Test (2010) 19: 537–557
DOI 10.1007/s11749-010-0185-3

ORIGINAL PAPER

A simple multiway ANOVA for functional data

J.A. Cuesta-Albertos · M. Febrero-Bande

Received: 30 July 2009 / Accepted: 10 February 2010 / Published online: 6 March 2010

Abstract We propose a procedure to test complicated ANOVA designs for functional data. The procedure is effective, flexible, easy to compute and does not require a heavy computational effort. It is based on the analysis of randomly chosen one-dimensional projections. The paper contains some theoretical results as well as some simulations and the analysis of some real data sets. Functional data include multidimensional data, so the paper contains a comparison between the proposed procedure and some usual MANOVA tests.

Keywords ANOVA · Functional data · Random projections · Testing · Two-way ANOVA

Mathematics Subject Classification (2000) 62H15 · 62J10

1 Introduction

The popularization of computers has allowed many processes to be continuously monitored, thus providing data which are functions, the so-called functional data.

Communicated by Domingo Morales.

J.A. Cuesta-Albertos has been partially supported by the Spanish Ministerio de Ciencia y Tecnología, grant MTM2008-0607-C02-02 and the Consejería de Educación y Cultura de la Junta de Castilla y León, grant PAPIJCL VA102/06.
M. Febrero-Bande has been partially supported by the Spanish Ministerio de Ciencia y Tecnología, grant MTM2008-03010.

J.A. Cuesta-Albertos
Departamento de Matemáticas, Estadística y Computación, Facultad de Ciencias, Universidad de Cantabria, Santander, Spain
e-mail: juan.cuesta@unican.es

M. Febrero-Bande (✉)
Departamento de Estadística e Inv. Operativa, Facultad de Matemáticas, Universidad Santiago de Compostela, Santiago de Compostela, Spain
e-mail: manuel.febrero@usc.es

Those include the changes in value of a particular share or index on the stock market, the temperatures in a given location, etc., but it is also possible to include, in this setting, data which are not functions but rather highly multidimensional like those handled in genomics. Usual references on functional data are Ramsay and Silverman (2002, 2005), Ferraty and Vieu (2006), and references therein.

A natural statistical problem is to decide on the existence or not of differences in the process of interest when some conditions which may affect it vary. We refer, for example, to the existence of differences between the temperatures in different locations, between the values of the stocks in American, European, and Japanese markets, ...

This problem is usually handled by employing a model which assumes the existence of an underlying function describing the typical evolution of the process under consideration, assuming that the data we have at hand have been obtained by adding random fluctuations to this typical function. Thus, the problem becomes a kind of functional ANOVA problem.

Let us state the problem more precisely. Assume that we have a real-valued random process $\mathbf{X} = \{\mathbf{X}(t), t \in [a, b]\}$, measured along the closed and bounded interval $[a, b]$. Assume that there are R different conditions which may affect the process and that $\mathbf{X}$ have been measured n_r times under each of those conditions, giving the sample $\mathbf{X}_i^r(t), i = 1, \ldots, n_r, r = 1, \ldots, R$. Assume also that for every $r = 1, \ldots, R$, there exists a (nonrandom) function $f^r : [a, b] \to \mathbb{R}$ such that

$$\mathbf{X}_i^r(t) = m(t) + f^r(t) + \epsilon_i^r(t), \quad t \in [a, b], \ i = 1, \ldots, n_r, \tag{1}$$

where ϵ_i^r are random functions centered in mean (i.e., $E[\epsilon_i^r(t)] = 0$ for every t). Our aim is to test the hypothesis $f^1 = f^2 = \cdots = f^r = 0$.

In Ramsay and Silverman (2005), as well as in many other papers on the subject, advantage is taken of the fact that the measurements are usually made in a finite set of values $\{t_1, \ldots, t_m\} \subset [a, b]$. The idea is to apply a real ANOVA analysis for every value t_j. Thus, we obtain m values for the F-test, one for each t_j. However, even in the case that $R = 1$ and that the distribution of $(\epsilon_i^1(t_1), \ldots, \epsilon_i^1(t_m))$ is standard Gaussian, it is shown in Fan (1996) that if we want to test the null hypothesis $f^1 = 0$, then the powers of the likelihood ratio test converge to the level of the test if $m \to \infty$, and $\|(f^1(t_1), \ldots, f^1(t_m))\| \to \infty$ with $\|(f^1(t_1), \ldots, f^1(t_m))\| = o(m^{1/2})$. To solve those problems, a kind of bootstrap procedure is proposed in Ramsay and Silverman (2005), and, in Fan and Lin (1998), the authors suggest a criterion to select a finite subset of t_j's.

Other approaches are based on the application of a dimension-reduction and/or smoothing technique (see, for instance, Abramovich et al. 2004; Abramovich and Angelini 2005; Angelini and Vidakovic 2002; Antoniadis and Sapatinas 2007; Brumback and Rice 1998; Faraway 1997; Ferraty et al. 2007; Gu 2002; Guo 2002). Different points of view are handled in Cuevas et al. (2004), Fujikoshi et al. (2004), and Schott (2007). In Cuevas et al. (2004), an F-test based on the squares of the norms of the between samples and within samples differences is proposed. In Fujikoshi et al. (2004) and Schott (2007), the authors analyze the behavior of the traces of the between and within covariance matrices. A common drawback of those solu-

tions (excepting Brumback and Rice 1998; Gu 2002, and Guo 2002) is that they are intended to solve only one-way problems.

The proposals in Brumback and Rice (1998), Gu (2002), and Guo (2002) are based on the use of the so-called smoothing splines (SS-ANOVA). Initially, the SS-ANOVA had the strong limitation that the time was considered as an additional factor, which implies that the values $\epsilon_i^r(t)$ and $\epsilon_i^r(s)$ are independent if $t \neq s$. This requirement is somewhat unreasonable because this independence would imply that the functions $\mathbf{X}_i^r$ are discontinuous in every point, while, in practice, the continuity of the observed functions seems to be customary (this limitation remains in the available version of the `gss` R-package (R Development Core Team 2007), which is a practical implementation of the SS-ANOVA).

Later, in Brumback and Rice (1998) and Gu (2002), this restriction was removed (we consider Brumback and Rice 1998 particularly interesting because it includes a discussion in which the authors apply several methods to analyze the same data). However, the technique underlying those papers is quite complicated because it relies on the so-called Reproductive Kernel Hilbert Spaces and the involved computations are quite demanding.

We can summarize the situation by saying that there exist several methods to handle ANOVA problems. Some of them only allow one-way designs. Some others admit complicated designs, but the computations are rather involved. There are also methods based on pointwise one-dimensional comparisons which could be blind against some alternatives even in simple situations. An additional limitation is that many of the previous methods are based on the normality assumption.

The goal of this paper is to propose an effective, flexible, and easy to compute technique able to deal with complicated ANOVA designs requiring no normality and as few additional hypotheses as possible. The key idea here is to use random projections to transform functional data into univariate data, solve the ANOVA problem in this simple situation, and obtain conclusions for the functional data by collecting the information from several projections.

The technique is based on Theorem 4.1 in Cuesta-Albertos et al. (2007c), which requires that the measurements belong to a separable Hilbert space. This result was generalized in Cuevas and Fraiman (2009) to general spaces without a Hilbert space structure. Thus, our technique also works (with some obvious modifications) in those general spaces. In spite of this, we have chosen in this paper the Hilbert framework in order to keep the exposition as simple as possible because we consider that this setting is general enough as to include many real-world problems: if the functions $t \to \mathbf{X}_i^r(t)$ are measured on a bounded interval $[a, b]$ and $\int_a^b (\mathbf{X}_i^r(t))^2\, dt$ do exist and are finite (which, obviously, happens if the functions are continuous), then the functions belong to the separable Hilbert space of the square-integrable functions.

