)_{\mathcal{H}}$ between $\theta_1 = \sum_{i=1}^N \theta_{1i} \phi_i(x)$, and $\theta_2 = \sum_{i=1}^N \theta_{2i} \phi_i(x)$ is of the form

$$\langle \theta_1, \theta_2 \rangle_{\mathcal{H}} = \sum_{i=1}^N \frac{\theta_{1i} \theta_{2i}}{\lambda_i}. \tag{4.7}$$

This induces the norm $\|\cdot\|_{\mathcal{H}}$, where $\|\theta\|_{\mathcal{H}}^2 = \sum_{i=1}^N \frac{\theta_i^2}{\lambda_i}$. A smoothness condition on the space is immediately imposed by requiring the norm to be finite – the eigenvalues must decay sufficiently fast.

The Hilbert space defined above is a unique RKHS with respect to κ , in that it satisfies the following reproducing property:

$$\langle \theta, \mathcal{K}(\cdot, x) \rangle = \sum_{i=1}^N \frac{\theta_i \lambda_i \phi_i(x)}{\lambda_i} = \theta(x). \tag{4.8}$$

Further, the kernel satisfies the following :

$$\langle \mathcal{K}(x, \cdot), \mathcal{K}(x', \cdot) \rangle = \sum_{i=1}^N \frac{\lambda_i^2 \phi_i(x)}{\lambda_i} = \mathcal{K}(x, x'). \quad (4.9)$$

Now, with reference to (4.6), observe that the square norm $\|\theta\|_{\mathcal{H}}^2 = \sum_{i=1}^N \theta_i^2 / \lambda_i$ and the quadratic form $\theta^T \kappa \theta$ have the same form if the latter is expressed in terms of the eigenvectors of κ , albeit the latter has n terms, while the square norm has N terms.

4.4. Regularization

The ill-posed-ness of inverse problems can be understood from the fact that for any given data set y_n , all functions that pass through the data set minimize any given measure of discrepancy $\mathbb{D}(y_n, \theta)$ between the data y_n and θ . To combat this, one considers minimization of the following regularized functional:

$$R(\theta) = \mathbb{D}(y_n, \theta) + \frac{\tau}{2} \|\theta\|_{\mathcal{H}}^2, \quad (4.10)$$

where the second term, which is the regularizer, controls smoothness of the function and τ is the appropriate Lagrange multiplier.

$$\pi(\theta(x^*) | y_n, \mathbf{x}_n) \equiv N(\hat{\mu}(x^*), \hat{\sigma}^2(x^*)), \quad (4.12)$$

where, for any $x^* \in \mathfrak{X}$,

$$\hat{\mu}(x^*) = \mathbf{s}^T(x^*) (\mathbf{K} + \sigma^2 \mathbb{I}_n)^{-1} \mathbf{y}_n; \quad (4.13)$$

$$\hat{\sigma}^2(x^*) = \mathcal{K}(x^*, x^*) - \mathbf{s}^T(x^*) (\mathbf{K} + \sigma^2 \mathbb{I}_n)^{-1} \mathbf{s}(x^*), \quad (4.14)$$

with $\mathbf{s}(x^*) = (\mathcal{K}(x^*, x_1), \dots, \mathcal{K}(x^*, x_n))^T$.

Observe that the posterior mean admits the following representation:

$$\hat{\mu}(x^*) = \sum_{i=1}^n \tilde{c}_i \mathcal{K}(x^*, x_i), \quad (4.15)$$

where \tilde{c}_i is the i -th element of $(\mathbf{K} + \sigma^2 \mathbb{I}_n)^{-1} \mathbf{y}_n$.

In other words, the posterior mean of the Gaussian process based model is consistent with the representer theorem.

5. Regularization using differential operators and connection with Gaussian process

For $x = (x_1, \dots, x_d)^T \in \mathbb{R}^d$, let

$$\|\mathcal{L}^m \theta\|^2 = \int \sum_{j_1 + \dots + j_d = m} \left(\frac{\partial^m \theta(x)}{\partial x_1^{j_1} \dots \partial x_d^{j_d}} \right)^2, \quad (5.1)$$

and

$$\|\mathcal{P}\theta\|^2 = \sum_{m=0}^M b_m \|\mathcal{L}^m \theta\|^2, \quad (5.2)$$

The well-known representer theorem (see, for example, Kimeldorf and Wahba [20], O'Sullivan *et al.* [28], Wahba [38], Schölkopf and Smola [32]) guarantees that each minimizer $\theta \in \mathcal{H}$ can be represented as $\theta(x) = \sum_{i=1}^n c_i \mathcal{K}(x, x_i)$, where κ is the corresponding reproducing kernel. If $\mathbb{D}(y_n, \theta)$ is convex, then there is a unique minimizer $\hat{\theta}$.

