



SAVE: An R Package for the Statistical Analysis of Computer Models

Jesús Palomo
Rey Juan Carlos University

Rui Paulo
Universidade de Lisboa

Gonzalo García-Donato
Universidad de
Castilla-La Mancha

Abstract

This paper introduces the R package **SAVE** which implements statistical methodology for the analysis of computer models. Namely, the package includes routines that perform emulation, calibration and validation of this type of models. The methodology is Bayesian and is essentially that of Bayarri, Berger, Paulo, Sacks, Cafeo, Cavendish, Lin, and Tu (2007). The package is available through the Comprehensive R Archive Network. We illustrate its use with a real data example and in the context of a simulated example.

Keywords: R package, computer models, calibration, emulation, Bayesian analysis, Gaussian processes.

1. The analysis of complex computer models

Complex computer models are implementations of sophisticated mathematical models that aim at reproducing a particular real process. The R (R Core Team 2014) package **SAVE** (statistical analysis and validation engine; Palomo, García-Donato, and Paulo 2015) implements statistical methodology developed for the analysis of this type of models, which is based on Craig, Goldstein, Seheult, and Smith (1996), Kennedy and O'Hagan (2001), Kennedy, O'Hagan, and Higgins (2002), Higdon, Kennedy, Cavendish, Cafeo, and Ryne (2004), and most directly on Bayarri *et al.* (2007). The package is available through the Comprehensive R Archive Network at <http://CRAN.R-project.org/package=SAVE>.

The following aspects of the statistical analysis of a computer model are addressed in **SAVE**:

- *Emulation.* A crucial characteristic of these models is that they are often computationally very demanding and a single run may take several minutes to complete. It is then important to produce fast approximations to the output of these models, and these approximations are referred to as emulators.

- *Calibration.* Computer models usually depend on a vector of unknown inputs that needs to be specified before the model can be run. Calibration refers to the process of determining estimates of these calibration parameters based on field observations of the real process.
- *Validation.* Ultimately, we want to assess the degree to which the computer model is an effective surrogate for the real process. We do so by producing predictions of reality and associated tolerance bounds, which measure the degree of the accuracy of said predictions.

Related R packages include **BACCO** (Hankin 2005) and the suite of **Dice** packages: **DiceKriging** and **DiceOptim** (Roustant, Ginsbourger, and Deville 2012), **DiceEval** (Dupuy and Helbert 2014) and **DiceDesign** (Franco, Dupuy, Roustant, Damblin, and Iooss 2014). **DiceKriging** computes estimates of Gaussian process parameters, and our package takes advantage of its functionalities, while **DiceOptim** is dedicated to optimization of complex computer models based on kriging models. **DiceDesign** facilitates the construction of various space-filling designs for computer experiments. **DiceEval** tackles the problem of validation but follows an approach that is quite different from ours; in particular, it is not Bayesian.

BACCO is a package that is similar to ours in its goals. However, it implements the methodology in Kennedy and O’Hagan (2001) which, although Bayesian, is distinct from that of Bayarri *et al.* (2007), the one we implement (**BACCO** also implements the methods in Kennedy and O’Hagan 2000, which is a topic we do not cover). The most important distinction between **BACCO** and **SAVE** is that we explore the posterior distribution of the parameters of the statistical model using simulation-based techniques, namely Markov chain Monte Carlo, whereas **BACCO** relies either on analytical or on numerical integration. As a consequence, with **SAVE** one can for instance explore the posterior distribution of calibration parameters, which often have a physical meaning, and take advantage of all the benefits that come with simulation-based inference. This will be illustrated in Section 4.

We should note that the **SAVE** package relies on C code to perform computer intensive calculations. Additionally, in order to maintain numerical stability as much as possible, we make from our own C code extensive calls to numerical routines written in Fortran, notably those available from **BLAS** and **LAPACK**.

The rest of this paper is organized as follows. In the next section, we describe the statistical methodology. In Section 3, we describe the structure of the package, and in Section 4 we illustrate its use in the context of a real example. In Section 5, we use it in a situation where the truth is known. Technical details are, whenever possible, relegated to the appendices.

2. Introducing the statistical framework

This section is subdivided into three subsections, each corresponding to a different stage of our modeling/fitting process.

2.1. Stage I: Emulation

Denote the output of the computer model by $y^M(\mathbf{x}, \mathbf{v})$, where \mathbf{x} is a vector of controllable inputs and \mathbf{v} is a vector of unknown calibration and/or tuning parameters in the model. We

have access to the data obtained by evaluating the computer model at a design set consisting of N points $D^M = \{(\mathbf{x}_1, \mathbf{v}_1), \dots, (\mathbf{x}_N, \mathbf{v}_N)\}$. We denote by $\mathbf{y}^M = (y^M(\mathbf{x}_1, \mathbf{v}_1), \dots, y^M(\mathbf{x}_N, \mathbf{v}_N))$ the vector of model evaluations, and we refer to these data as computer model data, or simply as model data.

A preliminary but central question in the analysis is the construction of an emulator of the computer model, that is, a method to produce estimates of the output at untested configurations along with an associated measure of uncertainty. This is stage I of the analysis of a computer model. For this task, we follow the popular strategy (cf. Sacks, Welch, Mitchell, and Wynn 1989; Kennedy and O’Hagan 2000; Bayarri *et al.* 2007) of using a Gaussian process-based response-surface approximation to the model output. This approach results in that, conditional on \mathbf{y}^M and on a set of parameters specifying the Gaussian process, $y^M(\cdot)$ follows a Gaussian process with mean and covariance functions which are available in closed form. The approach that **SAVE** currently implements for emulating $y^M(\cdot)$ estimates the unknown parameters by maximum likelihood using the R package **DiceKriging** (Roustant *et al.* 2012).

Analytic expressions for the mean and covariance functions can be found in Appendix A, along with a description of the parameters $\boldsymbol{\theta}^{M,L}$ and $\boldsymbol{\theta}^{M,C}$ which specify, respectively, the linear trend and the covariance kernel of the Gaussian process prior. In practical terms, the output of the computer model at a set of untested configurations (given \mathbf{y}^M and the parameter estimates) follows a multivariate normal distribution with known mean vector and covariance matrix.

2.2. Stage II: Calibration

Field data consists of noisy observations of the real process, possibly with replicates. To be more precise, we have a design set $D^F = \{\mathbf{x}_1^*, \dots, \mathbf{x}_n^*\}$ and we observe

$$y_k^F(\mathbf{x}_j^*) = y^R(\mathbf{x}_j^*) + \varepsilon_{kj}, \quad j = 1, \dots, n; \quad k = 1, \dots, n_j \quad (1)$$

where $y^R(\cdot)$ represents the real process and ε_{kj} are independent and identically distributed $N(0, 1/\lambda^F)$ random variables. Notice that the field data may contain replicates, that is, independent measurements of the experiment using the same configuration of the controllable inputs. We denote the complete set of field observations by \mathbf{y}^F .

Calibrating a model stands for finding estimates of the vector of calibration parameters based on field observations of the real phenomenon. This is achieved by postulating a statistical model relating the output of the model and the real phenomenon which introduces the notion of model discrepancy (cf. Craig, Goldstein, Seheult, and Smith 1997; Kennedy and O’Hagan 2001; Goldstein 2010), namely,

$$y^R(\mathbf{x}) = y^M(\mathbf{x}, \mathbf{u}) + b(\mathbf{x}), \quad (2)$$

where $b(\mathbf{x})$ stands for the bias or discrepancy function, and \mathbf{u} is the unknown value of the calibration vector \mathbf{v} which we are ultimately interested in estimating. In Figure 1 we have included a diagram which schematically characterizes the framework we have just described.