Theorem 4.1 in Cuesta-Albertos et al. (2007c), roughly speaking, states that if we have two different distributions, one of them determined by their moments (see Shohat and Tamarkin 1963), and we randomly choose a marginal of them, those marginals are almost surely different. We state this result here for further reference.

Theorem 1.1 *Let $\mathcal{H}$ be a separable Hilbert space with norm $\|\cdot\|$. Let μ be a Gaussian distribution on $\mathcal{H}$ such that each of its one-dimensional projections is non-degenerate.*

Let P and Q be two probability distributions on $\mathcal{H}$ such that P is determined by their moments. If $P \neq Q$, then

$$\mu\left\{v \in \mathcal{H} : P \circ \Pi_v^{-1} = Q \circ \Pi_v^{-1}\right\} = 0,$$

where Π_v is the orthogonal projection onto the one-dimensional subspace generated by v.

Several conditions have been proposed to assure that a distribution P is determined by their moments like, for instance, the so-called Carleman condition: $\sum_k m_k^{-1/k} = \infty$, where $m_k := \int \|x\|^k \, dP(x)$ are assumed to be finite (see Shohat and Tamarkin 1963). This condition is satisfied for compact-supported or Gaussian distributions and, more generally, for distributions with light tails.

According to Theorem 1.1, if we want to test the null hypothesis

$$H_0 : f^1 = \cdots = f^R = 0$$

in (1), we only need to randomly take a one-dimensional subspace, project onto this subspace the functions $f^1, \ldots, f^R$, and test whether those one-dimensional projections coincide.

We have presented this reasoning in the one-way ANOVA merely to simplify the exposition, but the procedure can be extended to cover more complicated linear models. Roughly speaking, it so happens that our procedure covers every functional design which can be handled in the one-dimensional case.

The test that we propose here is conditional on the obtained one-dimensional projection which has to be chosen independently of the sample at hand. Conditional tests are not new in statistics. In fact, most bootstrap procedures are conditional (on the obtained sample), but we can trace conditional procedures, at least, to the exact Fisher test for contingency tables.

We also want to remark that since Theorem 1.1 holds if we take $\mathcal{H} = \mathbb{R}^d$, the proposed procedure can be used as a simple alternative to the usual MANOVA. In this setting, there exist procedures to handle complicated designs, some of them enjoying optimality properties. The comparison between the proposed procedure and some available standard alternatives is interesting in order to see the existence or not of differences in power.

The paper follows the following scheme: In Sect. 2 we develop the procedure. In Sect. 3 we present some simulations. In Sect. 4 we analyze three real data sets. The first one contains the results of a two-way physiological experiment. The second one is related to the temperatures in different locations in Spain; this is a two-way ANOVA plus a continuous covariable. The third one is a nested two-factor model on the progesterone levels in some women; it is the example analyzed in Brumback and Rice (1998). Section 5 is devoted to comparing the proposed procedure with usual MANOVA tests; in order to make this comparison more interesting, we have chosen a very unfavorable situation: since our procedure is intended for functional data, it seems that the lower the dimension, the worse for the procedure; thus, we have chosen a three-dimensional example to make the comparison. The paper ends with a Discussion.

2 The procedure

In this section $\mathcal{H}$ will denote a separable Hilbert space endowed with scalar product $\langle \cdot, \cdot \rangle$. Without loss of generality we will assume that the closed interval in which we measure the process is $[0, 1]$.

2.1 The model. Assumptions

We begin generalizing the model (1) to a two-way ANOVA model with a covariable. Obviously, other models can be treated similarly.

Let $R, S \in \mathbb{N}$ and assume that for all $r = 1, \ldots, R$ and $s = 1, \ldots, S$, there exist $\mathbf{X}_i^{r,s}, i = 1, \ldots, n_{r,s} \in \mathbb{N}$, random functions in $\mathcal{H}$, such that

$$\mathbf{X}_i^{r,s}(t) = m(t) + f^r(t) + g^s(t) + h^{r,s}(t) + \gamma(t) Y_i^{r,s} + \epsilon_i^{r,s}(t), \quad t \in [0, 1], \qquad (2)$$

where:

1. The function m is nonrandom and describes the overall shape of the process.
2. The (nonrandom) functions $f^r, g^s, h^{r,s}$ belong to $\mathcal{H}$ and account, respectively, for the main effect of the first and second factors and for the interaction between them.

   In order to make the model identifiable, we will assume, as usual, that, for all $t \in [0, 1]$, $r_0 = 1, \ldots, R$, and $s_0 = 1, \ldots, S$,

$$\sum_r f^r(t) = \sum_s g^s(t) = \sum_r h^{r,s_0}(t) = \sum_s h^{r_0,s}(t) = 0.$$

3. The $Y_i^{r,s} \in \mathbb{R}$ are both random and known quantities that influence the process according to the weights given by the (nonrandom and unknown) function $\gamma \in \mathcal{H}$. They play the same role as the covariables in the one-dimensional case.
4. The random trajectories $\epsilon_i^{r,s}$ are assumed to be $\mathcal{H}$-valued, independent, and centered in mean. Moreover, for each r, s fixed, $\epsilon_i^{r,s}, i = 1, \ldots, n_{r,s}$, are identically distributed.

We are interested in testing the following null hypotheses:

$$\begin{aligned}
&H_0^A : f^1 = \cdots = f^R = 0 && \text{(the first factor has no influence)},\\
&H_0^B : g^1 = \cdots = g^S = 0 && \text{(the second factor has no influence)},\\
&H_0^I : h^{1,1} = \cdots = h^{R,S} = 0 && \text{(there is no interaction between factors)},\\
&H_0^C : \gamma = 0 && \text{(the covariable has no influence)},
\end{aligned}$$

against the alternatives that their negations hold.

2.2 The procedure

The procedure that we propose is a direct application of the following result.

Theorem 2.1 *Let us assume the model and assumptions stated in Sect. 2.1. Let μ be a Gaussian distribution on $\mathcal{H}$ such that each of its one-dimensional projections is nondegenerate. Then*

1. *If there exist* r_1, r_2 *such that* $f^{r_1} \neq f^{r_2}$, *then*

$$\mu\left\{v \in \mathcal{H} : \textit{ such that } \langle v, f^1\rangle = \cdots = \langle v, f^R\rangle\right\} = 0.$$

2. *If there exist* s_1, s_2 *such that* $g^{s_1} \neq g^{s_2}$, *then*

$$\mu\left\{v \in \mathcal{H} : \textit{ such that } \langle v, g^1\rangle = \cdots = \langle v, g^S\rangle\right\} = 0.$$

3. *If there exist* $(r_1, s_1), (r_2, s_2)$ *such that* $h^{r_1,s_1} \neq h^{r_2,s_2}$, *then*

$$\mu\left\{v \in \mathcal{H} : \textit{ such that } \langle v, h^{1,1}\rangle = \cdots = \langle v, h^{R,S}\rangle\right\} = 0.$$

4. *If* $\gamma \neq 0$, *then* $\mu\{v \in \mathcal{H} : \textit{ such that } \langle v, \gamma\rangle = 0\rangle\} = 0$.