4.5. Gaussian process modeling of the unknown function θ

For simplicity, let us consider the model

$$y_i = \theta(x_i) + \epsilon_i \quad (4.11)$$

for $i = 1, \dots, n$, where $\epsilon_i \stackrel{iid}{\sim} N(0, \sigma^2)$, where we assume σ to be known for simplicity of illustration. Let $\theta(\chi)$ be modeled by a Gaussian process with mean function $\mu(\chi)$ and covariance kernel κ associated with the RKHS. In other words, for any $\chi \in \mathcal{X}$, $E[\theta(x)] = \mu(x)$ and for any $x_1, x_2 \in \mathfrak{X}$, $Cov(\theta(x_1), \theta(x_2)) = \mathcal{K}(x_1, x_2)$.

Assuming for convenience that $\mu(\chi) = 0$ for all $\chi \in \mathcal{X}$, it follows that the posterior distribution of $\theta(\chi^*)$ for any $\chi^* \in \mathcal{X}$ is given by

for some $M > 0$, where the co-efficients $b_m \geq 0$. In particular, we assume for our purpose that $b_0 > 0$. It is clear that $\|\mathcal{P}\theta\|^2$ is translation and rotation invariant. This norm penalizes θ in terms of its derivatives up to order M .

5.1. Relation to RKHS

It can be shown, using the fact that the complex exponentials $\exp(2\pi i s^T \chi)$ are eigen functions of the differential operator, that

$$\|\mathcal{P}\theta\|^2 = \int \sum_{m=0}^M b_m (4\pi^2 s^T s)^m |\tilde{\theta}(s)|^2 ds, \tag{5.3}$$

where $\hat{\theta}(s)$ is the Fourier transform of $\theta(s)$. Comparison of (5.3) with (4.7) yields the power spectrum of the form $[\sum_{m=0}^M b_m (4\pi^2 s^T s)^m]^{-1}$ which yields the following kernel by Fourier inversion:

$$\mathcal{K}(x, x') = \mathcal{K}(x - x') = \int \frac{\exp(2\pi i s^T (x - x'))}{\sum_{m=0}^M b_m (4\pi^2 s^T s)^m} ds. \tag{5.4}$$

Calculus of variations can also be used to minimize $R(\theta)$ with respect to θ , which yields (using the Euler-Lagrange equation)

$$\theta(x) = \sum_{i=1}^n b_i \mathcal{G}(x - x_i), \tag{5.5}$$

with

$$\sum_{i=1}^m (-1)^m b_m \nabla^m \mathcal{G} = \delta_{x-x'}, \tag{5.6}$$

where G is known as the Green’s function. Using Fourier transform on (5.6) it can be shown that the Green’s function is nothing but the kernel K given by (5.4). Moreover, it follows from (5.6) that

$$G(x, x') = G(x - x') = \begin{cases} c_{m,d} |x - x'|^{2m-d} \log |x - x'| & \text{if } 2m > d \text{ and } d \text{ even} \\ c_{m,d} |x - x'|^{2m-d} & \text{otherwise} \end{cases} \tag{5.9}$$

where $c_{m,D}$ are constants (see Wahba [38] for the explicit form).

We now specialize the above arguments to the spline set-up. As before, let us consider the model $y_i = \theta(x_i) + \epsilon_i$, where, for $i = 1, \dots, n$, $\epsilon_i \stackrel{iid}{\sim} N(0, \sigma^2)$. For simplicity, we consider the one-dimensional set-up, and consider the cubic spline smoothing problem that minimizes

$\sum_{i=1}^m (-1)^m b_m \nabla^m$ and \mathcal{K} are inverses of each other.

Examples of kernels derived from differential operators are as follows. For $d = 1$, setting $b_0 = b_2, b_1 = 1$ and $b_m = 0$ for $m \geq 2$, one obtains $\kappa(\chi, \chi') = \kappa(\chi - \chi')$
 $= \frac{1}{2b} \exp(-b|\chi - \chi'|)$, which is the covariance of the Ornstein-Uhlenbeck process. For general d dimension, setting $b_m = b^{2m} / (m!2^m)$, yields $\kappa(\chi, \chi') = \kappa(\chi - \chi') = \frac{1}{(2\pi b^2)^{d/2}} \exp[-\frac{1}{2b^2}(x - x')^T(x - x')]$.

Considering a grid x_n , note that

$$\|\mathcal{P}\theta\|^2 \approx \sum_{m=0}^M b_m (D_m \theta)^T (D_m \theta) = \theta^T \left(\sum_{m=0}^M D_m^T D_m \right) \theta, \tag{5.7}$$

where D_m is a suitable finite-difference approximation of the differential operator. Note that such finite-difference approximation has been explored in section 3, which we now investigate in a rigorous setting. Also, since (5.7) is quadratic in θ , assuming a prior for θ , the logarithm of which has this form, and further assuming that $\log[\mathbb{D}(y_n, \theta)]$ is a log-likelihood quadratic in θ , a Gaussian posterior results.

