Emulating dynamic non-linear simulators using Gaussian processes

Hossein Mohammadi [a,b,*], Peter Challenor [a,b], Marc Goodfellow [a,b,c,d]

[a] College of Engineering, Mathematics and Physical Sciences, University of Exeter, Exeter, United Kingdom
[b] EPSRC Centre for Predictive Modelling in Healthcare, University of Exeter, Exeter, United Kingdom
[c] Centre for Biomedical Modelling and Analysis, University of Exeter, Exeter, United Kingdom
[d] Living Systems Institute, University of Exeter, Exeter, United Kingdom

ABSTRACT

The dynamic emulation of non-linear deterministic computer codes where the output is a time series, possibly multivariate, is examined. Such computer models simulate the evolution of some real-world phenomenon over time, for example models of the climate or the functioning of the human brain. The models we are interested in are highly non-linear and exhibit tipping points, bifurcations and chaotic behaviour. However, each simulation run could be too time-consuming to perform analyses that require many runs, including quantifying the variation in model output with respect to changes in the inputs. Therefore, Gaussian process emulators are used to approximate the output of the code. To do this, the flow map of the system under study is emulated over a short time period. Then, it is used in an iterative way to predict the whole time series. A number of ways are proposed to take into account the uncertainty of inputs to the emulators, after fixed initial conditions, and the correlation between them through the time series. The methodology is illustrated with two examples: the highly non-linear dynamical systems described by the Lorenz and van der Pol equations. In both cases, the predictive performance is relatively high and the measure of uncertainty provided by the method reflects the extent of predictability in each system.

© 2019 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Computer models, e.g. numerical simulators, are sophisticated mathematical representations of some real-world phenomenon implemented in computer programs (O'Hagan, 2006). Such models are widely used in many fields of science and technology to aid our understanding of physical processes or because conducting physical experiments is too costly, highly time-consuming or even impossible in some cases (Sacks et al., 1989). Often, simulators are available as commercial packages and the underlying functions are unknown to the user. In most applications, it is crucial to understand the sensitivity of model outputs to variation or uncertainty in inputs (O'Hagan, 2006). Performing such quantitative studies requires a large number of simulation runs, see for example Ferrat et al. (2018). It becomes impractical if each simulation run is time-consuming.

Emulators, also known as *surrogate models*, *metamodels* or *response surfaces* (Kleijnen, 2009) provide a "fast" approximation of complex simulation models using a limited number of training runs. The most popular classes of emulators are

neural networks, splines, regression models, etc. We refer the reader to Jin et al. (2000), Chen et al. (2006) and Forrester and Keane (2009) for more information on different types of emulators and their properties. Among the diverse types of emulators, Gaussian processes (GPs) have become increasingly popular over the last two decades in the field of the design and analysis of computer experiments (Sacks et al., 1989; Santner et al., 2003; Jones and Johnson, 2009). Also known as Kriging, especially in geostatistics (Cressie, 1993), GPs have been effectively used in many real-world applications including wireless communication (Schwaighofer et al., 2004), metallurgy (Bailer-Jones et al., 1997), biology (Swain et al., 2016; Kalaitzis and Lawrence, 2011; Lawrence et al., 2006), environmental science (Lee et al., 2011; Challenor, 2004), and sensor placements (Krause et al., 2008).

There are several reasons for the popularity of GPs. Firstly, they can be used to fit any smooth (with different degrees of smoothness), continuous function thanks to the variety of covariance kernels available (Neal, 1998). See Section 2 for more details on kernels. Secondly, GPs are non-parametric models, i.e., no strong assumptions about the form of the underlying function are required, unlike polynomial regression (Rawlings et al., 2006). Moreover, the prediction performance of GPs is comparable to (if not better than) other methods such as neural networks (Rasmussen, 1997; Kamath et al., 2018). The limit of a single layer neural network as the number of neurons tends to infinity is a Gaussian process (MacKay, 1998; Neal, 1996). The main advantage of GPs is that they provide not only a mean predictor but also a quantification of the associated uncertainty. This uncertainty reflects the prediction accuracy and can serve as a criterion to enhance prediction capability of the emulator (Jin et al., 2002).

This paper deals with the emulation of dynamic computer models that simulate phenomena evolving with time. The output of a dynamic simulator is a time series for each input. The time series represents the values of the state variables at each time step. Such models are often expressed by a system of differential equations.

Dynamic simulators appear in many applications. For instance, Stommel's box model (Stommel, 1961) simulates the evolution of temperature and salinity to determine the ocean density. In Birrell et al. (2011) a dynamic model is developed whose output is a time series of general practice consultations for the 2009 A/H1N1 influenza epidemic in London. Since this model is computationally expensive, a GP emulator is developed for calibration (Farah et al., 2014). Another real-world example of dynamic computer models is presented in Kuczera et al. (2006) where a saturated path hydrology model simulates the movement of water at catchment scales. In Williamson and Blaker (2014) large climate models with time series output that exhibit chaotic behaviour are emulated using Bayesian dynamic linear model Gaussian processes. We refer to Conti and O'Hagan (2010) for more examples on such simulators.

There are many different proposed approaches for emulating dynamic simulators. According to Reichert et al. (2011), these approaches can be divided into four categories:

1. One method is to use a multi-output emulator for predicting time series output (Conti and O'Hagan, 2010). In this case, the dimension of output space is $q = T$ where T is the number of time steps the simulator is run for. However, when T is large, the efficiency will reduce or may cause numerical problems. In addition, prediction is possible only for a fixed time horizon and one needs to repeat the prediction procedure for different time horizons. Building q separate emulators for q time points has the drawback of losing some information, as the correlation between various outputs (which we expect to be high) is not considered. Such correlation is taken into account in Fricker et al. (2013) and Rougier (2008) within multivariate emulators. However, as mentioned earlier, multivariate emulators are not efficient when the simulator's output is highly multivariate. A common approach to alleviate this problem is to perform dimension reduction techniques on the output space such as principal components analysis (Higdon et al., 2008) and wavelet decomposition (Bayarri et al., 2007). A potential drawback of these techniques is that we may lose information by leaving out some components.
2. A second approach is to treat time as an additional model input (Kennedy and O'Hagan, 2001). Gaussian processes have a computational complexity of $\mathcal{O}(n^3)$ where n is the number of sample points. Considering time as an extra parameter increases the computational cost to $\mathcal{O}(n^3 + T^3)$ using a separable covariance function (Plumlee, 2014). As a result, the method can be burdensome when T is large. Moreover, it is shown in Conti and O'Hagan (2010) that the performance of multi-output emulators exceeds emulators with time as an extra input.
3. One-step ahead emulations are another example in which the basic assumption is that the model output at a given time depends only on the previous output in time. Then, the transition function needs to be approximated. This method is reported to be efficient (Conti et al., 2009).
4. Finally, methods have been described that combine stochastic dynamic models with innovation terms in the form of GPs. For example, in Liu and West (2009) a time-varying auto regression time series, which is a type of dynamic linear model, combined with GPs is used to emulate a dynamic computer code in a hydrological system. Similar work is carried out in Williamson and Blaker (2014) with application to climate models.

We propose a methodology based on iterative one-step ahead predictions. Given that simulating a long time series from the system is computationally expensive, our strategy is to emulate the flow map of the system over a short period of time. Then, we use the estimated flow map, which is computationally cheaper, to approximate the whole time series in an iterative way similar to the work in Conti et al. (2009). However, our method is different from that work in several ways. First, we build separate emulators to approximate each state variable that allows to have different covariance properties. Second, we propose a methodology to incorporate the uncertainty of inputs to the emulators at time t and the correlation between them through the time series, starting from a fixed initial condition. This is an important aspect of one-step ahead

predictions because input to the GP model is uncertain after the first time step. Besides, it can be used as a criterion to estimate the predictability horizon of an emulator. Third, we emulate the flow map which is novel to our knowledge and has not yet been pursued. The last item is the main difference between our work and the method presented in Girard et al. (2003) where multi-step ahead forecasting of a given time series is performed by doing repeated one-step ahead prediction conditional on the observed time series.

2. Gaussian processes as emulators

Let f be the underlying function of an expensive simulator we want to approximate (or predict) defined as $f : \mathcal{X} \longmapsto \mathcal{F}$. Here, $\mathcal{X} \subset \mathbb{R}^d$ and $\mathcal{F} \subset \mathbb{R}^q$ are the input and output space respectively. We further assume that f is a "black-box" function; there is no analytic expression for it and additional information such as gradients are not available. Also throughout this paper we assume the simulator to be *deterministic* (vs. *stochastic*); i.e. if it is run twice with the same inputs, the outputs will be identical.

A GP defines a distribution over functions which can be regarded as a generalization of the normal distribution to infinite dimensions. Formally, a GP indexed by $\mathcal{X}$ is a collection of random variables $\{Z_{\mathbf{x}} : \mathbf{x} \in \mathcal{X}\}$ such that for any $N \in \mathbb{N}$ and any $\mathbf{x}^1, \ldots, \mathbf{x}^N \in \mathcal{X}$, $\left(Z_{\mathbf{x}^1}, \ldots, Z_{\mathbf{x}^N}\right)^{\top}$ follows a multivariate Gaussian distribution (Rasmussen and Williams, 2005). GPs are fully characterized by their mean function $\mu(.)$ and covariance kernel $k(.,.)$, which are defined as

$$\mu : \mathcal{X} \longmapsto \mathbb{R} \; ; \; \mu(\mathbf{x}) = \mathbb{E}[Z_{\mathbf{x}}] \tag{1}$$

$$k : \mathcal{X} \times \mathcal{X} \longmapsto \mathbb{R} \; ; \; k(\mathbf{x}, \mathbf{x}') = \mathbb{C}\text{ov}\left(Z_{\mathbf{x}}, Z_{\mathbf{x}'}\right). \tag{2}$$

The mean function reflects our prior belief about the form of f. That is why $\mu(.)$ is also called the "prior" mean within the Bayesian framework. While $\mu(.)$ could be any function, $k(.,.)$ must be symmetric positive definite. The most commonly used kernel is the squared exponential (SE) which has the form

$$k\left(\mathbf{x}, \mathbf{x}'\right) = \sigma^2 \prod_{l=1}^{d} \exp\left(-\frac{\mid x_l - x_l' \mid^2}{2\theta_l^2}\right). \tag{3}$$

In the above equation, the parameter σ^2 is referred to as the *process variance* and controls the scale of the amplitude of sample paths. The parameter θ_l is called the *characteristic length-scale* and controls the degree of smoothness of sample paths along the coordinate l, $1 \leq l \leq d$. The Matérn family of covariance functions is also widely used (Rasmussen and Williams, 2005). Usually the kernel parameters are unknown and need to be estimated. Choosing appropriate kernel parameters has a huge impact on the accuracy of emulators. Maximum likelihood, cross validation or Bayesian estimation are common methods for this purpose.

Covariance kernels play an important role in GP modelling. They customize the structure of sample paths of GPs. As an example, three different kernels (exponential, Matérn 3/2, and SE, see Rasmussen and Williams, 2005 for more information) and the associated sample paths are illustrated in Fig. 1. While in a process incorporating the SE kernel the sample paths are smooth (infinitely differentiable), they are only continuous (not differentiable) when the exponential kernel is used. Herein, we consider *stationary* covariance kernels that are translation invariant. The value of a stationary kernel depends only on the difference between input vectors. In other words, $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} + \mathbf{h}, \mathbf{x}' + \mathbf{h})$ for any $\mathbf{h} \in \mathbb{R}^d$.

To fit a GP, the true function f is evaluated at n locations $\mathbf{X}_n = \left\{\mathbf{x}^1, \ldots, \mathbf{x}^n\right\}$ with the corresponding outputs (observations) $\mathbf{y} = \left(f(\mathbf{x}^1), \ldots, f(\mathbf{x}^n)\right)^{\top}$. Together, $\mathbf{X}_n$ and $\mathbf{y}$ form the set of *training* samples/data denoted by $D = \{\mathbf{X}_n, \mathbf{y}\}$. Then the conditional distribution of $Z_{\mathbf{x}}$ is calculated as: $\left\{Z_{\mathbf{x}} \mid D : Z_{\mathbf{x}^1} = f(\mathbf{x}^1), \ldots, Z_{\mathbf{x}^n} = f(\mathbf{x}^n)\right\}$. If the mean function $\mu(.)$ is known, the prediction (conditional mean, $m(.)$) and its uncertainty (conditional variance, $s^2(.)$) at a generic location $\mathbf{x}$ are of the form

$$m(\mathbf{x}) = \mu(\mathbf{x}) + \mathbf{k}(\mathbf{x})^{\top}\mathbf{K}^{-1}(\mathbf{y} - \mu(\mathbf{X}_n)) \tag{4}$$

$$s^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x})^{\top}\mathbf{K}^{-1}\mathbf{k}(\mathbf{x}), \tag{5}$$

where $\mathbf{k}(\mathbf{x}) = \left(k(\mathbf{x}, \mathbf{x}^i)\right)_{1 \leq i \leq n}$ is the vector of covariances between the observation at $\mathbf{x}$ and the outputs at the $\mathbf{x}^i$s and $\mathbf{K} = \left(k(\mathbf{x}^i, \mathbf{x}^j)\right)_{1 \leq i,j \leq n}$ denotes the matrix of covariances between sample outputs. Also, $\mu(\mathbf{X}_n)$ is the vector of mean function values at the training samples. The mean predictor obtained by Eq. (4) interpolates the points in the training data. Moreover, the prediction uncertainty vanishes at the training points and grows as we get further from them. An illustrative example is shown in Fig. 2.