Proof Given $r = 1, \ldots, R$, let P^r be the probability distribution on $\mathcal{H}$ concentrated on f^r (i.e., P^r satisfies that $P^r[f^r] = 1$). Obviously, every probability distribution concentrated on a single point satisfies Carleman's condition. Thus, we can apply Theorem 1.1 to every pair of probability distributions P^{r_1} and P^{r_2}, and the result in item 1 follows.

The proofs of the remaining statements are identical. □

To explain the procedure, let us focus our attention on the testing of hypothesis H_0^A. Let $v \in \mathcal{H}$ be a vector chosen at random employing the distribution μ. If H_0^A holds, then for every $v \in \mathcal{H}$, the (conditional on v) null hypothesis

$$H_0^{A,v} : \langle v, f^1\rangle = \cdots = \langle v, f^R\rangle = 0$$

also holds. Also, according to 1 in Theorem 2.1, if H_0^A fails, then for μ-almost every $v \in \mathcal{H}$, $H_0^{A,v}$ also fails. Thus, a statistical test at level α to test $H_0^{A,v}$ is a statistical test at the same level to test H_0^A.

Therefore, if Φ is a statistical test at level α to test the null hypothesis H_0^A under the model (2) if the data are real, it is then possible to test H_0^A in the $\mathcal{H}$-valued case by applying Φ to test $H_0^{A,v}$.

The same can be said with respect the null hypotheses H_0^B, H_0^I, and H_0^C and the corresponding (random) counterparts $H_0^{A,v}$, $H_0^{B,v}$, and $H_0^{C,v}$.

Let us present a toy example of application of the method. With this example, it is not our aim to show all of its possibilities as neither is it to carry out a full analysis of the example. In fact, to be perfectly honest, it has to be said that this example was selected as the most suitable in a series of 15 trials.

Example 2.2 Let us assume that we have two samples (see left graph in Fig. 1) of size 15 each, $X_1, \ldots, X_{15}$ (in blue) and $Y_1, \ldots, Y_{15}$ (in green) of functions measured on the interval $[0, 1]$.

Assume that we know that there exist two functions m_1 and m_2 such that

$$X_i(t) = m_1(t) + \epsilon_i^1(t), \quad i = 1, \ldots, 15,$$
$$Y_j(t) = m_2(t) + \epsilon_j^2(t), \quad j = 1, \ldots, 15,$$

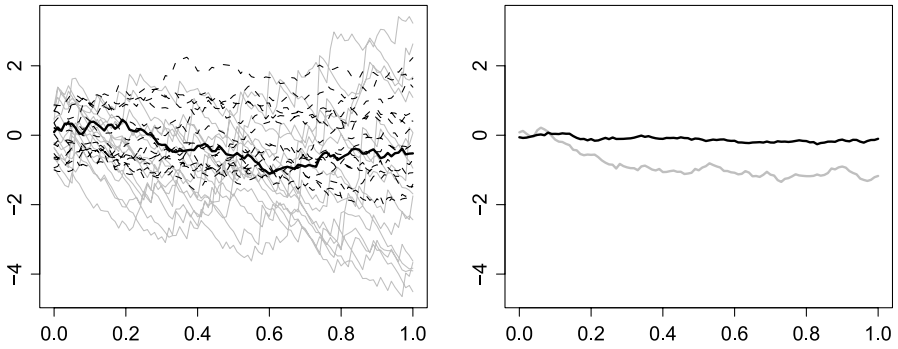

Fig. 1 *Left-hand graph*: Samples of trajectories (*gray* and *dashed black curves*) and randomly chosen vector (*solid black curve*). *Right-hand graph*: Means of the samples

where, for every t, $E[X_i(t)] = m_1(t)$ and $E[Y_j(t)] = m_2(t)$, and the random functions $\epsilon_1^1, \ldots, \epsilon_{15}^1$ (respectively $\epsilon_1^2, \ldots, \epsilon_{15}^2$) are identically distributed. Assume that we are interested in testing the null hypothesis

$$H_0 : m_1 = m_2.$$

The sample means of both samples appear in the right-hand graph in Fig. 1. There are many procedures in the literature to test H_0 in this functional setting, but, perhaps it is not too evident which one is more suitable for this particular problem. In particular, some of the available procedures require that the data be Gaussian, and it is not easy either to decide whether this assumption holds or not.

On the other hand, the involved functions are continuous, so, $\int_0^1 X_i^2(t)\,dt$ and $\int_0^1 Y_j^2(t)\,dt$ are finite, and we can apply the previous theory. To do this, we need to choose, with a Gaussian distribution, a vector v in the space of the square-integrable functions. In this case, we use the distribution of a standard Brownian motion. To select the precise v, we take advantage of the fact that the curves in both samples were only measured in the points $0, .01, .02, \ldots, .99, 1$. Therefore, we only need to know the values of v in those points. For this, we can employ the usual approximation of the Brownian motion by a sequence of partial sums of independent normal variables as follows:

1. Let $Z_1, \ldots, Z_{100}$ be independent, identically distributed standard normal variables.
2. Take $v(0) = 0$ and for every $h = 1, \ldots, 100$, let $t_h = .01 \times h$ and

$$v(t_h) = v(t_{h-1}) + .01 \times Z_h.$$

The trajectory v is represented in red in Fig. 1. To apply Theorem 2.1, let us consider the integrals

$$m_i^v = \int_0^1 m_i(t)v(t)\,dt, \quad i = 1, 2. \tag{3}$$

If H_0 holds, obviously, $m_1^v = m_2^v$. On the other hand, the real numbers

$$\frac{m_1^v}{(\int_0^1 v^2(t)\,dt)^{1/2}} \quad \text{and} \quad \frac{m_2^v}{(\int_0^1 v^2(t)\,dt)^{1/2}}$$

are, respectively, the projections of the functions m_1 and m_2 onto the one-dimensional subspace generated by v. Therefore, according to Theorem 2.1, if, in fact, $m_1 \neq m_2$, then it is almost surely impossible that $m_1^v = m_2^v$, and, in consequence, it is equivalent either directly testing H_0 or the (conditional on v) null hypothesis

$$H_0^v : m_1^v = m_2^v.$$

Notice that m_1^v and m_2^v are real numbers and that they are the means of the (also real) random variables $\int_0^1 X_i(t)v(t)\,dt$ and $\int_0^1 Y_j(t)v(t)\,dt$, respectively. Therefore, to test H_0^v, we only need to compute those integrals for all the members in both samples (obtaining two samples of integrals) and, then, apply a test to check if it is reasonable to assume that the mean of the distributions which generated those samples was the same or not. The integrals in (3) can be approximated by the sums

$$p_i^1 = .01\sum_{h=1}^{101} X_i(t_h)v(t_h) \quad \text{and} \quad p_j^2 = .01\sum_{h=1}^{101} Y_j(t_h)v(t_h), \quad i, j = 1, \ldots, 15. \quad (4)$$

Now, we can employ the (real) samples p_i^1 and p_j^2 to test H_0^v. Here the situation is more clearer because we are dealing with real data. For instance, Fig. 2 shows the box-plot of those values.