5.2. Spline models and connection with Gaussian process

Let us consider the penalty function to be $\|\mathcal{L}^m \theta\|^2$. Then polynomials up to degree $m - 1$ are not penalized and so, are in the null space of the regularization operator. In this case, it can be shown that a minimizer of $R(\theta)$ is of the form

$$\theta(x) = \sum_{j=1}^k d_j \psi_j(x) + \sum_{i=1}^n c_i G(x, x_i), \tag{5.8}$$

where $\{\psi_1, \dots, \psi_k\}$ are polynomials that span the null space and the Green’s function G is given by (see Duchon [15], Meinguet [23])

$$R(\theta) = \sum_{i=1}^n (y_i - \theta(x_i))^2 + \tau \int_0^1 [\theta''(x)]^2 dx, \tag{5.10}$$

where $0 < \chi_1 < \dots < \chi_n < 1$. The solution to this minimization problem is given by

$$\theta(x) = \sum_{j=0}^1 d_j x^j + \sum_{i=1}^n c_i (x - x_i)_+^3, \tag{5.10}$$

where, for any χ , $(\chi)_+ = \chi$ if $\chi > 0$ and zero otherwise.

Following Wahba [37], let us consider

$$f(x) = \sum_{j=0}^1 \beta_j x^j + \theta(x), \quad (5.12)$$

where $\beta = (\beta_0, \beta_1)^T \sim N(\mathbf{0}, \sigma_\beta^2 \mathbb{I}_2)$, and θ is a zero mean Gaussian process with covariance

$$\sigma_\theta^2 \mathcal{K}(x, x') = \int_0^1 (x-u)_+(x'-u)_+ du = \sigma_\theta^2 \left(\frac{|x-x'|v^2}{2} + \frac{v^3}{3} \right), \quad (5.13)$$

where $v = \min\{\chi, \chi'\}$.

Taking $\sigma_\beta \frac{2}{\beta} \rightarrow \infty$ makes the prior of β vague, so that penalty on the polynomial terms in the null space is effectively washed out. It follows that

$$E[\theta(x^*) | \mathbf{y}_n, \mathbf{x}_n] = \mathbf{h}(x^*)^T \hat{\beta} + \mathbf{s}(x^*)^T \hat{\mathbf{K}}^{-1} (\mathbf{y}_n - \mathbf{H}^T \hat{\beta}), \quad (5.14)$$

where, for any χ , $h(\chi) = (1, \chi)^T$, $\mathbf{H} = (h(\chi_1), \dots, h(\chi_n))$, $\hat{\mathbf{K}}$ is the covariance matrix corresponding to

$$\sigma_\theta^2 \mathcal{K}(x_i, x_j) + \sigma^2 \delta_{ij}, \text{ and } \hat{\beta} = (\mathbf{H} \hat{\mathbf{K}}^{-1} \mathbf{H})^{-1} \mathbf{H} \hat{\mathbf{K}}^{-1} \mathbf{y}_n.$$

Since the elements of $s(\chi^*)$ are piecewise cubic polynomials, it is easy to see that the posterior mean (5.14) is also a piecewise cubic polynomial. It is also clear that (5.14) is a first order polynomial on $[0, \chi_1]$ and $[\chi_n, 1]$.

5.2.1. Connection with the ℓ -fold integrated Wiener process

Shepp [33] considered the ℓ -fold integrated Wiener process, for $\ell = 0, 1, 2, \dots$, as follows:

$$W_\ell(x) = \int_0^1 \frac{(x-u)_+^\ell}{\ell!} Z(u) du, \quad (5.15)$$

where Z is a Gaussian white noise process with covariance $\delta(u-u')$. As a special case, note that W_0 is the standard Wiener process. In our case, note that

$$\mathcal{K}(x, x') = \text{Cov}(W_1(x), W_1(x')). \quad (5.16)$$

The above ideas can be easily extended to the case of the regularizer $\int [f^{(m)}(x)]^2 dx$, for $m \geq 1$ by replacing $(\chi-u)_+$ with $(\chi-1)_+^{m-1} / (m-1)!$ and letting $h(\chi) = (1, \chi, \dots, \chi^{m-1})^T$

6. The Bayesian approach to inverse problems in Hilbert spaces

We assume the following model

$$y = G(\theta) + \epsilon, \quad (6.1)$$

where y , θ and ϵ are in Banach or Hilbert spaces.

6.1. Bayes theorem for general inverse problems

We will consider the model stated by equation (6.1). Let \mathcal{Y} and Θ denote the sample spaces for y and θ , respectively. Let us first assume that both are separable Banach spaces. Assume μ_0 to be the prior measure for θ . Assuming well-defined joint distribution for (y, θ) , let us denote the posterior of θ given y as μ_y . Let $\epsilon \sim Q_0$ where Q_0 such that ϵ and θ are independent. Let Q_0 be the distribution of ϵ . Let us denote the conditional distribution of y given θ by Q_θ , obtained from a translation of Q_0 by $G(\theta)$.

Assume that $Q_\theta \ll Q_0$. Thus, for some potential $\Phi : \Theta \times \mathcal{Y} \mapsto \mathbb{R}$,

$$\frac{dQ_\theta}{dQ_0} = \exp(-\Phi(\theta, y)). \quad (6.2)$$

Thus, for fixed θ , $\Phi(\theta, \cdot) : \mathcal{Y} \mapsto \mathbb{R}$ is measurable and $E_{Q_0}[\exp(-\Phi(\theta, y))] = 1$. Note that $-\Phi(\cdot, y)$ is nothing but the log-likelihood.