The approach that this package implements is (partially) Bayesian and therefore requires the specification of a prior for all the unknown quantities that appear in the statistical model, namely, $b(\cdot)$, \mathbf{u} , and λ^F . In line with Bayarri *et al.* (2007), we specify these priors in a fashion that requires very little input from the user. An exception is the prior on \mathbf{u} which should reflect

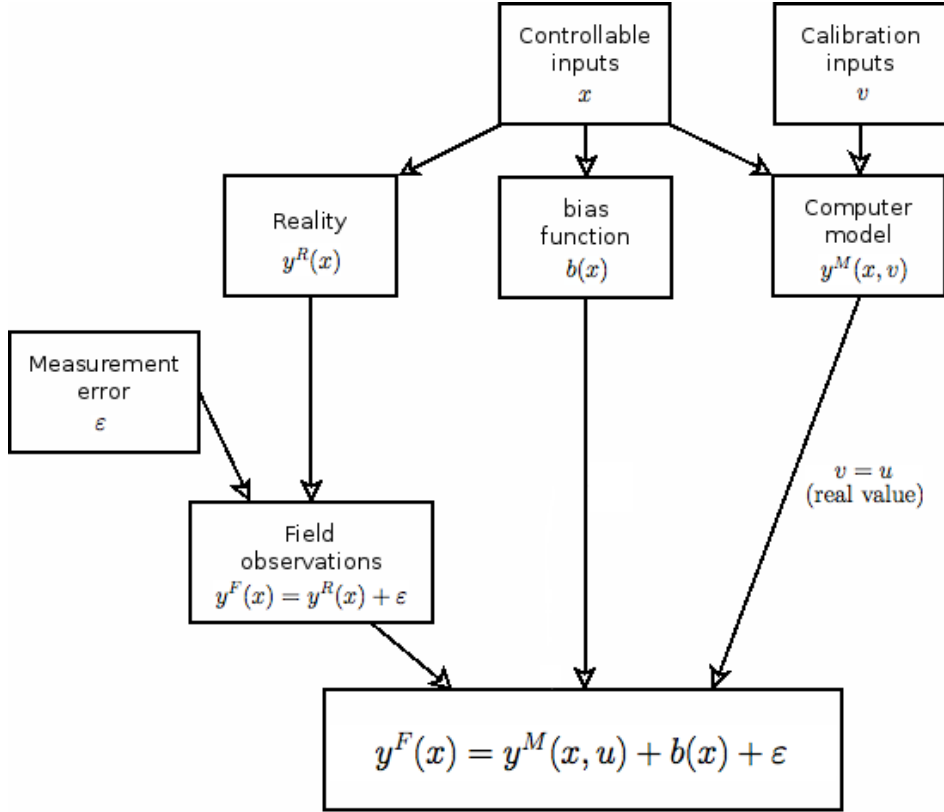


Figure 1: Schematic representation of the framework.

expert opinion about the calibration parameter. Details of the prior specification are available in Appendix B. Let $(\mathbf{u}, \lambda^b, \boldsymbol{\beta}^b, \lambda^F) \equiv (\mathbf{u}, \boldsymbol{\theta}^F)$ denote the vector of unknown parameters at this stage of the analysis, which we refer to as stage II. (The bias function $b(\cdot)$ gets a Gaussian process prior; λ^b denotes the precision and the vector $\boldsymbol{\beta}^b$ controls the correlation structure of this process.)

The posterior distribution $\pi(\mathbf{u}, \boldsymbol{\theta}^F \mid \mathbf{y}^M, \mathbf{y}^F)$ is obtained using Markov chain Monte Carlo (MCMC) methods and is ultimately represented by a sample of correlated draws, which we denote by $\{(\mathbf{u}_i, \boldsymbol{\theta}_i^F), i = 1, \dots, M\}$. Details on the sampling method used are described in Appendix C.

2.3. Stage III: Validation

Once the posterior distribution of all the unknowns in the statistical model is obtained, we can proceed to the validation step. For that, we need to select a set of configurations for the controllable inputs, D^V – e.g., $D^V = \{\mathbf{x}_1^\dagger, \dots, \mathbf{x}_r^\dagger\}$ – at which we are interested in assessing the accuracy of the computer model. For a particular value \mathbf{u} of the calibration parameter, let $D_{\mathbf{u}}^V = \{(\mathbf{x}_1^\dagger, \mathbf{u}), \dots, (\mathbf{x}_r^\dagger, \mathbf{u})\}$ be the design that results from augmenting each of the configurations in D^V with the vector \mathbf{u} for the calibration parameters. In general, we denote by $h(D)$ the vector that results from evaluating the function h at the elements of D .

We must obtain draws from the distribution

$$\begin{aligned} f(y^M(D^V), b(D^V) \mid \mathbf{y}^M, \mathbf{y}^F) &= \\ &= \int f(y^M(D_{\mathbf{u}}^V), b(D^V) \mid \mathbf{y}^M, \mathbf{y}^F, \mathbf{u}, \boldsymbol{\theta}^F) \pi(\mathbf{u}, \boldsymbol{\theta}^F \mid \mathbf{y}^M, \mathbf{y}^F) d\mathbf{u} d\boldsymbol{\theta}^F, \end{aligned} \quad (3)$$

where the notation D^V reflects that the \mathbf{u} 's are integrated out (with respect to the posterior). These are obtained by drawing the vectors $\mathbf{y}_i^M, \mathbf{b}_i$ from $f(y^M(D_{\mathbf{u}_i}^V), b(D^V) \mid \mathbf{y}^M, \mathbf{y}^F, \mathbf{u}_i, \boldsymbol{\theta}_i^F)$ for every $(\mathbf{u}_i, \boldsymbol{\theta}_i^F)$ in the previously constructed MCMC sample drawn from the posterior distribution. (The distribution of $y^M(D_{\mathbf{u}}^V), b(D^V) \mid \mathbf{y}^M, \mathbf{y}^F, \mathbf{u}, \boldsymbol{\theta}^F$ is multivariate normal, cf. Bayarri *et al.* 2007, for further details.)

Having obtained the samples $\{(\mathbf{y}_i^M, \mathbf{b}_i), i = 1, \dots, M\}$, we can compute several quantities which will aid us in the validation task:

- The bias-corrected prediction of the real process (i.e., reality) at D^V

$$\hat{\mathbf{y}}^R = \frac{1}{M} \sum_{i=1}^M (\mathbf{y}_i^M + \mathbf{b}_i).$$

- The tolerance bars measuring the accuracy of $\hat{\mathbf{y}}^R$ as a predictor of $y^R(D^V)$ are computed as follows: pick $\gamma \in (0, 1)$; then, compute $\boldsymbol{\tau} = (\tau(\mathbf{x}) : \mathbf{x} \in D^V)$ such that $(1 - \gamma) \times 100\%$ of the samples satisfy

$$|\hat{\mathbf{y}}^R - (\mathbf{y}_i^M + \mathbf{b}_i)| \leq \boldsymbol{\tau},$$

with the inequality interpreted in a component-wise fashion. We can then state that, for each $\mathbf{x} \in D_V$, $\mathbb{P}(|y^R(\mathbf{x}) - \hat{\mathbf{y}}^R(\mathbf{x})| < \boldsymbol{\tau}(\mathbf{x}) \mid \mathbf{y}^F, \mathbf{y}^M) = 1 - \gamma$.