3. Emulating dynamical simulators

3.1. One-step ahead emulation: general methodology

We wish to predict the output of a computationally expensive dynamical simulator relying on a d-dimensional autonomous system of ordinary differential equations (ODEs) of which the state variable is given by the real-valued vector
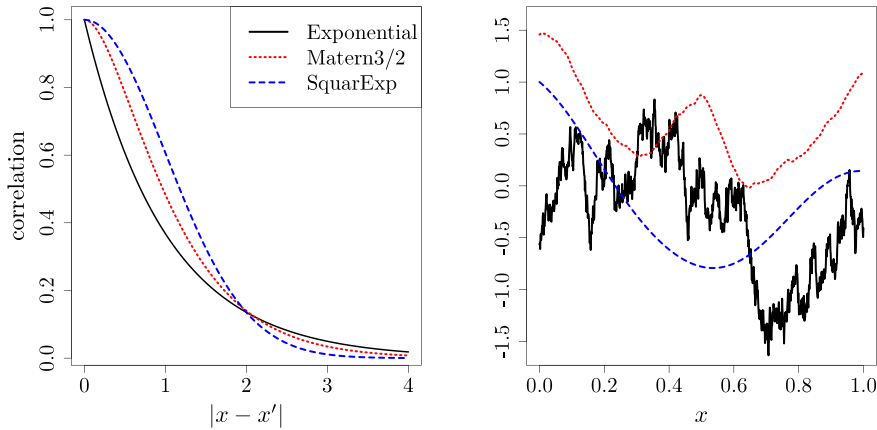
Fig. 1. The structure of GP sample paths is determined by the covariance kernel. Left: Graphs of three (stationary) kernels: exponential (solid), Matérn 3/2 (dotted), and squared exponential (dashed). Right: Three sample paths corresponding to the covariance kernels shown on the left picture. The process with squared exponential kernel is infinitely differentiable, whilst with Matérn 3/2 the process is only once differentiable. The process with exponential kernel is not differentiable.
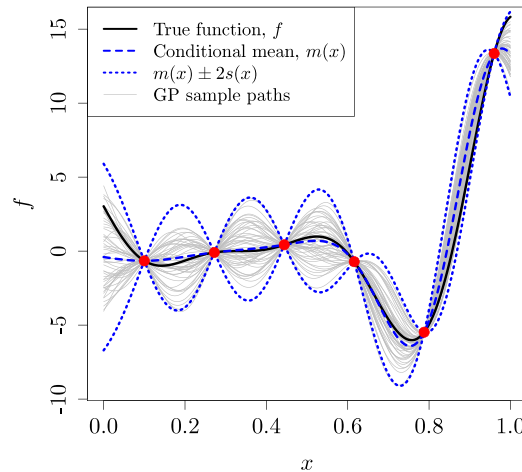

Fig. 2. Gaussian process mean (thick blue dashed line) conditional on 6 training samples (red bullets), which is also known as the conditional mean (denoted by $m(x)$), along with confidence intervals (thick dotted blue lines) equal to $m(x) \pm 2s(x)$. The true function is the thick solid line. The thin lines are 50 sample paths of the GP.

$\mathbf{x}(t) = (x_1(t), \ldots, x_d(t))^\top$. This system gives rise to the flow map $\Phi : \mathbb{R}^d \times \mathbb{R} \longmapsto \mathbb{R}^d$ such that $\mathbf{x}(t_0 + \Delta t) = \Phi(\mathbf{x}(t_0), \Delta t)$ for any $\mathbf{x}(t_0)$ and Δt. We are interested in a "short" *fixed* time step Δt to give $\mathbf{x}(t_1)$ at time $t_1 = t_0 + \Delta t$. Since Δt is fixed, we consider the flow map as a function of $\mathbf{x}(t_0)$ only.

To predict $\mathbf{x}(t)$ over time, we assume that $\Phi(.)$ consists of d components given by

$$\Phi(.) := (f_1(.), \ldots, f_d(.))^\top , \quad f_l : \mathbb{R}^d \longmapsto \mathbb{R}, \ 1 \leq l \leq d, \tag{6}$$

such that each f_l maps $\mathbf{x}(t_0)$ to $x_l(t_1)$ the lth component of $\mathbf{x}(t_1)$. A 2D example is illustrated in Fig. 3 to clarify our assumption. Then, f_ls are treated as black-box functions that are replaced with their emulators denoted by $\hat{f}_l$s which are iteratively used for one-step ahead predictions over the time horizon T.

The training set consists of n initial conditions with the corresponding outputs, which are the solution of the system at time t_1, obtained by running the simulator over the short time horizon Δt. This training set is then used to approximate each function f_l by a GP. The instructions are summarized in Algorithm 1.

Note that in Algorithm 1, only the initial input to the emulators is certain. Thereafter, inputs are actually outputs of the emulators in the previous step. For example, to predict $\mathbf{x}$ at $t_2 = t_0 + 2\Delta t$, the input is $\hat{\mathbf{x}}^*(t_1) = \left(\hat{f}_1(\mathbf{x}(t_0)), \ldots, \hat{f}_d(\mathbf{x}(t_0))\right)^\top$ in which $p\left(\hat{f}_l(\mathbf{x}(t_0))\right) \sim \mathcal{N}\left(m_l(\mathbf{x}(t_0)), s_l^2(\mathbf{x}(t_0))\right), \ 1 \leq l \leq d$, see Eqs. (4) and (5). So, we need to incorporate the input uncertainty in our modelling which is discussed below. Propagating such uncertainty, which is neglected in Conti et al. (2009), results in a more accurate representation of the uncertainty in the emulator over the time horizon.
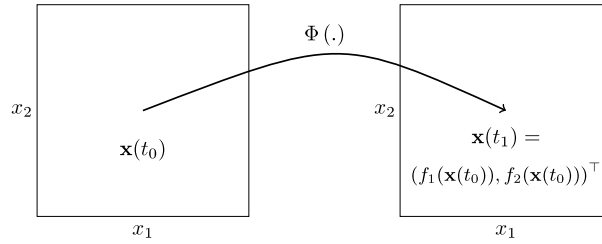
Fig. 3. Left: space of the initial conditions $\mathbf{x}(t_0) = (x_1(t_0), x_2(t_0))^\top$. Right: space of the solution of the system at time $t_1 = t_0 + \Delta t$, i.e. $\mathbf{x}(t_1)$. The flow map $\Phi(.)$ maps from $\mathbf{x}(t_0)$ to $\mathbf{x}(t_1) : \mathbf{x}(t_1) = \Phi(\mathbf{x}(t_0)) = (f_1(\mathbf{x}(t_0)), f_2(\mathbf{x}(t_0)))^\top$ where f_1 and f_2 are unknown functions approximated by GPs.

3.2. Emulation with uncertain input: uncorrelated emulators

GPs with uncertain inputs have been studied in Girard et al. (2003), Candela et al. (2003) and Kuß (2006). Suppose $\mathbf{x}^*$ is drawn from a distribution that has mean $\boldsymbol{\mu}^*$ and variance $\boldsymbol{\Sigma}^*$. The probability distribution of the prediction at $\mathbf{x}^*$ with the GP emulator $\hat{f}$ is determined by

$$p\left(\hat{f}(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*, D\right) = \int p\left(\hat{f}(\mathbf{x}^*)|\mathbf{x}^*, D\right) p(\mathbf{x}^*)\, d\mathbf{x}^*, \tag{7}$$

where $p\left(\hat{f}(\mathbf{x}^*)|\mathbf{x}^*, D\right)$ has a normal distribution whose mean and variance are specified by Eqs. (4) and (5). The integral in (7) is analytically intractable (Girard et al., 2003). However, it can be approximated by different techniques which are divided into two groups: Monte Carlo-based methods and deterministic techniques such as Laplace's approximation. In this work, the former approach is used because it is simple, we only need to sample from $\mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$, and the approximated distribution will converge to the true distribution as the number of samples grows (Girard et al., 2003; Rasmussen, 1997). We refer the reader to MacKay (2002) for more information on the deterministic techniques.

Let $\mathbf{x}^* \sim \mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$, the first and second moments of $p\left(\hat{f}(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*, D\right)$ using the law of iterated expectations and conditional variance are given by

$$\mathbb{E}\left[\hat{f}(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right] = \mathbb{E}_{\mathbf{x}^*}\left[\mathbb{E}_{\hat{f}(\mathbf{x}^*)}\left[\hat{f}(\mathbf{x}^*)|\mathbf{x}^*\right]\right]$$

$$= \mathbb{E}_{\mathbf{x}^*}\left[m(\mathbf{x}^*)\right] \tag{8}$$

$$\mathbb{V}\mathrm{ar}\left[\hat{f}(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right] = \mathbb{E}_{\mathbf{x}^*}\left[\mathbb{V}\mathrm{ar}_{\hat{f}(\mathbf{x}^*)}\left[\hat{f}(\mathbf{x}^*)|\mathbf{x}^*\right]\right]$$

$$+ \mathbb{V}\mathrm{ar}_{\mathbf{x}^*}\left[\mathbb{E}_{\hat{f}(\mathbf{x}^*)}\left[\hat{f}(\mathbf{x}^*)|\mathbf{x}^*\right]\right]$$

$$= \mathbb{E}_{\mathbf{x}^*}\left[s^2(\mathbf{x}^*)\right] + \mathbb{V}\mathrm{ar}_{\mathbf{x}^*}\left[m(\mathbf{x}^*)\right]. \tag{9}$$

Algorithm 1 Emulation of dynamic non-linear computer models

1: Select n samples of initial conditions: $\mathbf{X}_n = \left\{\mathbf{x}^1(t_0), \ldots, \mathbf{x}^n(t_0)\right\}$
2: Run the simulator to obtain $\left\{\mathbf{x}^1(t_1), \ldots, \mathbf{x}^n(t_1)\right\}$
3: **for** $l = 1$ to d **do**
4:     $\mathbf{y} \leftarrow \left(x_l^1(t_1), \ldots, x_l^n(t_1)\right)^\top$
5:     Build the l-th emulator $\hat{f}_l$, based on $D = \{\mathbf{X}_n, \mathbf{y}\}$
6: **end for**
7: Predict over time horizon T given initial condition $\mathbf{x}(t_0)$ as follow

$$\hat{\mathbf{x}}^*(t_1) = \left(\hat{f}_1(\mathbf{x}(t_0)), \ldots, \hat{f}_d(\mathbf{x}(t_0))\right)^\top \quad \triangleright \ \mathbf{x} \ \text{with a superscript} * \text{indicates that it is uncertain.}$$

$$\hat{\mathbf{x}}^*(t_2) = \left(\hat{f}_1\left(\hat{\mathbf{x}}^*(t_1)\right), \ldots, \hat{f}_d\left(\hat{\mathbf{x}}^*(t_1)\right)\right)^\top$$

$$\vdots \qquad\qquad\qquad \vdots$$

$$\hat{\mathbf{x}}^*(t_T) = \left(\hat{f}_1\left(\hat{\mathbf{x}}^*(t_{T-1})\right), \ldots, \hat{f}_d\left(\hat{\mathbf{x}}^*(t_{T-1})\right)\right)^\top$$
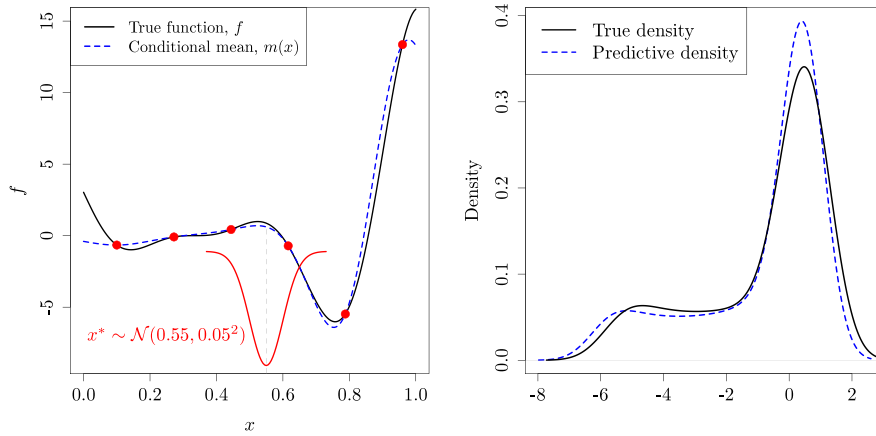

Fig. 4. Estimating the probability distribution of the output given a distribution of inputs using the Monte Carlo method. Left: the model input is uncertain with the probability distribution: $x^* \sim \mathcal{N}(0.55, 0.05^2)$, as shown by the red line. Right: the estimated output distributions for the true function (solid) and the emulator (blue dashed). These estimates are based on $n_{MC} = 10\,000$ samples drawn from x^* that are propagated in $f(.)$ and $m(.)$. Performing such a large number of function evaluations is impractical if f is computationally expensive, but in our case this step is performed using the emulator.

