

Contents

1	Introduction	2
1.1	Outline	2
1.2	Further details and advantages of I-prior modelling	3
1.3	Relation with other work	4
1.4	Overview of paper	4
2	Function spaces with reproducing kernels	5
2.1	Definitions and basic properties	5
2.2	Some useful RKHSs	6
2.2.1	RKHS of constant functions	6
2.2.2	RKHSs over finite sets	7
2.2.3	RKHSs over a Hilbert space	7
2.3	Centering of an RKKS	7
3	Construction of RKKSs over product spaces	8
3.1	ANOVA RKHSs	8
3.2	Scale parameters for kernels	9
3.3	Application in regression modelling	9
3.3.1	Models with one covariate	10
3.3.2	Models with two covariates	11
3.3.3	Functional response model	11
3.3.4	Multi-class classification	12
4	The I-prior	12
5	Maximum marginal likelihood estimation of I-prior models	14
5.1	Marginal likelihood and posterior distribution of parameter estimates	14
5.2	Estimation of parameters using the EM algorithm	15
5.3	Asymptotic theory	17
5.4	Significance tests	19
5.4.1	HSIC	19
5.4.2	Likelihood ratio test	20
6	Application to data	22
6.1	Motivation for use of FBM RKHS	23
6.2	Regression with a functional covariate	23
6.3	Multilevel models	25
6.4	Multi-class classification	26
6.5	Longitudinal data analysis	27
7	Discussion and topics for further research	30
	References	30
A	ANOVA kernel construction	32
B	OLD	33
B.1	Asymptotic distribution of MLEs 1	33
B.2	Asymptotic distribution of MLEs 2	33
B.3	Test for significance of one covariate keeping other constant	34
B.3.1	OLD	35
B.4	old	36
B.4.1	Comparison with existing tests	37

Regression modelling with I-priors

With applications to functional, multilevel, and longitudinal data

Wicher Bergsma and Haziq Jamil

February 24, 2020

Abstract

We introduce a methodology with the aim of unifying and generalizing a variety of regression methods and models, including multilevel, varying coefficient, and longitudinal models, and models with functional covariates and responses. As we show, a natural space for the regression functions pertaining to such models is the reproducing kernel Krein space (RKKS). We introduce the I-prior over the RKKS of the model, defined as the maximizer of entropy subject to a suitable constraint involving the Fisher information on the regression function.

The I-prior is Gaussian with covariance kernel proportional to the Fisher information on the regression function, and the regression function is estimated by its posterior distribution under the I-prior. The I-prior has the intuitively appealing property that the more information is available on a linear functional of the regression function, the larger the prior variance, and the smaller the influence of the prior mean on the posterior distribution.

The methodology we introduce has some advantages compared to commonly used methods in terms of ease of estimation and model comparison. Firstly, a single methodology can be used for a variety of models, for which previously a different methods were used. Secondly, an EM algorithm with a simple E and M step for estimating the scale parameter of each covariate is available, facilitating estimation for complex models. Thirdly, we propose a novel parsimonious model formulation, requiring a single scale parameter for each covariate and no further parameters for interaction effects, allowing a semi-Bayes approach to the selection of interaction effects.

An R-package implementing our methodology is available (Jamil, 2019).

Keywords: reproducing kernel, RKHS, RKKS, Fisher information, objective prior, g -prior, empirical Bayes, regression, supervised learning, statistical learning, longitudinal data analysis, nonparametric regression, functional data analysis, multilevel models, classification, regularization, random effects.

1 Introduction

1.1 Outline

Consider a sample $(x_1, y_1), \dots, (x_n, y_n)$, where y_i is a real-valued measurement on unit i , and $x_i = (x_{i1}, \dots, x_{ip})$ is a row vector of p covariates, where each x_{ik} belongs to some set \mathcal{X}_k and may for example be real, categorical, multidimensional, or functional. To describe the dependence of the y_i on the x_i , we consider the regression model

$$y_i = f(x_i) + \varepsilon_i, \quad f \in \mathcal{F} \tag{1}$$

where \mathcal{F} is a space of functions. We assume the errors have a multivariate normal distribution, i.e.,

$$(\varepsilon_1, \dots, \varepsilon_n) \sim \text{MVN}(0, \Psi^{-1}), \tag{2}$$

where $\Psi = (\psi_{ij})$ is an $n \times n$ positive definite precision matrix. Here, Ψ is taken to be known up to a low dimensional parameter, e.g., $\Psi = \psi I_n$ ($\psi > 0$, I_n the $n \times n$ identity matrix), reflecting iid errors.

The function f is assumed to be partitioned into a sum of main effects and possible interactions. An example for $p = 2$ is

$$f(x) = f(x_1, x_2) = f_\emptyset + f_1(x_1) + f_2(x_2) + f_{12}(x_1, x_2), \tag{3}$$

and for $p = 3$,

$$f(x) = f(x_1, x_2, x_3) = f_\emptyset + f_1(x_1) + f_2(x_2) + f_3(x_3) + f_{12}(x_1, x_2) + f_{23}(x_2, x_3).$$

Here, each of the x_k may be, for example, scalar, categorical, Euclidean, or functional. As we explain, this setup includes multilevel, varying coefficient, and longitudinal models.

The intercept f_\emptyset is constant and we assume that the functions f_k are in a reproducing kernel Hilbert space (RKHS) of functions over a covariate space \mathcal{X}_k . The functions f_{kl} , f_{klm} , etc. describing interaction effects are assumed to lie in the tensor product space of the corresponding main effect function spaces. An RKHS possesses a positive definite kernel $h_k(x, x')$, where $x, x' \in \mathcal{X}_k$, and this kernel is multiplied by a scale parameter $\lambda_k \in \mathbb{R}$ which may be negative. If one or more of the scale parameters are negative, the resulting kernel for the space of regression functions \mathcal{F} is indefinite, so that \mathcal{F} is a *reproducing kernel Krein space* (RKKS). If $p = 1$, however, the model only has an intercept and a main effect, so there is only one scale parameter and the RKHS framework suffices.

The RKHS is the natural choice of Hilbert space for the main effect functions. In particular, in an RKHS point evaluation is continuous, implying that a normwise consistent estimator is also pointwise consistent. It is hard to envisage pointwise consistent estimation from a finite data set if point evaluation is discontinuous. Also from a practical point of view, RKHSs are convenient because the kernel allows us to do computations. However, for regression with more than one covariate, the RKHS framework is arbitrarily restrictive in that the aforementioned scale parameters are required to be positive, and the natural extension which doesn't suffer from this restriction is the RKKS framework. An RKKS can be viewed as an RKHS with the positive definiteness restriction relaxed, which turns out to simplify computations and statistical significance testing.

The I-prior for $f \in \mathcal{F}$ where \mathcal{F} is an RKKS is a Gaussian prior whose covariance kernel is the Fisher information for f . If \mathcal{F} has reproducing kernel h , then the Fisher information

between $f(x)$ and $f(x')$ is given as

$$\mathcal{I}[f(x), f(x')] = \sum_{i=1}^n \sum_{j=1}^n \psi_{ij} h(x, x_i) h(x', x_j).$$

Hence, f follows an I-prior distribution if it can be written in the form

$$f(x) = f_0(x) + \sum_{i=1}^n h(x, x_i) w_i, \quad (4)$$

where $f_0 \in \mathcal{F}$ is the prior mean (typically set to zero), and

$$(w_1, \dots, w_n) \sim \text{MVN}(0, \Psi) \quad (5)$$

As we show, the I-prior has a maximum entropy interpretation. The I-prior methodology for the case $p = 1$, when as mentioned the RKHS framework suffices, was described in some detail by Bergsma (2019).

An intuitively attractive property of the I-prior is that if much information about a linear functional of f (e.g., a regression coefficient) is available, its prior variance is large, and the data have a relatively large influence on the posterior, while if little information about a linear functional is available, the posterior will be largely determined by the prior mean, which serves as a ‘best guess’ of f . The I-prior methodology consists of estimation of the regression function by its posterior distribution under the I-prior, where we take the posterior mean as the summary measure.

1.2 Further details and advantages of I-prior modelling

As mentioned, each covariate, which may be categorical, multidimensional, or functional, is assigned a scale parameter. In this paper we estimate these using maximum marginal likelihood. An advantage of the I-prior methodology is the availability of an EM algorithm with simple E and M steps for estimating these scale parameters, which can greatly facilitate estimation, especially if there are many scale parameters. To compare, for Gaussian process regression there is no simple EM algorithm for estimating scale parameters.

Furthermore, in our approach scale parameters are only needed for each covariate, and no further parameters are needed for interaction effects. This is in contrast with the usual approach in the regularization and Gaussian process regression literature, where separate scale parameters are assigned to each interaction effect. Our parsimonious approach yields a simpler likelihood and hence simplifies estimation. In addition, this allows a semi-Bayes approach to the selection of interaction effects, potentially able to detect effects with smaller sample sizes than with existing approaches. The use of RKKSs implies there are no range restrictions for the (scalar) scale parameters. As a consequence, simple chi-square tests can be done for the significance of covariates, simplifying the chi-bar tests usually done in random effects modelling.

For simplicity three main classes of RKHSs will be used in this paper, allowing linear and smooth effects of Euclidean and functional covariates on a response, as well as the incorporation of categorical covariates: the *canonical* RKHS, consisting of linear functions of the covariates; the *fractional Brownian motion* (FBM) RKHS, consisting of smooth functions of the covariates; and the *canonical* or the *Pearson* RKHS for nominal categorical covariates. The FBM RKHS has smoothness parameter $\gamma \in (0, 1)$, called the Hurst coefficient.

Note that any RKHS has an associated Gaussian process which has the reproducing kernel as its covariance kernel. The functions in an RKHS are typically smoother than the paths of

the associated Gaussian process. For this reason, the FBM RKHS is quite useful in regression, containing functions of an appropriate smoothness range, whereas the FBM Gaussian process may have paths which are too rough for regression (see Bergsma (2019) for a detailed discussion).

The I-prior methodology has a theoretical advantage compared to Tikhonov regularization, described in more detail in Bergsma (2019). In particular, as is well-known, regularization can *systematically undersmooth*, whereas the I-prior does not. For example, with mean squared error loss and a penalty proportional to $\int \ddot{f}(x)^2 dx$, the regularizer is a cubic spline smoother, which is also the posterior mean under an integrated Brownian motion prior. However, with probability one, an integrated Brownian motion path does not have two derivatives, and as a result the regularizer is an inadmissible estimator (with respect to squared error loss) for the set of functions with two derivatives (Chakraborty & Panaretos, 2019).

For some commonly used models our methodology has some additional advantages. For example, in the standard approach, multilevel models require estimation of a latent covariance matrix for the random effects. If there are more than a few covariates this can be problematic due to the positive definiteness constraint and number of parameters to be estimated. In our approach, a smaller number of unrestricted scale parameters needs to be estimated.

1.3 Relation with other work

The present paper complements Bergsma (2019), which covers the case of a single, possibly multidimensional covariate. As mentioned, in this case there is a single scale parameter, and the RKHS framework suffices. In that paper, more details are given on the I-prior derivation, and generalization of I-priors to a broad class of statistical models is given. The relation with competing methods is outlined, including g -priors, Jeffreys and reference priors, and Fisher kernels. A detailed comparison with Tikhonov regularization is given, with particular detail on the relation with cubic spline smoothing. It is explained in detail how I-priors work when the regression functions are linear, or when they are assumed to lie in the FBM RKHS, which is a particularly attractive RKHS for I-prior modelling.

Jamil (2018) provides a number of extensions to the present methodology, including probit and logit models using a fully Bayes approach, Bayesian variable selection using I-priors, and Nyström approximations for speeding up the I-prior methodology. Furthermore, he contributed a user friendly R package `iprior` (Jamil, 2019), further described in Jamil and Bergsma (2019).

Ong, Mary, Canu, and Smola (2004) previously used RKKs in the context of regularization. In particular, they considered a regularization framework, where the usual RKHS squared penalty norm $\|f\|_{\mathcal{F}}^2$ is replaced by the RKKs indefinite inner product $\langle f, f \rangle_{\mathcal{F}}$. As the latter may be negative, it does not make sense to minimize the “penalized” loss function, and instead they sought a saddle point. Their approach is very different from ours, firstly in that they considered very different RKKs, and secondly by constructing a Gaussian prior over the RKKs the indefiniteness of the inner product becomes irrelevant.

1.4 Overview of paper

In Section 2, a summary of existing theory of RKHSs and RKKs is given as needed for this paper. In Section 3, we describe the construction of RKKs over product spaces. We describe how a number of well-known models, such as the varying intercept model, one-dimensional smoothing, and multidimensional (or functional) response models can be described using the RKKs framework. In Section 4, the I-prior is defined and its representation (4) is derived for model (1) with multivariate normal errors. In Section 5, the EM algorithm for estimating scale

parameters is described. In Section 6, we apply the I-prior methodology to a number of data examples in the respective areas of multilevel modelling, functional data analysis, classification and longitudinal data analysis, illustrating some possible advantages over existing techniques, and showing competitive predictive performance.

2 Function spaces with reproducing kernels

This section summarizes existing theory as needed for this paper. In Section 2.1 we give the definition and some well-known basic properties of RKHSs and RKKSs. Section 2.2 briefly lists the RKHSs used in this paper. These RKHSs are used as building blocks to construct RKKSs over product spaces, called ANOVA RKKSs, which is the topic of the next Section 3. The RKHSs we use in this paper will normally be *centered*, i.e., the functions in the RKHS have zero mean, which is formally described in Section 2.3.

2.1 Definitions and basic properties

The first comprehensive treatment of RKHSs was given by Aronszajn (1950), and their usefulness for statistics was initially demonstrated by Parzen (1961) and further developed by Kimeldorf and Wahba (1970). Some more recent overviews of RKHS theory with a view to application in statistics and machine learning are Wahba (1990b), Berlinet and Thomas-Agnan (2004), Steinwart and Christmann (2008, Chapter 4) and Hofmann, Schölkopf, and Smola (2008). Schwartz (1964) developed a general theory of Hilbertian subspaces of topological vector spaces which includes the theory of RKKSs. The first applications to statistics and machine learning of RKKSs were given by Ong et al. (2004) and Canu, Ong, and Mary (2009). A recent technical survey of the theory of RKKSs is given by Gheondea (2013). Below, we give a very brief overview of the theory as needed for this paper, more details can be found in the aforementioned literature.

We begin with the definition of the (possibly indefinite or negative definite) inner product.

Definition 1. *Let \mathcal{F} be a vector space over the reals. A function $\langle \cdot, \cdot \rangle_{\mathcal{F}} : \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}$ is called an inner product on \mathcal{F} if, for all $f, f', f'' \in \mathcal{F}$,*

- (symmetry) $\langle f, f' \rangle_{\mathcal{F}} = \langle f', f \rangle_{\mathcal{F}}$
- (linearity) $\langle \alpha f + f', f'' \rangle_{\mathcal{F}} = \alpha \langle f, f'' \rangle_{\mathcal{F}} + \langle f', f'' \rangle_{\mathcal{F}}$
- (nondegeneracy) $(\forall g \in \mathcal{F} : \langle f, g \rangle_{\mathcal{F}} = 0) \Rightarrow f = 0$

If $\langle f, f \rangle_{\mathcal{F}} \geq 0$ for all $f \in \mathcal{F}$, the inner product is called positive definite and $\|f\|_{\mathcal{F}} := \langle f, f \rangle_{\mathcal{F}}$ is called a norm on \mathcal{F} . If $\langle f, f \rangle_{\mathcal{F}} \leq 0$ for all $f \in \mathcal{F}$, the inner product is called negative definite. An inner product which is neither positive definite nor negative definite is called indefinite.

Recall that a Hilbert space is a complete inner product space with a positive definite inner product. The more general notion of Krein space is defined as follows.

Definition 2. *A vector space \mathcal{F} equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{F}}$ is called a Krein space if there are two Hilbert spaces \mathcal{F}_+ and \mathcal{F}_- spanning \mathcal{F} such that*

- All $f \in \mathcal{F}$ can be decomposed as $f = f_+ + f_-$ where $f_+ \in \mathcal{F}_+$ and $f_- \in \mathcal{F}_-$.
- For all $f, f' \in \mathcal{F}$, $\langle f, f' \rangle_{\mathcal{F}} = \langle f_+, f'_+ \rangle_{\mathcal{F}_+} - \langle f_-, f'_- \rangle_{\mathcal{F}_-}$

\mathcal{X}	RKHS	Functions $f(x)$	Kernel $h(x, x')$	Centered kernel
Any set	Constant	Constant functions	1	-
Finite set	Canonical	All functions	$\delta_{xx'}$	$\delta_{xx'} - p(x)p(x')$
Finite set	Pearson	All zero mean functions	$\delta_{xx'}/p(x) - 1$	$\delta_{xx'}/p(x) - 1$
Hilbert space	Canonical	$\langle x, \beta \rangle_{\mathcal{X}}$	$\langle x, x' \rangle_{\mathcal{X}}$	$\langle x - \bar{x}, x' - \bar{x} \rangle_{\mathcal{X}}$
\mathbb{R}^p	Mahalanobis	$x^\top \beta$	$x^\top S^{-1} x'$	$(x - \bar{x})^\top S^{-1} (x' - \bar{x})$
\mathbb{R}	Brownian motion	$\int_{-\infty}^x \beta(t) dt$	$\frac{1}{2}(x + x' - x - x')$	Eq. (7) ($\gamma = 1/2$)
Hilbert space	Brownian motion	Hölder $\geq 1/2$	$\frac{1}{2}(\ x\ + \ x'\ - \ x - x'\)$	Eq. (7) ($\gamma = 1/2$)
Hilbert space	FBM- γ	Hölder $\geq \gamma$	$\frac{1}{2}(\ x\ ^{2\gamma} + \ x'\ ^{2\gamma} - \ x - x'\ ^{2\gamma})$	Eq. (7)

Table 1: List of RKHSs

Note that any Hilbert space is a Krein space, which can be seen by taking $\mathcal{F}_- = \{0\}$.