- The pure-model prediction of reality at D^V is obtained by selecting an estimate of \mathbf{u} , $\hat{\mathbf{u}}$, say, which can be, for instance, its posterior mean or median. Then, output of the model at $D_{\hat{\mathbf{u}}}^V$ is computed by either actually running the model or by exercising the emulator. The mean vector of the emulator (Equation 6, in Appendix A, with unknown parameters replaced by estimates), can then be used as an estimate of the output of the model. Denote the pure-model prediction of reality by $\hat{\mathbf{y}}^M$.
- The tolerance bars measuring the accuracy of $\hat{\mathbf{y}}^M$ as a predictor of $y^R(D^V)$ are computed in a similar fashion: pick $\gamma \in (0, 1)$; then, compute $\boldsymbol{\tau} = (\tau(\mathbf{x}) : \mathbf{x} \in D^V)$ such that $(1 - \gamma) \times 100\%$ of the samples satisfy

$$|\hat{\mathbf{y}}^M - (\mathbf{y}_i^M + \mathbf{b}_i)| \leq \boldsymbol{\tau},$$

with the inequality interpreted in a component-wise fashion. We can then state that, for each $\mathbf{x} \in D_V$, $\mathbb{P}(|y^R(\mathbf{x}) - \hat{\mathbf{y}}^M(\mathbf{x})| < \boldsymbol{\tau}(\mathbf{x}) \mid \mathbf{y}^F, \mathbf{y}^M) = 1 - \gamma$.

- It is also possible to estimate the bias associated with the pure-model prediction, $\mathbf{b}_{\hat{\mathbf{u}}} = y^R(D^V) - \hat{\mathbf{y}}^M$: samples from its posterior predictive distribution can be obtained by computing $\{\mathbf{y}_i^M + \mathbf{b}_i - \hat{\mathbf{y}}^M\}$ so that a point estimate is $\hat{\mathbf{b}}_{\hat{\mathbf{u}}} = \hat{\mathbf{y}}^R - \hat{\mathbf{y}}^M$ and $1 - \gamma$ pointwise credible intervals can be determined by computing the associated $\gamma/2 \times 100\%$ and $(1 - \gamma/2) \times 100\%$ sample quantiles.

3. An overview of SAVE

There are three high-level functions in **SAVE** which allow the user to perform all the tasks described in the previous section: `SAVE`, `bayesfit` and `validate`. In general,

- `SAVE` creates an object of the class ‘`SAVE`’ and essentially sets up the problem by filling out a number of slots of this object. The data structures are set up (more on this below). Maximum likelihood calculations are performed using **DiceKriging** (Roustant *et al.* 2012). These calculations serve two purposes: fit the emulator of the computer model, and aid in the specification of the prior of θ^F (cf. Appendix B).
- `bayesfit` produces a sample from the posterior distribution of the parameters λ^F , λ^b and \mathbf{u} (β^b is fixed at an estimate throughout the analysis, cf. Appendix B). It takes as an argument an object of the class ‘`SAVE`’ that has been created using `SAVE`, and returns a copy of this object but with several additional slots filled out. These slots pertain to the MCMC sample obtained.
- `validate` ultimately produces the bias-corrected prediction, the pure-model prediction, associated tolerance bounds and estimated bias function for any set of configurations for the controllable inputs and a posterior estimate of the vector of calibration inputs. It performs this task based on the information contained in an object of the class ‘`SAVE`’ which has been produced by a call of the function `bayesfit`.

Two additional functions are available in **SAVE**, but these can be considered low-level routines. The function `predictcode` produces i.i.d. draws from the emulator evaluated at a set of design points, along with its mean vector and covariance matrix. It expects as an argument an object of the class ‘`SAVE`’. The function `predictreality` expects as an argument an object of the class ‘`SAVE`’ which has been produced by the function `bayesfit`, i.e., containing an MCMC sample. It outputs draws from the distribution in (3) for a design set of configurations for the controllable inputs of the problem. These functions are internally called by `validate`, but can be utilized to further explore the problem.

The output of each of these functions can be appropriately summarized by calls to `print`, `summary`, `plot` and `show`, where suitable methods are provided for the objects returned.

The data is handled in the following way. **SAVE** assumes that there are two R data frames loaded: one containing all the field data, and another containing all the model data. The response and the input variables in the designs are identified by the names associated with each of these data frames, so they must be consistent. For illustrative purposes, consider the synthetic example in Table 1. The data frame at the top of the figure contains the field data and the one at the bottom contains the model data. If we decide to analyze the problem where the response is the variable `expand`; controllable inputs are `temp`, `press` and `weight`; and calibration inputs are `delta1` and `shift`, then we must call the function `SAVE` with arguments, `field.data = field`, `model.data = model`, `response.name = "expand"`, `controllable.names = c("temp", "press", "weight")`, `calibration.names = c("delta1", "shift")`. Notice that not all the columns present in the data frames are incorporated in this analysis.

If in order to obtain the field data the controllable inputs (if any are present) have not been varied but have instead been fixed at a single value, you have to specify `controllable.names = NULL`.

Data frame called <code>field</code> containing the field data					
	temp	press	weight	expand	height
1	35.1	2.65	600	84.1	5.1
2	35.1	2.75	600	90.6	9.2
3	35.1	2.65	600	80.4	6.1
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	30.3	2.65	700	83.4	7.1

Data frame called <code>model</code> containing the model data							
	temp	press	weight	delta1	shift	expand	delta2
1	23.2	2.12	629.1	0.22	-1.1	99.1	0.22
2	43.7	2.11	711.0	0.84	-0.9	70.1	0.83
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
N	34.4	2.12	700.1	0.49	-0.8	67.6	0.33

Table 1: Synthetic example: The data frame at the top of the figure is called `field` and contains the field data; the data frame at the bottom of the figure is called `model` and contains the model data.

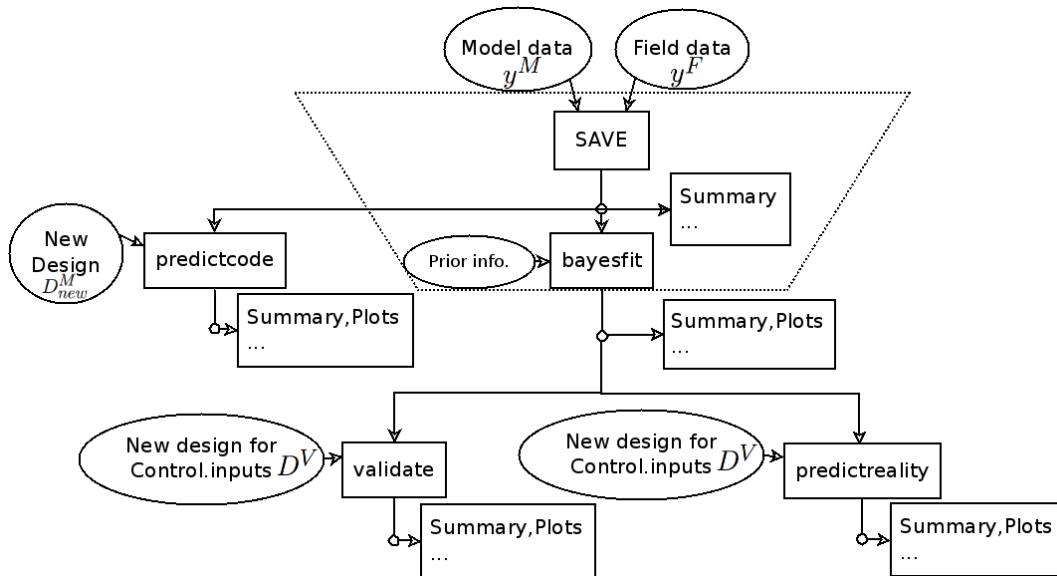


Figure 2: Schematic representation of the structure of the package.