We can even apply the Kolmogorov–Smirnov test to check their fit to the normal family. This test rejects the normality of the p_i^1's at the usual .05 level (p-value equals .0255) but not that of the p_j^2's (p-value equals .3038). Under those circumstances, the t-test does not seem appropriate to test H_0^v and, the Wilcoxon–Mann–Whitney test may be preferable. Anyway, the p-value obtained with the t-test is .0392 (thus leading to the rejection of H_0^v (and then to that of H_0) at the .05 level, while the

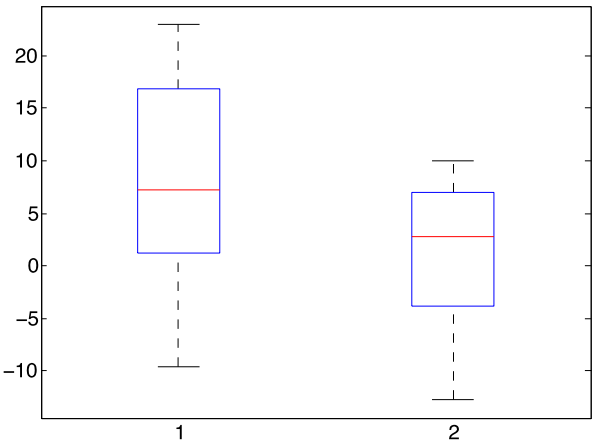

Fig. 2 Box-plot of the quantities computed in (4)

second test gives $p = .1057$, which is not enough to reject H_0^v. In conclusion, we should reject the normality hypothesis, declare the t-test inappropriate, and conclude that there is not enough evidence to reject H_0.

An interesting conclusion of the Kolmogorov–Smirnov test is that the functional data are not Gaussian (if they were Gaussian, then every projection would be Gaussian, and we have one which is not). Therefore, if we had chosen to analyze the functional data using a method requiring Gaussianity, then, the analysis would have been flawed.

In this case, the null hypothesis holds with $m_1 = m_2 = 0$. The distribution of ϵ_i^1 is a mixture of a non-Gaussian distribution plus a 10% of outliers, while the distribution of ϵ_j^2 is that of the standard Brownian motion plus a random fixed quantity chosen in the interval $[-1, 1]$.

2.2.1 Model assumptions

It is well known that, in classical ANOVA, hypotheses like the homoscedasticity or the Gaussianity are crucial. Thus, it is worth paying some attention to the question of which hypotheses are required in our procedure.

The key here is that the problem is split into two: First project and, then, apply ANOVA. Thus, it is possible to decide which is the right ANOVA to be applied looking at the projections. Take, for instance, Example 2.2. After computing the projections, we saw that the t-test was not appropriate, and, then, we selected a nonparametric test.

The same happens with the ANOVA: after computing the projections, we can select the right (one-dimensional) ANOVA technique to be applied, depending on when the data are Gaussian, homoscedastic, ... or not. In particular, in random functions, homoscedasticity does not look reasonable because, very often, the oscillation of a process strongly depends on their values.

Therefore, it is good to have at hand a one-dimensional ANOVA test which works under heteroscedastic and/or non-Gaussian conditions. There are several available procedures in the literature which fulfil those requirements. Here, we have chosen the one proposed in Brunner et al. (1997) which will be the procedure used unless otherwise stated.

2.2.2 Stability of the procedure

According to the previous reasoning, the exposed procedure is consistent in the sense that the rejection of any null hypothesis $H^{\cdot,v}$ is almost surely equivalent to rejecting the corresponding null hypothesis $H_0^{\cdot}$. However, it is obvious that this procedure has two main drawbacks.

On one hand, we are losing some information because we are replacing a function by just one real number and this should bring some loss of power. On the other hand, when employing a randomly chosen projection, we are accepting some random instability in the problem because what may happen is that if we repeat the procedure twice, a null hypothesis could first be rejected and later accepted.

In order to reduce these inconveniences, it has been proposed (for instance, in Cuesta-Albertos et al. 2007a, 2007b) to take $k > 1$ random projections, test the null

hypothesis for each projection, and, then, adjust the obtained p-values by bootstrap. Bonferroni's method is the selection in Cuesta-Albertos et al. (2008). Alternatively, in descriptive studies where the goal is to provide an idea about how far from acceptable the null hypothesis under consideration is, reporting, for instance, the mean value of the obtained p-values could be enough.

It is very well known that Bonferroni's correction is quite conservative. On the other hand, the bootstrap is quite time consuming, and the proofs to show that it works in specific problems are quite involved. Therefore, our advice here is to employ a correction directed at controlling the False Discovery Rate (FDR). If k different statistical hypotheses are tested, then the FDR is the expected proportion of wrongly rejected hypotheses. Since in our situation all the hypotheses that we are testing are equivalent, the FDR coincides with the level of the procedure.

The FDR was introduced in Benjamini and Hochberg (1995) for independent tests. Here, we will employ the improvement proposed in Benjamini and Yekutieli (2001).

Given a set $D \subset \mathbb{R}^k$, we say that D is *increasing* if $x \in D$ and $y \geq x$ implies that $y \in D$ (here we handle the coordinate-wise order). Now, let $Z = (Z_1, \ldots, Z_k)'$ be a random k-dimensional vector. We say that Z is *positively regression dependent* (PRD) if, for every increasing set D and $i = 1, \ldots, k$, the map $x \to P[Z \in D | Z_i = x]$ is increasing in x. Notice that we can safely assume that the tests that we are going to apply are PRD because they are based on the same initial data set.

Proposition 2.3 (Benjamini and Yekutieli 2001) *Let us assume that we apply k PRD tests to check the same null hypothesis. Assume that the ordered p-values that we obtain are $p_{(1)}, \ldots, p_{(k)}$ where $p_{(1)} \leq \cdots \leq p_{(k)}$. Let $\alpha \in (0, 1)$. The level of the test which rejects if the set*

$$\left\{ i : p_{(i)} \leq \frac{i}{k} \alpha \right\}$$

is not empty is, at most, α.

This result gives us a procedure to compute corrected p-values as follows: given the ordered p-values $p_{(1)}, \ldots, p_{(k)}$, Proposition 2.3 allows us to reject the null hypothesis for every level $\alpha \geq \inf\{\frac{k}{i} p_{(i)}, i = 1, \ldots, k\}$. Therefore, we will take as a corrected p-value the quantity $\inf\{\frac{k}{i} p_{(i)}, i = 1, \ldots, k\}$.

If Bonferroni's correction leads to rejection at a given level α (which happens when $p_{(1)} \leq \alpha/k$), this procedure also rejects. On the other hand, the correction based on the FDR remains conservative because the test is *at most* at the desired level; but the experience shows that it is far less conservative than Bonferroni's correction. We will return to this question in the next subsection.

The only remaining problem is the selection of the number k of random projections, which according to Theorem 4.1 in Cuesta-Albertos et al. (2007c), must be independent with respect to the data. For this, we refer the interested reader to Cuesta-Albertos and Nieto-Reyes (2007, 2008) where this problem is studied. In Sect. 3 in Cuesta-Albertos and Nieto-Reyes (2008) some quite conservative upper bounds for k are given. This conservatism is shown in Cuesta-Albertos and Nieto-Reyes (2008), where a not so large value for k is taken when analyzing practical problems. The

functional case is analyzed in Sect. 4.1 in Cuesta-Albertos and Nieto-Reyes (2007) in connection with a classification by sex problem. Here the authors handle $k \in \{1, 2\}$ (in Sect. 4.1.1) and $k \in \{10, 20\}$ (in Sect. 4.1.2). This leads us to use in this paper up to $k = 30$ which is considerably above the number of projections effectively handled in those papers.