We next define the notion of a *reproducing kernel*:

Definition 3. Let \mathcal{F} be a Krein space of functions over a set \mathcal{X} . A symmetric function $h : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a reproducing kernel of \mathcal{F} if and only if

(a) $h(x, \cdot) \in \mathcal{F}$ for all $x \in \mathcal{X}$

(b) $f(x) = \langle f, h(x, \cdot) \rangle_{\mathcal{F}}$ for all $f \in \mathcal{F}$ and $x \in \mathcal{X}$.

A Hilbert space resp. Krein space is called a *reproducing kernel Hilbert space* (RKHS) resp. *reproducing kernel Krein space* (RKKS) if it possesses a reproducing kernel. Sometimes in this paper we will use the shorthand ‘kernel’ to refer to ‘reproducing kernel’.

A function $h : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is said to be positive definite on \mathcal{X} if $\sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j h(x_i, x_j) \geq 0$ for all scalars $\alpha_1, \dots, \alpha_n$ and all $x_1, \dots, x_n \in \mathcal{X}$. From the definition of positive definite inner products it follows that the reproducing kernel of an RKHS is symmetric and positive definite. The reproducing kernel of an RKKS can be shown to be the difference of two positive definite kernels so need not be positive definite. It follows that an inner product in an RKKS $(\mathcal{F}, \langle \cdot, \cdot \rangle_{\mathcal{F}})$ is the difference of two positive definite inner products

$$\langle f, f' \rangle_{\mathcal{F}} = \langle f, f' \rangle_{\mathcal{F}^+} - \langle f, f' \rangle_{\mathcal{F}^-}$$

We define the norm of $f \in \mathcal{F}$ as the norm in the corresponding RKHS, i.e.,

$$\|f\|_{\mathcal{F}}^2 = \|f\|_{\mathcal{F}^+}^2 + \|f\|_{\mathcal{F}^-}^2 \quad (6)$$

The Moore-Aronszajn theorem states that every symmetric positive definite function defines a unique RKHS. Every RKSS also has a unique kernel, but a given kernel may have more than one RKKS associated with it (e.g., Alpay, 1991).

2.2 Some useful RKHSs

Below we describe some RKHSs that we will use in this paper. A summary is given in Table 1.

2.2.1 RKHS of constant functions

The RKHS of constant functions with reproducing kernel given by $h(x, x') = 1$. For a constant function f with $f(x) = c$, $\|f\|_{\mathcal{F}} = |c|$. (The RKHS of constant functions will be an essential component in the construction of RKKSs over product spaces in Section 3.1.)

2.2.2 RKHSs over finite sets

Let \mathcal{X} be a finite set. The *canonical* RKHS over \mathcal{X} is the RKHS whose kernel is the Kronecker delta function, i.e., $h(x, x') = \delta_{xx'}$ consists of the set of all functions $f : \mathcal{X} \rightarrow \mathbb{R}$, with squared norm

$$\|f\|_{\mathcal{F}}^2 = \sum_{x \in \mathcal{X}} f(x)^2.$$

Note that, viewing f as a $|\mathcal{X}|$ -dimensional vector, the canonical RKHS over \mathcal{X} is just standard Euclidean space.

Alternatively, the Pearson RKHS over a finite probability space (\mathcal{X}, p) , defined as the RKHS with reproducing kernel $h(x, x') = p(x)^{-1}\delta_{xx'} - 1$, consisting of all functions with $\sum p(x)f(x) = 0$ and

$$\|f\|_{\mathcal{F}}^2 = \sum p(x)f(x)^2$$

2.2.3 RKHSs over a Hilbert space

Let \mathcal{X} be a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{X}}$. The *canonical RKHS* over \mathcal{X} is defined as its continuous dual space whose reproducing kernel is given by

$$h(x, x') = \langle x, x' \rangle_{\mathcal{X}}.$$

Functions in this space are of the form $f(x) = \langle x, \beta \rangle_{\mathcal{X}}$, with norm $\|f\|_{\mathcal{F}} = \|\beta\|_{\mathcal{X}}$.

A special case is the *Mahalanobis RKHS*, defined as the canonical RKHS over $\mathcal{X} = \mathbb{R}^p$ equipped with the Mahalanobis inner product; for a covariance matrix S , it is defined as

$$\langle x, x' \rangle_{\text{Mah}} = x^{\top} S^{-1} x'$$

The *Brownian motion RKHS* is defined as the RKHS over \mathcal{X} whose reproducing kernel is the generalized Brownian motion covariance kernel

$$h(x, x') = -\frac{1}{2} \left(\|x - x'\|_{\mathcal{X}} - \|x\|_{\mathcal{X}} - \|x'\|_{\mathcal{X}} \right)$$

Functions in the Brownian motion RKHS are Hölder of degree at least 1/2 (see Bergsma (2019) for a proof). In the simplest nontrivial case, $\mathcal{X} = \mathbb{R}$, and the RKHS consists of functions with a square integrable derivative, whose norm is the L^2 norm of the derivative, i.e., every $f \in \mathcal{F}$ can be written as $f(x) = \int_{-\infty}^x \beta(t) dt$ for some square integrable β , and has norm $\int_{\mathbb{R}} \beta(t)^2 dt$.

The fractional Brownian motion (FBM) RKHS is the RKHS whose reproducing kernel is the generalized FBM covariance kernel

$$h(x, x') = -\frac{1}{2} \left(\|x - x'\|_{\mathcal{X}}^{2\gamma} - \|x\|_{\mathcal{X}}^{2\gamma} - \|x'\|_{\mathcal{X}}^{2\gamma} \right)$$

Functions in the FBM- γ RKHS are Hölder of degree at least γ (see Bergsma (2019) for a proof).

2.3 Centering of an RKKS

We say a function space \mathcal{F} over \mathcal{X} is *centered* with respect to a data set $\{x_1, \dots, x_n\} \subset \mathcal{X}$ if

$$\sum_{i=1}^n f(x_i) = 0 \quad \forall f \in \mathcal{F}$$

It can be verified that an RKKS \mathcal{F} is centered if and only if its kernel h is centered, in the sense that $\sum_{i=1}^n h(x, x_i) = 0$ for all $x \in \mathcal{X}$. If h is a kernel, then h_{cent} defined as follows is centered:

$$h_{\text{cent}}(x, x') = h(x, x') - \frac{1}{n} \sum_{j=1}^n h(x, x_j) - \frac{1}{n} \sum_{i=1}^n h(x_i, x') + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n h(x_i, x_j)$$

Table 1 gives a list of kernels discussed in Section 2.2 and their centered versions. The centered FBM RKKS has kernel

$$h_{\text{cent}}(x, x') = -\frac{1}{2} \left(\|x - x'\|_{\mathcal{X}}^{2\gamma} - \frac{1}{n} \sum_{j=1}^n \|x - x_j\|_{\mathcal{X}}^{2\gamma} - \frac{1}{n} \sum_{i=1}^n \|x_i - x'\|_{\mathcal{X}}^{2\gamma} + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \|x_i - x_j\|_{\mathcal{X}}^{2\gamma} \right) \quad (7)$$

where the Brownian motion RKHS is obtained if $\gamma = 1/2$.

3 Construction of RKKSs over product spaces

ANOVA constructions of RKHSs over product spaces are a natural tool for formulating regression models and were introduced for this purpose by Wahba (1990a) and Gu and Wahba (1993). In Section 3.1 we describe ANOVA RKKSs, an immediate extension of ANOVA RKHSs which are needed in this paper. In Section 3.2, we describe what as far as we are aware is a novel approach to use scale parameters parsimoniously in the ANOVA construction. In Section 3.3 we show how the framework is useful in regression, as, for example, it can be used to easily formulate multilevel and varying coefficient models.

3.1 ANOVA RKHSs

An ANOVA decomposition of a function f over product space $\mathcal{X}_1 \times \mathcal{X}_2$ is given by

$$f(x_1, x_2) = f_{\emptyset} + f_1(x_1) + f_2(x_2) + f_{12}(x_1, x_2)$$

where the components are orthogonal in some way. To formalize this, let us first define the tensor product of RKHSs. Let \mathcal{F}_1 and \mathcal{F}_2 be two RKHSs over \mathcal{X}_1 resp. \mathcal{X}_2 . For $f_1 \in \mathcal{F}_1$ and $f_2 \in \mathcal{F}_2$, the tensor product $f_{12} = f_1 \otimes f_2$ is defined by $f_{12}(x_1, x_2) = f_1(x_1)f_2(x_2)$. The tensor product of \mathcal{F}_1 and \mathcal{F}_2 is denoted as $\mathcal{F}_1 \otimes \mathcal{F}_2$ and is defined as the closure of the set of functions $\{f_1 \otimes f_2 | f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2\}$ equipped with the inner product

$$\langle f_1 \otimes f_2, f_1' \otimes f_2' \rangle_{\mathcal{F}_1 \otimes \mathcal{F}_2} = \langle f_1, f_1' \rangle_{\mathcal{F}_1} \langle f_2, f_2' \rangle_{\mathcal{F}_2}.$$

The tensor product of RKKSs is defined analogously, the closure being defined with respect to the corresponding positive definite inner product.

Let \mathcal{C}_k be the RKHS of constant functions over \mathcal{X}_k with kernel $c_k(x, x') = 1$ and let \mathcal{F}_k be an RKKS over \mathcal{X}_k with kernel h_k ($k = 1, 2$). An ANOVA RKKS over $\mathcal{X}_1 \times \mathcal{X}_2$ is given as

$$\mathcal{F} = \mathcal{C}_1 \otimes \mathcal{C}_2 + \mathcal{F}_1 \otimes \mathcal{C}_2 + \mathcal{C}_1 \otimes \mathcal{F}_2 + \mathcal{F}_1 \otimes \mathcal{F}_2$$

with kernel h given by

$$h((x_1, x_2), (x_1', x_2')) = 1 + h_1(x_1, x_1') + h_2(x_2, x_2') + h_1(x_1, x_1')h_2(x_2, x_2')$$

In this paper we assume the components h_k are centered relative to data $\{x_{1k}, \dots, x_{nk}\} \subset \mathcal{X}_k$ (Section 2.3). Hence, the \mathcal{C}_k and \mathcal{F}_k are orthogonal in the sense that

$$\sum_{i=1}^n c_k(x_{1k}) f_k(x_{1k}) = \sum_{i=1}^n f_k(x_{1k}) = 0$$

for any $c_k \in \mathcal{C}_k$ and $f_k \in \mathcal{F}_k$.

With p covariates, the ANOVA model with all interactions can be written succinctly as

$$\mathcal{F} = \bigoplus_{k=1}^p (\mathcal{C}_k \otimes \mathcal{F}_k) \quad (8)$$

with reproducing kernel

$$h(x, x') = \prod_{k=1}^p (1 + h_k(x_k, x'_k)) \quad (9)$$

More general ANOVA kernels are described in Appendix A.

3.2 Scale parameters for kernels

In practice, the length of a vector in an RKHSs is measured on an arbitrary scale, and this can be taken into account by multiplying the kernel by a real-valued scale parameter which is to be estimated. In the ANOVA case, we can use component kernels $\mu_k h_k$ and $\tau_k c_k$ for real-valued μ_k and τ_k , giving the ANOVA kernel

$$h_{\mu, \tau}((x_1, x_2), (x'_1, x'_2)) = \tau_1 \tau_2 + \mu_1 \tau_2 h_1(x_1, x'_1) + \tau_1 \mu_2 h_2(x_2, x'_2) + \mu_1 \mu_2 h_1(x_1, x'_1) h_2(x_2, x'_2)$$

This expression is overparameterized, and setting $\lambda_0 = \tau_1 \tau_2$ and $\lambda_k = \mu_k / \tau_k$ (assuming $\tau_k \neq 0$), we obtain the identified parameterization

$$h_\lambda((x_1, x_2), (x'_1, x'_2)) = \lambda_0 \{1 + \lambda_1 h_1(x_1, x'_1) + \lambda_2 h_2(x_2, x'_2) + \lambda_1 \lambda_2 h_1(x_1, x'_1) h_2(x_2, x'_2)\} \quad (10)$$

Typically, the kernels h_k will be positive definite, so that the corresponding \mathcal{F}_k are RKHSs. Then if at least one of the lambda parameters is negative, a function space with h_λ as its kernel will be an RKKS.

In the literature, a less parsimonious construction than (10) has been used, namely

$$h_v((x_1, x_2), (x'_1, x'_2)) = v_0 + v_1 h_1(x_1, x'_1) + v_2 h_2(x_2, x'_2) + v_{12} h_1(x_1, x'_1) h_2(x_2, x'_2) \quad (11)$$

(e.g., Wahba, 1990b, Section 10.2, Berlinet & Thomas-Agnan, 2004, Section 10.2, Gu, 2013, Section 2.4.5). We refer to the corresponding RKKS as the *extended ANOVA RKKS*. Here, each of the four terms has a separate scale parameter, and is thus less parsimonious than our approach which only requires three scale parameters. For models with all interactions, our approach has $p + 1$ scale parameters, while 2^p parameters are required if every interaction is assigned a separate parameter.

3.3 Application in regression modelling

We show how some well-known models which can be written in the form (1) can be formulated using the ANOVA function space construction. In Sections 12 and 14 we consider models with one and two covariates respectively (see also Tables 2 and 3). In Section 3.3.3 we consider functional responses, and in Section 3.3.4 we consider multi-class classification models.

Model: $y_i = \alpha + f_1(x_i) + \varepsilon_i$, $x_i \in \mathcal{X}_1$, $f_1 \in \mathcal{F}_1$			
\mathcal{X}_1	RKHS \mathcal{F}_1	Model name	Usual notation
A finite set	Pearson	One-way ANOVA/Varying intercept model	$y_{ij} = \alpha + \beta_j + \varepsilon_{ij}$
\mathbb{R}	Canonical	Simple regression	$y_i = \alpha + x_i\beta + \varepsilon_i$
\mathbb{R}	FBM	Smoothing spline model	$y_i = \alpha + \int_0^{x_i} \beta(t)dt + \varepsilon_i$
An RKHS	Canonical	Functional linear regression	$y_i = \alpha + \int x_i(t)\beta(t)d\mu(t) + \varepsilon_i$
An RKHS	FBM	Smooth functional regression	$y_i = \alpha + f(x_i) + \varepsilon_i$

Table 2: Some models with one, possibly multidimensional covariate

3.3.1 Models with one covariate

We consider examples of regression models (1) where \mathcal{F} is of the form (8) with $p = 1$. We can write

$$f(x) = \alpha + f_1(x), \quad f_1 \in \mathcal{F}_1, x \in \mathcal{X}_1 \quad (12)$$

where \mathcal{F}_1 is in the centered RKHS over a set \mathcal{X}_1 with kernel h_1 . Recall that the centering means $\sum_{i=1}^n f(x_i) = 0$ for all $f \in \mathcal{F}_1$.

If \mathcal{X}_1 is a finite set, then (12) is known as a one-way ANOVA model or varying intercept model. The usual notation for this model is

$$y_{jk} = \alpha + f_1(k) + \varepsilon_{jk} \quad f_1 \in \mathcal{F}_1, k \in \mathcal{X}_1 \quad (13)$$

where y_{jk} is the k th y_i -value for which $x_i = j$. Suitable RKHSs \mathcal{F}_1 are the canonical for which $\|f\|_{\mathcal{F}}^2 = \sum f(j)^2$ or the Pearson for which $\|f\|_{\mathcal{F}}^2 = \sum p(j)f(j)^2$, where $p(j)$ is the number of x_i s equal to j (Section 2.3).

If \mathcal{X}_1 is a subset of a Hilbert space, a flexible range of models is obtained by taking \mathcal{F}_1 to be the canonical RKHS, the Brownian motion RKHS, or more generally the FBM RKHS. In the first case, we obtain

$$y_i = \alpha + \langle x_i, \beta \rangle_{\mathcal{X}} + \varepsilon_i$$

If if $\mathcal{X} = \mathbb{R}^q$ we obtain the special case

$$y_i = \alpha + \sum_{k=1}^q x_{ik}\beta_k + \varepsilon_i$$

or if $\mathcal{X} = L^2(\mathcal{X}, \mu)$ we obtain the functional linear model (e.g., Yao, Müller, and Wang (2005))

$$y_i = \alpha + \int_{\mathcal{X}} x_i(t)\beta(t)d\mu(t) + \varepsilon_i$$

A smooth dependence model is obtained if \mathcal{F}_1 is the Brownian motion or FBM RKHS. A special case is the smoothing spline model

$$y_i = \alpha + \int_{-\infty}^{x_i} \beta(t)dt + \varepsilon_i$$

when $\mathcal{X}_1 = \mathbb{R}$. However, there are many other potentially useful linear or smooth dependence models, such as the model where \mathcal{X} is a Brownian motion RKHS.