We have included in Figure 2 a diagram which schematically describes the structure of the package, its different parts and interconnections.

When putting together the package, we have taken into account the fact that many of the calculations implemented by the functions described above can be computationally demanding and numerically unstable. For example, it is a well-known fact that covariance matrices arising in Gaussian process likelihoods can be close to singular even for relatively small sample sizes. As such, when evaluating the likelihood we never invert matrices and instead perform the Cholesky decomposition (which is also useful to compute the determinant) and solve

triangular systems. In some situations, e.g., when performing predictions, we resort to the pivotal Cholesky decomposition because the regular Cholesky decomposition tends to break down. All the linear algebra calculations are performed with the help of the **BLAS** and **LAPACK** numerical libraries, which are known to be quite stable and accurate. The size of the field and model datasets determines the size of the matrices that need to be manipulated, and hence represent the biggest bottleneck when it comes to the speed of the calculations. As we have already stressed, in order to make the package as fast as possible we have implemented all the computationally intensive computations in C. Nevertheless, large sample sizes could certainly be an issue. So far, the package has only been tested in problems where the sample sizes can be considered small, say below 50.

The area of statistical analysis of computer models is quite active, both on the methodological and on the computational side. For instance, the problem of ill-conditioned covariance matrices we referred to above has been recently addressed in [Ranjan, Haynes, and Karsten \(2011\)](#). These authors propose an iterative regularization approach to construct the emulators that considerably improves on standard methods (see also the corresponding implementation in the **GPfit** package, [MacDonald, Ranjan, and Chipman 2015](#)). Given the modular structure of **SAVE**, it is conceptually easy to incorporate this and other advances into the package, and we plan on doing that in future versions.

One other aspect of **SAVE** that could be improved pertains to the kernel used when modeling the correlation structure of the Gaussian processes involved in the analysis. Appendices [A](#) and [B](#) make it clear that at the emulation stage **SAVE** is currently using a separable power exponential kernel, and that for the bias function the package uses a separable Gaussian correlation kernel. It is our goal to, in the future, allow the user to choose from a variety of kernels, as it is the case, for instance, in **DiceKriging**. The greater difficulty with that is at the MCMC level. We have to make sure that the strategy we follow with the Gaussian kernel still works with other kernels. Most likely it will not, and what we really need is to develop a more general strategy. In the meanwhile, the Gaussian/power exponential kernels are still among the most commonly used in computer modeling, although it is well-known that they possess several shortcomings (both methodological and computational).

4. A real data example

In this section we illustrate the use of the package with an analysis of a real example. It is the so-called spotweld example originally analyzed in [Bayarri *et al.* \(2007\)](#). We refer the interested reader to that paper for complete details on the application.

In [Figure 3](#) you can find a schematic representation of the spot welding process. Two sheets of metal of a particular thickness (**thickness**) are compressed by two electrodes under a certain applied load (**load**). Electric current of certain magnitude (**current**) is passed through said electrodes and the heat produced by the current flow causes the surfaces under pressure to melt. After cooling, a weld nugget is formed and as a result the two metal sheets are welded together. The scientists are interested in the diameter of this nugget (**diameter**).

Included in the package are two data sets, **spotweldfield** and **spotweldmodel** that pertain, respectively, to field experiments and computer model experiments associated with this problem. After loading the package (`library("SAVE")`), the data frames can be loaded using the commands

	<code>data.frame</code>	Response (y)	Controllable (\boldsymbol{x})	Calibration (\boldsymbol{v})
Model names	<code>spotweldmodel</code>	<code>\$diameter</code>	<code>\$current</code> , <code>\$load</code> , <code>\$thickness</code>	<code>\$tuning</code>
Field names	<code>spotweldfield</code>	<code>\$diameter</code>	<code>\$current</code> , <code>\$load</code> , <code>\$thickness</code>	–

Table 2: Correspondence between the notation in Section 2 and the variable names in the `spotweld` example.

```
R> data("spotweldfield", package = "SAVE")
```

and

```
R> data("spotweldmodel", package = "SAVE")
```

Notice that the columns of the data frames are appropriately named, and that the computer model features an additional input, named `tuning`, which is a calibration input related to contact resistance.

We start this analysis by setting up the problem. This is accomplished by creating `sw` – an object of the class ‘SAVE’ – using the function `SAVE`:

```
R> sw <- SAVE(response.name = "diameter",
+   controllable.names = c("current", "load", "thickness"),
+   calibration.names = "tuning", field.data = spotweldfield,
+   model.data = spotweldmodel, mean.formula = ~ 1),
+   bestguess = list(tuning = 4.0))
```

Here we are specifying which columns correspond to the response and which correspond to the controllable and the calibration inputs. Additionally, we are also

- setting the mean function of the Gaussian process approximation to the output of the computer model as a constant (with the option `mean.formula = ~ 1`) and
- providing an estimate of the vector of calibration inputs as a list (`bestguess`), which will be used in specifying the prior for $\boldsymbol{\theta}^F$ – see Appendix B for further details.

The object `sw` has now been created and several of its slots have been filled. The easiest way of accessing that information is by means of the command `summary(sw)`.

Since at this point the emulator has been fitted, we could now use the function `predictcode` to predict the output of the computer model at a set of input configurations – more on this later. Instead, we will proceed to fit the Bayesian model (2) relating reality to computer model. To do so, we use the function `bayesfit`:

```
R> set.seed(0)
R> swbayes <- bayesfit(object = sw, prior = c(uniform("tuning", upper = 8,
+   lower = 0.8)), n.iter = 20000, n.burnin = 100, n.thin = 2)
```

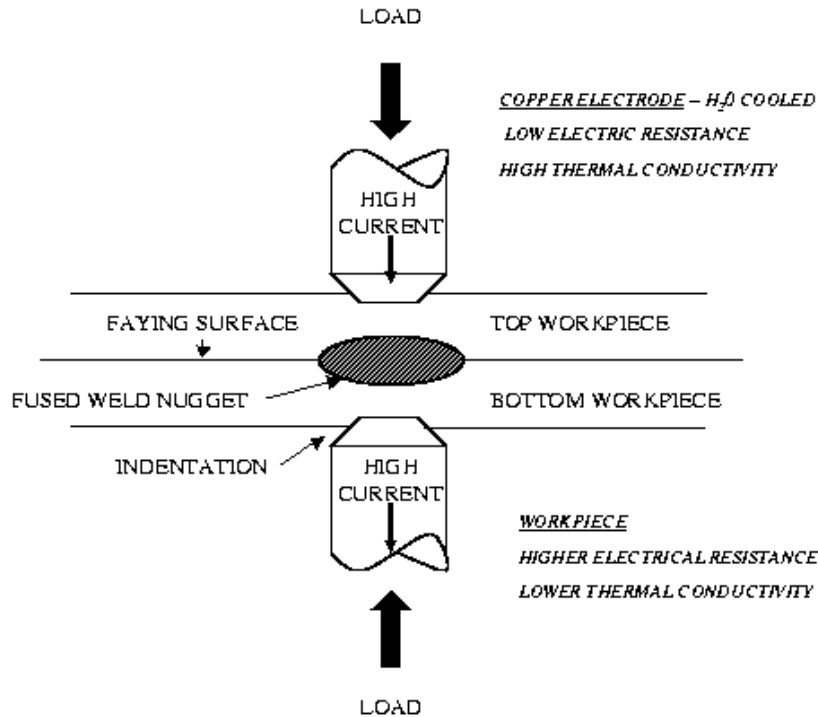


Figure 3: Schematic representation of the spotwelding process.