Let ν_0 denote the product measure

$$\nu_0(d\theta, dy) = \mu_0(d\theta)Q_0(dy), \tag{6.3}$$

and let us assume that Φ is ν_0 -measurable. Then $(\theta, y) \in \Theta \times Y$ is distributed according to the measure $\nu(d\theta, dy) = \mu_0(d\theta)Q_\theta(dy)$. It then also follows that $\nu \ll \nu_0$, with

$$\frac{d\nu_\theta}{d\nu_0}(\theta, y) = \exp(-\Phi(\theta, y)). \tag{6.4}$$

Then we have the following statement of Bayes' theorem for general inverse problems:

Theorem 3 (Bayes theorem for general inverse problems). Assume that $\Phi : \Theta \times Y \mapsto \mathbb{R}$ is ν_0 -measurable and

$$C = \int_{\Theta} \exp(-\Phi(\theta, y)) \mu_0(dy) > 0, \tag{6.5}$$

for Q_0 -almost surely all y . Then the posterior of θ given y , which we denote by μ^y , exists under ν . Also, $\mu^y \ll \mu_0$ and for all y ν_0 -almost surely,

$$\frac{d\mu_\theta^y}{d\mu_0}(\theta) = \frac{1}{C} \exp(-\Phi(\theta, y)). \tag{6.6}$$

Now assume that Θ and Y are Hilbert spaces. Suppose $\varepsilon \sim N(0, \Gamma)$. Then the following theorem holds:

Theorem 4 (Vollmer [36]).

$$\frac{d\mu^y}{d\mu_0} \propto \exp\left(-\frac{1}{2}\|G(\theta)\|_\Gamma^2 + \langle y, G(\theta) \rangle_\Gamma\right), \tag{6.7}$$

where $\langle \cdot, \cdot \rangle_\Gamma = \langle \Gamma^{-1} \cdot, \cdot \rangle$, and $\|\cdot\|_\Gamma$ is the norm induced by $\langle \cdot, \cdot \rangle_\Gamma$.

For the model $y_i = \theta(x_i) + \varepsilon_i$ for $i = 1, \dots, n$, with $\varepsilon_i \stackrel{iid}{\sim} N(0, \sigma^2)$, the posterior is of the form

$$\frac{d\mu^y}{d\mu_0} \propto \exp\left(-\sum_{i=1}^n \frac{(y_i - \theta(x_i))^2}{2\sigma^2}\right). \tag{6.8}$$

6.2. Connection with regularization methods

It is not immediately clear if the Bayesian approach in the Hilbert space setting has connection with the deterministic regularization methods, but Vollmer [36] prove consistency of the posterior assuming certain stability results which are used to prove convergence of regularization methods; see Engl *et al.* [16].

We next turn to inverse regression.

7. Inverse regression

We first provide some examples of inverse regression, mostly based on Avenhaus *et al.* [4].

7.1. Examples of inverse regression

7.1.1. Example 1: Measurement of nuclear materials

Measurement of the amount of nuclear materials such as plutonium by direct chemical means is an extremely difficult exercise. This motivates model-based methods. For instance, there are physical laws relating heat production or the number of neutrons emitted (the dependent response variable y) to the amount of material present, the latter being the independent variable x . But any measurement instrument based on the physical laws first needs to be calibrated. In other words, the unknown parameters of the model needs to be learned, using known inputs and outputs. However, the independent variables are usually subject to measurement errors,

motivating a statistical model. Thus, conditionally on x and parameter(s) θ , $y \sim P(\cdot|x, \theta)$, where $P(\cdot|x, \theta)$, denotes some appropriate probability model. Given y_n and x_n , and some specific \tilde{y} , the corresponding \tilde{x} needs to be predicted.

7.1.2. Example 2: Estimation of family incomes

Suppose that it is of interest to estimate the family incomes in a certain city through public opinion poll. Most of the population, however, will be unwilling to provide reliable answers to the questionnaires. One way to extract relatively reliable figures is to consider some dependent variable, say, housing expenses (y), which is supposed to strongly depend on family income (χ); see Muth [26], and such that the population is less reluctant to divulge the correct figures related to y . From past survey data on χ_n and y_n , and using current data from families who may provide reliable answers related to both χ and y , a statistical model may be built, using which the unknown family incomes may be predicted, given their household incomes.

7.1.3. Example 3: Missing variables

In regression problems where some of the covariate values χ_i are missing, they may be estimated from the remaining data and the model. In this context, Press and Scott [29] considered a simple linear regression problem

in a Bayesian framework. Under special assumptions about the error and prior distributions, they showed that an optimal procedure for estimating the linear parameters is to first estimate the missing χ_i from an inverse regression based only on the complete data pairs.

7.1.4. Example 4: Bioassay

It is usual to investigate the effects of substances (y) given in several dosages on organisms (χ) using bioassay methods. In this context it may be of interest to determine the dosage necessary to obtain some interesting effect, making inverse regression relevant (see, for example, Rasch *et al.* [30]).