Model: $y_i = \alpha + f_1(x_{1i}) + f_2(x_{2i}) + f_{12}(x_{1i}, x_{2i}) + \varepsilon_i$, $x_{ki} \in \mathcal{X}_k$, $f_k \in \mathcal{F}_k$, $f_{12} \in \mathcal{F}_1 \otimes \mathcal{F}_2$					
\mathcal{X}_1	RKHS \mathcal{F}_1	\mathcal{X}_2	RKHS \mathcal{F}_2	Model name	Usual notation
Finite set	Pearson	\mathbb{R}	Canonical	Varying slope model	$y_{ij} = \alpha + \alpha_j + x_{ij}\beta + x_{ij}\beta_j + \varepsilon_{ij}$
\mathbb{R}^p	FBM	\mathbb{R}^p	Canonical	Varying coefficient model	$y_i = \alpha + x_{2i}^\top \beta(x_{1i}) + \varepsilon_i$
\mathbb{R}^p	FBM	\mathbb{R}^p	Canonical	Multivariate regression	$y_{ij} = \alpha + x_i \beta_j + \varepsilon_{ij}$
\mathbb{R}^p	FBM	\mathbb{R}^p	Canonical	Functional response model	$y_i(t) = f(x_i, t) + \varepsilon_i$

Table 3: Some models with two, possibly multidimensional covariates

3.3.2 Models with two covariates

We next consider examples of regression models (1) where \mathcal{F} is of the form (8) with $p = 2$. We can write

$$f(x_1, x_2) = \alpha + f_1(x_1) + f_2(x_2) + f_{12}(x_1, x_2), \quad f_k \in \mathcal{F}_k, x_k \in \mathcal{X}_k, k = 1, 2, \quad (14)$$

where \mathcal{F}_k is in the centered RKHS over a set \mathcal{X}_k with kernel h_k , and $f_{12} \in \mathcal{F}_1 \otimes \mathcal{F}_2$.

First consider the case that $\mathcal{X}_1 = \mathbb{R}$ and \mathcal{X}_2 is a finite set. Taking \mathcal{F}_1 to be the canonical RKHS yields the *varying slope* model, that is, for each element of \mathcal{X}_2 , we have a linear dependence model. The usual representation of this model is as a two-level regression model asserting that the y_{ij} depend both on the covariate x_{ij} and on the cluster j ,

$$y_{ij} = f(j, x_{ij}) + \varepsilon_{ij} \quad i = 1, \dots, n_j, j = 1, \dots, m, x_{ij} \in \mathcal{X}. \quad (15)$$

where

$$f(j, x_{ij}) = \alpha + \beta_{1,j} + x_{ij}\beta_2 + x_{ij}\beta_{12,j}, \quad (16)$$

Here, $\alpha + \beta_{1,j}$ is called the intercept for cluster j and $\beta_2 + \beta_{12,j}$ its slope.

Next consider the case $\mathcal{X}_1 = \mathcal{X}_2 = \mathbb{R}$. The canonical RKHSs give the model

$$y_i = \alpha + x_{1i}\beta_1 + x_{2i}\beta_2 + x_{1i}x_{2i}\beta_{12} + \varepsilon_i$$

Note that this approach is suitable if the x_{1i} and x_{2i} are measured on different scales, such as height and weight. If the x_{ik} are measured on the same scale, e.g., height measured at two different time points, it may be better to consider the pairs (x_{i1}, x_{2i}) as a single covariate in \mathbb{R}^2 and use the approach in Section 3.3.1.

If \mathcal{F}_1 is the canonical RKHS and \mathcal{F}_2 the Brownian motion RKHS, we obtain

$$y_i = \alpha + x_{1i}\beta_1 + f_2(x_{2i}) + x_{1i}\beta(x_{2i}) + \varepsilon_i$$

This has been called the *varying coefficient model* (Hastie & Tibshirani, 1993), where $\beta(x_{2i})$ is the varying regression coefficient for the x_{1i} .

3.3.3 Functional response model

We now consider the case that, rather than scalars, the y_i are real-valued functions over a set \mathcal{T} . A regression model then be formulated as

$$y_i(t) = f(x_i, t) + \varepsilon_{it} \quad i = 1, \dots, n, t \in \mathcal{T} \quad (17)$$

where

$$f(x_i, t) = \alpha + f_1(x_i) + f_2(t) + f_{12}(x_i, t) \quad f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2, f_{12} \in \mathcal{F}_1 \otimes \mathcal{F}_2$$

If \mathcal{T} is a finite set, y_i can be viewed as a vector in $\mathbb{R}^{|\mathcal{T}|}$. Then if additionally \mathcal{F}_1 is the canonical RKHS over \mathbb{R}^p , we obtain the usual multivariate regression model

$$y_{it} = \alpha + \alpha_t + x_i^\top \beta_t + \varepsilon_{it}$$

where $y_{it} = y_i(t)$ and $\alpha_t = f_2(t)$.

In practice, we do not observe y_i entirely but rather a finite set of evaluations at index points u_{i1}, \dots, u_{im_i} , i.e., we observe $y_i(u_{i1}), \dots, y_i(u_{im_i})$. For example, in a repeated measurements setting, both the number of measurements m_i and the times of measurement may be different for different units i . Then (17) becomes

$$y_i(u_{is}) = f(x_i, u_{is}) + \varepsilon_{is} \quad i = 1, \dots, n, s = 1, \dots, m_i$$

Note that this can be viewed as an instance of model (1).

In Section 6.5 we apply this model, as well as an extension with an extra covariate, to a longitudinal data set, taking \mathcal{F}_2 the FBM RKHS.

3.3.4 Multi-class classification

Consider a multi-class classification problem where, with \mathcal{C} a finite set of classes, we have observations $(x_1, c_1), \dots, (x_n, c_n)$ for $x_i \in \mathcal{X}$ and $c_i \in \mathcal{C}$. The aim is to find a prediction function to predict the class $c \in \mathcal{C}$ for a future observation $x \in \mathcal{X}$. We can use the present framework as follows. Let $y_{ij} = 1$ if $c_i = j$ and let $y_{ij} = 0$ otherwise. We may now consider the model

$$y_{ij} = f(x_i, j) + \varepsilon_{ij}, \quad i = 1, \dots, n, j \in \mathcal{C}$$

where the y_{ij} satisfy the restriction

$$\sum_{j \in \mathcal{C}} y_{ij} = 1 \quad \forall i$$

Assuming $\sum_{j \in \mathcal{C}} \varepsilon_{ij} = 0$, we then have $\sum_{j \in \mathcal{C}} f(x_i, j) = 1$ and we may consider a decomposition

$$f(x_i, j) = |\mathcal{C}|^{-1} + f_2(j) + f_{12}(x_i, j) + \varepsilon_{ij} \quad f_2 \in \mathcal{F}_2, f_{12} \in \mathcal{F}_1 \otimes \mathcal{F}_2 \quad (18)$$

subject to the identifying restrictions $\sum_{j \in \mathcal{C}} f_2(j) = 0$, $\sum_{j \in \mathcal{C}} f_{12}(x_i, j) = 0$ for all i , and $\sum_{i=1}^n f_{12}(x_i, j) = 0$ for all j . Hence, we may assume \mathcal{F}_2 is the centered canonical RKHS, and \mathcal{F}_1 is any appropriate RKKS. Note that a main effect for x_i is not needed in (18).

Of course, more than one covariate can be incorporated, for example with $p = 2$ we may take $\mathcal{F}_1 = \mathcal{F}_{11} + \mathcal{F}_{12} + \mathcal{F}_{11} \otimes \mathcal{F}_{12}$ so that

$$f_{12}(x_i, j) = f_{12}(x_{1i}, x_{2i}, j) = g_1(x_{1i}, j) + g_2(x_{2i}, j) + g_{12}(x_{1i}, x_{2i}, j)$$

In this paper we naively assume the errors are iid normal conditional on $\sum_j \varepsilon_{ij} = 0$. Though this is unrealistic, the real data examples in Section 6 show competitive performance of this approach.

4 The I-prior

Consider model (1) subject to (2), where \mathcal{F} is an RKKS with reproducing kernel h . In this section we derive a prior for the regression function f based on the Fisher information on f .

As shown by Bergsma (2019), the Fisher information on f is given by

$$I[f](x, x') = \sum_{i=1}^n \sum_{j=1}^n \psi_{ij} h(x, x_i) h(x, x_j)$$

$I[f]$ is positive definite and hence induces an RKHS over \mathcal{X} , which we denote by \mathcal{F}_n . Note that \mathcal{F}_n is a finite dimensional subspace of \mathcal{F} , consisting of functions of the form $f(x) = \sum_{i=1}^n h(x, x_i) w_i$. With \hat{f} an unbiased estimator of the true regression function f , the Crámer-Rao inequality implies that for any $g \in \mathcal{F}$

$$\text{var}(\langle g, \hat{f} \rangle_{\mathcal{F}}) \geq \|g\|_{\mathcal{F}_n}^2$$

By standard weighted least squares theory, equality is achieved if \hat{f} is a maximum likelihood estimator of f .

We can write any $f \in \mathcal{F}$ as $f = f_n + r_n$, where $f_n \in \mathcal{F}_n$ and $r_n(x_1) = \dots = r_n(x_n) = 0$. Then $r_n \in \mathcal{F}_n^\perp$, where \mathcal{F}_n^\perp is the orthogonal complement of \mathcal{F}_n in \mathcal{F} . The likelihood for f does not depend on r_n , i.e., the data contain no information on r_n , and we can replace r_n by a ‘best guess’. In this paper, we set $r_n = 0$.

We define the I-prior as a maximum entropy prior as follows. Let ν be volume measure induced by $\|\cdot\|_{\mathcal{F}_n}$. The entropy of a prior π over \mathcal{F}_n relative to ν is

$$\mathcal{E}(\pi) = - \int_{\mathcal{F}_n} \pi(f) \log \pi(f) \nu(df).$$

We define the I-prior for f as the prior π maximizing entropy subject to the constraint

$$E_{g \sim \pi} \|g\|_{\mathcal{F}_n}^2 = \text{constant}$$

Variational calculus shows that an I-prior for f is the Gaussian variable with mean $r_0 = 0$ and covariance kernel proportional to the Fisher information on f , i.e.,

$$\text{cov}_\pi(f(x), f(x')) = \tau^2 \sum_{i=1}^n \sum_{j=1}^n \psi_{ij} h(x, x_i) h(x', x_j)$$

for some $\tau > 0$. Equivalently, under the I-prior, f can be written in the form

$$f(x) = \tau \sum_{i=1}^n h(x, x_i) w_i, \quad (w_1, \dots, w_n) \sim \text{MVN}(0, \Psi), \quad (19)$$

Note that the ANOVA kernel described in Section 3 has a scale parameter λ_0 which makes the parameter τ superfluous, and we will omit τ in further developments.

Let us consider some examples, described in more detail in Bergsma (2019). Firstly, if \mathcal{F} consists of functions $f(x) = x^\top \beta$ ($x, \beta \in \mathbb{R}^p$) with norm $\|f\|_{\mathcal{F}} = \|\beta\|_{\mathbb{R}^p}$, then under the I-prior, $\beta \sim \text{MVN}(0, \lambda X^\top \Psi X)$, where X is the $n \times p$ matrix with i th row x_i . If instead of with the Euclidean norm we equip \mathbb{R}^p with the Mahalanobis norm, $\beta \sim \text{MVN}(0, \lambda(X^\top \Psi X)^{-1})$, which is the g -prior (Zellner, 1986) with $g = \lambda$.

Secondly, if \mathcal{F} is the Brownian motion RKHS, $f \in \mathcal{F}$ are of the form

$$f(x) = \int_{-\infty}^x \beta(t) dt,$$

and $\|f\|_{\mathcal{F}}^2 = \int_{\mathbb{R}} \beta(t)^2 dt$. If $f \in \mathcal{F}_n \subset \mathcal{F}$, then it can be shown that

$$\beta(t) = \sum_{i:x_i \leq t} w_i = \frac{f(x_{i_t+1}) - f(x_{i_t})}{x_{i_t+1} - x_{i_t}},$$

with $i_t = \max_{x_i \leq t} i$. Note that $\sum w_i = 0$ and hence $\lim_{t \rightarrow \pm\infty} \beta(t) = 0$. Under the I-prior, (5) holds, hence the I-prior for f is an generalized integrated discrete Brownian bridge (generalized as the normal increments need not be independent). If the errors are iid, β is an ordinary Brownian bridge. The posterior mean is therefore similar to a cubic spline smoother, which is the posterior mean based on an integrated Brownian motion. Note however that the I-prior assumes the true regression function has only a single derivative, while the cubic spline smoother, based on a penalty $\int \dot{f}(t)^2 dt$, assumes it has two.

5 Maximum marginal likelihood estimation of I-prior models

In this section we describe an EM algorithm for estimating the hyperparameters of model (1) subject to (2). In this paper we use ANOVA kernels as described in Section 3, so the hyperparameters are the lambda scale parameters (of which there are $p + 1$ for p possibly multi-dimensional covariates) and the parameters of the error precision matrix Ψ . For example, if errors are iid, $\Psi = \psi I$ for some $\psi > 0$ which is to be estimated, or if the errors are AR(1) two real-valued parameters need to be estimated.

In Section 5.1, we give the marginal likelihood of the hyperparameters, and in Section 5.2 we describe an efficient EM algorithm for estimating them. In particular, the E step is in closed form and the M steps is either also in closed form or, with d scale parameters, requires d polynomials in d unknowns to be solved, which can be done very efficiently. We may compare this with Gaussian process regression, for which it is well-known that the M step is as complex as maximizing the marginal likelihood directly.

In this section, for readability, we switch to boldface notation for finite-dimensional vectors and matrices. Hence the matrix Ψ in (2) will be denoted by the boldface Ψ .

5.1 Marginal likelihood and posterior distribution of parameter estimates

Denote $\mathbf{y} = (y_1, \dots, y_n)^\top$, $\mathbf{f} = (f(x_1), \dots, f(x_n))^\top$, $\mathbf{f}_0 = (f_0(x_1), \dots, f_0(x_n))^\top$, $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_n)^\top$, $\mathbf{w} = (w_1, \dots, w_n)^\top$. Then (1) implies $\mathbf{y} = \mathbf{f} + \boldsymbol{\varepsilon}$. Let \mathbf{H}_λ be the $n \times n$ matrix with (i, j) th coordinate $h_\lambda(x_i, x_j)$, where h_λ is an ANOVA reproducing kernel with scale parameter vector λ (see Section 3).

Under the I-prior, $\mathbf{f} \sim \text{MVN}(\mathbf{f}_0, \mathbf{H}_\lambda \Psi \mathbf{H}_\lambda)$, so the marginal distribution of \mathbf{y} is

$$\mathbf{y} \sim \text{MVN}(\mathbf{f}_0, \mathbf{V}_\mathbf{y}) \tag{20}$$

where the marginal covariance is given as

$$\mathbf{V}_\mathbf{y} = \mathbf{H}_\lambda \Psi \mathbf{H}_\lambda + \Psi^{-1} \tag{21}$$

Thus, the marginal log likelihood of (λ, Ψ) is

$$L(\lambda, \Psi | y) = -\frac{n}{2} \log(2\pi) - \frac{1}{2} \log |\mathbf{V}_\mathbf{y}| - \frac{1}{2} (\mathbf{y} - \mathbf{f}_0)^\top \mathbf{V}_\mathbf{y}^{-1} (\mathbf{y} - \mathbf{f}_0). \tag{22}$$

The maximum likelihood (ML) estimate $(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\Psi}})$ of $(\boldsymbol{\lambda}, \boldsymbol{\Psi})$ maximizes $L(\boldsymbol{\lambda}, \boldsymbol{\Psi}|\mathbf{y})$, and its asymptotic distribution can be found from the Fisher information. In particular, assume $\boldsymbol{\lambda} = \boldsymbol{\lambda}(\boldsymbol{\theta})$ and $\boldsymbol{\Psi} = \boldsymbol{\Psi}(\boldsymbol{\theta})$ are sufficiently smooth functions of a vector $\boldsymbol{\theta}$. Then straightforward calculations give the well-known result that the Fisher information matrix \mathbf{U} for $\boldsymbol{\theta}$ has (i, j) th coordinate

$$u_{ij} = \frac{1}{2} \text{tr} \left(\mathbf{V}_{\mathbf{y}}^{-1} \frac{\partial \mathbf{V}_{\mathbf{y}}}{\partial \theta_i} \mathbf{V}_{\mathbf{y}}^{-1} \frac{\partial \mathbf{V}_{\mathbf{y}}}{\partial \theta_j} \right),$$

where the derivatives are applied to each coordinate of the matrix. Now under well-known conditions, $\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$ has an asymptotic multivariate normal distribution with mean zero and covariance matrix \mathbf{U}^{-1} .

The next lemma gives the posterior distribution of f in (1) under the I-prior. A proof is provided in Bergsma (2019).

Lemma 1. *The posterior distribution of f in (1) subject to (2) given \mathbf{y} under the I-prior π is Gaussian with mean given by*

$$E_{\pi}[f(x)|\mathbf{y}] = f_0(x) + \sum_{i=1}^n h(x, x_i) \hat{w}_i$$

where

$$\tilde{\mathbf{w}} = \boldsymbol{\Psi} \mathbf{H}_{\boldsymbol{\lambda}} \mathbf{V}_{\mathbf{y}}^{-1} (\mathbf{y} - \mathbf{f}_0) \tag{23}$$

and covariance kernel given by

$$\text{cov}_{\pi}(f(x), f(x')|y_1, \dots, y_n) = \sum_{i=1}^n \sum_{j=1}^n h(x, x_i) h(x', x_j) (\mathbf{V}_{\mathbf{y}}^{-1})_{ij}$$

5.2 Estimation of parameters using the EM algorithm

We now describe the EM algorithm for estimating the scale parameter $\boldsymbol{\lambda}$ of $\mathbf{H}_{\boldsymbol{\lambda}}$, as well as parameters of the precision matrix $\boldsymbol{\Psi}$ (often it is just assumed the errors are iid $N(0, \psi^{-1})$, i.e., $\boldsymbol{\Psi} = \psi \mathbf{I}_n$). For estimating these parameters, EM turns out to be particularly efficient. There can be other unknown parameters as well, e.g., the Hurst coefficient for the FBM RKHS, but for this parameter EM is computationally much less attractive, and we will not go into this.