2.2.3 Finite-dimensional case

It is obvious that finite-dimensional Euclidean spaces can be considered as particular cases of Hilbert spaces. Then, in principle, there is no need for modifications to apply the procedure in those spaces. However, there are some practical points worth considering. The first one is related to the distribution to be employed to select the random directions. Stationarity is often a reasonable assumption in random processes, and, consequently, there is no need to distinguish between values of t. However, in the finite-dimensional case there is no reason to assume that the marginals of the underlying distribution are the same. Thus, it may not be particularly wise to select μ in Theorem 2.1, for instance, as the standard Gaussian distribution.

The second problem is related to the fact that if the dimension of the space is low as compared with k, then excessive dependencies between marginals may damage the FDR correction, which becomes a little conservative. We have sometimes observed this phenomenon, mostly in the low-dimensional applications (see Sect. 5). Our proposal is to choose k not greater than the dimension of the data being analyzed.

Concerning the first question, our proposal is, first, to apply an affine transformation to the data in order to make them centered in mean and its covariance matrix the identity. Then, we choose as μ the standard Gaussian distribution. The proposed affine transformation can be seen as a way to easily choose the random projections by taking into account the marginal distributions and their dependences. In other words, this transformation plus μ equal to the standard Gaussian distribution is equivalent to no transformation on the data and to μ being a Gaussian distribution with a covariance matrix depending on the marginal distributions and their dependences.

Moreover, from the point of view of the comparisons that we carry out in Sect. 5, since the affine transformations do not affect the usual MANOVA tests, this transformation has no influence on them.

3 Simulations

A Monte Carlo study has been carried out in order to evaluate the performance of the procedure proposed in Sect. 2. To this end, we have simulated several examples following the model (2). We have taken $R = S = 2$ and $n_{r,s} = 100, r, s = 1, 2$, in all of them. The underlying distributions depend on two parameters α, β as follows:

- $m(t) := 30(1 - t)t$.
- $f^r(t) := \alpha(-1)^r |\sin(4\pi t)|, r \in \{1, 2\}, \alpha \in \{0, .025, .05, .075, .10, .15, .25, .35, .50\}$.
- $g^s(t) := \beta(-1)^s I_{\{t>0.5\}}, s \in \{1, 2\}, \beta \in \{0, .025, .05, .075, .10, .15, .25, .35, .50\}$.
- $h^{r,s}(t) := -f^r(t)g^s(t)I_{\{\alpha \geq 0.25\}}$.

- $\epsilon_i^{r,s}(t)$ is a Gaussian process with mean zero and covariance function

$$\mathrm{Cov}\big(\epsilon_i^{r,s}(t_1), \epsilon_i^{r,s}(t_2)\big) := \sigma_{r,s}^2 \exp\big(-|t_1 - t_2|/.3\big), \quad \sigma_{r,s}^2 = .2 + r \times s \times .1,$$

where $t \in [0, 1]$ was discretized in $N = 101$ points, i.e., the processes were observed at the points $t_l = l/100, l = 0, \ldots, 100$. We wish to remark that we are in a heteroscedastic setting with $\sigma_{r,s}^2 \in \{.3, .4, .5\}$ depending on the particular cell we are considering.

The parameter α (resp. β) controls the strength of the factor $f^r(t)$ (resp. $g^s(t)$). In fact, depending on α, the interaction is present or not.

We have done $K = 500$ replications for every possible combination of the values of α and β whose results appear in Tables 1, 2, and 3. In order to save space, when some cases are statistically identical, we have summarized all the results in a single cell. For instance, when analyzing the null hypothesis H_0^A (resp. H_0^B), the values of β (resp. α) have no influence. Then, since we have 9 different values for α and β, every cell in Tables 1 and 2 summarizes the result of 4500 replications.

We have taken $k = 1, 5, 15, 30$ random projections, and each cell shows the proportion of rejections at level .05 after correcting the p-values with the procedure described following Proposition 2.3.

The results are quite satisfactory. The first column in Table 1 corresponds to the situation of no effect of the factor A, and the behavior is nice. The same can be said about Table 2 for factor B. As α is growing in Table 1 (resp. β in Table 2), the proportion of rejections quickly increases.

Table 3 has the information about interaction which is only present for $\alpha \geq .25$. Here, the first four rows are the summary for $\alpha \in \{0, .025, .05, .075, .10, .15\}$, where no interaction is present, and so, every cell is the result of $K = 3000$ replications. The following rows correspond to those values of α when interaction is present, and here, the number of replications is $K = 500$. Since the amount of interaction depends on

Table 1 Proportion of rejections for H_0^A at the level $p = .05$

RP	$\alpha = .00$	$\alpha = .025$	$\alpha = .05$	$\alpha = .075$	$\alpha = .10$	$\alpha = .15$	$\alpha = .25$	$\alpha = .35$	$\alpha = .50$
1	.051	.089	.175	.303	.438	.605	.766	.827	.886
5	.049	.093	.268	.544	.769	.955	.999	1	1
15	.044	.099	.320	.671	.902	.998	1	1	1
30	.037	.094	.351	.745	.944	1	1	1	1

Table 2 Proportion of rejections for H_0^B at the level $p = .05$

RP	$\beta = .00$	$\beta = .025$	$\beta = .05$	$\beta = .075$	$\beta = .10$	$\beta = .15$	$\beta = .25$	$\beta = .35$	$\beta = .50$
1	.048	.077	.182	.330	.432	.620	.787	.835	.889
5	.049	.082	.266	.545	.786	.969	.998	1	1
15	.039	.089	.306	.659	.905	.998	1	1	1
30	.036	.091	.322	.697	.932	1	1	1	1

Table 3 Proportion of rejections for H_0^I at the level $p = .05$

α	RP	$\beta = .00$	$\beta = .025$	$\beta = .05$	$\beta = .075$	$\beta = .10$	$\beta = .15$	$\beta = .25$	$\beta = .35$	$\beta = .50$
$\leq .15$	1	.052	.057	.049	.059	.050	.054	.049	.051	.049
	5	.048	.043	.047	.051	.046	.044	.043	.045	.050
	15	.043	.042	.045	.046	.045	.040	.039	.043	.042
	30	.040	.040	.048	.045	.038	.039	.037	.036	.043
.25	1	.042	.066	.048	.052	.052	.060	.110	.176	.300
	5	.054	.070	.068	.038	.052	.066	.136	.240	.510
	15	.032	.062	.056	.040	.046	.086	.190	.328	.632
	30	.024	.052	.042	.034	.042	.080	.200	.350	.732
.35	1	.048	.050	.062	.072	.072	.102	.178	.282	.454
	5	.048	.042	.046	.060	.062	.114	.280	.454	.754
	15	.050	.048	.052	.050	.082	.130	.322	.614	.896
	30	.056	.036	.042	.064	.086	.118	.342	.666	.954
.50	1	.042	.066	.060	.066	.104	.158	.266	.424	.584
	5	.046	.060	.076	.072	.130	.248	.554	.770	.926
	15	.030	.048	.080	.082	.148	.272	.658	.918	.994
	30	.032	.038	.062	.084	.134	.294	.730	.956	1

the product $\alpha \times \beta$, the proportion of rejections increases when we approach the lower right corner and goes near one only when both α and β are large.

Concerning the effect of the number of projections we wish to mention that the conservative character of the procedure proposed in Proposition 2.3 clearly appears under the null: in the three tables, in practically every cell corresponding to the null, the proportions decrease with k. Interestingly, the opposite happens under the alternative, excluding some cells in Table 3 for low values of α and β. The reason for this could be that, as k increases, the chance of finding a direction in which the alternative appears more strongly increases, and this increment overcomes the conservativeness of the procedure. However, this chance becomes smaller when k is great in comparison with the real dimension of the data. This is what led us to limit k being, as most, equal to the dimension of the data (see Sect. 5).