7.1.5. Example 5: Learning the Milky Way

The modelling of the Milky Way galaxy is an integral step in the study of galactic dynamics; this is because knowledge of model parameters that define the Milky Way directly influences our understanding of the evolution of our galaxy. Since the nature of the Galaxy’s phase space, in the neighbourhood of the Sun, is affected by distinct Milky Way features, measurements of phase space coordinates of individual stars that live in this neighbourhood of the Sun, will bear information about the influence of such features. Then, inversion of such measurements can help us learn the parameters that describe such MilkyWay features. In this regard, learning about the location of the Sun with respect to the center of the galaxy, given the two-component velocities of the stars in the vicinity of the Sun, is an important problem. For k such stars, Chakrabarty *et al.* [13] model the $k \times 2$ -dimensional velocity matrix V as a function of the galactocentric location (S) of the Sun, denoted by $V = \xi(S)$. For a given observed value V^* of V , it is then of interest to obtain the corresponding S^* . Since ξ is unknown, Chakrabarty *et al.* [13] model ξ as a matrix-variate Gaussian process, and consider the Bayesian approach to learning about S , given data $\{(S_i, V_i) : i = 1, \dots, n\}$ simulated from established astrophysical models, and the observed velocity matrix V^* .

We now provide a brief overview of of the methods of inverse linear regression, which is the most popular among inverse regression problems. Our discussion is generally based on Hoadley [19] and Avenhaus *et al.* [4].

7.2. Inverse linear regression

Let us consider the following simple linear regression model: for $i = 1, \dots, n$,

$$y_i = \sigma + \beta\chi_i + \sigma\epsilon_i, \tag{7.1}$$

where $\epsilon_i \stackrel{iid}{\sim} N(0, 1)$.

For simplicity, let us consider a single unknown $\tilde{\chi}$, associated with a further set of m responses $\{\tilde{y}_1, \dots, \tilde{y}_m\}$, related by

$$\tilde{y}_i = \alpha + \beta\tilde{x} + \tau\tilde{\epsilon}_i, \tag{7.2}$$

for $i = 1, \dots, m$, where $\tilde{\epsilon}_i \stackrel{iid}{\sim} N(0, 1)$ and are independent of the ϵ_i ’s associated with (7.1).

The interest in the above problem is inference regarding the unknown χ . Based on (7.1), first least squares estimates of α and β are obtained as

$$\hat{\beta} = \frac{\sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x})}{\sum_{i=1}^n (x_i - \bar{x})^2}; \tag{7.3}$$

$$\hat{\alpha} = \bar{y} - \hat{\beta}\bar{x}, \tag{7.4}$$

where $\bar{y} = \sum_{i=1}^n y_i/n$ and $\bar{x} = \sum_{i=1}^n x_i/n$. Then, letting $\bar{\tilde{y}} = \sum_{i=1}^m \tilde{y}_i/n$, ‘classical’ estimator of χ is given by

$$\hat{x}_C = \frac{\bar{\tilde{y}} - \hat{\alpha}}{\hat{\beta}}, \tag{7.5}$$

which is also the maximum likelihood estimator for the likelihood associated with (7.1) and (7.2), assuming known σ and τ . However,

$$E \left[(\hat{x}_C - x)^2 \mid \alpha, \beta, \sigma, \tau, x \right] = \infty, \tag{7.6}$$

which prompted Krutchkoff [22] to propose the following ‘inverse’ estimator:

$$\hat{x}_I = \hat{\gamma} + \hat{\delta}\bar{\tilde{y}}, \tag{7.7}$$

where,

$$\hat{\delta} = \frac{\sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x})}{\sum_{i=1}^n (y_i - \bar{y})^2}; \tag{7.8}$$

$$\hat{\gamma} = \bar{x} - \hat{\delta}\bar{\tilde{y}}, \tag{7.9}$$

are the least squares estimators of the slope and intercept when the χ_i are regressed on the y_i . It can be shown that the mean square error of this inverse estimator is finite. However, Williams [39] showed that if $\sigma^2 = \tau^2$ and if the sign of β is known, then the unique unbiased estimator of x has infinite variance. Williams advocated the use of confidence limits instead of point estimators.

Hoadley [19] derive confidence limits setting $\sigma = \tau$ and assuming without loss of generality that $\sum_{i=1}^n x_i = 0$. Under these assumptions, the maximum likelihood estimators of σ^2 with x_n and y_n only, $\tilde{\mathbf{y}}_n = (\tilde{y}_1, \dots, \tilde{y}_n)^T$ only, and with the entire available data set are, respectively,

$$\hat{\sigma}_1^2 = \frac{1}{n-2} \sum_{i=1}^n \left(y_i - \hat{\alpha} - \hat{\beta}x_i \right)^2; \tag{7.10}$$

$$\hat{\sigma}_2^2 = \frac{1}{m-1} \sum_{i=1}^m (\tilde{y}_i - \bar{\tilde{y}})^2; \tag{7.11}$$

$$\hat{\sigma}^2 = \frac{1}{n-2+m-1} \left[(n-2)\sigma_1^2 + (m-1)\sigma_2^2 \right]. \tag{7.12}$$