The EM algorithm has, as is well-known, guaranteed convergence under conditions often satisfied in practice. In the present case, the E-step is in closed form while the M-step is typically not, but in practice the M-step is computationally easy to carry out.

With g denoting the density function related to its argument and using (4), the complete

data log likelihood is

$$\begin{aligned}
L(\boldsymbol{\lambda}, \boldsymbol{\Psi} | \mathbf{y}, \mathbf{w}) &= \log g(\mathbf{y} | \mathbf{w}, \boldsymbol{\lambda}, \boldsymbol{\Psi}) + \log g(\mathbf{w} | \boldsymbol{\Psi}) \\
&= c + \frac{1}{2} \log |\boldsymbol{\Psi}| - \frac{1}{2} (\mathbf{y} - \mathbf{f})^\top \boldsymbol{\Psi} (\mathbf{y} - \mathbf{f}) - \frac{1}{2} \log |\boldsymbol{\Psi}| - \frac{1}{2} \mathbf{w}^\top \boldsymbol{\Psi}^{-1} \mathbf{w} \\
&= c - \frac{1}{2} (\mathbf{y} - \mathbf{f})^\top \boldsymbol{\Psi} (\mathbf{y} - \mathbf{f}) - \frac{1}{2} \mathbf{w}^\top \boldsymbol{\Psi}^{-1} \mathbf{w} \\
&= c - \frac{1}{2} (\mathbf{y} - \mathbf{f}_0 - \mathbf{H}_\lambda \mathbf{w})^\top \boldsymbol{\Psi} (\mathbf{y} - \mathbf{f}_0 - \mathbf{H}_\lambda \mathbf{w}) - \frac{1}{2} \mathbf{w}^\top \boldsymbol{\Psi}^{-1} \mathbf{w} \\
&= c - \frac{1}{2} (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} (\mathbf{y} - \mathbf{f}_0) - \frac{1}{2} \mathbf{w}^\top \mathbf{H}_\lambda^\top \boldsymbol{\Psi} \mathbf{H}_\lambda \mathbf{w} + (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} \mathbf{H}_\lambda \mathbf{w} - \frac{1}{2} \mathbf{w}^\top \boldsymbol{\Psi}^{-1} \mathbf{w} \\
&= c - \frac{1}{2} (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} (\mathbf{y} - \mathbf{f}_0) - \frac{1}{2} \mathbf{w}^\top \mathbf{V}_y \mathbf{w} + (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} \mathbf{H}_\lambda \mathbf{w} \\
&= c - \frac{1}{2} (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} (\mathbf{y} - \mathbf{f}_0) - \frac{1}{2} \text{tr} [\mathbf{V}_y \mathbf{w} \mathbf{w}^\top] + (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} \mathbf{H}_\lambda \mathbf{w},
\end{aligned}$$

where c is a constant. Write

$$\tilde{\mathbf{W}} = E(\mathbf{w} \mathbf{w}^\top | \mathbf{y}, \boldsymbol{\lambda}, \boldsymbol{\Psi}) = \tilde{\mathbf{V}}_w + \tilde{\mathbf{w}} \tilde{\mathbf{w}}^\top,$$

where $\tilde{\mathbf{w}}$ and $\tilde{\mathbf{V}}_w$ are given by (23) and (21). Let $\tilde{\mathbf{w}}^{(0)}$ and $\tilde{\mathbf{W}}^{(0)}$ be $\tilde{\mathbf{w}}$ and $\tilde{\mathbf{W}}$ with $\boldsymbol{\Psi}$ and $\boldsymbol{\lambda}$ replaced by $\boldsymbol{\Psi}^{(0)}$ and $\boldsymbol{\lambda}^{(0)}$. The E-step consists of computing

$$\begin{aligned}
Q(\boldsymbol{\lambda}, \boldsymbol{\Psi}) &= E \left\{ L(\boldsymbol{\lambda}, \boldsymbol{\Psi} | \mathbf{y}, \mathbf{w}) | \mathbf{y}, \boldsymbol{\lambda}^{(0)}, \boldsymbol{\Psi}^{(0)} \right\} \\
&= c - \frac{1}{2} (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} (\mathbf{y} - \mathbf{f}_0) - \frac{1}{2} \text{tr} [\mathbf{V}_y \tilde{\mathbf{W}}^{(0)}] + (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} \mathbf{H}_\lambda \tilde{\mathbf{w}}^{(0)}. \quad (24)
\end{aligned}$$

The M-step entails maximizing $Q(\boldsymbol{\lambda}, \boldsymbol{\Psi})$. We assume the global maximum can be found by differentiating, equating to zero, and solving. Supposing $\boldsymbol{\Psi}$ but not \mathbf{H}_λ depends on a parameter ψ and \mathbf{H}_λ but not $\boldsymbol{\Psi}$ depends on a parameter λ , the derivatives are given by

$$\begin{aligned}
\frac{\partial Q(\boldsymbol{\lambda}, \boldsymbol{\Psi})}{\partial \lambda} &= -\text{tr} \left[\frac{\partial \mathbf{H}_\lambda}{\partial \lambda} \boldsymbol{\Psi} \mathbf{H}_\lambda \tilde{\mathbf{W}}^{(0)} \right] + (\mathbf{y} - \mathbf{f}_0)^\top \boldsymbol{\Psi} \frac{\partial \mathbf{H}_\lambda}{\partial \lambda} \tilde{\mathbf{w}}^{(0)} \\
\frac{\partial Q(\boldsymbol{\lambda}, \boldsymbol{\Psi})}{\partial \psi} &= -\frac{1}{2} (\mathbf{y} - \mathbf{f}_0)^\top \frac{\partial \boldsymbol{\Psi}}{\partial \psi} (\mathbf{y} - \mathbf{f}_0) - \frac{1}{2} \text{tr} \left[\frac{\partial \mathbf{V}_y}{\partial \psi} \tilde{\mathbf{W}}^{(0)} \right] + (\mathbf{y} - \mathbf{f}_0)^\top \frac{\partial \boldsymbol{\Psi}}{\partial \psi} \mathbf{H}_\lambda \tilde{\mathbf{w}}^{(0)}.
\end{aligned}$$

For the examples in Section 6, the errors are iid, so $\boldsymbol{\Psi} = \psi \mathbf{I}_n$ for a scalar ψ , and \mathbf{H}_λ represents an ANOVA kernel as described in Section 3, so it is of the form $\mathbf{H}_\lambda = \sum_{s=1}^k g_s(\lambda) \mathbf{H}_s$, where g_s is a polynomial function. Then

$$\frac{\partial \boldsymbol{\Psi}}{\partial \psi} = \mathbf{I}_n \quad \frac{\partial \mathbf{H}_\lambda}{\partial \lambda} = \sum_{s=1}^k \frac{\partial g_s(\lambda)}{\partial \lambda} \mathbf{H}_s$$

The partial derivatives of Q set to zero can then normally be solved very quickly numerically for the purposes of this paper. For example, if $g_s(\lambda) = \lambda_s$ (which is the case, e.g., if there are no interactions in the model) then the equations have a closed form solution. In general, with iid errors, and \mathcal{V} the set of (possibly multidimensional) variables involved, $2|\mathcal{V}| + 1$ polynomial equations in $2|\mathcal{V}| + 1$ unknowns need to be solved, which we found can be done very rapidly using built in solvers in R and Mathematica. The computational bottleneck is not the M step, but the E step which is $O(n^3)$.

We did sometimes encounter accuracy problems, making it impossible to obtain full convergence of the EM algorithm (this problem was encountered and mentioned in Section 6.4). This is not a problem just for EM, as the marginal likelihood is difficult to estimate as well, see Bergsma (2019) for a visualization of this problem. For prediction purposes, the lack of convergence did not seem to matter, but the lack of an accurate maximal value of the marginal likelihood can make model comparison difficult.

5.3 Asymptotic theory

Let \hat{f} have the posterior distribution of f given \mathbf{y} under the I-prior. Let $\mathbf{Q} = \lambda^2 \mathbf{H} \Psi \mathbf{H} (\lambda^2 \mathbf{H} \Psi \mathbf{H} + \mathbf{I})^{-1}$. We have $\hat{\mathbf{f}} = \mathbf{H} \hat{\mathbf{w}}$ where $\hat{\mathbf{w}} = \lambda \Psi \mathbf{H} \mathbf{V}_{\mathbf{y}}^{-1} \mathbf{y} + \boldsymbol{\xi}$, where $\boldsymbol{\xi} \sim N(\mathbf{0}, \mathbf{V}_{\mathbf{y}}^{-1})$ is independent of $\mathbf{y} \sim \text{MVN}(\mathbf{0}, \mathbf{V}_{\mathbf{y}})$. The next lemma gives sufficient conditions on λ as $n \rightarrow \infty$ for the posterior of f to contract to the true value.

Lemma 2. Suppose $n\lambda \rightarrow \infty$ and $\lambda \rightarrow 0$ as $n \rightarrow \infty$. Then $E\|\hat{f} - f\|_{\mathcal{F}}^2 \rightarrow 0$ as $n \rightarrow \infty$.

Proof. We have $\mathbf{w} - \hat{\mathbf{w}} = \mathbf{w} - \lambda \Psi \mathbf{H} \mathbf{V}_y^{-1} \mathbf{y} + \boldsymbol{\xi} = (\mathbf{I} - \lambda^2 \mathbf{H} \Psi \mathbf{H} \mathbf{V}_y^{-1}) \mathbf{w} + \lambda \Psi \mathbf{H} \mathbf{V}_y^{-1} \boldsymbol{\varepsilon} + \boldsymbol{\xi}$. Set $f = \sum_{i=1}^n h(\cdot, x_i) w_i + r$ where $r \in \mathcal{F}_n^\perp$. Then

$$\begin{aligned}
E\|\hat{f} - f\|_{\mathcal{F}}^2 &= E \left\| \sum_{i=1}^n h(\cdot, x_i) (\hat{w}_i - w_i) - r \right\|_{\mathcal{F}}^2 \\
&= E \left\| \sum_{i=1}^n h(\cdot, x_i) (\hat{w}_i - w_i) \right\|_{\mathcal{F}}^2 + \|r\|_{\mathcal{F}}^2 \\
&= E \sum_{i=1}^n \sum_{j=1}^n h(x_i, x_j) (\hat{w}_i - w_i) (\hat{w}_j - w_j) + \|r\|_{\mathcal{F}}^2 \\
&= E(\lambda \hat{\mathbf{w}} - \mathbf{w})^\top \mathbf{H} (\lambda \hat{\mathbf{w}} - \mathbf{w}) + E\|r\|_{\mathcal{F}}^2 \\
&= E \left\| \lambda^2 \Psi \mathbf{H} \mathbf{V}_y^{-1} (\mathbf{H} \mathbf{w} + \boldsymbol{\varepsilon}) + \lambda \boldsymbol{\xi} - \mathbf{w} \right\|_{\mathbf{H}}^2 + \|r\|_{\mathcal{F}}^2 \\
&= E \left\| (\lambda^2 \Psi \mathbf{H} \mathbf{V}_y^{-1} \mathbf{H} - \mathbf{I}) \mathbf{w} + \lambda^2 \Psi \mathbf{H} \mathbf{V}_y^{-1} \boldsymbol{\varepsilon} + \lambda \boldsymbol{\xi} \right\|_{\mathbf{H}}^2 + \|r\|_{\mathcal{F}}^2 \\
&= E \left\| -\Psi^{-1} \mathbf{V}_y^{-1} \mathbf{w} + \lambda^2 \Psi \mathbf{H} \mathbf{V}_y^{-1} \boldsymbol{\varepsilon} + \lambda \boldsymbol{\xi} \right\|_{\mathbf{H}}^2 + E\|r\|_{\mathcal{F}}^2 \\
&= \left\| \Psi^{-1} \mathbf{V}_y^{-1} \mathbf{w} \right\|_{\mathbf{H}}^2 + \lambda^4 E \left\| \Psi \mathbf{H} \mathbf{V}_y^{-1} \boldsymbol{\varepsilon} \right\|_{\mathbf{H}}^2 + \lambda^2 E \left\| \boldsymbol{\xi} \right\|_{\mathbf{H}}^2 + \|r\|_{\mathcal{F}}^2 \\
&= \mathbf{w}^\top \mathbf{V}_y^{-1} \Psi^{-1} \mathbf{H} \Psi^{-1} \mathbf{V}_y^{-1} \mathbf{w} + \lambda^4 \text{tr}(\mathbf{H} \Psi \mathbf{H} \mathbf{V}_y^{-1} \Psi^{-1} \mathbf{V}_y^{-1} \mathbf{H} \Psi) + \lambda^2 \text{tr}(\mathbf{H} \mathbf{V}_y^{-1}) + \|r\|_{\mathcal{F}}^2 \\
&= \mathbf{w}^\top \mathbf{V}_y^{-1} \Psi^{-1} \mathbf{H} \Psi^{-1} \mathbf{V}_y^{-1} \mathbf{w} + \lambda^2 \text{tr}((\mathbf{I} - \Psi^{-1} \mathbf{V}_y^{-1}) \Psi^{-1} \mathbf{V}_y^{-1} \mathbf{H} \Psi) + \lambda^2 \text{tr}(\mathbf{H} \mathbf{V}_y^{-1}) + \|r\|_{\mathcal{F}}^2 \\
&= \mathbf{w}^\top \mathbf{V}_y^{-1} \Psi^{-1} \mathbf{H} \Psi^{-1} \mathbf{V}_y^{-1} \mathbf{w} - \lambda^2 \text{tr}(\mathbf{V}_y^{-1} \Psi^{-1} \mathbf{V}_y^{-1} \mathbf{H}) + 2\lambda^2 \text{tr}(\mathbf{H} \mathbf{V}_y^{-1}) + \|r\|_{\mathcal{F}}^2 \\
&= \sum_{k=1}^n \frac{na_k (\mathbf{g}_k^\top \mathbf{w})^2}{(n^2 \lambda^2 a_k^2 + 1)^2} - \sum_{k=1}^n \frac{n\lambda^2 a_k}{(n^2 \lambda^2 a_k^2 + 1)^2} + 2 \sum_{k=1}^n \frac{n\lambda^2 a_k}{n^2 \lambda^2 a_k^2 + 1} + \|r\|_{\mathcal{F}}^2 \\
&= \sum_{k=1}^n \frac{na_k (\mathbf{g}_k^\top \mathbf{w})^2}{(n^2 \lambda^2 a_k^2 + 1)^2} + \sum_{k=1}^n \frac{n\lambda^2 a_k}{n^2 \lambda^2 a_k^2 + 1} + \sum_{k=1}^n \frac{n\lambda^2 a_k}{(n^2 \lambda^2 a_k^2 + 1)^2} + \|r\|_{\mathcal{F}}^2 \\
&=: A + B + C + \|r\|_{\mathcal{F}}^2
\end{aligned}$$

By definition, $\|r\|_{\mathcal{F}}^2 \rightarrow 0$ as $n \rightarrow \infty$. Also, $C \leq B$. Now

$$B = \sum_{k=1}^{\lfloor n\lambda \rfloor} \frac{n\lambda^2 a_k}{n^2 \lambda^2 a_k^2 + 1} + \sum_{k=\lfloor n\lambda \rfloor}^n \frac{n\lambda^2 a_k}{n^2 \lambda^2 a_k^2 + 1} \leq \lambda + 2n\lambda^2 \sum_{k=\lfloor n\lambda \rfloor}^n a_k \rightarrow 0$$

$$A = \sum_{k=1}^{\lfloor (n\lambda)^q \rfloor} \frac{na_k (\mathbf{g}_k^\top \mathbf{w})^2}{(n^2 \lambda^2 a_k^2 + 1)^2} + \sum_{k=\lfloor (n\lambda)^q \rfloor}^n \frac{na_k (\mathbf{g}_k^\top \mathbf{w})^2}{(n^2 \lambda^2 a_k^2 + 1)^2}$$

Taking q sufficiently small, the first term converges to zero if $n\lambda \rightarrow \infty$. The second term

Also show $B \rightarrow \infty$ if $\lambda \not\rightarrow 0$ and $A \rightarrow a > 0$ if $n\lambda \rightarrow \ell < \infty$. \square

Lemma 3. With probability 1, $n\hat{\lambda} \rightarrow \infty$ and $\hat{\lambda} \rightarrow \infty$ as $n \rightarrow \infty$.