In this Monte Carlo study, the hypothesis H_0^C cannot be checked because under a heteroscedastic setting, as far as we know, there is no univariate method that provides a good approximation of p-values with continuous covariates. In order to check the hypothesis H_0^C, we simulated examples of model (2) with a functional covariate in a homoscedastic setting, as follows:

- $\alpha, \beta, r, s, t, m(t), f^r(t), g^s(t), h^{r,s}(t)$ as above.
- $\gamma(t) := (t - 0.5)$, and the distribution of the real random variable Y is $U[-.5, .5]$.
- $\epsilon_i^{r,s}(t)$ is a Gaussian process with mean zero and covariance function

$$\mathrm{Cov}\big(\epsilon_i^{r,s}(t_1), \epsilon_i^{r,s}(t_2)\big) := \sigma^2 \exp\big(-|t_1 - t_2|/.3\big), \quad \sigma^2 = .3.$$

The results for this model are quite similar to those shown above for the principal factors and the interaction. With respect to the hypothesis H_0^C, the proportion of rejections is, not surprisingly, near one for every cell. Obviously, under the assumed hypotheses, we have that the one-dimensional projections are Gaussian and homoscedastic. Thus, we have handled the unidimensional projections with a standard ANOVA.

4 Three real data sets

In this section we illustrate the proposed procedure by using three real datasets. On each problem, the number of random projections is the minimum between 30 and the dimension of the data.

4.1 Orthosis data

Those data have been previously studied in Abramovich et al. (2004), Abramovich and Angelini (2005), and Antoniadis and Sapatinas (2007). We want to thank Prof. Anestis Antoniadis who kindly provided us the data that, according to Antoniadis and Sapatinas (2007), were obtained at the Laboratoire Sport et Performance Motrice, Grenoble University, France.

The goal of the study was to analyze how muscle copes with an external perturbation. For this, seven male volunteers where measured while stepping on the spot using a spring-loaded orthosis under four conditions: without orthosis (Control), with the orthosis only (Orthosis), and two conditions (Spring1, Spring2) in which a spring-loaded orthosis was fitted onto the right knee joint. All subjects tried all conditions 10 times for 20 seconds each; but only the central 10 seconds were used in the study in order to avoid possible perturbations in the initial and final parts of the exercise. The resultant moment at the knee was derived by means of body segment kinematics recorded with a sampling frequency of 200 Hz. For each replication, 64 measurements, equally spaced, were taken and scaled so that a time interval corresponds to an individual gait cycle.

Therefore, the dataset consists of 280 separate runs and involves the seven subjects over four described experimental conditions, replicated 10 times for each subject. The graphical representation of the data appears in Fig. 3(a) in Antoniadis and Sapatinas (2007).

We have a two-way functional ANOVA (subjects and treatments being the qualitative factors) with 10 observations per cell. The associated problem does not seem particularly easy. In fact, in the analysis carried out in Abramovich et al. (2004) and Abramovich and Angelini (2005), the authors found significative differences between treatments, but not between Control and Orthosis, and not between Spring1 and Spring2 either. Then, in Antoniadis and Sapatinas (2007), they treated subjects as random effects that allowed them to carry out a more powerful analysis that led them to find significant differences in the analyses Spring1 + Spring2 vs. Control + Orthosis conditions (p-value $= 0$), Control vs. Orthosis (p-value $= .001$), and Spring1 vs. Spring2 (p-value $= .020$).

Table 4 p-values for the considered null hypothesis. Spr1&2 vs. Co&Or refers to (Spring1 + Spring2) vs. (Control + Orthosis)

Null hypothesis being tested					
Subject	Treatm.	Interact.	Spr1&2 vs. Co&Or	Cont vs. Orth	Spr1 vs. Spr2
0	0	0	0	.0007	.0280

Here, we will consider that both factors are fixed-effects. This allows us to analyze the existence of interactions between them which is an aspect not previously considered. We have applied our test choosing 30 random directions. Table 4 shows the p-values obtained.

The results in Table 4 show a strong evidence on the influence of both factors and on the existence of interaction. The data also very strongly support the existence of differences between wearing some kind of spring against not and between Control and Orthosis measurements. The evidence favoring differences between the springs is less. In fact, the test fails to find them if we take only 5 random projections (p-value equals .6208), but we obtain p-values similar to the one obtained in Antoniadis and Sapatinas (2007) when we increase this number. Therefore, the random ANOVA procedure allows all considered null hypotheses to be rejected in accordance with the results in Antoniadis and Sapatinas (2007) and overcoming the difficulties found in Abramovich et al. (2004) and Abramovich and Angelini (2005).

4.2 Spanish temperature data

Those data were downloaded from http://clima.meteored.com in September, 2007. This site contains climatological data from many meteorological stations. We consider the daily mean temperature as the functional datum for certain locations in Spain and certain months in the latest annual cycle. More precisely, we have downloaded the daily mean temperature in A Coruña, Avilés, Bilbao, San Sebastián, Santander, Vigo, Burgos, León, Madrid, Salamanca, Segovia, Soria, Valladolid, Vitoria, and Zamora. The last nine locations are in the interior of Spain and have cold winters and hot summers, whereas the first six locations are situated on the North Coast of Spain (Atlantic Ocean) and have temperatures that are less extreme. Also, in order to take into account the variability of seasons, we have selected the most representative available months in the considered period: October 2006, January 2007, May 2007, and July 2007.

Indeed, we have also taken into account the Monthly Total Amount of Rainfall as a covariate that can influence the temperature. Thus, this is the real known random variable Y in (2) which multiplies the unknown and nonrandom function γ, which measures the influence of this quantity each day in the month. It seems to us that in this particular example, Y should be constant, but this is not important in our model.

In Fig. 3 we have represented the curves of the mean temperatures in each month and zone. The influence of the seasonal factor becomes obvious from those graphs. But, in our opinion, the differences between the corresponding curves are not so great as to make the influence of the location factor so obvious. The same happens with the

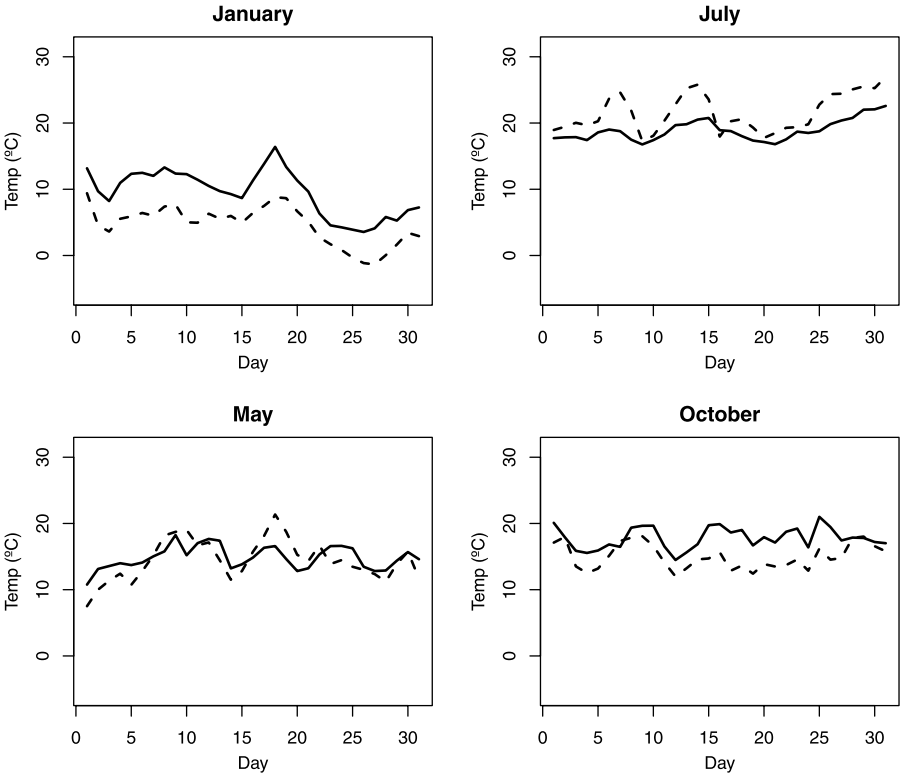

Fig. 3 Mean temperatures in coastal (*thick line*) and inland (*dotted*) locations by months