Now consider the F -statistic $F = \frac{n\hat{\beta}^2}{\hat{\sigma}^2}$ for testing the hypothesis $\beta = 0$. Note that under the null hypothesis this statistic has the F distribution with 1 and $n + m$ degrees of freedom. For $m = 1$,

$$\hat{\beta}(\hat{x}_C - x) \sqrt{\frac{n}{\sigma^2(n+1+x^2)}}$$

has a t distribution with $n - 2$ degrees of freedom. Letting $F_{\alpha,1,\nu}$ denote the upper α point of the F distribution with 1 and ν degrees of freedom, a confidence set S can be derived as follows:

$$S = \begin{cases} \{x : x_L \leq x \leq x_U\} & \text{if } F > F_{\alpha;1,n-2}; \\ \{x : x \leq x_L\} \cup \{x \geq x_U\} & \text{if } \frac{n+1}{n+1+\hat{x}_C^2} F_{\alpha;1,n-2} \leq F < F_{\alpha;1,n-2}; \\ (-\infty, \infty) & \text{if } F < \frac{n+1}{n+1+\hat{x}_C^2} F_{\alpha;1,n-2}, \end{cases} \quad (7.13)$$

where x_L and x_U are given by

$$\frac{F\hat{x}_C}{F - F_{\alpha;1,n-2}} \pm \frac{\{F_{\alpha;1,n-2} [(n+1)(F - F_{\alpha;1,n-2}) + F\hat{x}_C^2]\}^{\frac{1}{2}}}{F - F_{\alpha;1,n-2}}.$$

Hence, if $F < \frac{n+1}{n+1+\hat{x}_C^2} F_{\alpha;1,n-2}$, then the associated confidence interval is $S = (-\infty, \infty)$, which is of course useless.

Hoadley [19] present a Bayesian analysis of this problem, presented below in the form of the following two theorems.

Theorem 5 (Hoadley [19]). Assume that $\sigma = \tau$, and let χ be independent of $(\alpha, \beta, \sigma^2)$ a priori. With any prior $\pi(\chi)$ on χ and the prior

$$\pi(\alpha, \beta, \sigma^2) \propto \frac{1}{\sigma^2}$$

on $(\alpha, \beta, \sigma^2)$, the posterior density of x given by

$$\pi(x|\mathbf{y}_n, \mathbf{x}_n, \tilde{\mathbf{y}}_n) \propto \pi(x)L(x),$$

where,

$$L(x) = \frac{(1 + \frac{n}{m} + x^2)^{\frac{m+n-3}{2}}}{\left[1 + \frac{n}{m} + R\hat{x}_C^2 + \left(\frac{F}{m+n-3} + 1\right)(x - R\hat{x}_C)^2\right]^{\frac{m+n-2}{2}}},$$

where,

$$R = \frac{F}{F + m + n - 3}.$$

For $m = 1$, Hoadley [19] present the following result characterizing the inverse estimator \hat{x}_I :

Theorem 6 (Hoadley [19]). Consider the following informative prior on χ :

$$x = t_{n-3} \frac{n+1}{n-3},$$

where t_ν denotes the t distribution with ν degrees of freedom. Then the posterior distribution of χ given y_n, χ_n and \tilde{y}_n has the same distribution as

$$\hat{x}_I + t_{n-2} \sqrt{\frac{n+1 + \frac{\hat{x}_I^2}{R}}{F+n-2}}.$$

In particular, it follows from Theorem 6 that the posterior mean of x is \hat{x}_I when $m = 1$. In other words, the inverse estimator \hat{x}_I is Bayes with respect to the squared error loss and a particular informative prior distribution for χ .

Since the goal of Hoadley [19] was to provide a theoretical justification of the inverse estimator, he had to choose a somewhat unusual prior so that it leads to \hat{x}_I as the posterior mean. In general it is not necessary to confine ourselves to any specific prior for Bayesian analysis of inverse regression. It is also clear that the Bayesian framework is appropriate for any inverse regression problem, not just linear inverse regression;

indeed, the palaeoclimate reconstruction problem (Haslett *et al.* [18]) and the MilkyWay problem (Chakrabarty *et al.* [13]) are examples of very highly non-linear inverse regression problems.

7.3. Connection between inverse regression problems and traditional inverse problems

Note that the class of inverse regression problems includes the class of traditional inverse problems. The Milky Way problem is an example where learning the unknown, matrix-variate function ξ (inverse problem) was required, even though learning about S , the galactocentric location of the sun (inverse regression problem) was the primary goal. The Bayesian approach allowed learning both S and ξ simultaneously and coherently.

In the palaeoclimate models proposed in Haslett *et al.* [18], Bhattacharya [7] and Mukhopadhyay and Bhattacharya [25], although species assemblages are modeled conditionally on climate variables, the functional relationship between species and climate are not even approximately known. In all these works, it is of interest to learn about the functional relationship as well as to predict the unobserved climate values, the latter being the main aim. Again, the Bayesian approach facilitated appropriate learning of both the unknown quantities.