Computing quantities in (8) and (9) is not straightforward because they are functions of the random variable $\mathbf{x}^*$. In this work, $m(\mathbf{x}^*)$ and $s^2(\mathbf{x}^*)$ are approximated using a Monte Carlo (MC) method which relies on samples repeatedly drawn from a probability distribution and statistical analysis to infer the results (Raychaudhuri, 2008). For example, to approximate $\mathbb{E}_{\mathbf{x}^*}[m(\mathbf{x}^*)]$ in Eq. (8), samples are repeatedly drawn from the random variable $\mathbf{x}^*$. Then, they are propagated through the function $m(.)$ defined in Eq. (4). Finally, the desired quantity is approximated using

$$\mathbb{E}_{\mathbf{x}^*}\left[m(\mathbf{x}^*)\right] \approx \frac{1}{n_{MC}} \sum_{i=1}^{n_{MC}} m(\mathbf{x}^{*i}), \tag{10}$$

where n_{MC} denotes the number of MC samples. To shed more light on the MC method, an illustrative example is demonstrated in Fig. 4.

3.3. Emulation with uncertain input: correlated emulators

In the previous section, each element of $\mathbf{x}(t) = (x_1(t), x_2(t), \ldots, x_d(t))^\top$ is emulated separately; d different GP emulators denoted by $\hat{f}_1, \hat{f}_2, \ldots, \hat{f}_d$ are employed independently such that the lth emulator $\hat{f}_l$ emulates the transition function f_l defined as $f_l : \mathbf{x}(t_0) \longmapsto x_l(t_1)$. However, we may lose some information if correlation between emulators is neglected.

Let $\mathbf{x}^* \sim \mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$ be an uncertain input to the d emulators. As a result, $\mathbf{x}^{**} = \left(\hat{f}_1(\mathbf{x}^*), \ldots, \hat{f}_d(\mathbf{x}^*)\right)^\top$ is a random vector whose mean is determined by

$$\boldsymbol{\mu}^{**} = \left(\mathbb{E}\left[\hat{f}_1(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right], \ldots, \mathbb{E}\left[\hat{f}_d(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right]\right)^\top. \tag{11}$$

The elements in $\boldsymbol{\mu}^{**}$ are approximated using (10). Notice that $\mathbf{x}^{**}$ is not necessarily a random normal variable, see Fig. 4 as an example in which the underlying function is highly non-linear. However, we approximate it by a Gaussian which has been used in similar works such as Girard et al. (2003). It is convenient to generate samples from a Gaussian distribution as we use the Monte Carlo method to approximate the unknown quantities. Note that if we use a very small time step relative to the size of the vector field ($\Delta t \to 0$), the change in $\mathbf{x}(t_1)$ under the flow map for Δt is very small. Therefore, the function will be approximately $f_l(\mathbf{x}(t_0)) = \mathbf{x}(t_0)$ (this is the limit) and the assumption of normality on $\mathbf{x}^{**}$ is quite reasonable, see Fig. 5.

Let $\boldsymbol{\Sigma}^{**}$ be the covariance matrix of $\mathbf{x}^{**}$. In order to include the correlation between emulators, $\boldsymbol{\Sigma}^{**}$ must be of the form:

$$\begin{bmatrix} \mathbb{V}\mathrm{ar}\left[\hat{f}_1(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right] & \cdots & \mathbb{C}\mathrm{ov}\left[\hat{f}_1(\mathbf{x}^*), \hat{f}_d(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right] \\ \vdots & \ddots & \vdots \\ \mathbb{C}\mathrm{ov}\left[\hat{f}_d(\mathbf{x}^*), \hat{f}_1(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right] & \cdots & \mathbb{V}\mathrm{ar}\left[\hat{f}_d(\mathbf{x}^*)|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*\right] \end{bmatrix}.$$
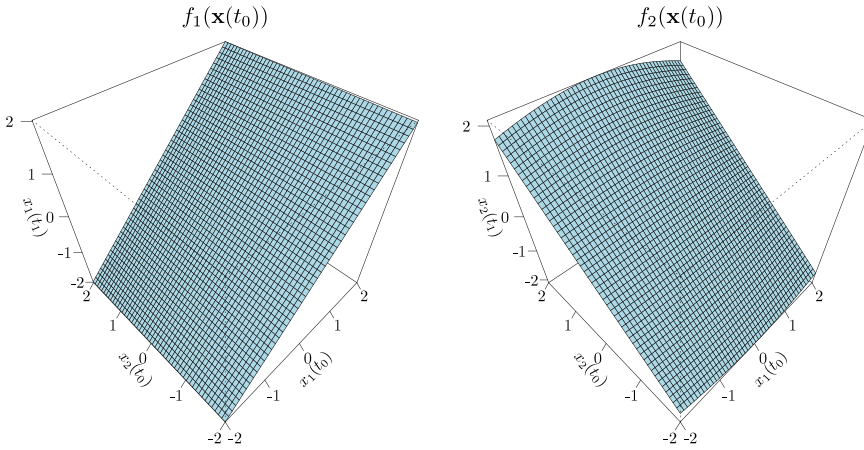
Fig. 5. The flow map of the van der Pol model as a function of the initial condition $\mathbf{x}(t_0)$ for the time step $\Delta t = 0.01$. Since the time step is small, the flow map is approximately linear. The two functions f_1 and f_2 are the components of the flow map defined as $\mathbf{x}(t_1) = \Phi(\mathbf{x}(t_0)) = (f_1(\mathbf{x}(t_0)), f_2(\mathbf{x}(t_0)))^\top$.

The diagonal elements of $\mathbf{\Sigma}^{**}$ are calculated using Eq. (9) which is approximated by the MC method. The off-diagonal elements, i.e. cross covariances, are given by:

$$
\mathbb{Cov}\left[\hat{f}_l(\mathbf{x}^*), \hat{f}_j(\mathbf{x}^*)|\boldsymbol{\mu}^*, \mathbf{\Sigma}^*\right] = \mathbb{E}\left[\hat{f}_l(\mathbf{x}^*)\hat{f}_j(\mathbf{x}^*)|\boldsymbol{\mu}^*, \mathbf{\Sigma}^*\right]
$$
$$
- \mathbb{E}\left[\hat{f}_l(\mathbf{x}^*)|\boldsymbol{\mu}^*, \mathbf{\Sigma}^*\right]\mathbb{E}\left[\hat{f}_j(\mathbf{x}^*)|\boldsymbol{\mu}^*, \mathbf{\Sigma}^*\right], \ 1 \le l, j \le d, l \ne j,
\tag{12}
$$

which are approximated by the MC method as below

$$
\mathbb{Cov}\left[\hat{f}_l(\mathbf{x}^*), \hat{f}_j(\mathbf{x}^*)|\boldsymbol{\mu}^*, \mathbf{\Sigma}^*\right] \approx \frac{1}{n_{MC}}\sum_{i=1}^{n_{MC}}\hat{f}_l(\mathbf{x}^{*i})\hat{f}_j(\mathbf{x}^{*i})
$$
$$
- \left(\frac{1}{n_{MC}}\sum_{i=1}^{n_{MC}}\hat{f}_l(\mathbf{x}^{*i})\right)\left(\frac{1}{n_{MC}}\sum_{i=1}^{n_{MC}}\hat{f}_j(\mathbf{x}^{*i})\right).
\tag{13}
$$

For example, the mean and covariance of $\hat{\mathbf{x}}^*(t_2)$ are obtained analogous to the way that $\boldsymbol{\mu}^{**}$ and $\mathbf{\Sigma}^{**}$ are computed. The uncertain input is $\hat{\mathbf{x}}^*(t_1)$ with the following distribution

$$
\mathcal{N}\left(\begin{bmatrix} \mathbb{E}\left[\hat{f}_1(\mathbf{x}(t_0))\right] = m_1\left(\mathbf{x}(t_0)\right) \\ \vdots \\ \mathbb{E}\left[\hat{f}_d(\mathbf{x}(t_0))\right] = m_d\left(\mathbf{x}(t_0)\right) \end{bmatrix}, \begin{bmatrix} \mathbb{Var}\left[\hat{f}_1(\mathbf{x}(t_0))\right] = s_1^2\left(\mathbf{x}(t_0)\right) & \dots & 0 \\ \vdots & \vdots & \vdots \\ 0 & \dots & \mathbb{Var}\left[\hat{f}_d(\mathbf{x}(t_0))\right] = s_d^2\left(\mathbf{x}(t_0)\right) \end{bmatrix}\right),
$$

where the mean and covariance are equivalent to $\boldsymbol{\mu}^*$ and $\mathbf{\Sigma}^*$ in Eqs. (11)–(13). The same rules apply to predict the state variable at times $t = 3, 4, \dots, T$.

It is worth mentioning that by applying a Gaussian distribution as a "prior" on the input at each time step ($t \ge 2$), the process can be considered as part of a family of "Deep Gaussian Processes" in which some properties of a GP are themselves another GP, see Dunlop et al. (2018). In that paper, they show that for the squared exponential kernel and certain parameter values, the process can become degenerate. We have not encountered this problem in any of our examples.

To shed more light on this method, the mean and the covariance matrix of the input in one-step ahead predictions are presented in Table 1. If cross covariances are not calculated and $\mathbf{\Sigma}^{**}$ is assumed to be a diagonal matrix, then it means the emulators are independent. Note that the GP emulators $\hat{f}_1, \dots, \hat{f}_d$ are independent and hence, the cross covariances are zero if the input to the emulators is deterministic, e.g. at $t = 1$. However, when the input is uncertain ($t \ge 2$), the predictive distribution of two independent GP becomes correlated (Kuß, 2006; Deisenroth et al., 2009).

4. Application to nonlinear dynamical systems

In this section we first describe the emulator which is applied for predicting dynamic models. We then examine the prediction capability of the emulator on two well studied dynamical systems: the Lorenz and the van der Pol systems, which are described in subsequent sections.

The GP emulator we use in our experiments consists of a first order polynomial regression for the mean function in Eq. (1) (i.e., $\mu(\mathbf{x}) = \beta_0 + \beta_1\mathbf{x}$) and a squared exponential kernel, given in Eq. (3), for the covariance kernel k. These choices

Table 1
Mean and covariance matrix of the input in one-step ahead predictions.