Table 5 p-values obtained with the random ANOVA

	Location	Month	Interaction	Monthly Rainfall
p-value	$2.98 \cdot 10^{-5}$	0.00	$7.8 \cdot 10^{-10}$	.089

existence of the interaction, in spite of the graphs suggesting its existence, since the mean temperature in the interior is lower than in the coastal area in winter and greater in summer.

Table 5 contains the p-values obtained for each null hypothesis with the proposed test using 30 random projections. The proposed procedure detects the existence of influence of both factors and, also, the existence of interaction between them. The influence of the monthly rainfall is rejected at the usual .05 level but not at the .10 level.

4.3 Urinary progesterone data

In this study the levels of urinary progesterone of several women with a healthy reproductive function enrolled in an artificial insemination program were measured. They had undergone insemination attempts in every menstrual cycle. The objective

was to determine the existence or not of differences in the daily progesterone levels depending on when a woman is in a conceptive or nonconceptive menstrual cycle. The duration of those cycles is 24 days aligned by the day of ovulation, and thus data outside days $[-8, 15]$ were excluded. A second question of interest is the existence or not of differences in the levels between those groups in days 2 to 8.

The levels of urinary progesterone were measured in 22 conceptive and 69 nonconceptive menstrual cycles. An interesting characteristic of this data set is that there were 22 women in the conceptive group and 29 in the nonconceptive because there were some cycles in this group measured on the same woman. Therefore, we have a two-way nested ANOVA.

The paper (Brumback and Rice 1998) focuses on the analysis of those data. A graphical representation of some data appears there. This paper contains a discussion, and, then, the data are analyzed using a linear mixed-effects model for curved data employing the SS structure (Brumback and Rice), a pointwise ANOVA followed by a smoothing step (Fan and Zhang), and a pointwise ANOVA (Ramsay and Silverman, RS). Globally speaking, the conclusions of those studies are similar.

There are some problems in this data set. First, there are some missing data. In particular, only 36 cycles contain the measurement of the last day, and only 56 contain that of the penultimate day (the other days are more complete). In order to handle this problem, we have followed the suggestion given by RS, and we have deleted the last day from the computations. Therefore, our data are 23-dimensional. Concerning the other missing data, we again follow RS using a cubic spline smoother with 9 degrees of freedom to assign absent data.

The results of our analysis appear in Table 6. To analyze the 8-day period starting on day 2, we have considered the sample restricted to this period. In other words, we have applied two independent analyses to the same sample: the first one to a set of 23-dimensional data and a second one to the 7-dimensional data obtained by restricting the sample to the shorter period.

The results in this table are quite similar to those in Brumback and Rice (1998): they show strong evidence of differences in the levels along the cycle, and they do not find any evidence of differences in the period between days 2 and 8.

However, there is an important difference between all the analyses in Brumback and Rice (1998) and our analyses. The analyses carried out in Brumback and Rice (1998) are based on the Gaussianity hypotheses. We have applied the Kolmogorov–Smirnov test to the 23 projections of the full data, and, on applying the FDR correction, we find that we can reject the null hypotheses stating that the initial sample was produced by a Gaussian distribution (the obtained p-value was $14 \cdot 10^{-13}$). This suggests that the analyses in Brumback and Rice (1998) could be flawed.

Table 6 p-values obtained with the random ANOVA for urinary progesterone data

Interval	Projections	group	group: subj
days -8 to 14	$k = 23$	$2.96 \cdot 10^{-6}$	$2.3 \cdot 10^{-6}$
days 2 to 8	$k = 7$	.5820	.0003

5 A comparison with MANOVA

In this section we compare our procedure to some usual MANOVA tests. We make the comparison by analyzing one of the examples proposed in `R` (R Development Core Team 2007), as an illustration of the `manova` procedure. The results are similar to those of the other proposed example.

As stated, this example was selected to illustrate the `manova` procedure which is based on Gaussianity assumptions. Thus, we have decided to assume this hypothesis too and employ the usual one-dimensional ANOVA test to analyze the projections.

The data appear in Krzanowski (1988, p. 81) and are related to the production of plastic film. These data are three-dimensional. They served to analyze how some characteristics of the plastic film ("tear", "gloss", and "opacity") vary depending on two factors ("rate" and "additive"), each with two levels ("Low" and "High"). Five measurements were taken under each set of production conditions.

Thus, we are in a three-dimensional, two-way MANOVA problem with two levels in each factor and five observations in each cell. Therefore, we will assume the model (2) with $Y_i^{r,s} = 0$, with t restricted to belonging to the finite set $\{1, 2, 3\}$, and $\mathcal{H}$ being the Euclidean space $\mathbb{R}^3$ with the usual norm and scalar product.

In this example we have tried $k = 3, 5, 15, 30$ random directions to see the effect of this number in this very low-dimensional problem. Table 7 contains the p-values obtained with the application of the random ANOVA and with the Pillai test. We see that both factors (especially the rate) are found to be significant and also that there is no statistical evidence on the existence of interaction between them. Moreover, the p-values obtained with the random ANOVA when $k = 3$ are similar to those provided by the Pillai test.

The role of k here looks similar to the one we found in Sect. 3: if we are in a situation with low p-values (as when testing the factor rate), the power of the test increases with k. However, if we are in a situation with large or, even moderate p-values (as when testing the interaction or the rate), then, the power decreases with k. Obviously, we are interested in detecting those situations in which the p-value is moderate, and we are not particularly interested into lowering the p-value as much as possible in those situations in which the p-value is already low. This is the reason for advising that k be taken lower or equal to the dimension of the analyzed data.

Table 7 p-values obtained with the random ANOVA and usual MANOVA for Krzanowski's data

	Number of projections				Pillai
	$k = 3$	$k = 5$	$k = 15$	$k = 30$	test
Rate	.0039	.0029	.0011	.0013	.0030
Additive	.0271	.0301	.0391	.0520	.0247
Interaction	.2561	.2846	.4869	.5160	.3018

6 Discussion

In this paper a procedure to test complicated linear designs in the functional setting is proposed. The procedure is random since it is based on the replacement of the given functional data by randomly chosen one-dimensional projections.

The consideration of a number of projections is proposed as a way to stabilize the results, using the False Discovery Rate to correct the obtained p-values.

The results obtained in the simulations carried out in Sect. 3 and the analysis of the real data sets in Sect. 4 cause us to be optimistic about its capability to discover the existence of influence of factors, interactions, covariables, ..., the only limitation being the existence or not of one-dimensional methods able to handle the model under consideration.