7.4. Consistency of inverse regression problems


In the above linear inverse regression, notice that if $\tau > 0$, then the variance of the estimator of χ can not tend to zero, even as the data size tends to infinity. This

shows that no estimator of χ can be consistent. The same argument applies even to Bayesian approaches; for any sensible prior on χ that does not give point mass to the true value of χ , the posterior of χ will not converge to the point mass at the true value of χ as the data size increases indefinitely. The arguments remain valid for any inverse regression problem where the response variable y probabilistically depends upon the independent variable χ . Not only in inverse regression problems, even in forward regression problems where the interest is in prediction of y given x , any estimate of y or any posterior predictive distribution y will be inconsistent.

To give an example of inconsistency in non-linear and non-normal inverse problem, consider the following set-up: $y_i \stackrel{iid}{\sim} \text{Poisson}(\theta_{\chi_i})$, for $i = 1, \dots, n$, where $\theta > 0$ and $\chi_i > 0$ for each i . Let us consider the prior $\pi(\theta) \equiv 1$ for all $\theta > 0$. For some $i^* \in \{1, \dots, n\}$ let us assume the leave-one-out cross-validation set-up in that we wish to learn $\chi = \chi_{i^*}$ assuming it is unknown, from the rest of the data. Putting the prior $\pi(\chi) \equiv 1$ for $\chi > 0$, the posterior of χ is given by (see Bhattacharya and Haslett [9], Bhattacharya [8])

$$\pi(x|\mathbf{x}_n \setminus x_i, \mathbf{y}_n) \propto \frac{x^{y_i}}{(x + \sum_{j \neq i} x_j)^{(\sum_{j=1}^n y_j + 1)}} \tag{7.14}$$

Figure 7.1 displays the posterior of χ when $i = 10$, for increasing sample size. Observe that the variance of the posterior does not decrease even with sample size as large as 100,000, clearly demonstrating inconsistency.



Hence, special, innovative priors are necessary for consistency in such cases.

8. Conclusion

In this review article, we have clarified the similarities and dissimilarities between the traditional inverse problems and the inverse regression problems. In particular, we have argued that only the latter class of problems qualify as authentic inverse problems in they have significantly different goals compared to the corresponding forward problems. Moreover, they include the traditional inverse problems on learning unknown

Figure 7.1: Demonstration of posterior inconsistency in inverse regression problems. The vertical line denotes the true value. functions as a special case, as exemplified by our palaeoclimate and Milky Way examples. We advocate the Bayesian paradigm for both classes of problems, not only because of its inherent flexibility, coherency and posterior uncertainty quantification, but also because the prior acts as a natural penalty which is very important to regularize the so-called ill-posed inverse problems. The well-known Tikhonov regularizer is just a special case from this perspective.

It is important to remark that the literature on inverse function learning problems and inverse regression problems is still very young and a lot of research is necessary to develop the fields. Specifically, there is hardly any well-developed, consistent model adequacy test or model comparison methodology in either of the two fields, although Mohammad-Djafari [24] deal with some specific inverse problems in this context, and Bhattacharya [8] propose a test for model adequacy in the case of inverse regression problems. Moreover, as we have demonstrated, inverse regression problems are inconsistent in general. The general development in these respects will be provided in the PhD thesis of the first author.