- $t = 0$, the input is deterministic:

$$\mathbf{x}(t_0) \sim \mathcal{N}\left(\mathbf{x}(t_0), \begin{bmatrix} 0 & \ldots & 0 \\ \vdots & \vdots & \vdots \\ 0 & \ldots & 0 \end{bmatrix}\right)$$

- $t = 1$, the cross covariances are zero because the emulators are independent:

$$\hat{\mathbf{x}}^*(t_1) \sim \mathcal{N}\left(\begin{bmatrix} \mathbb{E}\left[\hat{f}_1(\mathbf{x}(t_0))\right] \\ \vdots \\ \mathbb{E}\left[\hat{f}_d(\mathbf{x}(t_0))\right] \end{bmatrix}, \begin{bmatrix} \mathbb{V}\mathrm{ar}\left[\hat{f}_1(\mathbf{x}(t_0))\right] & \ldots & 0 \\ \vdots & \vdots & \vdots \\ 0 & \ldots & \mathbb{V}\mathrm{ar}\left[\hat{f}_d(\mathbf{x}(t_0))\right] \end{bmatrix}\right)$$

- $t = 2$, the input to emulators are no longer deterministic and the method described in this section should be applied:

$$\hat{\mathbf{x}}^*(t_2) \overset{app.}{\sim} \mathcal{N}\left(\begin{bmatrix} \mathbb{E}\left[\hat{f}_1(\hat{\mathbf{x}}^*(t_1))\right] \\ \vdots \\ \mathbb{E}\left[\hat{f}_d(\hat{\mathbf{x}}^*(t_1))\right] \end{bmatrix}, \begin{bmatrix} \mathbb{V}\mathrm{ar}\left[\hat{f}_1(\hat{\mathbf{x}}^*(t_1))\right] & \ldots & \mathbb{C}\mathrm{ov}\left[\hat{f}_1(\hat{\mathbf{x}}^*(t_1)), \hat{f}_d(\hat{\mathbf{x}}^*(t_1))\right] \\ \vdots & \vdots & \vdots \\ \mathbb{C}\mathrm{ov}\left[\hat{f}_d(\hat{\mathbf{x}}^*(t_1)), \hat{f}_1(\hat{\mathbf{x}}^*(t_1))\right] & \ldots & \mathbb{V}\mathrm{ar}\left[\hat{f}_d(\hat{\mathbf{x}}^*(t_1))\right] \end{bmatrix}\right)$$

$\quad \vdots \qquad\qquad \vdots \qquad\qquad\qquad \vdots$

- $t = T$:

$$\hat{\mathbf{x}}^*(t_T) \overset{app.}{\sim} \mathcal{N}\left(\begin{bmatrix} \mathbb{E}\left[\hat{f}_1(\hat{\mathbf{x}}^*(t_{n-1}))\right] \\ \vdots \\ \mathbb{E}\left[\hat{f}_d(\hat{\mathbf{x}}^*(t_1))\right] \end{bmatrix}, \begin{bmatrix} \mathbb{V}\mathrm{ar}\left[\hat{f}_1(\hat{\mathbf{x}}^*(t_{n-1}))\right] & \ldots & \mathbb{C}\mathrm{ov}\left[\hat{f}_1(\hat{\mathbf{x}}^*(t_{n-1})), \hat{f}_d(\hat{\mathbf{x}}^*(t_{n-1}))\right] \\ \vdots & \vdots & \vdots \\ \mathbb{C}\mathrm{ov}\left[\hat{f}_d(\hat{\mathbf{x}}^*(t_{n-1})), \hat{f}_1(\hat{\mathbf{x}}^*(t_{n-1}))\right] & \ldots & \mathbb{V}\mathrm{ar}\left[\hat{f}_d(\hat{\mathbf{x}}^*(t_{n-1}))\right] \end{bmatrix}\right).$$

of μ and k are recommended in Conti et al. (2009). A set of training samples of size $n = 12d$, as recommended in Jones et al. (1998) and Loeppky et al. (2009), is drawn over the space of initial conditions. Our training sample is constrained to lie in a cube. For example, a suitable boundary for the Lorenz attractor, as described in Afraimovich et al. (1977) and Williams (1979), is given by the unstable manifold of the origin. To define a bounding box that contains the attractor, we therefore simulated the system with initial conditions close to the origin and chose as boundaries in each coordinate the extremes of the simulation. Note we do assume that the system is constrained to lie in the same volume, although clearly we would like to capture most of its variation. The points of the training samples should be selected based on a space-filling sampling scheme, and we therefore use a Latin hypercube (Stein, 1987; Pronzato and Müller, 2012). The goal in a space-filling design is to spread the points *evenly* within the input space. No attempt is made to have the points lie along the stable manifold, we simply try to 'fill' space.

To build each emulator $\hat{f}_l$, $1 \leq l \leq d$, the training data consists of $\mathbf{X} = \{\mathbf{x}^1(t_0), \ldots, \mathbf{x}^n(t_0)\}$ with the corresponding outputs $\mathbf{y} = \left(x_l^1(t_1), \ldots, x_l^n(t_1)\right)^\top$. The R package DiceKriging (Roustant et al., 2012) is employed to fit the GP emulator. The unknown parameters of the SE kernel k (i.e. σ and θ_ls) and the mean function μ (i.e. β_0 and β_1) are estimated by maximum likelihood implemented in DiceKriging. After building the emulators, their accuracies are assessed by the leave-one-out cross-validation mean squared error (MSE_{LOO}) defined as

$$MSE_{LOO} = \frac{1}{n} \sum_{i=1}^{n} \left(\hat{f}_{l,-i}\left(\mathbf{x}^i(t_0)\right) - x_l^i(t_1)\right)^2. \tag{14}$$

In the above equation, $\hat{f}_{l,-i}\left(\mathbf{x}^i(t_0)\right)$ is the prediction obtained by the GP emulator $\hat{f}_l$ based on all the data points in $\mathbf{X}$ except the ith one.

The ODEs are solved by the default solver of the R package deSolve (Soetaert et al., 2010) which is called "ode". It is based on a variable order method to integrate the system over the next step ahead, i.e. $t_1 = t_0 + \Delta t$. More precisely, it uses the LSODA (Livermore solver for ordinary differential equations with automatic switching between stiff and nonstiff methods) method (Petzold, 1983). A full Jacobian matrix is used which is calculated internally by LSODA. In these two examples, we use a fixed time step equal to $\Delta t = 0.01$. In the following sections, we first apply the method of uncorrelated emulators and subsequently examine the method of correlated emulators on two dynamical systems, i.e. the Lorenz and van der Pol models.

4.1. Lorenz system: uncorrelated emulators

The Lorenz system was first proposed by Edward Lorenz in 1963 (Lorenz, 1963) as a mathematical representation of atmospheric convection. It is a three-dimensional system of ordinary differential equations. Under certain choices of
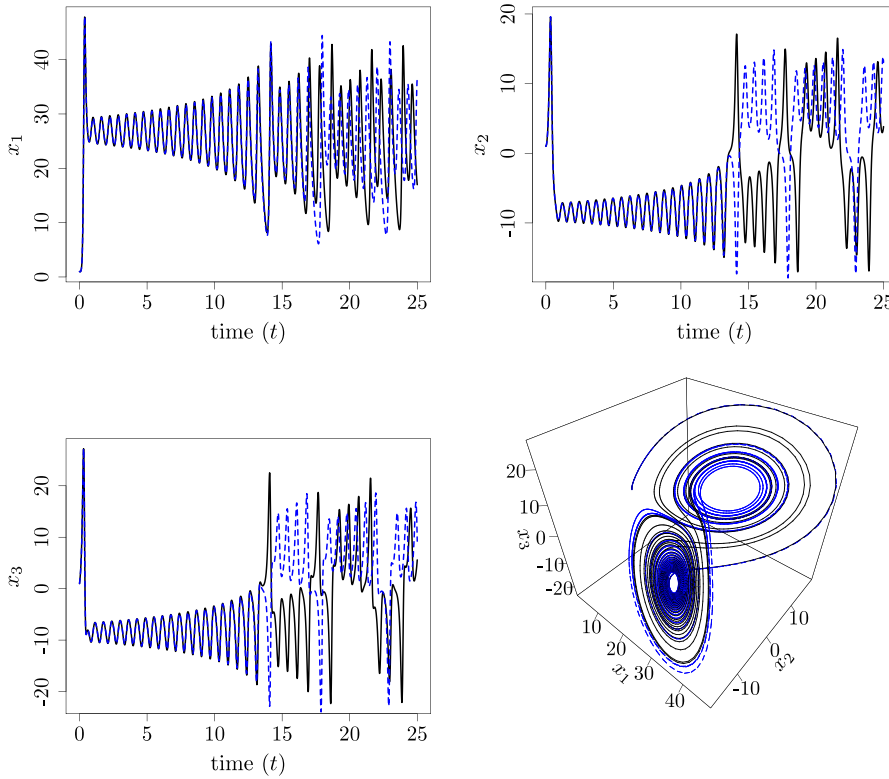
Fig. 6. The Lorenz system (solid black) and its emulation (blue dashed) considering the input uncertainty, but neglecting correlation between emulators, as described in Algorithm 1. The 1-D pictures illustrate evolution of state variables x_1, x_2, x_3 and their predictions. The emulators built based on iterative one-step ahead predictions are able to well predict up to about $t = 14$. The 3-D picture shows the evolution and prediction of the whole system.

parameters it can display chaotic behaviour, i.e. its behaviour is highly sensitive to initial conditions. The evolution of three state variables is described by Soetaert et al. (2010)

$$\begin{cases} \dfrac{dx_1}{dt} = ax_1 + x_2 x_3 \\[2mm] \dfrac{dx_2}{dt} = b(x_2 - x_3) \\[2mm] \dfrac{dx_3}{dt} = -x_1 x_2 + cx_2 - x_3, \end{cases} \tag{15}$$

where a, b and c are parameters. Here, we assume $a = -8/3$, $b = -10$ and $c = 28$. We focus on the case with initial conditions $\mathbf{x}(t_0) = (x_1(t_0) = 1, x_2(t_0) = 1, x_3(t_0) = 1)^\top$. The accuracy of the emulators is high based on the MSE_{LOO} criterion which is given below.

	$\hat{f}_1$	$\hat{f}_2$	$\hat{f}_3$
MSE_{LOO}	1.946×10^{-4}	3.533×10^{-7}	1.288×10^{-4}

Emulation of the Lorenz model using the iterative one-step ahead predictions considering the input uncertainty, but neglecting correlation between emulators, as described in Algorithm 1 is demonstrated in Fig. 6. We show the evolution of predictions for each system variable over time, as well as a three-dimensional picture showing the evolution of the whole system, $(x_1(t), x_2(t), x_3(t))^\top$. The solid line represents the true model and the blue dashed line is the GP prediction. It can be seen that the prediction precision is high at the beginning of the time course, for example $t \leq 14$. However, the emulator deviates from the true model as time progresses. Fig. 6 suggests that the emulator is well suited to describing the evolution of the system within a "wing" of the Lorenz attractor, but that predictions break down upon switching to the other part of the attractor.

Fig. 7 shows the uncertainties (solid black) associated with the predictions illustrated in Fig. 6. The uncertainties are compared with the case in which the input uncertainty is not considered (red dashed line). Generally, if emulation is
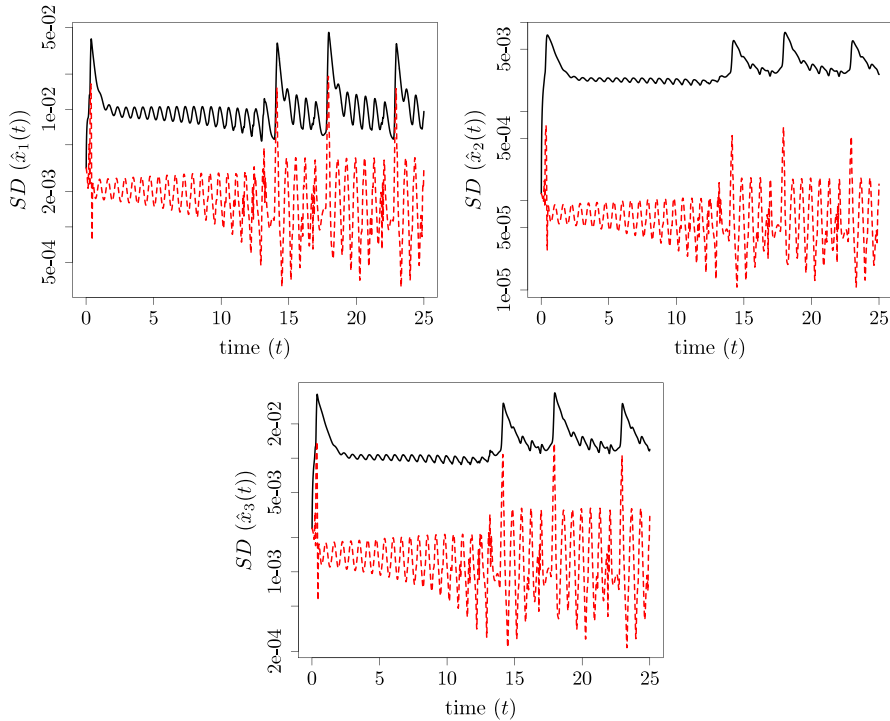
Fig. 7. Standard deviation (SD) associated with the predictions of state variables of Lorenz system (Fig. 6) with and without considering input uncertainty, solid black and red dashed lines respectively. Incorporating input uncertainty augments the prediction SD. Contrary to our expectations, the uncertainties do not grow with time.
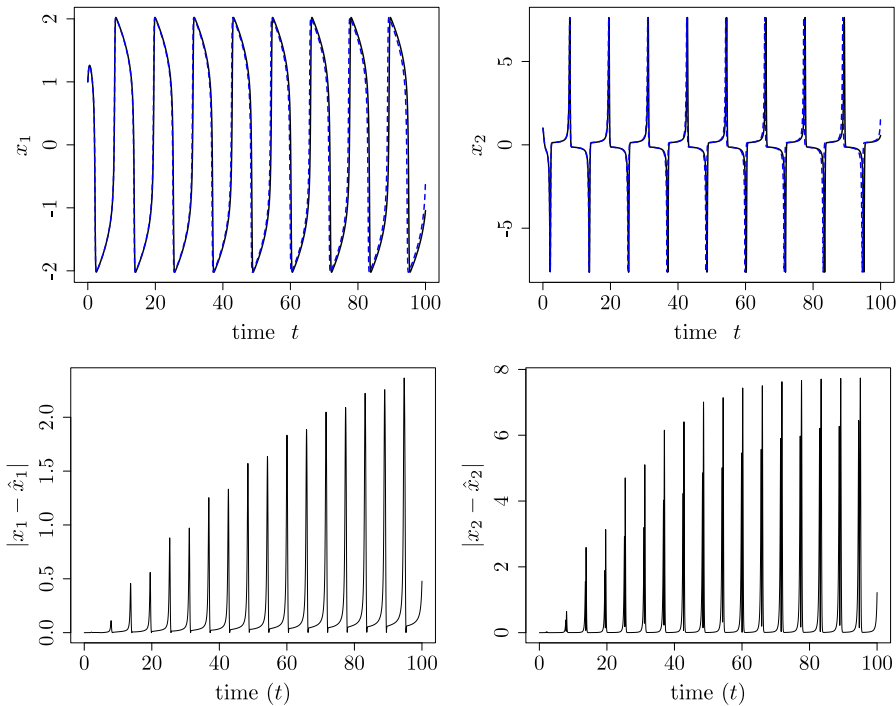