Proof. The marginal distribution is $\mathbf{y} \sim \text{MVN}(\mathbf{0}, \mathbf{V}_y)$ where $\mathbf{V}_y = \lambda^2 \mathbf{H} \Psi \mathbf{H} + \Psi^{-1}$ with marginal log-likelihood

$$L(\nu | \mathbf{y}) = -\frac{1}{2} \log |\mathbf{V}_y| - \frac{1}{2} \mathbf{y}^\top \mathbf{V}_y^{-1} \mathbf{y}$$

The derivative is

$$\begin{aligned}
\frac{dL(\nu|\mathbf{y})}{d\lambda^2} &= -\frac{1}{2} \text{tr}(\mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H}) + \frac{1}{2} \mathbf{y}^\top \mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H} \mathbf{V}_y^{-1} \mathbf{y} \\
&= -\frac{1}{2} \text{tr}(\mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H}) + \frac{1}{2} (\mathbf{H} \mathbf{w} + \boldsymbol{\varepsilon})^\top \mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H} \mathbf{V}_y^{-1} (\mathbf{H} \mathbf{w} + \boldsymbol{\varepsilon}) \\
&= -\frac{1}{2} \text{tr}(\mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H}) + \frac{1}{2} \mathbf{w}^\top \mathbf{H} \mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H} \mathbf{V}_y^{-1} \mathbf{H} \mathbf{w} + \mathbf{w}^\top \mathbf{H} \mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H} \mathbf{V}_y^{-1} \boldsymbol{\varepsilon} + \boldsymbol{\varepsilon}^\top \mathbf{V}_y^{-1} \mathbf{H} \Psi \mathbf{H} \mathbf{V}_y^{-1} \boldsymbol{\varepsilon} \\
&= -\frac{1}{2} \sum_{k=1}^n \frac{n^2 a_k^2}{n^2 a_k^2 \lambda^2 + 1} + \frac{1}{2} \sum_{k=1}^n \frac{n^4 a_k^4 (\mathbf{g}_k^\top \mathbf{w})^2}{(n^2 a_k^2 \lambda^2 + 1)^2} + \sum_{k=1}^n \frac{n^3 a_k^3 (\mathbf{g}_k^\top \mathbf{w})(\mathbf{g}_k^\top \boldsymbol{\varepsilon})}{(n^2 a_k^2 \lambda^2 + 1)^2} + \frac{1}{2} \sum_{k=1}^n \frac{n^2 a_k^2 (\mathbf{g}_k^\top \boldsymbol{\varepsilon})^2}{(n^2 a_k^2 \lambda^2 + 1)^2}
\end{aligned}$$

If $n\lambda \rightarrow c < \infty$, or if $\lambda \rightarrow c' > 0$, then this diverges from zero with probability 1... \square

5.4 Significance tests

5.4.1 HSIC

Consider the model

$$y_i = \alpha + f_1(x_i) + \varepsilon_i \quad f_1 \in \mathcal{F}_1, x_i \in \mathcal{X}, i = 1, \dots, n,$$

where \mathcal{F}_1 is an RKHS over \mathcal{X}_1 with centered kernel h_1 , i.e., $\sum_{i=1}^n h_1(x, x_i) = 0$ for all $x \in \mathcal{X}$. Consider the null and alternative hypotheses

$$H_0 : f_1 = 0 \quad H_1 : f_1 \neq 0$$

The *Hilbert-Schmidt independence criterion* (HSIC) was introduced by Gretton, Bousquet, Smola, and Schölkopf (2005). With centered kernel matrices $\mathbf{H}_1 = (h_1(x_i, x_j))$ and $\mathbf{H}_2 = (h_2(y_i, y_j))$, $\text{HSIC} = \sum_{i=1}^n \sum_{j=1}^n h_1(x_i, x_j) h_2(y_i, y_j)$. Taking $h_2(y, y') = (y - \bar{y})(y' - \bar{y})$, we obtain

$$\text{HSIC} = \sum_{i=1}^n \sum_{j=1}^n (y_i - \bar{y})(y_j - \bar{y}) h_1(x_i, x_j)$$

This quantity can be used as a test statistic for H_0 . The permutation p -value based on HSIC is defined as the proportion of permutations of y_1, \dots, y_n which give test statistic value larger than T^2 . Clearly under H_0 all permutations have equal probability so the test controls type I error.

To gain some insight into the test statistic, suppose \mathbf{H} has decomposition

$$h(x_i, x_j) = \sum_{k=1}^n \alpha_k g_k(x_i) g_k(x_j)$$

for an orthonormal system in the RKHS with kernel Ψ^{-1} , i.e., $\sum_{i=1}^n \sum_{j=1}^n \psi_{ij} a_k(x_i) a_l(x_j) = \delta_{kl}$ and $\sum_{i=1}^n g_k(x_i) = 0$. Then

$$T^2 = n^2 \sum_{k=1}^n \alpha_k \widehat{\text{cov}}(g_k(X), Y)^2$$

where

$$\widehat{\text{cov}}(g_k(X), Y) = \frac{1}{n} \sum_{i=1}^n g_k(x_i) (y_i - \bar{y})$$

Thus, the modified score test statistic downweights covariances which the data have less information on, in particular, the Fisher information (under H_1) on $\text{cov}(g_k(X), Y)$ equals α_k^2 . In contrast, the usual score statistic is

$$S^2 = \sum_{k=1}^q \widehat{\text{cov}}(g_k(X), Y)^2$$

where q is the rank of \mathbf{H}_1 , which gives equal weight to all covariances. Note that

$$S^2 \sim \sum_{k=1}^q Z_k^2 \quad T^2 \sim \sum_{i=1}^n \alpha_k Z^2$$

where Z_1, Z_2, \dots is a sequence of independent standard normal random variables. Davies (1980) shows how to evaluate p -values for mixture chi-squares. However, the permutation procedure gives exact conditional p -values, conditionally on the sufficient statistics which are the observed marginal distributions of X and Y .

5.4.2 Likelihood ratio test

With $\varepsilon_1, \dots, \varepsilon_n$ independent $N(0, \psi^{-1})$ random variables, consider the model

$$y_i = f_1(x_{1i}) + f_2(x_{2i}) + \varepsilon_i \quad f_k \in \mathcal{F}_k, x_{ki} \in \mathcal{X}_k, k = 1, 2, i = 1, \dots, n$$

where \mathcal{F}_k is the RKHS with reproducing kernel h_k ($k = 1, 2$). Consider the null and alternative hypotheses

$$H_0 : f_1 = 0 \quad H_1 : f_1 \neq 0$$

Let

$$\mathbf{Q}_0 = \tau_0^2 \mathbf{H}_1^2 (\tau_0^2 \mathbf{H}_1^2 + \mathbf{I})^{-1} \quad \mathbf{Q}_1 = \tau_1^2 (\mathbf{H}_1 + \mathbf{H}_2)^2 (\tau_1^2 (\mathbf{H}_1 + \mathbf{H}_2)^2 + \mathbf{I})^{-1}$$

Then, using the I-prior, the posterior means of \mathbf{f}_2 and $\mathbf{f}_1 + \mathbf{f}_2$ under H_0 and H_1 , respectively, are given as

$$\tilde{\mathbf{f}}_2 = \mathbf{Q}_0 \mathbf{y} \quad \hat{\mathbf{f}}_1 + \hat{\mathbf{f}}_2 = \mathbf{Q}_1 \mathbf{y}$$

The estimated error precisions are then

$$\hat{\psi}_0 = \|\mathbf{y} - \tilde{\mathbf{f}}_2\|_2^{-2} \quad \hat{\psi}_1 = \|\mathbf{y} - \hat{\mathbf{f}}_1 - \hat{\mathbf{f}}_2\|_2^{-2}$$

and the likelihood ratio statistic is

$$\text{LR} = \left(\frac{\hat{\psi}_1}{\hat{\psi}_0} \right)^{n/2} = \left(\frac{\mathbf{y}^\top (\mathbf{I} - \mathbf{Q}_0)^2 \mathbf{y}}{\mathbf{y}^\top (\mathbf{I} - \mathbf{Q}_1)^2 \mathbf{y}} \right)^{n/2}$$

The distribution of LR under H_0 depends on the unknown \mathbf{f}_2 , so it is unclear how to compute p -values. The next lemma shows that asymptotically, the dependence on \mathbf{f}_2 disappears.

Lemma 4. *Suppose, for $\ell = 0, 1$, $\tau_\ell n \rightarrow \infty$ and $\tau_\ell \rightarrow 0$ as $n \rightarrow \infty$, and h_ℓ has the Mercer expansion*

$$h_\ell(x, x') = \sum_{i=1}^{\infty} \alpha_{\ell, i} g_{\ell, k}(x) g_{\ell, k}(x')$$

where $\alpha_{\ell,i} \geq 0$, and $\sum_{i=1}^{\infty} \alpha_{\ell,i} < \infty$. Then

$$\sup_{c>0} \left| P(\text{LR} > c) - P \left[\left(\frac{\boldsymbol{\varepsilon}^\top (\mathbf{I} - \mathbf{Q}_0)^2 \boldsymbol{\varepsilon}}{\boldsymbol{\varepsilon}^\top (\mathbf{I} - \mathbf{Q}_1)^2 \boldsymbol{\varepsilon}} \right)^{n/2} > c \right] \right| \rightarrow 0$$

as $n \rightarrow \infty$.

The lemma follows from the next lemma and its proof.

Lemma 5. Let $y_i = f(x_i) + \varepsilon_i$, where f is in the RKHS \mathcal{F} over \mathcal{X} with reproducing kernel h , and assume it has a Mercer expansion

$$h(x, x') = \sum_{i=1}^{\infty} \alpha_i g_k(x) g_k(x')$$

where $\alpha_i \geq 0$, and $\sum_{i=1}^{\infty} \alpha_i < \infty$. Let $\mathbf{Q} = \tau^2 \mathbf{H}^2 (\tau^2 \mathbf{H}^2 + \mathbf{I})^{-1}$. Suppose $\tau n \rightarrow \infty$ and $\tau \rightarrow 0$ as $n \rightarrow \infty$. Then

$$\sup_{c>0} \left| P(\mathbf{y}^\top (\mathbf{I} - \mathbf{Q})^2 \mathbf{y} > c) - P(\boldsymbol{\varepsilon}^\top (\mathbf{I} - \mathbf{Q})^2 \boldsymbol{\varepsilon} > c) \right| \rightarrow 0$$

as $n \rightarrow \infty$.

Proof. First note that we can write $f(x) = \sum h(x, x_i) w_i + r(x) \equiv f_n(x) + r(x)$, where $f_n \perp_{\mathcal{F}} r$. Pythagoras' theorem then gives

$$\mathbf{w}^\top \mathbf{H} \mathbf{w} = \|f_n\|_{\mathcal{F}}^2 \leq \|f_n\|_{\mathcal{F}}^2 + \|r\|_{\mathcal{F}}^2 = \|f\|_{\mathcal{F}}^2$$

Writing $\mathbf{H} = n \sum_{k=1}^n a_k \mathbf{g}_k \mathbf{g}_k^\top$ where $a_1 \geq \dots \geq a_n$ and $\mathbf{g}_1, \dots, \mathbf{g}_n$ is an orthonormal system, we have

$$\mathbf{I} - \mathbf{Q} = \sum_{k=1}^n \frac{1}{n^2 \tau^2 a_k^2 + 1} \mathbf{g}_k \mathbf{g}_k^\top$$

Then with $\mathbf{f} = \mathbf{H} \mathbf{w}$ and $\mathbf{v} = (\mathbf{w}^\top \mathbf{H} \mathbf{w})^{-1/2} \mathbf{H}^{1/2} \mathbf{w}$ so that $\mathbf{v}^\top \mathbf{v} = 1$,

$$\begin{aligned} \mathbf{f}^\top (\mathbf{I} - \mathbf{Q})^2 \mathbf{f} &= \mathbf{w}^\top \mathbf{H} \mathbf{w} \mathbf{v}^\top \mathbf{H}^{1/2} (\mathbf{I} - \mathbf{Q})^2 \mathbf{H}^{1/2} \mathbf{v} \\ &\leq \mathbf{w}^\top \mathbf{H} \mathbf{w} \max_{\mathbf{s}: \mathbf{s}^\top \mathbf{s} = 1} \mathbf{s}^\top \mathbf{H}^{1/2} (\mathbf{I} - \mathbf{Q})^2 \mathbf{H}^{1/2} \mathbf{s} \\ &\leq \|f\|_{\mathcal{F}}^2 \max_{\mathbf{s}: \mathbf{s}^\top \mathbf{s} = 1} \sum_{k=1}^n \frac{a_k}{(n^2 \tau^2 a_k^2 + 1)^2} (\mathbf{s}^\top \mathbf{g}_k)^2 \leq \frac{\|f\|_{\mathcal{F}}^2}{\sqrt{3} n^2 \tau^2} \end{aligned}$$

Furthermore,

$$\boldsymbol{\varepsilon}^\top (\mathbf{I} - \mathbf{Q})^2 \boldsymbol{\varepsilon} = \text{tr}((\mathbf{I} - \mathbf{Q})^2) = \sum_{k=1}^n \left(\frac{\boldsymbol{\varepsilon}^\top \mathbf{g}_k}{n^2 \tau^2 a_k^2 + 1} \right)^2$$

Then with $a_k \asymp k^{-1-\eta}$,

$$\begin{aligned}
E\boldsymbol{\varepsilon}^\top(\mathbf{I} - \mathbf{Q})^2\boldsymbol{\varepsilon} &= \sum_{k=1}^n \left(\frac{1}{n^2\tau^2 a_k^2 + 1} \right)^2 \\
&= \sum_{k=1}^{\lfloor \tau n \rfloor} \left(\frac{1}{n^2\tau^2 a_k^2 + 1} \right)^2 + \sum_{k=\lceil \tau n \rceil}^n \left(\frac{1}{n^2\tau^2 a_k^2 + 1} \right)^2 \\
&\geq \lfloor \tau n \rfloor \left(\frac{1}{n^2\tau^2 a_1^2 + 1} \right)^2 + (n - \lceil \tau n \rceil) \left(\frac{1}{n^2\tau^2 a_{\lceil \tau n \rceil}^2 + 1} \right)^2 \\
&= O(1/(\tau^3 n^3)) + (n - \lceil \tau n \rceil) \left(\frac{1}{n^2(\tau n)^{-2\eta} + 1} \right)^2 \\
&= O(n)
\end{aligned}$$

Since $\mathbf{y} = \mathbf{f} + \boldsymbol{\varepsilon}$ the triangle inequality gives

$$\sqrt{\boldsymbol{\varepsilon}^\top(\mathbf{I} - \mathbf{Q})^2\boldsymbol{\varepsilon}} \leq \sqrt{\mathbf{y}^\top(\mathbf{I} - \mathbf{Q})^2\mathbf{y}} \leq \sqrt{\mathbf{f}^\top(\mathbf{I} - \mathbf{Q})^2\mathbf{f}} + \sqrt{\boldsymbol{\varepsilon}^\top(\mathbf{I} - \mathbf{Q})^2\boldsymbol{\varepsilon}}$$

Hence

$$|P(\mathbf{y}^\top(\mathbf{I} - \mathbf{Q})^2\mathbf{y} > c) - P(\boldsymbol{\varepsilon}^\top(\mathbf{I} - \mathbf{Q})^2\boldsymbol{\varepsilon} > c)| \leq \dots = o(1)$$

□

6 Application to data

We reanalyze some well-known data sets in the respective areas of multilevel modelling, functional data analysis, classification, and longitudinal data analysis. Whereas in the literature different methods are typically used in different areas, and often more than method per area, we fit all models using the single method introduced in this paper. In all cases we obtain a performance competitive with existing techniques in terms of mean squared prediction error for test data. Furthermore, our methodology is flexible in that it is easy to incorporate extra covariates and interaction effects using the parsimonious ANOVA framework given in Section 3.

This section is organized as follows. In Section 6.3 we fit the standard varying intercept and varying slope models using the I-prior. The purpose of this section is mainly to illustrate the differences between the I-prior approach and the standard random effects approach. Here, the I-prior method has an estimation advantage in that there are no positive definiteness restriction on a latent covariance matrix. In Section 6.4 we look at multi-class classification, both with high and low dimensional covariates. The latter two sections illustrate the potentially good predictive performance of the I-prior methodology compared to other methods. Furthermore, they illustrate the ease with which high-dimensional smoothing can be done using the I-prior. In Section 6.5 we do a longitudinal data analysis with I-priors. Here, we treat the longitudinal response curves as ‘functional’ responses. In contrast to standard approaches, we do not need to specify a covariance structure for the longitudinal responses, instead we merely need to specify an appropriate class of functions, e.g., a class of smooth functions.

Throughout this section, we will use the parsimonious ANOVA approach described in Section 3.2 for which interaction effects do not require extra parameters. In Section 6.5, we will compare this approach with the ‘classical’ non-parsimonious extended ANOVA one (also under the I-prior) where each interaction effect does take an extra parameter.




Figure 1: Sample of spectrometric curves used to predict fat content of meat

6.1 Motivation for use of FBM RKHS

As explained in more detail in Bergsma (2019), the use of the I-prior methodology is particularly attractive if \mathcal{F} is a fractional Brownian motion (FBM) RKHS over a Euclidean space (which has as a special case the aforementioned centered Brownian motion). FBM *process paths* are non-differentiable and, having Hölder smoothness ranging between 0 and 1, an FBM process prior for the regression function may be too rough for many applications. In contrast, functions in the FBM *RKHS* with Hurst coefficient γ are (weakly) differentiable if the Hurst $\gamma \geq 1/2$ and have minimum Hölder smoothness 2γ . This wide range of smoothnesses make it an attractive general purpose function space for nonparametric regression. Another advantage is that it allows us to do multivariate smoothing with just one or two parameters to be estimated: either only the scale parameter λ , while using a default setting of, say, $1/2$ for the Hurst coefficient, or both the scale parameter and the Hurst coefficient. This is in contrast with standard kernel based smoothing methods, which require a scale parameter and at least one kernel hyperparameter to be estimated. For example, if we use the exponential kernel

$$r(x, x') = \exp\left(-\frac{\|x - x'\|^{2\xi}}{2\sigma^2}\right), \quad (25)$$

the scale parameter λ , the smoothness parameter ξ (somewhat analogous to the Hurst coefficient), and a ‘variance’ parameter σ^2 need to be estimated. Default settings $\xi = 1$ or $\xi = 2$ could be used to reduce the number of free parameters to two. Furthermore, the functions in the squared exponential RKHS are analytical, which is too smooth for many applications.