We have now created a new object – `swbayes` – but instead we could have updated the object we had already created, `sw`. Notice that we need to specify a prior for the calibration parameter `tuning`, and also options pertaining to the MCMC algorithm. We have set a uniform prior for the calibration parameter, 20000 iterations for the MCMC sampling with a burn-in period of 100 iterations, and a thinning of 2. Other options were left at the corresponding defaults – cf. Appendix C for additional details.

The object `swbayes` now contains not only the estimates previously computed using `SAVE` but also a sample from the posterior distribution of the calibration parameter `tuning` and of the field and bias precisions, λ^F and λ^b , respectively. To display this information we can again resort to the command `summary(swbayes)`. We can however also plot the samples obtained:

- `plot(swbayes, option = "trace")` will give us the traceplots;
- `plot(swbayes, option = "calibration")` will produce an histogram of the posterior samples of the calibration parameters and corresponding priors (see Figure 4) and,
- `plot(swbayes, option = "precision")` will plot histograms of the posterior samples of the field and bias precisions. These histograms also include a plot of the prior and of the estimates that are used in constructing the prior (see Figure 5).

We can of course access the raw data. By running the command `slotNames(swbayes)` we get a description of the names of all the slots of the object `swbayes`, and it is then clear how to obtain the MCMC samples: `swbayes@mcmc.samples`.

After fitting the Bayesian model, we can finally produce predictions of reality and also assess the quality of pure-model and bias-corrected predictions of reality. The package `SAVE` pro-




Figure 4: Posterior distribution of the calibration parameter `tuning` (spotweld example). The dotted line corresponds to the prior used.




Figure 5: Posterior distribution of the precisions (spotweld example). The dashed vertical lines indicate the estimates of the parameters; the dash-dotted lines indicate the priors used.

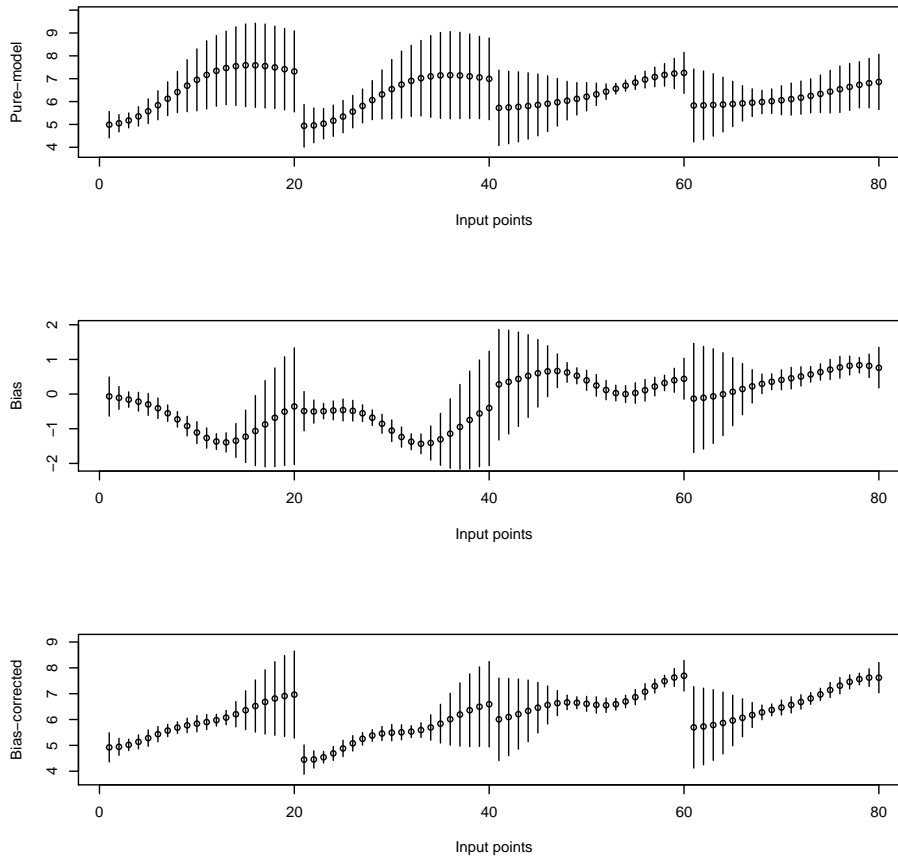


Figure 6: Default validation plot for the spotweld example. The plot at the top contains the pure-model predictions (circles) at each of the input configurations in `xnew` and associated 90% tolerance bounds. The middle plot depicts an estimate of the bias of these pure-model estimates. The bottom plot contains the bias-corrected prediction (circles) and associated 90% tolerance bounds.

vides a very convenient function that performs all these calculations. To illustrate its use in the present example, please consider the following R code:

```
R> load <- c(4.0, 5.3)
R> curr <- seq(from = 20, to = 30, length = 20)
R> g <- c(1, 2)
R> xnew <- expand.grid(current = curr, load = load, thickness = g)
R> set.seed(0)
R> valsw <- validate(object = swbayes, newdesign = xnew,
+   calibration.value = "mean", n.burnin = 100)
```

Above, we first construct the design of controllable inputs at which we want to predict reality. For four combinations of load and thickness of the metal plates, we want to predict the weld diameter as a function of current. As a remark, note that in the situation where `controllable.names = NULL`, the input `newdesign` must be left unspecified. This comment also applies when using the function `predictreality` in the same context.