Fig. 8. Top: Evolution of state variables of the van der Pol oscillator (solid black) and their predictions (blue dashed) considering the input uncertainty, but neglecting correlation between emulators, as described in Algorithm 1. Bottom: Difference between the state variables of the van der Pol system and their predictions, i.e. $|x_i - \hat{x}_i|$. The prediction capability of emulators is high up to approximately $t = 30$.
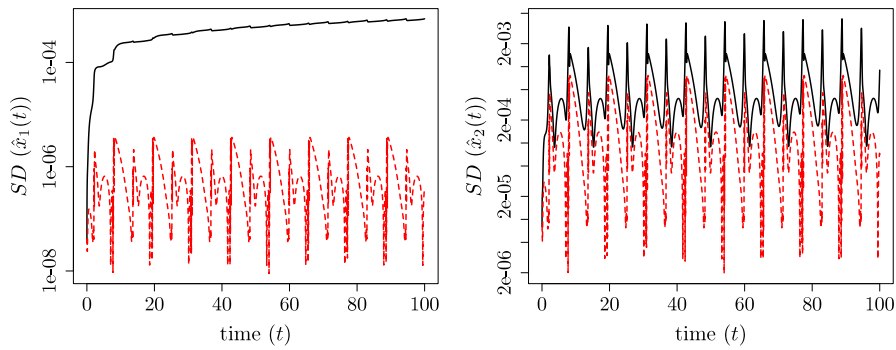
Fig. 9. Standard deviation (SD) of predicting the van der Pol system with and without considering input uncertainty, solid black and red dashed lines respectively.
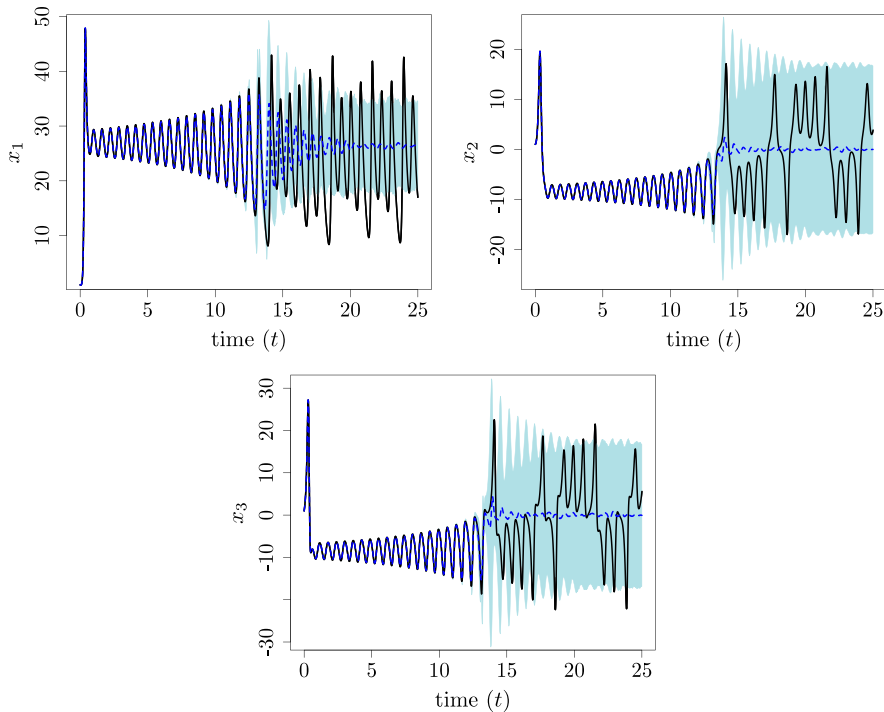

Fig. 10. The Lorenz system (solid black) and its prediction (blue dashed) considering input uncertainties and the correlation between emulators. The prediction accuracy is high at the beginning of the time course, say up to $t \approx 13$. From this point onwards, i.e. where the emulator is not able to predict the true model well, the prediction tends to the average of the process. However, the credible intervals (shaded) are large enough to contain the true model most of the time. Note when deviation from the true model occurs the prediction uncertainty reaches its maximum.

carried out with uncertain inputs, the magnitude of uncertainties is higher. Nevertheless, they are still too small and contrary to our expectations do not increase over time as the uncertainty builds up from step to step. The true model is not inside the credible intervals, which are defined as $m(\mathbf{x}) \pm 2s(\mathbf{x})$. Note the credible intervals are not shown in Fig. 6, but can easily be derived from Figs. 6 and 7. In particular, we would expect the uncertainty to "blow up" when we reach the point of switching between wings of the attractor (at about $t = 14$) where our emulator can be on a different wing to the true model but still has very small uncertainty.

4.2. van der Pol oscillator: uncorrelated emulators

The van der Pol model was first introduced by the Dutch electrical engineer Balthasar van der Pol in 1920. The van der Pol oscillator model expresses the behaviour of non-linear vacuum tube circuits. In its two-dimensional form, it is given
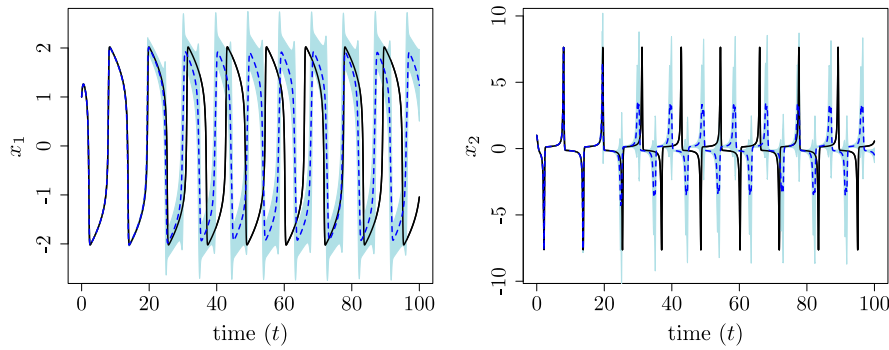
Fig. 11. The van der Pol model (solid black) and its prediction (blue dashed) considering input uncertainties and the correlation between emulators. The shaded area represents the credible intervals. The predictive capability of the emulator is high at the beginning of the time course, up to $t \approx 25$, but subsequently is a frequency miss-match and the prediction damps. From this point onwards, the prediction tends to the long term average of the system. However, the prediction uncertainty blows up when the emulator deviates from the true model.
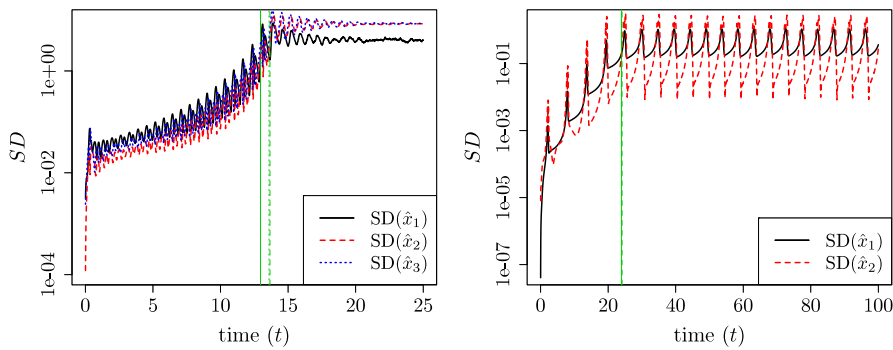
Fig. 12. Left: standard deviation (SD) of prediction associated with the three variables in the Lorenz model. Right: SD of prediction associated with the two variables in the van der Pol model. The y-axis is on logarithmic scale. In both cases, an SD reaches its maximum where the corresponding emulator cannot well predict the true model. This can be used as a criterion to detect the predictability horizon of dynamic emulators. The vertical green lines show the predictability horizon in the corresponding time series, change point in $\hat{x}_1(t)$ (solid), $\hat{x}_2(t)$ (dashed) and $\hat{x}_2(t)$ (dotted).
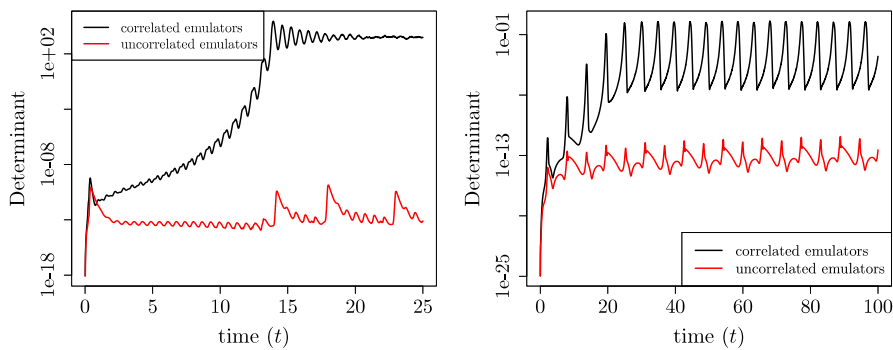