6.2 Regression with a functional covariate

We illustrate the prediction of a real valued response when one of the covariates is a function using a widely analysed data set used for quality control in the food industry. The data consist of measurements on a sample of 215 pieces of finely chopped meat. The response variable is fat content, and the covariate is light absorbance for 100 different wavelengths. The absorbance curve can be considered a ‘functional’ variable (see a sample of such curves plotted in Figure 1). For more details see <http://lib.stat.cmu.edu/datasets/tecator> and Thodberg (1996). Our aim is to predict fat content from the 100 measurements of absorbance. The first 172 observations in the data set are used as a training sample, and the remaining 43 observations are used as a test sample (following Thodberg’s original recommendation).

Many different methods have been applied in the literature to the data set, estimating a model using the training sample and evaluating its performance using the test sample. One of the best results was achieved early on by Thodberg (1996), who used neural networks on the first 10 principal components and achieved a test mean squared error of 0.36. The best test

Method	RMSE	
	Training	Test
Global constant model	12.50	13.3
Neural network (Vila et al., 2000)		0.34
Kernel smoothing (Ferraty & Vieu, 2006, Section 7.2)		1.85
Double index model (Chen, Hall, & Müller, 2011)		1.58
Single index model (Goia & Vieu, 2014)		1.18
Sliced inverse regression (Lian & Li, 2014)		0.90
MARS (Zhu, Yao, & Zhang, 2014)		0.88
Partial least squares (Zhu et al., 2014)		1.01
CSEFAM (Zhu et al., 2014)		0.85
Tikhonov regularization (linear)	3.32	3.54
Tikhonov regularization (FBM-1/2 kernel)	4.32	4.54
I-prior (linear)	2.82	3.15
I-prior (FBM RKHS with $\gamma = 0.5$)	0.00	0.67
I-prior (FBM RKHS with $\hat{\gamma} = 0.98$)	0.00	0.57
I-prior (squared exponential RKHS, $\hat{\sigma} = 0.0079$)	0.35	0.58

Table 4: RMSEs for predicting fat content from spectrometric functional covariate (see Figure 1): previously published results, Tikhonov regularization, and I-prior methodology.

error performance we found was by Vila, Wagner, and Neveu (2000) who achieved an error rate of 0.34, also using neural networks on the principal components. More recently various other statistical models have been tried on the data set, see Table 4 for a summary. In spite of their lesser performance compared to neural networks, the interest of these methods is that they do not rely on an a priori data reduction in terms of the main principal components.

The i th spectral curve is denoted x_i , with $x_i(t)$ denoting the absorbance for wavelength t . We assume $x_i \in \mathcal{X}$, where \mathcal{X} is a set of functions over \mathbb{R} and is equipped with an appropriate inner product. From Figure 1 it appears the curves are differentiable, so it seems reasonable to assume the \mathcal{X} is the Brownian motion RKHS over \mathbb{R} with squared norm

$$\|x\|_{\mathcal{X}}^2 = \int_{\mathbb{R}} \dot{x}(t)^2 dt.$$

Note that the L^2 inner product for a set of functions will rarely be appropriate.

Since \mathcal{X} is a Hilbert space, a linear effect of the spectral curve on fat content can now be modelled using the canonical RKHS over \mathcal{X} . We see in Table 4 that both Tikhonov regularization and the I-prior give a poor performance, with test RMSEs of 3.54 and 2.89, respectively. Next we fitted a smooth dependence of fat content on spectrometric curve using the FBM RKHS over \mathcal{X} . As seen in the table, Tikhonov regularization performs very poorly. We tried various values of the Hurst coefficient, but all give worse results than the linear model. On the other hand, the I-prior performs rather well for different RKHSs, including the FBM and the squared exponential ones. We had some convergence problems so could not get the ML estimator of γ , the Hurst coefficient for the FBM RKHS, so instead estimated it by minimizing the cross-validation error (10-fold cross-validation gave $\hat{\gamma} = 0.98$). For the squared exponential RKHS we did manage to find the ML estimator $\hat{\sigma}$ of σ , and it is given in Table 4.

Instead of fat content, protein content can be predicted from the spectral curve. With the I-prior based on a smooth dependence of protein content on the spectral curve we obtained an RMSE of 0.52, using a local (non-global) maximum likelihood estimate of the Hurst coefficient, $\hat{\gamma} = 0.997$. This improves on Zhu et al. (2014) who obtained an RMSE of 0.85.

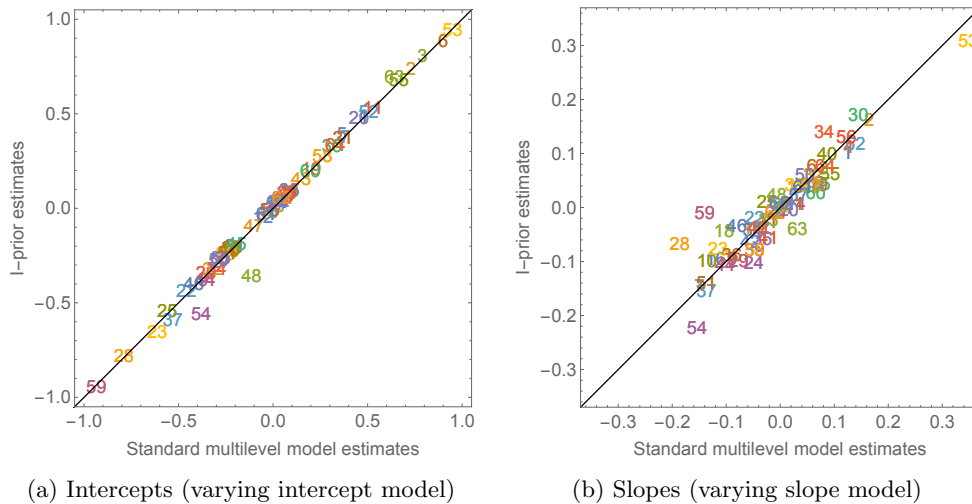


Figure 2: Estimated intercepts and slopes for school achievement data under varying intercept and varying slope model. The numbers are the school indices. In contrast to the I-prior, the standard random effects multilevel model assumes iid slopes and intercepts across schools. It is seen that I-prior estimated intercepts are bigger in absolute value than the standard ones for small schools (48 and 54; see text).

6.3 Multilevel models

The purpose of this section is to illustrate the differences between the standard random effects and the I-prior approaches for estimating varying intercept and varying slope models. We consider a data set which accompanies the MLwiN software (Rasbash, Steele, Browne, & Goldstein, 2012) on school achievement of 4059 pupils at 65 inner-London schools. The response variable is the GCSE score at age 16.

First we consider the varying intercept model (13). The ‘covariate’ is the nominal variable school (ranging from 1 to 65), which we assume to have a nominal effect on GCSE score, i.e., we assume that the regression function f in (13) is in the Pearson RKHS. A standard approach to fitting this model is the random intercept model, which is based on the assumption that the intercepts are iid normal with zero mean. In Figure 2(a), the posterior means of the intercepts are plotted for the standard (random effects) model and the I-prior. It can be seen the estimates are in broad agreement, with conspicuously different estimates for schools 48 (-0.11 vs. -0.36) and 54 (-0.38 vs. -0.56), the I-prior giving the largest estimate in absolute value in both cases. The reason for the relatively large I-prior estimates can be found from the prior variance formula ... which implies that the smaller the sample size for a particular school, the larger the prior variance of the intercept for that school. Indeed, schools 48 and 54 have the smallest sample sizes of all schools, namely 2 and 8 (the next smallest school is number 37, with 22 students).

Next we consider the varying slope model (16), which regresses, for each school, the GCSE score on the result of the London reading test (LRT), taken at age 11. A standard approach to fitting this model is the random intercept/slope model, which is based on the assumption that the intercept/slope pairs are iid bivariate normal with zero mean. To obtain an I-prior, we assume as above a nominal effect of school, and a linear effect of LRT (i.e., \mathcal{F}_2 in (??) is the canonical RKHS). In Figure 2(b), the posterior means of the slopes obtained using the




Figure 3: Gene expressions for 144 cancer patients with 14 different types of cancer. Blue indicates strong expression, yellow weak.

standard random effects model are plotted against the ones obtained using the I-prior. Again we see broad agreement of the estimates, but much less so than for the varying intercept model.

A limited cross-validation study yielded on average a small advantage of the standard random effects approach in terms of mean squared error, in the order of half a percent, indicating the iid assumption is reasonable. An advantage of the I-prior is that no a priori assumption about the distribution of the parameters need to be made. Furthermore, as discussed in Section 1, our approach is more parsimonious and allows potentially simpler estimation and testing.

6.4 Multi-class classification

We apply the model described in Section 3.3.4 to two data sets that have received widespread attention in the literature, one for which the most important covariate is low dimensional and one for which it is high dimensional. In both cases, the classes $1, \dots, T$ are unordered. The effect of x is first assumed to be linear, then smooth, and the results are compared. We show that the classifier obtained from the I-prior is competitive with the other classifiers that we were able to find in the literature. Bergsma (2019) previously considered I-prior modelling for two-class classification.

The first problem was originally presented by Ramaswamy et al. (2001), and concerns the prediction of cancer type based on 16,063 gene expression measurements (see Figure 3). A training set of 144 patients with 14 different types of cancer is available, as well as a test set of 54 patients which can be used to assess the performance of a classifier. Assuming a linear dependence on class of the covariate vector of 16,063 gene expressions, we obtained 0 training errors and test 12 errors, which competes well with other methods (see Table 5). In particular our training error rate far outperforms competing methods. Modelling a smooth dependence of class on the covariate vector gave some numerical problems and we were unable to obtain full convergence of the EM algorithm. Fortunately, the classification errors did not seem to be affected by this, so we are confident these are accurate (the maximum of the likelihood we obtained, not given here, was inaccurate). For $0.72 < \gamma < 0.86$ the training/test errors were 0/10, which was the best we could achieve.

The second problem concerns vowel recognition based on a 10-dimensional vector computed from a voice recording. The training and test sets are based on recordings of 8 resp. 7 people, each of whom spoke 11 different vowels 6 times. As seen in Table 6, if a linear effect of the covariates is assumed, the I-prior gives no advantage compared to ordinary least squares. The reason is that there are only a small number of predictors. Again we had some numerical difficulties fitting a smooth model, but the MLE seems to be $\hat{\gamma} = 0.652$ giving training/test

Method	Training errors Out of 144	Test errors Out of 54
Nearest neighbors	41	26
L^2 -penalized discriminant analysis	25	12
Support vector classifier	26	14
Lasso	30.7	12.5
L^1 penalized multinomial	17	13
Elastic net penalized multinomial	22	11.8
SCRDA (Guo, Hastie, & Tibshirani, 2007)	24	8
Scout (Witten & Tibshirani, 2011)	21	8
I-prior (linear)	0	12
I-prior (smooth, $\gamma = 0.8$)	0	10

Table 5: Comparison of classifier performance for the 14 classification problem. The covariate is high-dimensional, consisting of a vector of 16,063 gene expressions. The top six lines are taken from Hastie et al. (2009), Table 18.1. With the I-prior, high-dimensional smoothing is as straightforward as fitting a linear effect, while giving slightly better test performance. We did have some numerical problems with maximum likelihood estimation of the Hurst coefficient γ , which we took to be 0.8.

error rates 0/0.35, improving on the results given in Hastie, Tibshirani, and Friedman (2009).

6.5 Longitudinal data analysis

We consider a balanced longitudinal data set consisting of weights and 60 cows, 30 of which are randomly assigned to a treatment group A and 30 to a treatment group B . Weight was measured 11 times over a 133-day period, at two-week intervals, except for the last measurement, which was taken one week after the preceding measurement. As the response variable of interest we take the weight growth curve. In Figure 4 a sample of growth curves is shown. Of interest is investigating whether a treatment effect is present, and if it is, to assess its nature.

The usual approach to analyze a longitudinal data set such as this one is to assume that




Figure 4: Sample of 20 (out of available 60) growth curves of cows. The two colors green and orange represent the two treatments received. The question of interest is whether and how treatment affects growth.

Method	Error rates	
	Training	Test
Nearest neighbours		0.44
OLS regression (linear effects)	0.48	0.67
Linear discriminant analysis	0.32	0.56
Neural network (Gaussian nodes)		0.45
FDA/BRUTO	0.06	0.44
FDA/MARS (best reduced dimension)	0.13	0.39
I-prior (linear)	0.48	0.67
I-prior (smooth, $\hat{\gamma} \approx 0.652$)	0	0.35

Table 6: Comparison of classifier performance for the vowel classification problem (first six lines taken from Hastie et al. (2009)). Since the dimension of the covariate is low (equal to 10), the I-prior gives no advantage over ordinary least squares (OLS) if the covariate effect is linear.

the observed growth curves are realizations of a Gaussian process. For example, Kenward (1987) assumed a so-called ante-dependence structure of order k , which assumes an observation depends on the previous k observations, but given these is independent of any preceding observations. Various other process families have been considered (Núñez-Antón & Zimmerman, 2000; Pourahmadi, 2000; Pan & Mackenzie, 2003; Zhang, Leng, & Tang, 2014), see the latter for an overview. Steele (2008) points out the connections with multilevel and structural equation modelling.

Using the I-prior, it is not necessary to assume the growth curves were drawn randomly, instead, it suffices to assume they lie in an appropriate function class. In this section we assume this function class is the FBM RKHS, i.e., we assume a ‘smooth’ effect of time on weight (see Table 9). The growth curves form a multidimensional (or functional) response, so we can use the multidimensional/functional response model of Section 3.3.3. In the present case we have two covariates potentially influencing growth, namely cow index (C , indexed by i) and treatment (X), and we can write the regression model as

$$y_{it} = f(i, x_i, t) + \varepsilon_{it},$$

where y_{it} is weight of cow i at time t , and $x_i \in \{A, B\}$ is the treatment group of cow i . As discussed in Section 3.3.3, we can also interpret the model as a unidimensional response model, where the unidimensional response is weight, and covariates are cow index (C), treatment (X), and time (T). A main effect of C on (multidimensional) *growth* is then equivalent to an interaction effect of T and C on (unidimensional) *weight*.

Covariate	Range	Effect	RKHS
Time (T)	0 – 133 days	Smooth	Centered FBM ($\gamma = 0.3$)
Cow index (C)	$\{1, \dots, 60\}$	Nominal	Pearson
Treatment (X)	$\{A, B\}$	Nominal	Pearson

Table 7: Covariates used for modelling cow data

We assume iid errors and in addition to a smooth effect of time we assume a nominal effect of both C and X (see Table 7). We can estimate the Hurst coefficient γ of the FBM RKHS

using the maximum likelihood estimator, but a difficulty is that it varies a lot across models (for the present data set, we found estimates between 0.2 and 0.4). To make model comparison easier, we took as a compromise a fixed value $\gamma = 0.3$ for all models; we found that substantive conclusions were not greatly affected by the choice of γ . The model asserting that the growth curve does not vary with treatment or among cows is denoted $\{\}$ (no effect of either C or X on growth), and can be written as

$$y_{it} = \alpha + f(t) + \varepsilon_{it}.$$

The model asserting that the growth curve depends on cow but without a treatment effect is denoted $\{C\}$, and can be written as

$$y_{it} = \alpha + \beta_i + f(t) + f_i^C(t) + \varepsilon_{it},$$

subject to the identifying constraints $\sum_t f(t) = \sum_t f_i^C(t) = 0$ and $\sum_{i=1}^{60} f_i^C(t) = 0$ for all t . Here t ranges over $0, 14, \dots, 133$. Further, f and the f_i^C are functions in the centered FBM RKHS. The model which includes a treatment effect on growth, but without an interaction effect of treatment and cow index, is denoted $\{C, X\}$ and can be formulated as

$$y_{it} = \alpha + \beta_i + \xi_{x_i} + f(t) + f_i^C(t) + f_{x_i}^X(t) + \varepsilon_{it},$$

where $f_{x_i}^X$ also lies in the FBM RKHS. It can be seen that, since treatment x_i is determined by the cow index i , the model is not identified in the sense that model $\{C\}$ contains exactly the same functions as model $\{C, X\}$ (and similarly, model $\{CX\}$ also contains the same functions as model $\{C\}$). However, the inner products of these ANOVA RKHSs differ, so they also have different I-priors, allowing us to perform model selection.

In Section 3.2, the parsimonious ANOVA RKHS and the less parsimonious extended ANOVA RKHS were described. The former requires a parameter for each covariate which has at least a main effect, but no extra parameters for interaction effects; the latter requires one parameter for each main effect and each interaction effect. The two approaches are different, as can be seen in the summary in Table 8, but reassuringly yield the same substantive conclusions.