Pure-model predictions

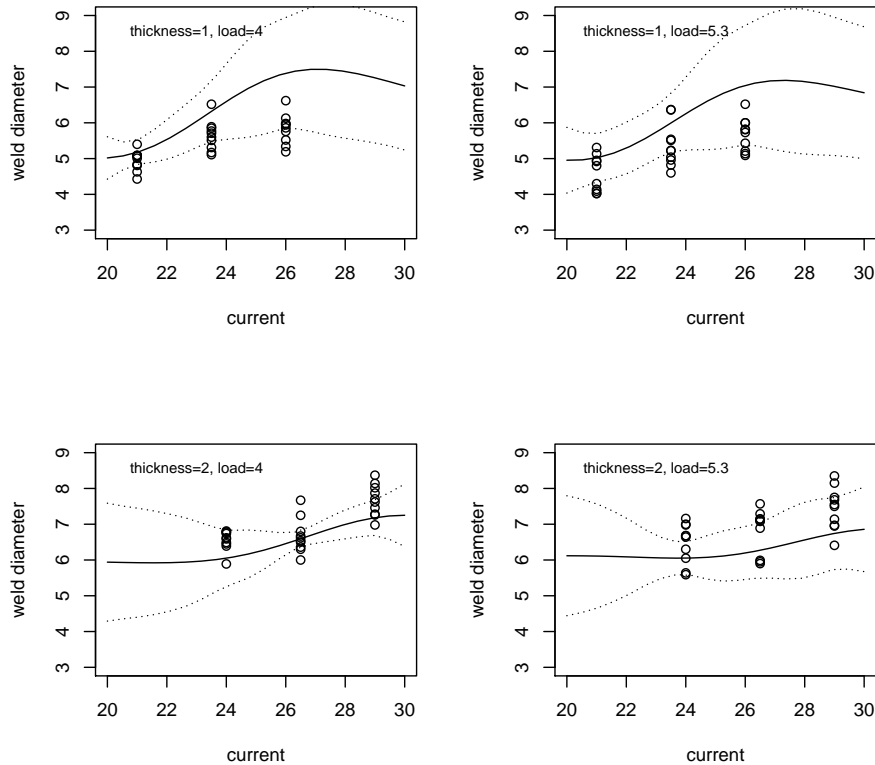


Figure 7: Customized validation plot for the spotweld example. Here we can see the pure-model predictions for 4 different combinations of load and thickness as a function of current. The dotted lines indicate the 90% tolerance bounds and the circles represent the observed field data corresponding to that particular combination of controllable inputs.

Regarding the pure-model prediction, we are setting the calibration parameters at the corresponding posterior mean. Note that we could set the calibration parameters at a particular value if we chose to. We would do it by specifying the input `calibration.value` to be a named `data.frame` containing those values.

The resulting object `valsw` contains a slot named `validate` where a matrix is stored. This matrix contains as columns the pure-model prediction of reality (`pure.model`) and associated tolerance bound (`tau.pm`); the estimate of the bias associated with the pure-model prediction and pointwise credible interval for that unknown (`bias.Lower` and `bias.Upper`); the bias-corrected prediction of reality (`bias.corrected`) and associated tolerance bounds (`tau.bc`). This information can be accessed using `summary(valsw)` but can also be plotted (`plot(valsw)`) – you can find the plot in Figure 6. Depending on the problem, this default plot will not always be the most appropriate way of displaying the estimates. Nevertheless, since we have access to the estimates, we can certainly construct customized plots for any problem at hand. As an example, in Figure 7 we can find a plot of the pure-model prediction and associated tolerance bounds as a function of current for the 4 different combinations of load and thickness. The circles correspond to the appropriate field observations.

It is also easy to incorporate a linear trend in the mean function for the emulator. This would be accomplished simply by starting the analysis with the command

```
R> sw2 <- SAVE(response.name = "diameter",
+   controllable.names = c("current", "load", "thickness"),
+   calibration.names = "tuning", field.data = spotweldfield,
+   model.data = spotweldmodel,
+   mean.formula = ~ current + load + thickness),
+   bestguess = list(tuning = 4.0))
```

which produces an new ‘SAVE’ object, `sw2`, and then we can proceed as before.

As stated above, `validate` is a function which calls two low-level functions, `predictcode` and `predictreality`. There may be situations where one must directly call these functions. We illustrate this in what follows.

Imagine that one is interested in understanding how the nugget `diameter` varies with `current`. One might assess this variation by looking at the derivative of `diameter` with respect to `current`. Let us do that for `thickness = 1` and `load = 4`. We start by predicting reality at an equally-spaced grid of `current` values between 20 and 30:

```
R> load <- 4
R> g <- 1
R> curr <- seq(from = 20, to = 30, length = 80)
R> xnew <- expand.grid(current = curr, load = load, thickness = g)
R> set.seed(0)
R> prsw <- predictreality(object = swbayes, newdesign = xnew)
```

The object `prsw` has slots named `"modelpred"` and `"biaspred"` where the draws from (3), which we have denoted in Section 2.3 by $(\mathbf{y}_i^M, \mathbf{b}_i)$, are stored. We now obtain the corresponding derivatives with respect to `current` of each of these draws, to finally obtain draws from the corresponding derivative of reality:

```
R> delta <- diff(curr)[1]
R> model <- prsw@modelpred
R> dmodel <- diff(t(model)) / delta
R> bias <- prsw@biaspred
R> dbias <- diff(t(bias)) / delta
R> dreal <- dmodel + dbias
```

These draws can be summarized by computing the corresponding mean and tolerance bounds as explained in Section 2.3. This is plotted in Figure 8.

Additionally, we can use `predictcode` to obtain draws from the emulator evaluated at a posterior estimate of the calibration input, `tuning`, say, `tuning = 3.2`. This allows us to produce the pure-model estimate of the derivative:

```
R> u <- 3.2; load <- 4; g <- 1; curr <- seq(from = 20, to = 30, length = 80)
R> xnewpure <- expand.grid(current = curr, load = load, thickness = g,
+   tuning = u)
```




Figure 8: Bias-corrected prediction of the derivative of `diameter` with respect to `current` – the estimate is the solid line; the dashed lines are the 90% tolerance bounds. The pure-model prediction is the dash-dotted line.

```
R> set.seed(0)
R> pmsw <- predictcode(object = swbayes, newdesign = xnewpure,
+   n.iter = 20000)
R> puremodel <- pmsw@samples
R> dpuremodel <- diff(t(puremodel)) / delta
```

The mean of the samples in `dpuremodel` is the pure-model prediction of the derivative. This is plotted in Figure 8. Notice how the two estimates, pure-model and bias-corrected, are in this case even qualitatively different.

5. A simulated example

In this section, we are going to illustrate the use of the package in a simulated data example: model and reality are known, and so is the true value of the calibration parameter and of the field data precision. Essentially the same setup has already been utilized in Bayarri *et al.* (2007) and in Bayarri (2010) to illustrate similar points.

In this example, we have a calibration parameter v whose true value is fixed at $u = 1.7$. The

controllable input is x ; reality is given by

$$y^R(x) = 3.5 \exp(-u \cdot x) + 1.5. \quad (4)$$

The computer model is given by

$$y^M(x, v) = 5 \exp(-v \cdot x), \quad (5)$$

but we are going to treat it as we would have a computationally expensive model: to obtain the model data, we evaluate it at a 6×6 equally-spaced grid on the (x, v) -space, which is set at $[0, 3] \times [0.5, 2]$. We use these data, available with the package as `synthmodel`, to produce the emulator; to load the data frame, use the command

```
R> data("synthmodel", package = "SAVE")
```

The field data consists of 3 replicates at each point in an equally-spaced grid of 10 points for x spanning the interval $[0.11, 3.01]$. The measurement error is $N(0, 0.3^2)$, that is $\lambda^F = 1/0.3^2 \approx 11.11$. These data are also available with the package; use the following command to load the corresponding data frame:

```
R> data("synthfield", package = "SAVE")
```

The R code to produce the ‘SAVE’ object associated with this analysis is

```
R> synth <- SAVE(response.name = "y", controllable.names = "x",
+   calibration.names = "v", field.data = synthfield, ,
+   model.data = synthmodel, mean.formula = ~ 1 + x,
+   bestguess = list(v = 1.5))
```

and the MCMC samples are obtained with

```
R> set.seed(0)
R> synth <- bayesfit(object = synth,
+   prior = uniform(var.name = "v", lower = 0, upper = 3), n.iter = 20000)
```

The plot of the posterior distribution on the calibration parameter can be found in Figure 9, and the posterior on the precisions is in Figure 10. Note how the posteriors include the true values of v and λ^F .