Fig. 13. The determinant of the covariance matrix, used in Table 1, over time in correlated (black) and uncorrelated (red) emulators in the Lorenz (left) and van der Pol (right) systems. In both examples, the determinant is larger in the correlated method which leads to a wider sampling of update directions for the next step. This is in line with the magnitude of uncertainties in the correlated and uncorrelated methods. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

by the following equations (Strogatz, 2007)

$$\begin{cases} \dfrac{dx_1}{dt} = x_2 \\ \dfrac{dx_2}{dt} = \alpha(1 - x_1^2)x_2 - x_1. \end{cases} \tag{16}$$
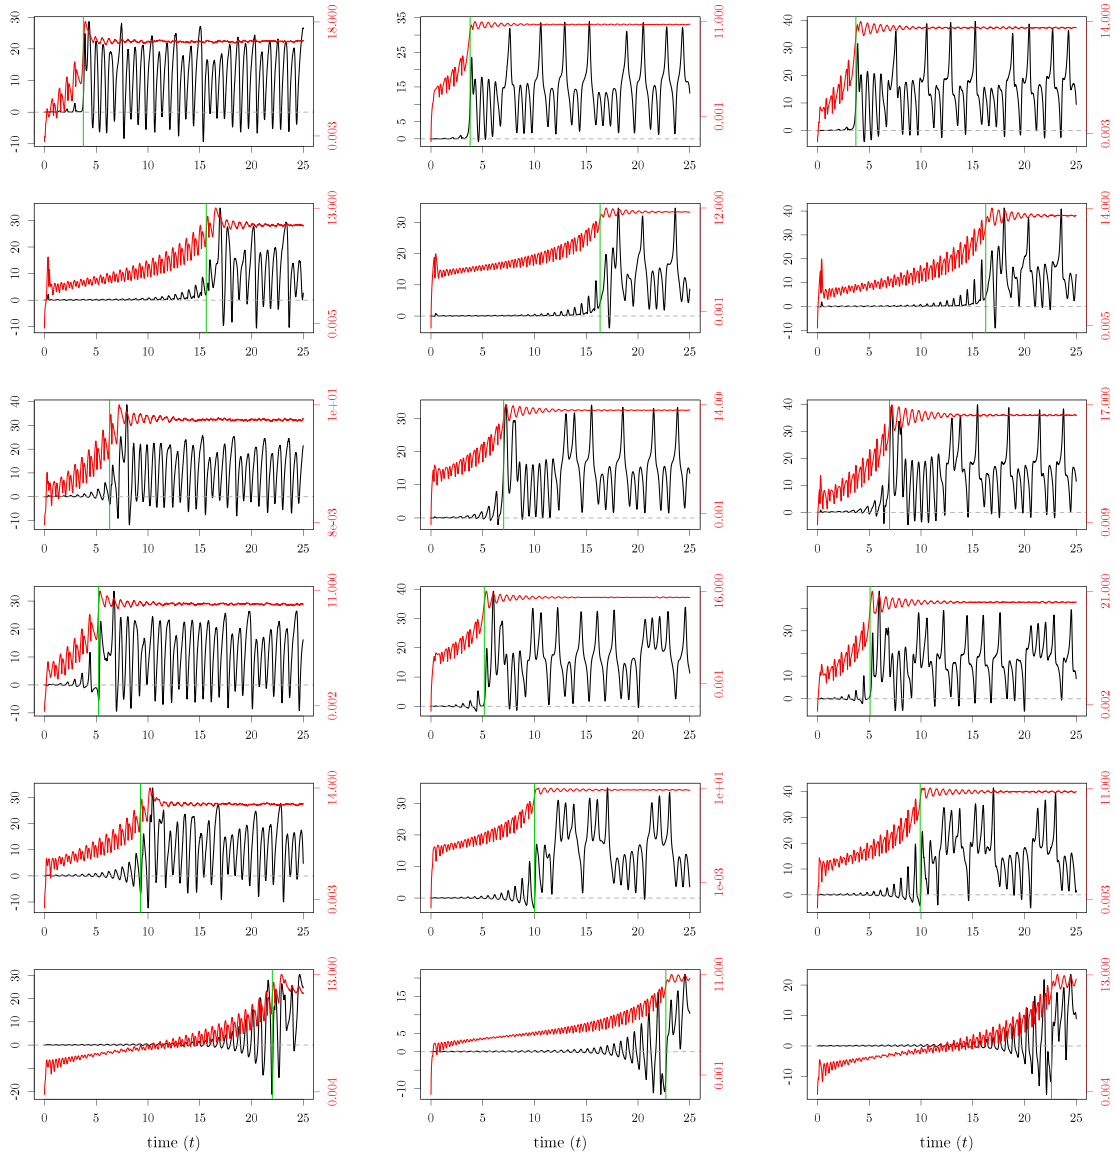
Fig. 14. Correlated method: prediction of the Lorenz system ($x_1(t)$: first panel, $x_2(t)$: second panel and $x_3(t)$: third panel) with six different initial conditions (corresponding to each row). The black line represents: $\hat{x}_i(t) + 2SD\left(\hat{x}_i(t)\right) - x_i(t)$ and the red line is $SD\left(\hat{x}_i(t)\right)$ with the (log scale) y-axis shown in red on the right side. The change point is shown by the vertical green line. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Here, the scalar $\alpha > 0$ determines the nonlinearity and the strength of damping. Here, we use the initial condition $\mathbf{x}(t_0) = (x_1 = 1, x_2 = 1)^\top$ and $\alpha = 5$. The accuracy of the emulators is given below.

	$\hat{f}_1$	$\hat{f}_2$
MSE_{LOO}	7.887×10^{-7}	2.580×10^{-7}

The results of predicting state variables of the van der Pol oscillator, neglecting correlations between emulators are illustrated in Fig. 8. The corresponding uncertainties are given by Fig. 9 where they are compared with the case that the input uncertainty is not considered. The difference between emulation and the true model is low up to approximately $t = 30$. Again, taking into account the input uncertainty augments prediction uncertainties everywhere. But, since they are small, the true model is not inside the credible intervals (= prediction $\pm 2 \times$ prediction standard deviation) when prediction accuracy declines. As is shown in the next section, considering both the input uncertainty and the correlation between emulators allows uncertainties to grow over time.
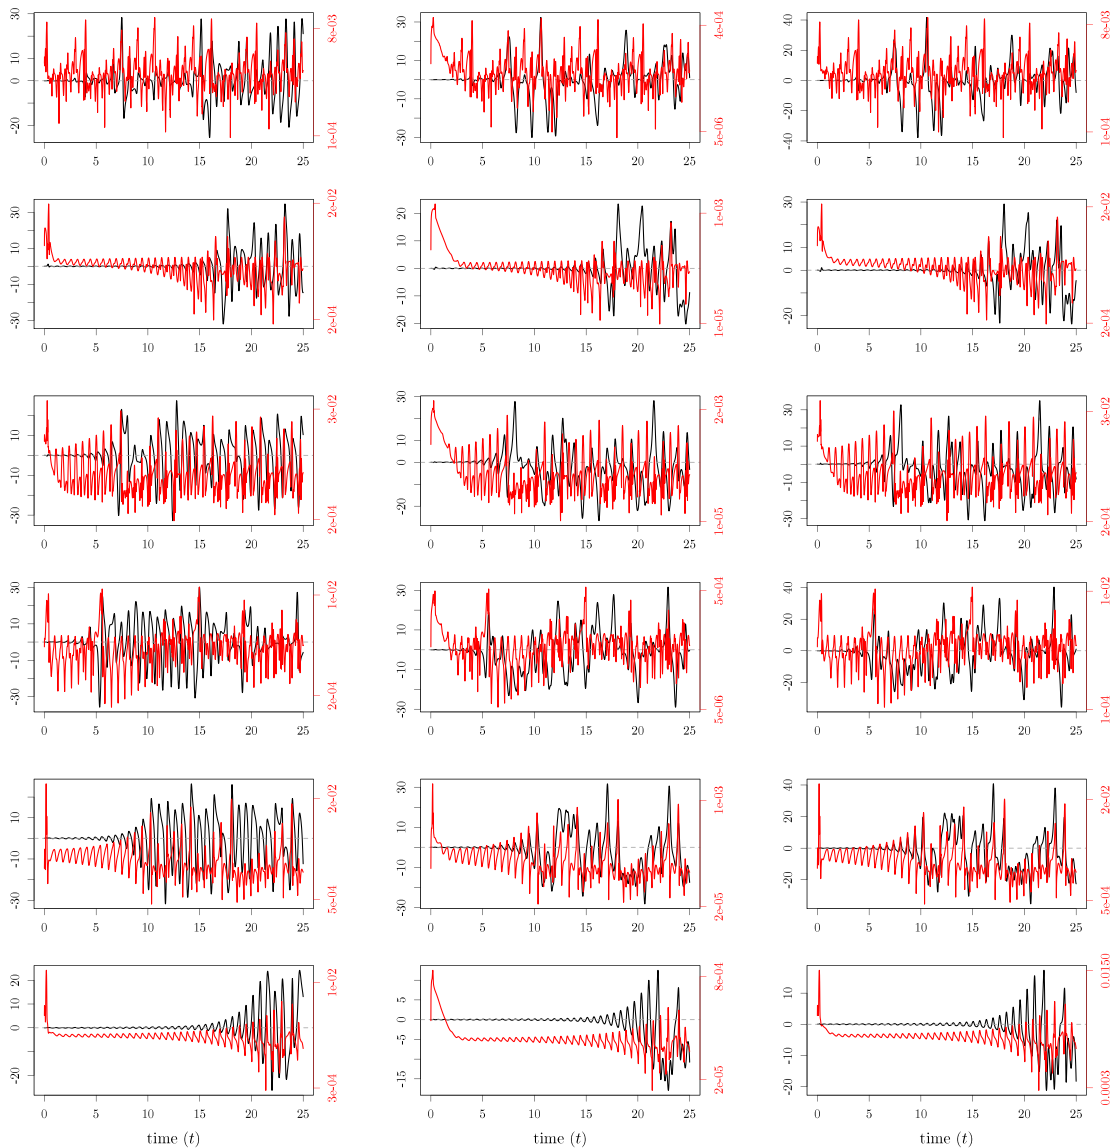
Fig. 15. Uncorrelated method: prediction of the Lorenz system ($x_1(t)$: first panel, $x_2(t)$: second panel and $x_3(t)$: third panel) with six different initial conditions (corresponding to each row). The black line represents: $\hat{x}_i(t) + 2SD\left(\hat{x}_i(t)\right) - x_i(t)$ and the red line is $SD\left(\hat{x}_i(t)\right)$ with the (log scale) y-axis shown in red on the right side. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.3. Application of correlated emulators to the Lorenz and van der Pol systems

The results of emulating the Lorenz and van der Pol systems considering input uncertainties together with the correlation between emulators are illustrated in Figs. 10 and 11, respectively. The predictive capability of these emulators is high at the beginning of the time course, say up to $t \approx 13$ for the Lorenz and $t \approx 25$ for the van der Pol models. In both cases, when the emulators deviate from the true models, the prediction uncertainties blow up which can be used to identify the time for which the prediction obtained by the emulator is reliable. This will be discussed later.

In the emulation of the Lorenz model, the prediction (blue dashed lines) tends to the average of the system after the emulator no longer predicts the true model well, i.e. $t \approx 13$. However, the uncertainty is large enough to encompass most values of the system such that the true model predominantly remains inside the credible intervals represented by the shaded area. Recall that the main drawback of the uncorrelated emulators method is that the prediction uncertainty is too small and the true model is not inside the credible intervals.

In the case of the van der Pol system, the prediction is accurate up to $t \approx 25$. From this time onwards, a frequency miss-match happens and deviation from the true trajectory grows such that the simulator output is rarely inside the
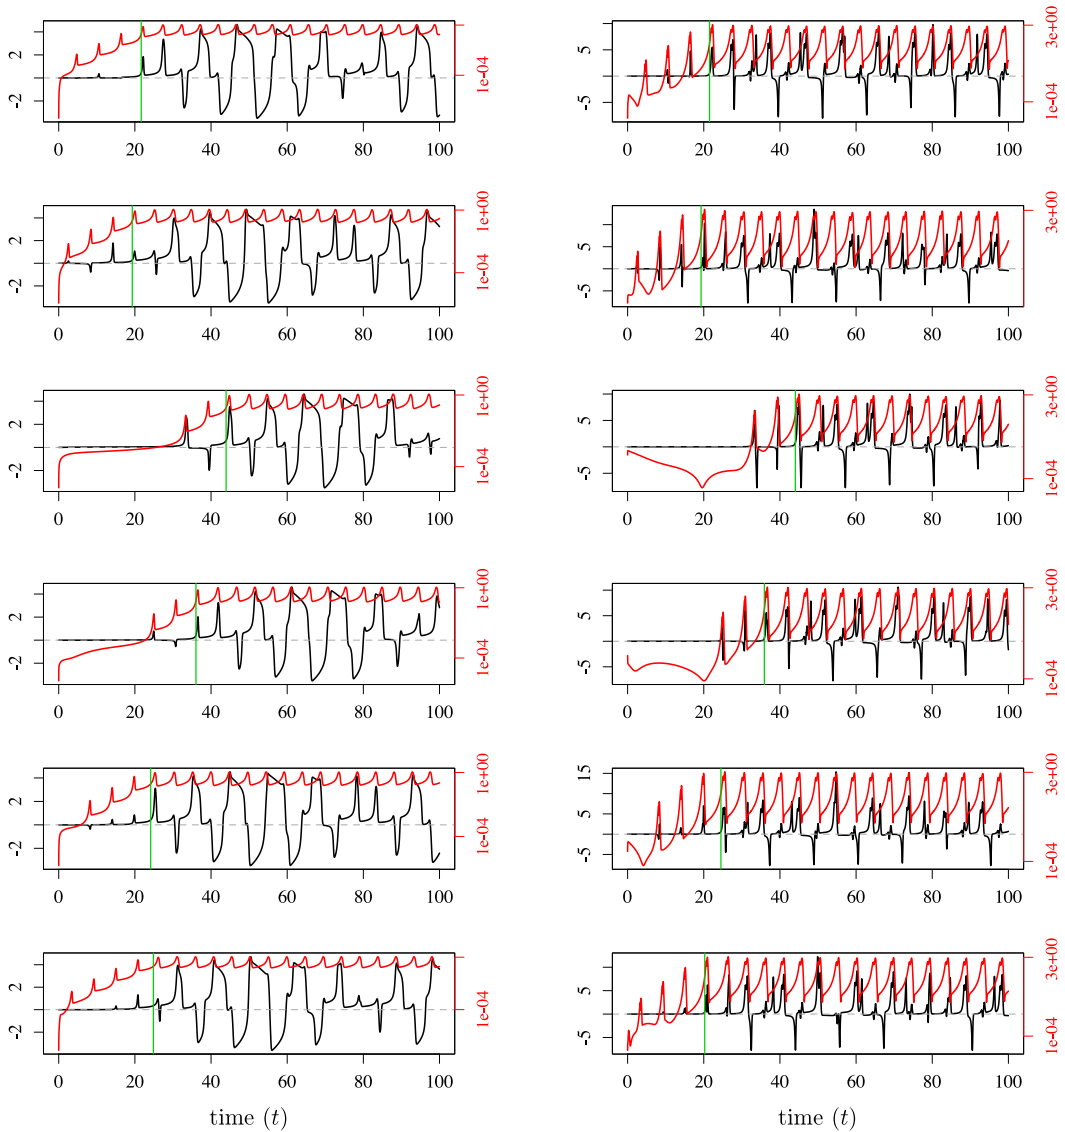
Fig. 16. Correlated method: prediction of the van der Pol model ($x_1(t)$: first panel and $x_2(t)$: second panel) with six different initial conditions (corresponding to each row). The black line represents: $\hat{x}_i(t) + 2SD\left(\hat{x}_i(t)\right) - x_i(t)$ and the red line is $SD\left(\hat{x}_i(t)\right)$. The change point is shown by the vertical green line. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

credible intervals after $t \approx 40$. Also, the amplitude of the prediction gradually damps which can be interpreted as the emulator in effect "giving up" on trying to emulate the value at a particular time t and instead falling back on a prediction for a random time. Such a prediction is useless in practice but statistically makes sense.