REFERENCES

- Arbogast, T., and Bona, J. L. [2008], "Methods of Applied Mathematics,". University of Texas at Austin.
- Aronszajn, N. [1950], "Theory of Reproducing Kernels," *Transactions of the American Mathematical Society*, **68** : 337-404.
- Aster, R. C., Borchers, B., and Thurber, C. H. [2013], *Parameter Estimation and Inverse Problems*, Oxford, UK: Academic Press.
- Avenhaus, R., H'opfinger, E., and Jewell, W. S. [1980], "Approaches to Inverse Linear Regression,". Technical Report. Available at <https://publikationen.bibliothek.kit.edu/270015256/3812158>.
- Baker, C. T. H. [1977], *The Numerical Treatment of Integral Equations*, Oxford: Clarendon Press.
- Berger, J. O. [1985], *Statistical Decision Theory and Bayesian Analysis*, New York: Springer-Verlag.
- Bhattacharya, S. [2006], "A Bayesian Semiparametric Model for Organism Based Environmental Reconstruction," *Environmetrics*, **17** : 763-76.
- Bhattacharya, S. [2013], "A Fully Bayesian Approach to Assessment of Model Adequacy in Inverse Problems," *Statistical Methodology*, **12** : 71-83.
- Bhattacharya, S., and Haslett, J. [2007], "Importance Resampling MCMC for Cross-Validation in Inverse Problems," *Bayesian Analysis*, **2** : 385-408.
- Bühlmann, P., and van de Geer, S. [2011], *Statistics for High-Dimensional Data*, New York : Springer.
- Bui-Thanh, T. [2012], "A Gentle Tutorial on Statistical Inversion Using the Bayesian Paradigm,". ICES Report 12-18. Available at <http://users.ices.utexas.edu/tanbui/PublishedPapers/BayesianTutorial.pdf>.
- Calvetti, D., and Somersalo, E. [2007], *Introduction to Bayesian Scientific Computing: Ten Lectures on Subjective Computing*, New York: Springer.
- Chakrabarty, D., Biswas, M., and Bhattacharya, S. [2015], "Bayesian Nonparametric Estimation of Milky Way Parameters Using Matrix-Variate Data, in a New Gaussian Process Based Method," *Electronic Journal of Statistics*, **9** : 1378-1403.
- Dashti, M., and Stuart, A. M. [2015], "The Bayesian Approach to Inverse Problems,". eprint: arXiv:1302.6989.
- Duchon, J. [1977], Splines Minimizing Rotation-Invariant Semi-norms in Sobolev Spaces, in *Constructive Theory of Functions of Several Variables*, eds. W. Schempp, and K. Zellner, Springer-Verlag, New York, pp. 85-100.
- Engl, H. W., Hanke, M., and Neubauer, A. [1996], *Regularization of Inverse Problems*, Dordrecht: Kluwer Academic Publishers Group. Volume 375 of Mathematics and its Applications.
- Giraud, C. [2015], *Introduction to High-Dimensional Statistics*, New York: Chapman and Hall.
- Haslett, J., Whitley, M., Bhattacharya, S., Salter-Townshend, M., Wilson, S. P., Allen, J. R. M., Huntley, B., and Mitchell, F. J. G. [2006], "Bayesian Palaeoclimate Reconstruction (with discussion)," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **169**, 395-438.
- Hoadley, B. [1970], "A Bayesian Look at Inverse Linear Regression," *Journal of the American Statistical Association*, **65** : 356-69.

- Kimeldorf, G., and Wahba, G. [1971], "Some Results on Tchebycheffian Spline Functions," *Journal of Mathematical Analysis and Applications*, **33** : 82-95.
- König, H. [1986], *Eigenvalue Distribution of Compact Operators*, : Birkhäuser.
- Krutchkoff, R. G. [1967], "Classical and Inverse Regression Methods of Calibration," *Technometrics*, **9** : 425-35.
- Meinguet, J. [1979], "Multivariate Interpolation at Arbitrary Points Made Simple," *Journal of the Applied Mathematics and Physics*, **30** : 292-304.
- Mohammad-Djafari, A. [2000], "Model Selection for Inverse Problems: Best Choice of Basis Function and Model Order Selection," Available at <https://arxiv.org/abs/math-ph/0008026>.
- Mukhopadhyay, S., and Bhattacharya, S. [2013], "Cross-Validation Based Assessment of a New Bayesian Palaeoclimate Model," *Environmetrics*, **24** : 550-568.
- Muth, R. F. [1960], The Demand for Non-Farm Housing, in *The Demand for Durable Goods*, ed. A. C. Harberger. The University of Chicago.
- O'Sullivan, F. [1986], "A Statistical Perspective on Ill-Posed Inverse Problems," *Statistical Science*, **1** : 502-12.
- O'Sullivan, F., Yandell, B. S., and Raynor, W. J. [1986], "Automatic Smoothing of Regression Functions in Generalized Linear Models," *Journal of the American Statistical Association*, **81**, 96-103.
- Press, S. J., and Scott, A. [1975], Missing Variables in Bayesian Regression, in *Studies in Bayesian Econometrics and Statistics*, eds. S. E. Fienberg, and A. Zellner, North-Holland, Amsterdam.
- Rasch, D., Enderlein, G., and Herrendorfer, G. [1973], "Biometrie," Deutscher Landwirtschaftsverlag, Berlin.
- Rasmussen, C. E., and Williams, C. K. I. [2006], *Gaussian Processes for Machine Learning*, Cambridge, Massachusetts: The MIT Press.
- Schölkopf, B., and Smola, A. J. [2002], *Learning with Kernels*, USA: MIT Press.
- Shepp, L. A. [1966], "Radon-Nikodym Derivatives of Gaussian Measures," *Annals of Mathematical Statistics*, **37** : 321-54.
- Tikhonov, A. [1963], "Solution of Incorrectly Formulated Problems and the Regularization Method," *Soviet Math. Dokl.*, **5** : 1035-38.
- Tikhonov, A., and Arsenin, V. [1977], *Solution of Ill-Posed Problems*, New York: Wiley.
- Vollmer, S. [2013], "Posterior Consistency for Bayesian Inverse Problems Through Stability and Regression Results," *Inverse Problems*, **29** : Article number 125011.
- Wahba, G. [1978], "Improper Priors, Spline Smoothing and the Problem of Guarding Against Model Errors in Regression," *Journal of the Royal Statistical Society B*, **40** : 364-72.
- Wahba, G. [1990], "Spline Functions for Observational Data," CBMS-NSF Regional Conference series, SIAM. Philadelphia.
- Williams, E. J. [1969], "A Note on Regression Methods in Calibration," *Technometrics*, **11** : 189-92.