Let us first consider the ANOVA approach. We can test for a treatment effect by testing model $\{C, X\}$ against the alternative that $\{C\}$ is true. The log-likelihood ratio is $(-2242.30) - (-2266.39) = 24.09$ and has an asymptotic chi-square distribution with $3 - 2 = 1$ degree of freedom if $\{C, X\}$ is true. The p -value is less than 10^{-6} so we conclude that model $\{C, X\}$ is significantly better than model $\{C\}$. We can next investigate whether the treatment effect differs among cows by comparing models $\{CX\}$ and $\{C, X\}$. As these two models have the same number of ANOVA (λ) parameters, we can simply choose the one with the highest likelihood, which is model $\{C, X\}$. We conclude that treatment has an effect on growth, but there is no reason to assume the effect of treatment differs among cows.

The extended ANOVA approach resembles classical least squares model selection since there are extra parameters for interaction effects. In this case model $\{C, X\}$ has five parameters, two more than model $\{C\}$, because X has a non-time-varying main effect and a time-varying effect on weight. Using likelihood ratio testing the best model is found to be $\{C, X\}$, giving the same result as in the parsimonious approach. Interestingly, $\{C, X\}$ and $\{CX\}$ give exactly the same fit for this data set, so we choose the simplest model $\{C, X\}$, which is the same result as for the ANOVA approach.

Time varying covariates were not available for the present data set but could easily be added. As the data set is balanced (the cows were weighted at the same time points), fitting can be done more efficiently without time varying covariates, however, we have 60 cows and 11 time points, giving $60 \times 11 = 660$ data points, which is easy to handle in any case.

Type of RKKS	Model	Log-likelihood	Error standard deviation	Number of λ parameters
ANOVA	$\{\}$	-2792.78	16.3	1
	$\{X\}$	-2792.73	16.3	2
	$\{C\}$	-2266.39	2.7	2
	$\{C, X\}$	-2242.30	2.5	3
	$\{CX\}$	-2251.30	3.3	3
Extended ANOVA	$\{\}$	-2792.78	16.3	2
	$\{X\}$	-2792.65	16.3	3
	$\{C\}$	-2250.31	3.6	3
	$\{C, X\}$	-2226.32	3.3	5
	$\{CX\}$	-2226.32	3.3	7

Table 8: Goodness of fit for cow data, see Table 7 for the covariates used. The model consists of the highest order effects on the growth curve, e.g., model $\{C, X\}$ means the growth curve depends on cow (C) and treatment (X), and there is no interaction meaning that the treatment effect is the same for all cows. In the ANOVA models, no extra parameters are needed for interaction effects, and among nested models with the same number of parameters the one with the highest likelihood can be chosen. For the extended ANOVA models, each interaction requires at least one extra parameter and model selection is via the likelihood ratio test.

7 Discussion and topics for further research

The most important reasons for using the I-prior methodology are that, as we aimed to show, it is relatively easy to use, flexible, and yields good predictive power. In particular, the I-prior methodology, consisting of an estimation method and a modelling method, can be used as a single solution for a wide range of applications where a large variety of methods have been used in the literature. For some well studied data sets we show that predictive performance is competitive with existing methods. The use of a full likelihood framework facilitates inference.

References

- Alpay, D. (1991). Some remarks on reproducing kernel Krein spaces. *Rocky Mountain J. Math.*
- Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3), 337–404.
- Bergsma, W. (2019). Regression with i-priors. *Econometrics and Statistics*.
- Berlinet, A., & Thomas-Agnan, C. (2004). *Reproducing kernel Hilbert spaces in probability and statistics*. Kluwer Academic.
- Canu, S., Ong, C. S., & Mary, X. (2009). Splines with non positive kernels. In W. S. Publishing (Ed.), *More progress in analysis* (pp. 163–173). Singapore.
- Chakraborty, A., & Panaretos, V. M. (2019, 08). Hybrid regularisation and the (in)admissibility of ridge regression in infinite dimensional Hilbert spaces. *Bernoulli*, 25(3), 1939–1976. Retrieved from <https://doi.org/10.3150/18-BEJ1041> doi: 10.3150/18-BEJ1041
- Chen, D., Hall, P., & Müller, H.-G. (2011). Single and multiple index functional regression models with nonparametric link. *Ann. Stat.*, 39(3), 1720–1747.

- Davies, R. B. (1980). Algorithm as 155: The distribution of a linear combination of χ^2 random variables. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 29(3), 323–333.
- Ferraty, F., & Vieu, P. (2006). *Nonparametric functional data analysis*. Springer Science + Business Media.
- Gheondea, A. (2013). A survey on reproducing kernel Krein spaces. *arXiv preprint arXiv:1309.2393*.
- Goia, A., & Vieu, P. (2014). Some advances on semi-parametric functional data modelling. *Contributions in infinite-dimensional statistics and related topics*, 135.
- Gretton, A., Bousquet, O., Smola, A., & Schölkopf, B. (2005). Measuring statistical dependence with Hilbert-Schmidt norms. *In: algorithmic learning theory: 16th international conference; Springer, 1*, 63–78.
- Gu, C. (2013). *Smoothing spline anova models*. Springer.
- Gu, C., & Wahba, G. (1993). Semiparametric analysis of variance with tensor product thin plate splines. *Journal of the Royal Statistical Society. Series B (Methodological)*, 353–368.
- Guo, Y., Hastie, T., & Tibshirani, R. (2007). Regularized linear discriminant analysis and its application in microarrays. *Biostatistics*, 8(1), 86–100.
- Hastie, T., & Tibshirani, R. (1993). Varying-coefficient models. *J. Roy. Statist. Soc. B*, 757–796.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: data mining, inference, and prediction*. New York: Springer.
- Hofmann, T., Schölkopf, B., & Smola, A. J. (2008). Kernel methods in machine learning. *Ann. Stat.*, 1171–1220.
- Jamil, H. (2018). *Regression modelling using priors depending on Fisher information covariance kernels (i-priors)* (Unpublished doctoral dissertation). The London School of Economics and Political Science (LSE).
- Jamil, H. (2019). iprior: Regression modelling using i-priors [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=iprior> (R package version 0.7.3)
- Jamil, H., & Bergsma, W. (2019). iprior: An R package for regression modelling using I-priors. <https://arxiv.org/abs/1912.01376>.
- Kenward, M. G. (1987). A method for comparing profiles of repeated measurements. *Applied Statistics*, 296–308.
- Kimeldorf, G. S., & Wahba, G. (1970). A correspondence between Bayesian estimation on stochastic processes and smoothing by splines. *Ann. Math. Stat.*, 41(2), 495–502.
- Lian, H., & Li, G. (2014). Series expansion for functional sufficient dimension reduction. *Journal of Multivariate Analysis*, 124, 150–165.
- Núñez-Antón, V., & Zimmerman, D. L. (2000). Modeling nonstationary longitudinal data. *Biometrics*, 56(3), 699–705.
- Ong, C. S., Mary, X., Canu, S., & Smola, A. J. (2004). Learning with non-positive kernels. In *ICML 21* (p. 639–646).
- Pan, J., & Mackenzie, G. (2003). On modelling mean-covariance structures in longitudinal studies. *Biometrika*, 90(1), 239–244.
- Parzen, E. (1961). An approach to time series analysis. *Ann. Math. Stat.*, 951–989.
- Pourahmadi, M. (2000). Maximum likelihood estimation of generalised linear models for multivariate normal covariance matrix. *Biometrika*, 87(2), 425–435.
- Ramaswamy, S., Tamayo, P., Rifkin, R., Mukherjee, S., Yeang, C.-H., Angelo, M., . . . others (2001). Multiclass cancer diagnosis using tumor gene expression signatures. *Proceedings*

- of the *National Academy of Sciences*, 98(26), 15149–15154.
- Rasbash, J., Steele, F., Browne, W. J., & Goldstein, H. (2012). A user’s guide to MLwiN, v2.26. *Centre for Multilevel Modelling, University of Bristol*.
- Schwartz, L. (1964). Sous-espaces Hilbertiens d’espaces vectoriels topologiques et noyaux associés (noyaux reproduisants). *Journal d’analyse mathématique*, 13(1), 115–256.
- Steele, F. (2008). Multilevel models for longitudinal data. *J. Roy. Statist. Soc. A*, 171(1), 5–19.
- Steinwart, I., & Christmann, A. (2008). *Support vector machines*. Springer.
- Thodberg, H. H. (1996). A review of Bayesian neural networks with an application to near infrared spectroscopy. *Neural Networks, IEEE Transactions on*, 7(1), 56–72.
- Vila, J.-P., Wagner, V., & Neveu, P. (2000). Bayesian nonlinear model selection and neural networks: a conjugate prior approach. *Neural Networks, IEEE Transactions on*, 11(2), 265–278.
- Wahba, G. (1990a). Multivariate model building with additive interaction and tensor product thin plate splines. In *Curves and surfaces* (pp. 491–504). Elsevier.
- Wahba, G. (1990b). *Spline models for observational data* (Vol. 59). Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM).
- Witten, D. M., & Tibshirani, R. (2011). Penalized classification using Fisher’s linear discriminant. *J. Roy. Statist. Soc. B*, 73(5), 753–772.
- Yao, F., Müller, H.-G., & Wang, J.-L. (2005). Functional linear regression analysis for longitudinal data. *The Annals of Statistics*, 33(6), 2873–2903.
- Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g -prior distributions. *Bayesian inference and decision techniques: Essays in Honor of Bruno De Finetti*, 6, 233–243.
- Zhang, W., Leng, C., & Tang, C. Y. (2014). A joint modelling approach for longitudinal studies. *J. Roy. Statist. Soc. B*.
- Zhu, H., Yao, F., & Zhang, H. H. (2014). Structured functional additive regression in reproducing kernel Hilbert spaces. *J. Roy. Statist. Soc. B*, 76(3), 581–603.

A ANOVA kernel construction

With \mathcal{V} a finite set of variables and any $v \in \mathcal{V}$, let \mathcal{X}_v be nonempty set, let \mathcal{F}_v be the RKKS over \mathcal{X}_v with kernel h_v , and let \mathcal{C}_v be the RKHS of constant functions over \mathcal{X}_v with kernel c_v defined by $c_v(x, x') = 1$. Denote the power set of a set A by $\mathbb{P}(A)$, and define the power set of a set of subsets \mathcal{A} as $\mathbb{P}(\mathcal{A}) = \cup_{A \in \mathcal{A}} \mathbb{P}(A)$. Let \mathcal{A} be a Sperner family of \mathcal{V} , i.e., \mathcal{A} consists of subsets of \mathcal{V} such that no set in \mathcal{A} contains another set in \mathcal{A} . Define the RKKS $\mathcal{F}_{\mathcal{A}}$ over $\times_{v \in \mathcal{V}} \mathcal{X}_v$ as

$$\mathcal{F}_{\mathcal{A}} = \sum_{A \in \mathbb{P}(\mathcal{A})} \bigotimes_{v \in A} \mathcal{F}_v \bigotimes_{v \in \mathcal{V} \setminus A} \mathcal{C}_v. \quad (26)$$

Then $f \in \mathcal{F}_{\mathcal{A}}$ are of the form

$$f(x) = \sum_{A \in \mathbb{P}(\mathcal{A})} f_A(x_A)$$

for $x \in \times_{v \in \mathcal{V}} \mathcal{X}_v$, with x_A the subvector of x which retains the coordinates corresponding to A . The reproducing kernel of $\mathcal{F}_{\mathcal{A}}$ is given as

$$h_{\mathcal{A}} = \sum_{A \in \mathbb{P}(\mathcal{A})} \bigotimes_{v \in A} h_v \bigotimes_{v \in \mathcal{V} \setminus A} c_v.$$

RKHS/RKKS	Name of effect
Centered canonical	Linear
Pearson	Nominal
Centered FBM	Smooth
Polynomial	Polynomial
ANOVA	Interaction

Table 9: Terminology for regression modelling, allowing a description of regression models without referring to RKKSs and RKHSs. For example, we say a (possibly multidimensional) covariate has a linear effect on the response if the corresponding regression function lies in the centered canonical RKKS.

B OLD

B.1 Asymptotic distribution of MLEs 1

Suppose $y \sim GP(0, G)$ where $G = \lambda A + B$ for positive definite $n \times n$ A and B , $\lambda \in \mathbb{R}$. The log likelihood is

$$l(\lambda) = -\frac{1}{2} \log|G| - \frac{1}{2} y^\top G^{-1} y$$

The derivative is

$$\dot{l}(\lambda) = -\frac{1}{2} \text{tr}(G^{-1} A) + \frac{1}{2} y^\top G^{-1} A G^{-1} y$$

and the second derivative is

$$\ddot{l}(\lambda) = \frac{1}{2} \text{tr}(G^{-1} A G^{-1} A) - y^\top G^{-1} A G^{-1} A G^{-1} y$$

Let us look at the distribution of $\hat{\lambda}$ for which $\dot{l}(\hat{\lambda}) = 0$ when the true value $\lambda_0 = 0$. Note that

$$E\dot{l}(\lambda) = 0 \quad \text{var}(\dot{l}(\lambda)) = \text{tr}(G^{-1} A G^{-1} A)$$

$$E\ddot{l}(\lambda) = \frac{1}{2} \text{tr}(G^{-1} A G^{-1} A) \quad \text{var}(\ddot{l}(\lambda)) = 2 \text{tr}(G^{-1} A G^{-1} A G^{-1} A G^{-1} A)$$

We have

$$\begin{aligned} \hat{\lambda} &= -\frac{\dot{l}(0)}{\ddot{l}(0)} + O_p(\hat{\lambda}^2) = -\frac{-\frac{1}{2} \text{tr}(B^{-1} A) + \frac{1}{2} y^\top B^{-1} A B^{-1} y}{\frac{1}{2} \text{tr}(B^{-1} A B^{-1} A) - y^\top B^{-1} A B^{-1} A B^{-1} y} + O_p(\hat{\lambda}^2) \\ &= -\frac{\frac{1}{2} \tilde{y}^\top B^{-1/2} A B^{-1/2} \tilde{y} - \frac{1}{2} \text{tr}(B^{-1/2} A B^{-1/2})}{\tilde{y}^\top B^{-1/2} A B^{-1} A B^{-1/2} \tilde{y} - \frac{1}{2} \text{tr}(B^{-1/2} A B^{-1} A B^{-1/2})} + O_p(\hat{\lambda}^2) \\ &= -\frac{\frac{1}{2} \tilde{y}^\top U \tilde{y} - \frac{1}{2} \text{tr}(U)}{\tilde{y}^\top U^2 \tilde{y} - \frac{1}{2} \text{tr}(U^2)} + O_p(\hat{\lambda}^2) \end{aligned}$$

where $\tilde{y} \sim N(0, I)$ and $U = B^{-1/2} A B^{-1/2}$. The numerator has mean zero and variance $\text{tr}(U^2)$, the denominator has mean $\frac{1}{2} \text{tr}(U^2)$ and variance $\text{tr}(U^4)$.

B.2 Asymptotic distribution of MLEs 2

Suppose $y \sim GP(0, G)$ where $G = \lambda^2 A + B$ for positive definite $n \times n$ A and B , $\lambda \in \mathbb{R}$. The log likelihood is

$$l(\lambda) = -\frac{1}{2} \log|G| - \frac{1}{2} y^\top G^{-1} y$$

The derivate is

$$\dot{l}(\lambda) = -\lambda \text{tr}(G^{-1}A) + \lambda \mathbf{y}^\top G^{-1}AG^{-1}\mathbf{y} = -\lambda [\text{tr}(G^{-1}A) - \mathbf{y}^\top G^{-1}AG^{-1}\mathbf{y}]$$

and the second derivative is

$$\ddot{l}(\lambda) = -\text{tr}(G^{-1}A) + \mathbf{y}^\top G^{-1}AG^{-1}\mathbf{y} + 2\lambda^2 \text{tr}(G^{-1}AG^{-1}A) - 4\lambda^2 \mathbf{y}^\top G^{-1}AG^{-1}AG^{-1}\mathbf{y}$$

Let us look at the distribution of $\hat{\lambda}$ for which $\dot{l}(\hat{\lambda}) = 0$ when the true value $\lambda_0 = 0$. We have

$$\hat{\lambda} = -\frac{\dot{l}(0)}{\ddot{l}(0)} + O_p(\hat{\lambda}^2) = -\frac{\dot{l}(0)}{\ddot{l}(0)} + O_p(\hat{\lambda}^2)$$

B.3 Test for significance of one covariate keeping other constant

We consider the case of iid errors. Suppose

$$y_i = f_1(x_{1i}) + f_2(x_{2i}) + \varepsilon_i \quad f_k \in \mathcal{F}_k, x_{ki} \in \mathcal{X}_k, k = 1, 2, i = 1, \dots, n$$

where

$$\varepsilon_i \stackrel{\text{iid}}{\sim} N(0, \sigma^2)$$

and \mathcal{F}_k a function space. Consider the null and alternative hypotheses

$$H_0 : f_1 = 0 \quad H_1 : f_1 \neq 0$$

A difficulty is we cannot use the test of the previous subsection, because if an estimate of f_2 is substituted we cannot normally control type I error. Instead, we propose a likelihood ratio type test.