We note that both the Lorenz and van der Pol models were set in regimes that contain an attractor, and we consider initial conditions within their basins of attraction. Therefore, any trajectory converges to a bounded region. This means that the uncertainty is bounded above by the uncertainty of all flow maps initiated within the basin of attraction, propagated through the emulator. In each case we begin emulation with a single initial condition, i.e. certain input. As the variance of the input distribution grows, a greater volume of the phase space is used as input into the emulation of the flow map. As the variance of the input distribution grows, the mean of the emulator output tends to zero. We believe this is because the mean is taken over flows with many different directions, such that on the average there is not a preferred direction. This is in contrast to the case in which the variance of the input distribution is small, for example within the initial emulated time period. In this case, the input distribution covers a small volume in phase space such that there is a preferred direction of the sampled flows on average (for example when the distribution spans a small region of the attractor).
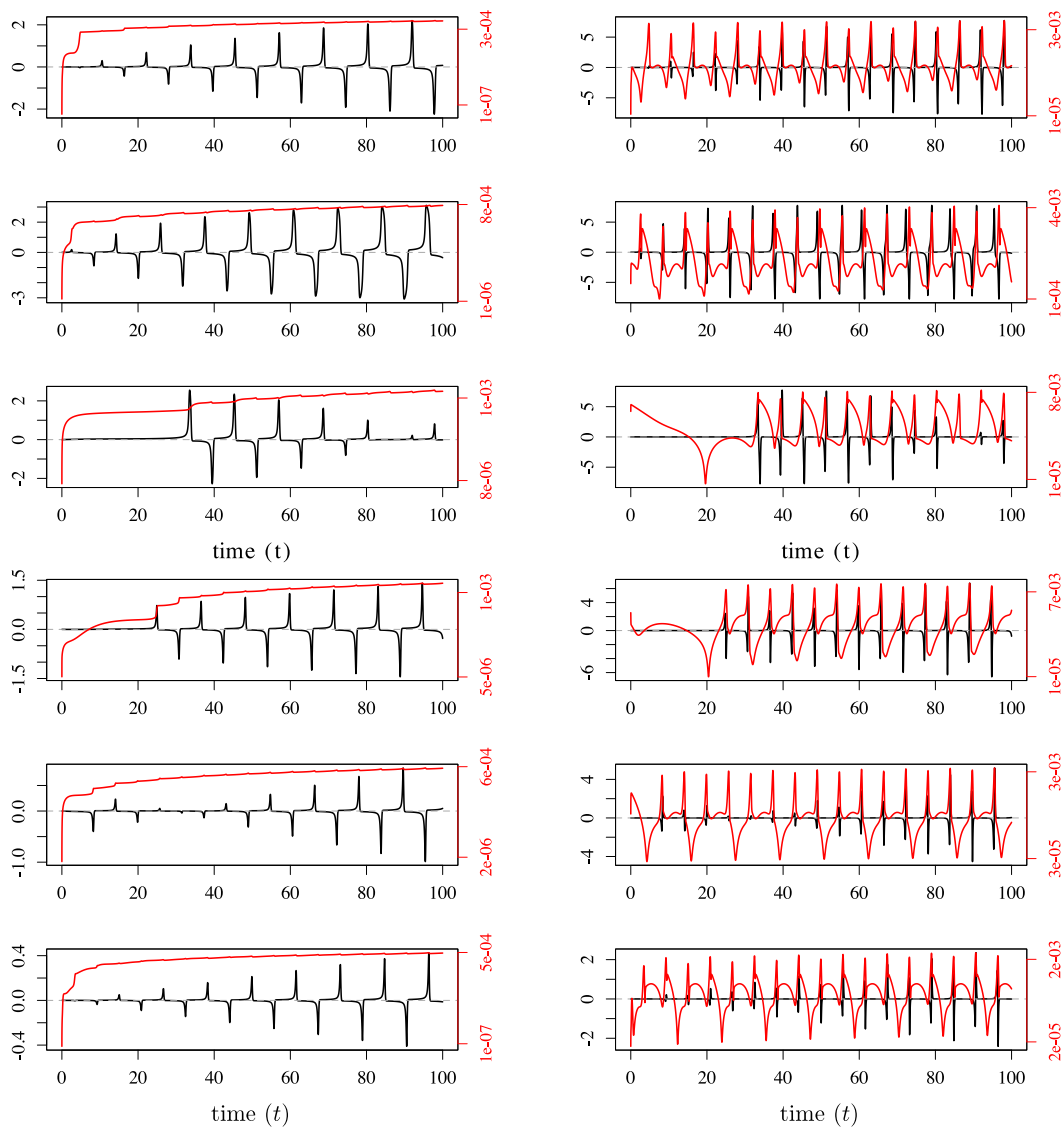
Fig. 17. Uncorrelated method: prediction of the van der Pol model ($x_1(t)$: first panel and $x_2(t)$: second panel) with six different initial conditions (corresponding to each row). The black line represents: $\hat{x}_i(t) + 2SD\left(\hat{x}_i(t)\right) - x_i(t)$ and the red line is $SD\left(\hat{x}_i(t)\right)$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Fig. 12 shows uncertainties associated with predictions of the variables in the Lorenz (left) and van der Pol systems (right) considering the correlation between the emulators through the time series. As can be seen, the uncertainty grows and reaches its maximum, interestingly, when deviation from the true model begins. From this point onwards, the expected value of the emulator is the long term average of the underlying model while the uncertainty of prediction is large. This point can be used as a measure for the *predictability horizon* of dynamic emulators. More precisely, the time at which the mean value of uncertainty changes significantly is considered as the predictability horizon. To identify the *change point* in the mean of time series, the `cpt.mean` function implemented in the `changepoint` R package (Killick and Eckley, 2014) is applied to the prediction uncertainties. The vertical green lines in Fig. 12 represent the change points in the time series. Similar results are obtained when emulating the Lorenz and van der Pol systems with different initial conditions using the correlated method. In Appendix A, the emulation of the two systems is demonstrated for additional initial conditions.

Using correlated emulators, the magnitude of uncertainty is generally larger than for the uncorrelated case, which would lead to a wider sampling of update directions for the next step, and hence greater deviation from the true underlying trajectory. This is especially the case of the van der Pol model where the prediction accuracy is higher in the uncorrelated method, see Figs. 11 and 8. However, in this approach the uncertainty is too small and the true model
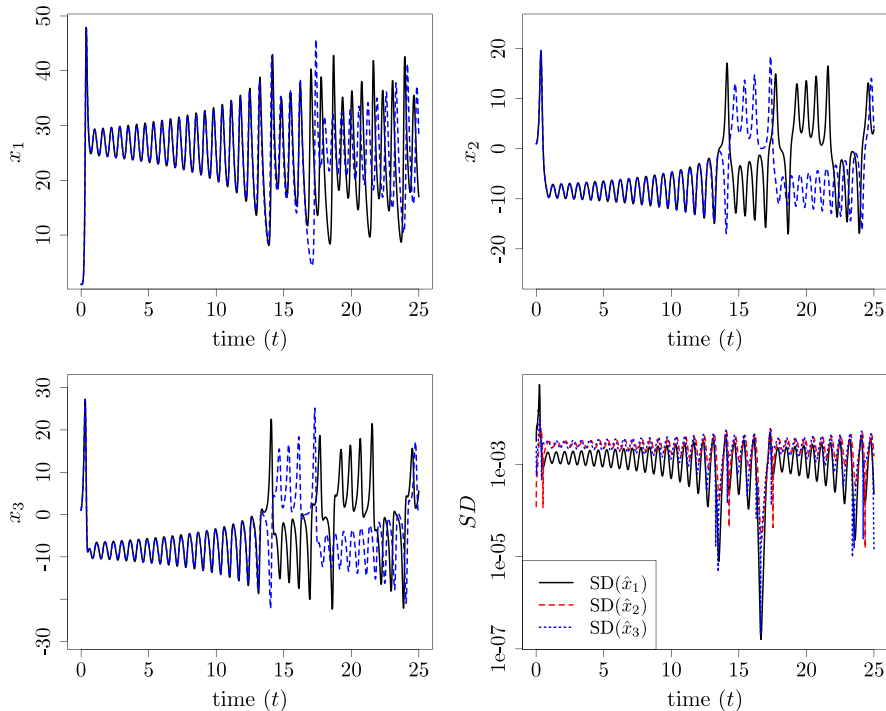
Fig. 18. Emulation of the Lorenz system in which, at each iteration, the output distribution is approximated by the Laplace's method.

is not inside the credible intervals. The larger uncertainty in the correlated method can be justified by looking at the determinant of the covariance matrix (Fig. 13) used in Table 1 to generate samples from the input distribution. As can be seen, the determinant of the covariance matrix is consistent with the prediction uncertainties. Note that the determinant measures overall dispersion of a multidimensional random variable and is referred to as generalized variance.

5. Conclusion

In this paper we develop a general framework for dynamically emulating highly non-linear functions with time series outputs using Gaussian processes. Such functions show the behaviour of phenomena evolving with time. One advantage of our method is that it is easy to implement in comparison to alternative methods; it uses a GP emulator to perform one-step ahead predictions in an iterative way over the whole time series. Moreover, we propose a number of ways to propagate uncertainty through the time series based on both the uncertainty of inputs to the emulators and the correlation between them. The capability of our method is illustrated in application to two non-linear dynamical systems: the Lorenz and van der Pol systems. In both examples, the prediction uncertainty obtained by the proposed method (i.e. correlated emulators) allows to measure a "predictability horizon", within which the prediction accuracy is high. It should be noted that in our two examples the computer model run time is small and we can compare the model and emulator directly. This is not the case for many applications where the model run time is large.

One can extend the idea of one-step to several-step ahead predictions. The main consideration is that the flow map becomes more non-linear if the number of steps increases. As a result, more training data is required to approximate the flow map well which increases the computational complexity of GPs. Recall that a GP has a computational complexity of $\mathcal{O}(n^3)$. With respect to the approximation of the distribution of inputs at subsequent time steps, as mentioned before, Eq. (7) can be approximated by the deterministic methods such as the Laplace's approximation. Another possible future research direction is to investigate such techniques in the framework of dynamic emulation. Fig. 18 (Appendix B) illustrates the results of emulating the Lorenz system in which, at each iteration, the output distribution is approximated by the Laplace's method. In this case, the prediction uncertainty does not grow over time, similarly to the case of uncorrelated emulators.

Acknowledgements

Appendix A. Correlated emulators with different initial conditions

The Lorenz and van der Pol systems are emulated (based on the correlated method) with six more different initial conditions selected randomly from $[-10, 10]^d$. These results are consistent with the previous observations in that the prediction uncertainty can be used to obtain the predictability horizon using the change point detection (see Figs. 14–17).

Appendix B. Laplace's approximation

See Fig. 18.