We next produce the bias-corrected prediction of reality at an equally-spaced grid of 25 points for x in the interval $[0.05, 3.05]$. For that, we have used the code

```
R> xnew <- data.frame(x = seq(from = 0.05, to = 3.05, length = 25))
R> valsynth <- validate(object = synth, newdesign = xnew, n.burnin = 100)
```

The bias-corrected prediction of reality is the solid black line in Figure 11. The dashed lines correspond to the tolerance bars, and the red solid line is the actual curve we are trying to estimate; the stars represent the simulated field data.

Using the fact that we know y^M , it is easy to find the least squares estimate of u , i.e., the value of v that minimizes the sum of the squared distance between the field observations




Figure 9: Posterior distribution of the calibration input for the simulated example.




Figure 10: Posterior distribution of the precisions for the simulated example.




Figure 11: Bias-corrected prediction for the simulated example. The solid black line corresponds to the bias-corrected prediction, the dashed lines represent the tolerance bars, and the stars are the simulated field observations. The red line corresponds to the actual reality curve, which is known in this case. The blue and green curves correspond, respectively, to the computer model evaluated at the posterior mean and at the least squares estimate of the calibration parameter.

and the model output. That estimate here turns out to be $\tilde{u} = 0.63$. This is a common strategy in calibration: not including a bias term. It has been argued extensively, starting with Kennedy and O'Hagan (2001) and Craig *et al.* (1997), that this approach often leads to over-fitting and to unreasonable estimates for the calibration parameter. In this example, we again see evidence of that. In Figure 11, the green curve corresponds the true model output evaluated at \tilde{u} . The blue curve represents the true model output evaluated at the posterior mean, which turns out to be $\hat{u} = 1.58$. The inclusion of the bias prevents over-fitting and allows for a more reasonable estimate of the calibration parameter. Additionally, it is clear that the bias-corrected prediction is a much better estimate of the true reality curve than the pure-model prediction (computer model evaluated at the posterior mean) and than the computer model evaluated at the least squares estimate.

Dedication

To Susie Bayarri, *in memoriam*.

Acknowledgments

This research was funded by the Portuguese National Funds through FCT – Fundação para a Ciência e a Tecnologia, project PTDC/MAT/105349/2008 and by the Ministerio de Economía y Competitividad, Spain, under grant MTM2013-42323-P. An initial version of this software was produced while the authors were postdoctoral researchers at the National Institute of Statistical Sciences (NISS), and at the time their research was supported in part by grants from General Motors and the National Science Foundation (Grant DMS-0073952). M.J. Bayarri (Valencia), James O. Berger (Duke), Jerry Sacks (NISS) were actively involved in that research, and so were J.A. Cafeo, J. Cavendish, C.H. Lin and J. Tu from General Motors.

References

- Bayarri MJ (2010). “A Methodological Review of Computer Models.” In MH Chen, P Müller, D Sun, K Ye, D Dey (eds.), *Frontiers of Statistical Decision Making and Bayesian Analysis: In Honor of James O. Berger*. Springer-Verlag.
- Bayarri MJ, Berger JO, Paulo R, Sacks J, Cafeo JA, Cavendish J, Lin CH, Tu J (2007). “A Framework for Validation of Computer Models.” *Technometrics*, **49**(2), 138–154.
- Craig PS, Goldstein M, Seheult AH, Smith JA (1996). “Bayes Linear Strategies for History Matching of Hydrocarbon Reservoirs.” In JM Bernardo, JO Berger, AP Dawid, D Heckerman, AFM Smith (eds.), *Bayesian Statistics 5*. Oxford University Press, London.
- Craig PS, Goldstein M, Seheult AH, Smith JA (1997). “Pressure Matching for Hydrocarbon Reservoirs: A Case Study in the Use of Bayes Linear Strategies for Large Computer Experiments.” In C Gatsonis, JS Hodges, RE Kass, R McCulloch, P Rossi, ND Singpurwalla (eds.), *Case Studies in Bayesian Statistics*, volume III, pp. 36–93.
- Dupuy D, Helbert C (2014). *DiceEval: Construction and Evaluation of Metamodels*. R package version 1.3, URL <http://CRAN.R-project.org/package=DiceEval>.
- Franco J, Dupuy D, Roustant O, Damblin G, Iooss B (2014). *DiceDesign: Designs of Computer Experiments*. R package version 1.6, URL <http://CRAN.R-project.org/package=DiceDesign>.
- Goldstein M (2010). “External Bayesian Analysis of Computer Simulators.” In JM Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, M West (eds.), *Bayesian Statistics 9*. Oxford University Press, London.
- Hankin RKS (2005). “Introducing **BACCO**, An R Bundle for Bayesian Analysis of Computer Code Output.” *Journal of Statistical Software*, **14**(16), 1–21. URL <http://www.jstatsoft.org/v14/i16/>.

- Higdon D, Kennedy MC, Cavendish J, Cafo J, Ryne RD (2004). “Combining Field Data and Computer Simulations for Calibration and Prediction.” *SIAM Journal on Scientific Computing*, **26**(2), 448–466.
- Kennedy MC, O’Hagan A (2000). “Predicting the Output from a Complex Computer Code When Fast Approximations Are Available.” *Biometrika*, **87**(1), 1–13.
- Kennedy MC, O’Hagan A (2001). “Bayesian Calibration of Computer Models.” *Journal of the Royal Statistical Society B*, **63**(3), 425–464.
- Kennedy MC, O’Hagan A, Higgins N (2002). “Bayesian Analysis of Computer Code Outputs.” In CW Anderson, V Barnett, PC Chatwin, AH El-Shaarawi (eds.), *Quantitative Methods for Current Environmental Issues*, pp. 227–243. Springer-Verlag.
- MacDonald B, Ranjan P, Chipman H (2015). “**GPfit**: An R Package for Fitting a Gaussian Process Model to Deterministic Simulator Outputs.” *Journal of Statistical Software*, **64**(12), 1–23. URL <http://www.jstatsoft.org/v64/i12/>.
- Palomo J, García-Donato G, Paulo R (2015). *SAVE: Bayesian Emulation, Calibration and Validation of Computer Models*. R package version 1.0, URL <http://CRAN.R-project.org/package=SAVE>.
- Ranjan P, Haynes R, Karsten R (2011). “A Computationally Stable Approach to Gaussian Process Interpolation of Deterministic Computer Simulation Data.” *Technometrics*, **53**(4), 366–378.
- R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>.
- Roustant O, Ginsbourger D, Deville Y (2012). “**DiceKriging**, **DiceOptim**: Two R Packages for the Analysis of Computer Experiments by Kriging-Based Metamodeling and Optimization.” *Journal of Statistical Software*, **51**(1), 1–55. URL <http://www.jstatsoft.org/v51/i01/>.
- Sacks J, Welch WJ, Mitchell TJ, Wynn HP (1989). “Design and Analysis of Computer Experiments.” *Statistical Science*, **4**(4), 409–423.