Under H_0 , let

$$\hat{\mathbf{f}}_2 = A(A + \tau I)^{-1}\mathbf{y}$$

Then

$$\mathbf{y} - \hat{\mathbf{f}}_2 = \mathbf{y} - A(A + \tau I)^{-1}\mathbf{y} = \tau(A + \tau I)^{-1}\mathbf{y} = \tau(A + \tau I)^{-1}(\mathbf{f}_2 + \boldsymbol{\varepsilon})$$

An estimator of σ^2 under H_0 is

$$\hat{\sigma}_0^2 = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{f}_2(x_{2i}))^2 = \frac{\tau^2}{n} \mathbf{y}^\top (A + \tau I)^{-2} \mathbf{y}$$

Similarly, defining under H_1

$$\widehat{\mathbf{f}}_1 + \widehat{\mathbf{f}}_2 = B(B + \tau I)^{-1}\mathbf{y}$$

we obtain

$$\hat{\sigma}_1^2 = \frac{1}{n} \sum_{i=1}^n (y_i - (\widehat{f}_1(x_{1i}) + \widehat{f}_2(x_{2i})))^2 = \frac{\tau^2}{n} \mathbf{y}^\top (B + \tau I)^{-2} \mathbf{y}$$

The likelihood ratio statistic is

$$\text{LR} = \frac{L(\hat{f}_1, \widehat{f}_2, \hat{\sigma}_1)}{L(0, \widehat{f}_2, \hat{\sigma}_0)} = \left(\frac{\hat{\sigma}_0}{\hat{\sigma}_1} \right)^n$$

We may also consider the F -like statistic

$$F = \frac{\hat{\sigma}_0^2 - \hat{\sigma}_1^2}{\hat{\sigma}_1^2}$$

With $\hat{f}_1 = A\mathbf{y}$ and $\widehat{f}_2 = B\mathbf{y}$, we obtain

$$\begin{aligned} \hat{\sigma}_1^2 &= n^{-1} \mathbf{y}^\top (I - A - B)^2 \mathbf{y} \\ \hat{\sigma}_0^2 &= n^{-1} \mathbf{y}^\top (I - B)^2 \mathbf{y} \end{aligned}$$

and

$$\hat{\sigma}_0^2 - \hat{\sigma}_1^2 = n^{-1} \mathbf{y}^\top [(I - B)^2 - (I - A - B)^2] \mathbf{y}$$

Lemma 6. Suppose $\hat{f}_1 = Ay$ and $\hat{f}_{2k} = B_k y$ are consistent, and H_0 is true. Then

$$E((\sigma - \hat{\sigma}_k)^2) \rightarrow 0$$

$$HSIC(\hat{\sigma}_0^2 - \hat{\sigma}_1^2, \hat{\sigma}_1^2) \rightarrow 0$$

B.3.1 OLD

Suppose

$$y_i = f(x_i) + g(z_i) + \varepsilon_i \quad f \in \mathcal{F}, x_i \in \mathcal{X}, z_i \in \mathcal{Z}, i = 1, \dots, n$$

where \mathcal{F} is an RKHS with kernel h and g is a given function over \mathcal{X} . Consider the null and alternative hypotheses

$$H_0 : f = 0 \quad H_1 : f \neq 0$$

Note that we have a conditional independence problem, that is, we wish to test whether $Y|f(X)$ is identically distributed to $Y|(f(X), g(X))$. We first consider the case that g is known, then that it needs to be estimated.

In practice, \mathbf{g} will often need to be estimated. Let $\hat{\mathbf{g}}$ be an estimator of g and define

$$\hat{T}^2 = \|s(0)\|_{\mathcal{F}}^2 = (\mathbf{y} - \hat{\mathbf{g}})^\top \Psi \mathbf{H} \Psi (\mathbf{y} - \hat{\mathbf{g}})$$

Recall that

interpretation???

$$I[\hat{g} - g] = \sum_{i=1}^n \sum_{j=1}^n \psi_{ij} (\hat{g}(x_i) - g(x_i)) (\hat{g}(x_j) - g(x_j)) = (\hat{\mathbf{g}} - \mathbf{g})^\top \Psi (\hat{\mathbf{g}} - \mathbf{g})$$

Lemma 7. It holds true that

$$(\hat{T} - T)^2 \leq \sqrt{\text{var}(T^2)/2} \times I[\hat{g} - g]$$

Proof. We have by Cauchy-Schwartz and Lemma 10,

$$\begin{aligned} \|\hat{s}(0) - s(0)\|_{\mathcal{F}}^2 &= (\hat{\mathbf{g}} - \mathbf{g})^\top \Psi \mathbf{H} \Psi (\hat{\mathbf{g}} - \mathbf{g}) = \text{tr}[\Psi \mathbf{H} \Psi (\hat{\mathbf{g}} - \mathbf{g})(\hat{\mathbf{g}} - \mathbf{g})^\top] \\ &\leq \sqrt{\text{tr}(\Psi \mathbf{H} \Psi \mathbf{H})} \sqrt{\text{tr}[\Psi (\hat{\mathbf{g}} - \mathbf{g})(\hat{\mathbf{g}} - \mathbf{g})^\top \Psi (\hat{\mathbf{g}} - \mathbf{g})(\hat{\mathbf{g}} - \mathbf{g})^\top]} \\ &= \sqrt{\text{var}(T^2)/2} I[\hat{g} - g] \end{aligned}$$

Hence, using Cauchy-Schwartz,

$$\begin{aligned} (\hat{T} - T)^2 &= \hat{T}^2 + T^2 - 2\hat{T}T \leq \|\hat{s}(0)\|_{\mathcal{F}}^2 + \|s(0)\|_{\mathcal{F}}^2 - 2\langle \hat{s}(0), s(0) \rangle_{\mathcal{F}} = \|\hat{s}(0) - s(0)\|_{\mathcal{F}}^2 \\ &\leq \sqrt{\text{var}(T^2)/2} I[\hat{g} - g] \end{aligned}$$

□

The following lemma shows that under weak conditions, \hat{T} can be used as an asymptotically equivalent test statistic as T .

Corollary 1. Suppose $I[\hat{g} - g] = o_P(\sqrt{n})$ and $\text{var}(T^2) = O(1)$. Then if H_0 is true,

$$\frac{1}{n} \hat{T}^2 \xrightarrow{P} \frac{1}{n} T^2$$

and if H_1 is true,

$$\frac{1}{\sqrt{n}} \hat{T}^2 \xrightarrow{P} \frac{1}{\sqrt{n}} T^2$$

Proof.

□

B.4 old

Let

$$r_2(x) = \sum_{i=1}^n \sum_{j=1}^n h_2(x, x_i) \psi_{ij}(\hat{f}_1(x_j) - f_1(x_j))$$

We then have

$$\hat{s}(0) = s(0) - r_2$$

and

$$\hat{T}^2 = T^2 - 2\langle r_2, s(0) \rangle_{\mathcal{F}_2} + \|r_2\|_{\mathcal{F}_2}^2$$

Now

$$\|r_2\|_{\mathcal{F}_2}^2 = (\hat{\mathbf{f}}_1 - \mathbf{f}_1)^\top \Psi \mathbf{H}_2 \Psi (\hat{\mathbf{f}}_1 - \mathbf{f}_1) \leq \sqrt{\text{tr}(\Psi \mathbf{H}_2)} (\hat{\mathbf{f}}_1 - \mathbf{f}_1) \Psi (\hat{\mathbf{f}}_1 - \mathbf{f}_1) = T \times o_P(\sqrt{n})$$

Furthermore,

$$\langle r_2, s_2(0) \rangle_{\mathcal{F}_2} \leq T \|r_2\|_{\mathcal{F}_2} = T^{3/2} o_P(n^{1/4})$$

where the inequality is via Cauchy-Schwartz.

In practice, \mathbf{f}_1 will often need to be estimated, say by $\hat{\mathbf{f}}_1$. Define

$$\hat{T}^2 = \|s_2(0)\|_{\mathcal{F}_2}^2 = (\mathbf{y} - \hat{\mathbf{f}}_1) \Psi \mathbf{H} \Psi (\mathbf{y} - \hat{\mathbf{f}}_1)$$

Lemma 8. *Suppose*

$$(\hat{\mathbf{f}}_1 - \mathbf{f}_1) \Psi \mathbf{H} \Psi (\hat{\mathbf{f}}_1 - \mathbf{f}_1) = o_P(\sqrt{n})$$

Then if H_0 is true,

$$\frac{1}{n} \hat{T}^2 \xrightarrow{P} \frac{1}{n} T^2$$

and if H_1 is true,

$$\frac{1}{\sqrt{n}} \hat{T}^2 \xrightarrow{P} \frac{1}{\sqrt{n}} T^2$$

Proof. We have the decomposition

$$\hat{T}^2 = T^2 - 2(\mathbf{y} - \mathbf{f}_1)^\top \Psi \mathbf{H} \Psi (\hat{\mathbf{f}}_1 - \mathbf{f}_1) + (\hat{\mathbf{f}}_1 - \mathbf{f}_1) \Psi \mathbf{H} \Psi (\hat{\mathbf{f}}_1 - \mathbf{f}_1) =: T^2 - 2\nu_1 + b$$

Then by the assumption, $n^{-1/2}b \xrightarrow{P} 0$. Furthermore, conditionally on \mathbf{X} , ν_1 has mean zero and

$$\text{var}(n^{-1/2}\nu_1 | \mathbf{X}) = n^{-1} (\hat{\mathbf{f}}_1 - \mathbf{f}_1) \Psi \mathbf{H} \Psi \mathbf{H} \Psi (\hat{\mathbf{f}}_1 - \mathbf{f}_1)$$

so $E(n^{-1}\nu_1^2 | \mathbf{X}) \xrightarrow{P} 0$ and it follows $P(n^{-1}\nu_1^2 > \epsilon) \rightarrow 0$ for any $\epsilon > 0$

□

Alternative approach:

With $\mathbf{A} = \mathbf{H}_1 \mathbf{\Psi} \mathbf{H}_1$ let

$$\begin{aligned} r_2(x) &= \sum_{i=1}^n \sum_{j=1}^n h_2(x, x_i) \psi_{ij} (\hat{f}_1(x_j) - f_1(x_j)) \\ &= \sum_{i=1}^n \sum_{j=1}^n h_2(x, x_i) \psi_{ij} (\mathbf{A}(\mathbf{A} + \mathbf{\Psi}^{-1})^{-1})_{ij} \varepsilon_i - \sum_{i=1}^n \sum_{j=1}^n h_2(x, x_i) ((\mathbf{A} + \mathbf{\Psi}^{-1})^{-1})_{ij} f_1(x_j) \end{aligned}$$

We then have

$$\hat{s}_2(0) = s(0) - r_2$$

and

$$\hat{T}^2 = T^2 - 2\langle r_2, s_2(0) \rangle_{\mathcal{F}_2} + \|r_2\|_{\mathcal{F}_2}^2$$

Now with $\mathbf{Q} = (\mathbf{A} + \mathbf{\Psi}^{-1})^{-1} \mathbf{A}$,

$$\|r_2\|_{\mathcal{F}_2}^2 = \varepsilon^\top \mathbf{Q} \mathbf{\Psi} \mathbf{H}_2 \mathbf{\Psi} \mathbf{Q}^\top \varepsilon + \mathbf{f}_1^\top (\mathbf{A} + \mathbf{\Psi}^{-1})^{-1} \mathbf{H}_2 (\mathbf{A} + \mathbf{\Psi}^{-1})^{-1} \mathbf{f}_1 - 2\varepsilon^\top \mathbf{Q} \mathbf{\Psi} \mathbf{H}_2 (\mathbf{A} + \mathbf{\Psi}^{-1})^{-1} \mathbf{f}_1$$

and

$$\begin{aligned} E(\|r_2\|_{\mathcal{F}_2}^2) &= \text{tr}(\mathbf{\Psi}^{-1} \mathbf{Q} \mathbf{\Psi} \mathbf{H}_2 \mathbf{\Psi} \mathbf{Q}^\top) + \mathbf{f}_1^\top (\mathbf{A} + \mathbf{\Psi}^{-1})^{-1} \mathbf{H}_2 (\mathbf{A} + \mathbf{\Psi}^{-1})^{-1} \mathbf{f}_1 \\ \text{var}(\|r_2\|_{\mathcal{F}_2}^2) &= \text{tr}(\mathbf{\Psi}^{-1} \mathbf{Q} \mathbf{\Psi} \mathbf{H}_2 \mathbf{\Psi} \mathbf{Q}^\top \mathbf{\Psi}^{-1} \mathbf{Q} \mathbf{\Psi} \mathbf{H}_2 \mathbf{\Psi} \mathbf{Q}^\top) + \dots \end{aligned}$$

Alternatively, with

$$\mathbf{r}_1 = (\mathbf{H}_1 \mathbf{\Psi} \mathbf{H}_1 + \mathbf{\Psi}^{-1})^{-1} \mathbf{f}_1 \quad \mathbf{R}_1 = (\mathbf{H}_1 \mathbf{\Psi} \mathbf{H}_1 + \mathbf{\Psi}^{-1})^{-1} \mathbf{\Psi}^{-1} (\mathbf{H}_1 \mathbf{\Psi} \mathbf{H}_1 + \mathbf{\Psi}^{-1})^{-1}$$

we have

$$E(T^2) = \text{tr}(\mathbf{\Psi} \mathbf{H}_2) \quad E(\hat{T}^2) = \text{tr}(\mathbf{R}_1 \mathbf{H}_2) + \mathbf{r}_1^\top \mathbf{H}_2 \mathbf{r}_1$$

and

$$\text{var}(T^2) = 2 \text{tr}(\mathbf{\Psi} \mathbf{H}_2 \mathbf{\Psi} \mathbf{H}_2) \quad \text{var}(\hat{T}^2) = 2 \text{tr}(\mathbf{R}_1 \mathbf{H}_2 \mathbf{R}_1 \mathbf{H}_2) + 4 \mathbf{r}_1^\top \mathbf{H}_2 \mathbf{R}_1 \mathbf{H}_2 \mathbf{r}_1$$

(NB: doesn't take into account that λ s need to be estimated!!!)

B.4.1 Comparison with existing tests

Note we can write

$$T^2 = \text{tr}(\mathbf{\Psi} \mathbf{H} \mathbf{\Psi} \mathbf{G})$$

where $\mathbf{G} = (\mathbf{y} - \mathbf{f}_1)(\mathbf{y} - \mathbf{f}_1)^\top$. We see that T^2 is a modification of HSIC = $\text{tr}(\mathbf{H} \mathbf{G})$ (Gretton et al., 2005), that is, if $\mathbf{\Psi}$ is proportional to the identity matrix (i.e., iid errors), T^2 is essentially HSIC based on the kernel h over \mathcal{X} and the canonical kernel over \mathbb{R} .

For example, if $f_1(x) = x_1 \beta_1$ and $f_2(x) = x_2 \beta_2$, we obtain the squared partial correlation coefficient

$$S^2 = \frac{(r_{X_2, Y} - r_{X_1, X_2} r_{X_1, Y})^2}{(1 - r_{X_1, X_2}^2)(1 - r_{X_1, Y}^2)}$$

while

$$T^2 =$$

The score function for the parameter f_1 is

$$s(f_1) = \sum_{i=1}^n \sum_{j=1}^n h(\cdot, x_i) \psi_{ij} (y_j - \alpha - f_1(x_j))$$

Note that $s(f_1) \in \mathcal{F}_1$. Substituting the ML estimator $\hat{\alpha} = \bar{y} = n^{-1} \sum_{i=1}^n y_i$ gives

$$\hat{s}(f) = \sum_{i=1}^n \sum_{j=1}^n h(\cdot, x_i) \psi_{ij} (y_j - \bar{y})$$

The usual score statistic is

$$S^2 = \|\hat{s}(0)\|_{\mathcal{F}_n}^2 = (\mathbf{y} - \bar{y}) \boldsymbol{\Psi} \mathbf{H} (\mathbf{H} \boldsymbol{\Psi} \mathbf{H})^{-1} \mathbf{H} \boldsymbol{\Psi} (\mathbf{y} - \bar{y})$$

where $(\cdot)^{-}$ represents a generalized inverse. Unfortunately, S^2 can only be used if \mathcal{F}_1 is low-dimensional. As an alternative, we propose a *modified score statistic*, defined as the \mathcal{F} -norm of the estimated score function evaluated at H_0 :

$$T^2 = \|\hat{s}(0)\|_{\mathcal{F}}^2 = (\mathbf{y} - \bar{y})^\top \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} (\mathbf{y} - \bar{y})$$

The permutation p -value based on T^2 is defined as the proportion of permutations of y_1, \dots, y_n which give test statistic value larger than T^2 .

Lemma 9. *It holds true that*

$$ET^2 = \mathbf{f}_1^\top \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} \mathbf{f}_1 + \text{tr}(\boldsymbol{\Psi} \mathbf{H}) + O_p(1)$$

Proof. We have

$$\begin{aligned} ET^2 &= E(\mathbf{y} - \bar{y})^\top \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} (\mathbf{y} - \bar{y}) \\ &= E(\alpha - \bar{y} + \mathbf{f}_1 + \boldsymbol{\varepsilon})^\top \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} (\alpha - \bar{y} + \mathbf{f}_1 + \boldsymbol{\varepsilon}) \\ &= E(\alpha - \bar{y} + \mathbf{f}_1)^\top \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} (\alpha - \bar{y} + \mathbf{f}_1) + E\boldsymbol{\varepsilon}^\top \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} \boldsymbol{\varepsilon} \\ &= \mathbf{f}_1^\top \boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} \mathbf{f}_1 + \text{tr}(\boldsymbol{\Psi} \mathbf{H}) + O_p(1) \end{aligned}$$

□

The next lemma will be used later. The proof follows immediately from the well-known formulas for the mean and variance of quadratic forms in multivariate normal variates.

Lemma 10. *If holds true that*

$$E(T^2) = \text{tr}(\boldsymbol{\Psi} \mathbf{H}) \quad \text{var}(T^2) = 2 \text{tr}(\boldsymbol{\Psi} \mathbf{H} \boldsymbol{\Psi} \mathbf{H})$$