References

Afraimovich, V.S., Bykov, V.V., Shilnikov, L.P., 1977. On the origin and structure of the Lorenz attractor. Akad. Nauk SSSR Dokl. 234, 336–339.

Bailer-Jones, C., Sabin, T.J., Mackay, D.J.C., Withers, P., 1997. Prediction of deformed and annealed microstructures using bayesian neural networks and gaussian processes. In: Proceedings of the Australasia Pacific Forum on Intelligent Processing and Manufacturing of Materials. pp. 913–919.

Bayarri, M.J., Berger, J.O., Cafeo, J., Garcia-Donato, G., Liu, F., Palomo, J., Parthasarathy, R.J., Paulo, R., Sacks, J., Walsh, D., 2007. Computer model validation with functional output. Ann. Statist. 35 (5), 1874–1906.

Birrell, P., Ketsetzis, G., Gay, N., Cooper, B., Presanis, A., Harris, R., Charlett, A., Zhang, X., White, P., Pebody, R., De, A.D., 2011. Bayesian Modeling to unmask and predict influenza a/h1n1 pdm dynamics in London. Proc. Natl. Acad. Sci. U.S.A. 108, 18238–18243. http://dx.doi.org/10.1073/pnas.1103002108.

Candela, J.Q., Girard, A., Larsen, J., Rasmussen, C.E., 2003. Propagation of uncertainty in bayesian kernel models - application to multiple-step ahead forecasting. In: International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 701–704.

Challenor, P., 2004. The probability of rapid climate change. Significance 1 (4), 155–158. http://dx.doi.org/10.1111/j.1740-9713.2004.00059.x.

Chen, V.C.P., Tsui, K.L., Barton, R.R., Meckesheimer, M., 2006. A review on design, modeling and applications of computer experiments. IIE Trans. 38 (4), 273–291.

Conti, S., Gosling, J.P., Oakley, J.E., O'Hagan, A., 2009. Gaussian process emulation of dynamic computer codes. Biometrika 96, 663–676.

Conti, S., O'Hagan, A., 2010. Bayesian emulation of complex multi-output and dynamic computer models. J. Statist. Plann. Inference 140 (3), 640–651.

Cressie, N., 1993. Statistics for Spatial Data. In: Wiley series in probability and mathematical statistics: Applied probability and statistics, J. Wiley.

Deisenroth, M.P., Huber, M.F., Hanebeck, U.D., 2009. Analytic moment-based gaussian process filtering. In: Proceedings of the 26th Annual International Conference on Machine Learning. In: ICML '09, ACM, New York, NY, USA, pp. 225–232. http://dx.doi.org/10.1145/1553374.1553403.

Dunlop, M.M., Girolami, M., Stuart, A.M., Teckentrup, A.L., 2018. How deep are deep Gaussian processes?. J. Mach. Learn. Res. 19 (54), 1–46.

Farah, M., Birrell, P., Conti, S., Angelis, D.D., 2014. Bayesian emulation and Calibration of a dynamic epidemic model for A/H1N1 influenza. J. Amer. Statist. Assoc. 109 (508), 1398–1411. http://dx.doi.org/10.1080/01621459.2014.934453.

Ferrat, L.A., Goodfellow, M., Terry, J.R., 2018. Classifying dynamic transitions in high dimensional neural mass models: a random forest approach. PLoS Comput. Biol. 14 (3), 1–27. http://dx.doi.org/10.1371/journal.pcbi.1006009.

Forrester, A.I., Keane, A.J., 2009. Recent advances in surrogate-based optimization. Prog. Aerosp. Sci. 45 (1–3), 50–79.

Fricker, T.E., Oakley, J.E., Urban, N.M., 2013. Multivariate gaussian process emulators with nonseparable covariance structures. Technometrics 55 (1), 47–56. http://dx.doi.org/10.1080/00401706.2012.715835.

Girard, A., Rasmussen, C.E., Candela, J.Q.n., Murray-Smith, R., 2003. Gaussian process priors with uncertain inputs - application to multiple-step ahead time series forecasting. In: Advances in Neural Information Processing Systems, Vol. 15. MIT Press, pp. 545–552.

Higdon, D., Gattiker, J., Williams, B., Rightley, M., 2008. Computer model Calibration using high-dimensional output. J. Amer. Statist. Assoc. 103 (482), 570–583. http://dx.doi.org/10.1198/016214507000000888.

Jin, R., Chen, W., Simpson, T.W., 2000. Comparative studies of metamodeling techniques under multiple modeling criteria. Struct. Multidiscip. Optim. 23, 1–13.

Jin, R., Chen, W., Sudjianto, A., 2002. On sequential sampling for global metamodeling in engineering design. In: Design Engineering Technical Conferences and Computers and Information in Engineering, Vol. 2. pp. 539–548.

Jones, B., Johnson, R.T., 2009. Design and analysis for the Gaussian process model. Qual. Reliab. Eng. Int. 25, 515–524. http://dx.doi.org/10.1002/qre.1044.

Jones, D.R., Schonlau, M., Welch, W.J., 1998. Efficient global optimization of expensive black-box functions. J. Global Optim. 13 (4), 455–492. http://dx.doi.org/10.1023/A:1008306431147.

Kalaitzis, A.A., Lawrence, N.D., 2011. A simple approach to ranking differentially expressed gene expression time courses through Gaussian process regression. BMC Bioinformatics 12 (1), 180. http://dx.doi.org/10.1186/1471-2105-12-180.

Kamath, A., Vargas-Hernández, R.A., Krems, R.V., Carrington, T., Manzhos, S., 2018. Neural networks vs Gaussian process regression for representing potential energy surfaces: A comparative study of fit quality and vibrational spectrum accuracy. J. Chem. Phys. 148 (24), 241702. http://dx.doi.org/10.1063/1.5003074.

Kennedy, M.C., O'Hagan, A., 2001. Bayesian calibration of computer models. J. R. Stat. Soc. Ser. B Stat. Methodol. 63 (3), 425–464. http://dx.doi.org/10.1111/1467-9868.00294.

Killick, R., Eckley, I., 2014. Changepoint: An r package for changepoint analysis. J. Stat. Softw. Articles 58 (3), 1–19. http://dx.doi.org/10.18637/jss.v058.i03.

Kleijnen, J.P., 2009. Kriging metamodeling in simulation: A review. European J. Oper. Res. 192 (3), 707–716.

Krause, A., Singh, A., Guestrin, C., 2008. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. J. Mach. Learn. Res. 9, 235–284.

Kuczera, G., Kavetski, D., Franks, S., Thyer, M., 2006. Towards a Bayesian total error analysis of conceptual rainfall-runoff models: Characterising model error using storm-dependent parameters. J. Hydrol. 331 (1), 161–177.

Kuß, M., 2006. Gaussian Process Models for Robust Regression, Classification, and Reinforcement Learning (Ph.D. thesis). Technische Universität Darmstadt, Darmstadt, Germany, Darmstadt.

Lawrence, N., Sanguinetti, G., Rattray, M., 2006. Modelling transcriptional regulation using Gaussian processes. In: NIPS 2006.

Lee, L.A., Carslaw, K.S., Pringle, K.J., Mann, G.W., Spracklen, D.V., 2011. Emulation of a complex global aerosol model to quantify sensitivity to uncertain parameters. Atmos. Chem. Phys. 11 (23), 12253–12273. http://dx.doi.org/10.5194/acp-11-12253-2011.

Liu, F., West, M., 2009. A dynamic modelling strategy for Bayesian computer model emulation. Bayesian Anal. 4 (2), 393–412. http://dx.doi.org/10.1214/09-BA415.

Loeppky, J.L., Sacks, J., Welch, W.J., 2009. Choosing the sample size of a computer experiment: A practical guide. Technometrics 51 (4), 366–376. http://dx.doi.org/10.1198/TECH.2009.08040.

Lorenz, E.N., 1963. Deterministic nonperiodic flow. Journal of Atmospheric Sciences 20, 130–148. http://dx.doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2.

MacKay, D.J.C., 1998. Introduction to Gaussian processes. In: Bishop, C.M. (Ed.), Neural Networks and Machine Learning. Springer-Verlag, pp. 133–166.

MacKay, D.J.C., 2002. Information Theory, Inference & Learning Algorithms. Cambridge University Press, New York, NY, USA.

Neal, R.M., 1996. Bayesian Learning for Neural Networks. In: Lecture Notes in Statistics, Springer-Verlag New York.

Neal, R.M., 1998. Regression and Classification Using Gaussian Process Priors. In: Bayesian Statistics, vol. 6, Oxford University Press, pp. 475–501.

O'Hagan, A., 2006. Bayesian analysis of computer code outputs: A tutorial. Reliab. Eng. Syst. Saf. 91 (10–11), 1290–1300. http://dx.doi.org/10.1016/j.ress.2005.11.025.

Petzold, L., 1983. Automatic selection of methods for solving stiff and nonstiff systems of ordinary differential equations. SIAM J. Sci. Stat. Comput. 4 (1), 136–148. http://dx.doi.org/10.1137/0904010.

Plumlee, M., 2014. Fast prediction of deterministic functions using sparse grid experimental designs. J. Amer. Statist. Assoc. 109 (508), 1581–1591. http://dx.doi.org/10.1080/01621459.2014.900250.

Pronzato, L., Müller, W.G., 2012. Design of computer experiments: space filling and beyond. Stat. Comput. 22 (3), 681–701. http://dx.doi.org/10.1007/s11222-011-9242-3.

Rasmussen, C.E., 1997. Evaluation of Gaussian Processes and Other Methods for Non-Linear Regression (Ph.D. thesis). University of Toronto, Toronto, Ont., Canada, Canada.

Rasmussen, C.E., Williams, C.K.I., 2005. Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning). The MIT Press.

Rawlings, J., Pantula, S., Dickey, D., 2006. Applied Regression Analysis: A Research Tool. In: Springer Texts in Statistics, Springer New York.

Raychaudhuri, S., 2008. Introduction to Monte Carlo simulation. In: Proceedings of the 40th Conference on Winter Simulation. pp. 91–100.

Reichert, P., White, G., Bayarri, M.J., Pitman, E.B., 2011. Mechanism-based emulation of dynamic simulation models: Concept and application in hydrology. Comput. Stat. Data Anal. 55 (4), 1638–1655. http://dx.doi.org/10.1016/j.csda.2010.10.011.

Rougier, J., 2008. Efficient emulators for multivariate deterministic functions. J. Comput. Graph. Statist. 17 (4), 827–843. http://dx.doi.org/10.1198/106186008X384032.

Roustant, O., Ginsbourger, D., Deville, Y., 2012. Dicekriging, diceoptim: Two r packages for the analysis of computer experiments by kriging-based metamodeling and optimization. J. Stat. Softw. 51 (1), 1–55. http://dx.doi.org/10.18637/jss.v051.i01.

Sacks, J., Welch, W.J., Mitchell, T.J., Wynn, H.P., 1989. Design and analysis of computer experiments. Statist. Sci. 4 (4), 409–423. http://dx.doi.org/10.1214/ss/1177012420.

Santner, T.J., Williams, B., Notz, W., 2003. The Design and Analysis of Computer Experiments. Springer-Verlag, p. 283.

Schwaighofer, A., Grigoras, M., Tresp, V., Hoffmann, C., 2004. In: Thrun, S., Saul, L., Schölkopf, B. (Eds.), GPPS: A Gaussian Process Positioning System for Cellular Networks. In: Advances in Neural Information Processing Sys, vol. 16, MIT Press.

Soetaert, K., Petzoldt, T., Setzer, R.W., 2010. Solving differential equations in R: Package desolve. J. Stat. Softw. 33 (9), 1–25. http://dx.doi.org/10.18637/jss.v033.i09.

Stein, M., 1987. Large sample properties of simulations using latin hypercube sampling. Technometrics 29 (2), 143–151.

Stommel, H., 1961. Thermohaline convection with two stable regimes of flow. Tellus 13, 224–230.

Strogatz, S., 2007. Nonlinear dynamics and chaos. In: Studies in nonlinearity, Sarat Book House.

Swain, P.S., Stevenson, K., Leary, A., Montano-Gutierrez, L.F., Clark, I.B.N., Vogel, J., Pilizota, T., 2016. Inferring time derivatives including cell growth rates using gaussian processes. Nature Commun. 7 (4), 273–291.

Williams, R.F., 1979. The structure of lorenz attractors. Publ. Math. Inst. Hautes Etudes Sci. 50, 73–99.

Williamson, D., Blaker, A.T., 2014. Evolving Bayesian emulators for structured chaotic time series, with application to large climate models. SIAM/ASA J. Uncertain. Quantif. 2 (1), 1–28. http://dx.doi.org/10.1137/120900915.