A. Details on the emulator

We assume that *a priori* $y^M(\cdot)$ follows a stationary Gaussian process with mean and covariance functions governed by unknown parameters $\boldsymbol{\theta}^{M,L}$ and $\boldsymbol{\theta}^{M,C} = (\lambda^M, \boldsymbol{\alpha}^M, \boldsymbol{\beta}^M)$, respectively. The mean function of the Gaussian process is assumed to be of the form $\boldsymbol{\Psi}^\top(\cdot)\boldsymbol{\theta}^{M,L}$ where $\boldsymbol{\Psi}(\mathbf{z})$ is a specified $k \times 1$ vector function of the input $\mathbf{z} = (\mathbf{x}, \mathbf{v})$ and $\boldsymbol{\theta}^{M,L}$ is a $k \times 1$ vector of unknown regression parameters. This mean function is specified through the argument `mean.formula` of the `SAVE` function. Note that one of the restrictions of `SAVE` is that $\boldsymbol{\Psi}$ can only be a function of \mathbf{x} , the vector of controllable inputs.

The parameter λ^M is the precision (the inverse of the variance) of the Gaussian process and the other parameters $(\boldsymbol{\alpha}^M, \boldsymbol{\beta}^M)$ control the correlation function of the Gaussian process, which we assume to be of the form

$$c^M(\mathbf{z}, \mathbf{z}^*) = \exp\left(-\sum_{j=1}^d \beta_j^M |z_j - z_j^*|^{\alpha_j^M}\right).$$

Here, d is the number of coordinates in $\mathbf{z} = (\mathbf{x}, \mathbf{v})$, the α_j^M are numbers between 0 and 2, and the β_j^M are positive parameters.

After observing \mathbf{y}^M , the conditional posterior distribution of y^M given the hyperparameters, $f(y^M(\cdot) | \mathbf{y}^M, \boldsymbol{\theta}^{M,L}, \boldsymbol{\theta}^{M,C})$, is a Gaussian process with updated mean and covariance functions given respectively by

$$E[y^M(\mathbf{z}) | \mathbf{y}^M, \boldsymbol{\theta}^{M,L}, \boldsymbol{\theta}^{M,C}] = \boldsymbol{\Psi}^\top(\mathbf{z})\boldsymbol{\theta}^{M,L} + \mathbf{r}_z^\top(\boldsymbol{\Gamma}^M)^{-1}(\mathbf{y}^M - \mathbf{X}\boldsymbol{\theta}^{M,C}) \quad (6)$$

$$\text{COV}[y^M(\mathbf{z}), y^M(\mathbf{z}^*) | \mathbf{y}^M, \boldsymbol{\theta}^{M,L}, \boldsymbol{\theta}^{M,C}] = \frac{1}{\lambda^M} c^M(\mathbf{z}, \mathbf{z}^*) - \mathbf{r}_z^\top(\boldsymbol{\Gamma}^M)^{-1}\mathbf{r}_{z^*}, \quad (7)$$

where $\mathbf{r}_z^\top = \frac{1}{\lambda^M} (c^M(\mathbf{z}, \mathbf{z}_1), \dots, c^M(\mathbf{z}, \mathbf{z}_N))$, \mathbf{X} is the matrix with rows $\boldsymbol{\Psi}^\top(\mathbf{z}_1), \dots, \boldsymbol{\Psi}^\top(\mathbf{z}_N)$, and $\boldsymbol{\Gamma}^M$ is a matrix with (i, j) entry given by $c^M(\mathbf{z}_i, \mathbf{z}_j)/\lambda^M$

To obtain an emulator for y^M , we replace in the formulae above the unknown parameter values by the corresponding maximum likelihood estimates.

B. Details on the stage II prior

The stage II unknowns are $b(\cdot)$, the bias function, \mathbf{u} , the vector of calibration parameters, and λ^F , the precision of the field measurement error. The prior for \mathbf{u} is specified using expert knowledge. Currently, the distributional choices are limited to uniform and normal, this last one truncated to an interval. This enters the function `bayesfit` through argument `prior`.

The prior for the bias is a stationary zero-mean Gaussian process with covariance λ^b and correlation function given by

$$c^b(\mathbf{x}, \mathbf{x}^*) = \exp\left(-\sum_{j=1}^p \beta_j^b |x_j - x_j^*|^2\right).$$

Here, p is the number of coordinates in \mathbf{x} , and the β_j^b are positive parameters. Let $\boldsymbol{\beta}^b = (\beta_1^b, \dots, \beta_p^b)$. We need to specify a prior for $\boldsymbol{\theta}^F = (\lambda^b, \boldsymbol{\beta}^b, \lambda^F)$, and we do so in a nearly

automatic fashion as follows: we start by selecting a best guess for the vector of calibration parameters, denoted by $\tilde{\mathbf{u}}$, which is the argument `bestguess` in function `SAVE`. Then, using the emulator, we predict the output of the computer model at $D_{\tilde{\mathbf{u}}}^F$, denoted $\mathbf{y}^M(D_{\tilde{\mathbf{u}}}^F)$. Next, treat $\mathbf{y}^F - \mathbf{y}^M(D_{\tilde{\mathbf{u}}}^F)$ as a realization of a Gaussian process with a nugget, namely as a realization of a multivariate normal with mean zero and covariance matrix $c^b(D^F)/\lambda^b + I/\lambda^F$, to get maximum likelihood estimates $\hat{\lambda}^b, \hat{\beta}^b, \hat{\lambda}^F$. Then,

- β^b is fixed throughout the analysis at $\hat{\beta}^b$;
- λ^b and λ^F are independent exponentially distributed quantities centered at a multiple of the corresponding estimates, $\hat{\lambda}^b$ and $\hat{\lambda}^F$. This multiple is set through the parameter `mcmcMultmle` in `bayesfit` and its purpose is to allow the user to specify a prior which is relatively flat in the region where the posterior distribution accumulates.

The function `SAVE` computes these maximum likelihood estimates (with the help of the package `DiceKriging`; Roustant *et al.* 2012) and stores these in the slot "mle" of the corresponding object of class 'SAVE'.

C. Details on the MCMC sampling

Full details on the sampling mechanism can be found in Bayarri *et al.* (2007). The algorithm implemented in `SAVE` requires very little input apart from the necessary length of the simulation, burn-in and thinning numbers. This is because all unknowns are sampled directly from their full conditionals with the exception of the vector of calibration parameters. This vector is sampled using a Metropolis-Hastings step, for which the user needs to decide on three aspects:

- The proposal distribution is a mixture between the prior and a local move. The user needs to specify the probability of sampling from the prior, which is argument `prob.prop` of the `bayesfit` function.
- The algorithm performs a fixed number of Metropolis-Hastings steps before deciding on a move; the user must set this number, and this is argument `nMH` of the `bayesfit` function.
- The package implements two alternative methods: `method = 2` specifies that computer model and bias are analytically integrated out before sampling \mathbf{u} ; this is the default and preferred method. If `method = 1`, these vectors are not integrated out.

Affiliation:

Jesús Palomo
 Department of Business Administration (Finance)
 Rey Juan Carlos University
 28032 Madrid, Spain
 E-mail: jesus.palomo@urjc.es

Rui Paulo
Department of Mathematics and CEMAPRE
ISEG, Universidade de Lisboa
1200-781 Lisboa, Portugal
E-mail: rui@iseg.ulisboa.pt

Gonzalo García-Donato
Department of Economics and Finance
Universidad de Castilla-La Mancha
02071 Albacete, Spain
E-mail: gonzalo.garciadonato@uclm.es