It is also noticeable that the method proposed produces similar results to those obtained with usual MANOVA tests. We also wish to stress the fact that the flexibility of the procedure makes it easy to deal with those situations in which the usual assumptions are not fulfilled because, through the projections, it is possible to apply a one-dimensional ANOVA test (like that proposed in Brunner et al. 1997) in which the assumptions are less stringent.

An important possibility (not analyzed in this paper) is that, in those cases in which a null hypothesis is rejected, the one-dimensional projections can be useful to determine the way in which the means depart from the null hypotheses, because usual one-dimensional ANOVA-related procedures allow us to identify, for instance, the subpopulations responsible for the null hypothesis failing. Also, those procedures can be applied to estimate the value of the departure in each random direction.

Moreover, the proposed method can be applied to every random process in which we have a Hilbert space that the trajectories belong to. Thus, it is not required that the random functions be defined on an interval or even on a one-dimensional subset. However, as stated in the introduction, the results in Cuevas and Fraiman (2009) also allow us to eliminate the assumption that the process belongs to a separable Hilbert space in our results.

Concerning the practical issues to apply our method, there are some questions that we would like to mention here:

1. According to Theorem 2.1, the random vectors can be chosen with any Gaussian distribution with nondegenerate one-dimensional marginals. In our simulations we have used Brownian motions and Ornstein–Uhlenbeck processes with no noticeable differences between them. However, an interesting open problem is the existence or not of an optimal distribution to chose the random projections depending on some particular alternatives which the researcher could have in mind.
2. With respect to the exact number of random projections to be used, our trials have shown that $k = 30$ is high enough in general and the gain in power obtained when, for example, $k = 60$ is used is negligible. So, our recommendation is to select $k = 30$ and just in those few cases where the practitioner is in doubt, select a higher number of projections.
3. An usual practical problem is to decide what to do when we have trajectories with some missing observations. We propose two possibilities: The first one is to

employ some smoothing procedure to fill in the gaps and, then, to compute the random projections of the filled observations.

The second possibility is to take advantage of the fact that the data we have at hand are, in every case, discrete approximations of continuous observations. For instance, let us assume that we have some observations $\mathbf{X}_i(t_h^i)$, $h = 1, \ldots, H_i$ and $i = 1, \ldots, n$, where $t_h^i \in [0, T]$. Let $v \in L_2[0, T]$ be a randomly chosen vector with unit norm. Theoretically, we should compute

$$\langle \mathbf{X}_i, v \rangle = \int_0^T \mathbf{X}_i(t) v(t)\, dt, \quad i = 1, \ldots, n,$$

which, in practice, is impossible. Therefore, we can make the approximations

$$\langle \mathbf{X}_i, v \rangle \approx \sum_{h=1}^{H_i} \mathbf{X}_i\big(t_h^i\big) v\big(t_h^i\big)\big(t_h^i - t_{h-1}^i\big), \quad i = 1, \ldots, n,$$

where we take $t_0^i = 0$. Notice that the procedure will work as long as those approximations are good; this does not require that the sets $\{t_h^i, h = 1, \ldots, H_i\}$ be identical. In fact, they may be even disjoint.

Acknowledgements The authors want to thank two anonymous referees whose comments improved considerably the paper.

References

Abramovich F, Angelini C (2005) Testing in mixed-effects FANOVA models. J Stat Plan Inference 136:4326–4348

Abramovich F, Antoniadis A, Sapatinas T, Vidakovic B (2004) Optimal testing in a fixed-effects functional analysis of variance model. Int J Wavelets Multiresolut Inf Process 2:323–349

Angelini C, Vidakovic B (2002) Some novel methods in wavelet data analysis: Wavelet ANOVA, F-test shrinkage and Γ-minimax wavelet shrinkage. Istituto per le Applicazioni del Calcolo "Mauro Picone." Rapporti Tecnici RT 256/02

Antoniadis A, Sapatinas T (2007) Estimation and inference in functional mixed-effects models. Comput Stat Data Anal 51:4793–4813

Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. J R Stat Soc Ser B 57:289–300

Benjamini Y, Yekutieli D (2001) The control of the false discovery rate in multiple testing under dependency. Ann Stat 29:1165–1188

Brumback BA, Rice JA (1998) Smoothing spline models for the analysis of nested and crossed samples of curves. J Am Stat Assoc 93:961–994

Brunner E, Dette H, Munk A (1997) Box-type approximations in nonparametric factorial designs. J Am Stat Assoc 92:1494–1502

Cuesta-Albertos JA, Nieto-Reyes A (2007) An infinite dimensional depth: The random Tukey depth. Technical Report

Cuesta-Albertos JA, Nieto-Reyes A (2008) The random Tukey depth. Comput Stat Data Anal 52:4979–4988

Cuesta-Albertos JA, del Barrio T, Fraiman R, Matrán C (2007a) The random projection method in goodness of fit for functional data. Comput Stat Data Anal 51:4814–4831

Cuesta-Albertos JA, Fraiman R, Galves A, García J, Svarc M (2007b) Identifying rhythmic classes of languages using their sonority: a Kolmogorov–Smirnov approach. J Appl Stat 34:749–761

Cuesta-Albertos JA, Fraiman R, Ransford T (2007c) A sharp form of the Cramér–Wold theorem. J Theor Probab 20:201–209
Cuesta-Albertos JA, Cuevas A, Fraiman R (2008) On projection-based tests for spherical and compositional data. Stat Comput 19:367–380
Cuevas A, Fraiman R (2009) On depth measures and dual statistics. A methodology for dealing with general data. J Multivar Anal 100:753–766
Cuevas A, Febrero M, Fraiman R (2004) An ANOVA test for functional data. Comput Stat Data Anal 47:111–122
Fan J (1996) Tests of significance based on wavelet thresholding and Neyman's truncation. J Am Stat Assoc 91:674–688
Fan J, Lin SK (1998) Test of significance when data are curves. J Am Stat Assoc 93:254–261
Faraway JJ (1997) Regression analysis for a functional response. J R Stat Soc B 64:887–898
Ferraty F, Vieu P (2006) Nonparametric functional data analysis: theory and practice. Springer series in statistics. Springer, New York
Ferraty F, Vieu P, Viguier-Pla S (2007) Factor-based comparison of groups of curves. Comput Stat Data Anal 51:4903–4910
Fujikoshi Y, Himeno T, Wakaki H (2004) Asymptotic results of a high dimensional MANOVA and power comparisons when the dimension is large compared to the sample size. J Jpn Stat Soc 34:19–26
Gu C (2002) Smoothing spline ANOVA models. Springer, New York
Guo W (2002) Inference in smoothing spline analysis of variance. J R Stat Soc B 64:887–898
Krzanowski WJ (1988) Principles of multivariate analysis. A user's perspective. Oxford University Press, Oxford
R Development Core Team (2007) R: a language and environment for statistical computing, Vienna, Austria. ISBN 3-900051-07-0, http://www.R-project.org
Ramsay JO, Silverman BW (2002) Applied functional data analysis. Methods and case studies. Springer, New York
Ramsay JO, Silverman BW (2005) Functional data analysis, 2nd edn. Springer, New York
Schott JR (2007) Some high-dimensional tests for a one-way ANOVA. J Multivar Anal 98:1825–1839
Shohat JA, Tamarkin JD (1963) The problem of moments. American Mathematical Society, Providence