

Research



Cite this article: Sapsis TP, Blanchard A. 2022 Optimal criteria and their asymptotic form for data selection in data-driven reduced-order modelling with Gaussian process regression. *Phil. Trans. R. Soc. A* **380**: 20210197. <https://doi.org/10.1098/rsta.2021.0197>

Received: 15 September 2021

Accepted: 7 January 2022

One contribution of 16 to a theme issue 'Data-driven prediction in dynamical systems'.

Subject Areas:

computational mechanics, applied mathematics, computational mathematics, statistics

Keywords:

optimal experimental design, data-driven modelling, Bayesian regression, optimal sampling, active learning

Author for correspondence:

Themistoklis P. Sapsis
e-mail: sapsis@mit.edu

Optimal criteria and their asymptotic form for data selection in data-driven reduced-order modelling with Gaussian process regression

Themistoklis P. Sapsis and Antoine Blanchard

Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

TPS, 0000-0003-0302-0691; AB, 0000-0001-6514-0619

We derive criteria for the selection of datapoints used for data-driven reduced-order modelling and other areas of supervised learning based on Gaussian process regression (GPR). While this is a well-studied area in the fields of active learning and optimal experimental design, most criteria in the literature are empirical. Here we introduce an optimality condition for the selection of a new input defined as the minimizer of the distance between the approximated output probability density function (pdf) of the reduced-order model and the exact one. Given that the exact pdf is unknown, we define the selection criterion as the supremum over the unit sphere of the native Hilbert space for the GPR. The resulting selection criterion, however, has a form that is difficult to compute. We combine results from GPR theory and asymptotic analysis to derive a computable form of the defined optimality criterion that is valid in the limit of small predictive variance. The derived asymptotic form of the selection criterion leads to convergence of the GPR model that guarantees a balanced distribution of data resources between probable and large-deviation outputs, resulting in an effective way of sampling towards data-driven reduced-order modelling.

This article is part of the theme issue 'Data-driven prediction in dynamical systems'.

1. Introduction

Reduced-order modelling has been a cornerstone of modern computational methods. The effectiveness of the reduced-order models relies both on their design, so they can capture the complexity of the underlying process, and also on the information they rely on. This fundamental information can have the form of (i) governing equations, which typically carry assumptions or simplifications of their own, and (ii) data, which is a more reliable but expensive source. The present work involves the development of criteria for the most effective selection of data—or associated experiments to generate these data—in order to perform data-driven reduced-order modelling.

The literature related to data-driven reduced-order modelling is vast and spans a great number of engineering and scientific fields ranging from fluid mechanics [1–5] to structural mechanics [6–9]. In the majority of these works, the assumption is plentiful data. While for many applications this is indeed the case, there are several important problems where plentiful data are not available, either because the associated physical or numerical experiments are too expensive or because of the nature of the problem, e.g. extreme events that occur rarely [10].

The scientific field that aims to tackle this issue is active learning [11–13]. A critical issue in active learning is the choice of acquisition function, i.e. the criterion used to select which sample to query next in an optimal manner. Acquisition functions come in various shapes and forms [11,14], but many popular criteria suffer from severe limitations, including high computational cost, intractability in high dimensions, and inability to discriminate between active and idle input variables [15]. Recently, a new class of active learning criteria was introduced, designed to take into account the effect of the output on the selection of samples [15,16]. This new class of output-weighted criteria has led to significantly improved performance in a variety of problems involving uncertainty quantification, Bayesian optimization [17,18], and decision-making [19,20]. However, there is no sound theoretical understanding of its favourable properties.

In this work, we rigorously show the optimal properties of output-weighted acquisition functions when it comes to the problem of selecting samples for data-driven reduced-order modelling in the Bayesian context. Specifically, we employ Gaussian process regression (GPR) as the building block to develop Bayesian surrogates or stochastic reduced-order models. Within the GPR framework, we derive asymptotically optimal acquisition functions, where optimality is defined in the sense of fastest convergence of the probability density function (pdf) describing the quantity of interest. We demonstrate the derived optimal acquisition functions in a mechanical oscillator subjected to high-dimensional stochastic forcing resulting in non-Gaussian heavy tails, as well as the reduced-order modelling of a beam under axial and transverse stochastic loads that result in buckling and bending.

2. Problem set-up

Our aim is to build a data-driven reduced-order model that will capture the behaviour of an output quantity $y \in \mathbb{R}$ (assumed scalar for simplicity) with respect to an input variable $x \in \mathbb{R}^n$. The input can represent, for example, initial conditions or parameters governing the evolution of a dynamical system, while the output is any quantity of interest that depends on the input variable. We also assume that the input variable has a prescribed probability distribution function, $p_x(x)$.

We are given a set of input datapoints $X = \{x_i\}_{i=1}^N$ and corresponding output values

$$Y = [y(x_1), \dots, y(x_N)]^T,$$

and the objective is to identify a criterion that allows us to select the next input point $x_{N+1} \triangleq h$ that will supplement—together with the corresponding output $y(x_{N+1})$ —the dataset $\mathcal{D} = \{X, Y\}$. There is a plethora of criteria in the active-learning field (see, e.g. [11] for a review), the majority of which are designed to select the next input h that either minimizes the uncertainty of the reduced-order model, i.e. the posterior variance $\sigma_y^2(x|\mathcal{D})$, or maximizes the information content between input and output variables. However, neither of these methods gives appropriate attention to

increasing the accuracy of the resulting model in the regions of the input space where it is needed the most: inputs that result in large deviations of the output variable y from its expected value.

To accomplish this goal, we will focus on accelerating the convergence of the model output pdf, $p_y(s|\mathcal{D})$. In the context of reduced-order modelling, this is the appropriate quantity to consider as it encodes the probabilistic information about the large deviations of y , while the input variables have been marginalized. Therefore, while criteria based on the posterior variance aim to minimize the error of the model for all or certain input variables, a criterion based on the output pdf explicitly targets regions of the input space that contribute to the output pdf. Moreover, to emphasize regions associated with large deviations of the output variable y , we will define convergence in terms of the *logarithm* of $p_y(s|\mathcal{D})$, using the below function to measure the discrepancy between two pdfs p_1 and p_2 :

$$\mathbb{D}(p_1, p_2) = \int_{S_y} |\log p_1(s) - \log p_2(s)| ds. \quad (2.1)$$

In the above, S_y is a finite domain of y (the output variable), over which we aim to build the reduced-order model. Note that this is different from the Kullback–Leibler divergence, which does not give the same emphasis to low probability events (associated with large deviations of y). Moreover, it can be easily shown that the above function defines a metric. In what follows, we first provide a quick review of GPR and its basic properties, and subsequently formulate the optimality condition that our sampling criteria will satisfy.

(a) Review of Gaussian process regression

We use GPR to build a surrogate for the unknown function $y(x) : \mathbb{R}^n \rightarrow \mathbb{R}$. The idea is to use a Gaussian prior on the function y , i.e.

$$y_0 \sim \text{GP}(m(x), k(x, x')), \quad (2.2)$$

with known mean $m(x) : \mathbb{R}^n \rightarrow \mathbb{R}$ and covariance function $k(x, x') : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ assumed to be positive-definite. We typically set the mean function m to be identically zero. As for the covariance kernel k , we often use the Gaussian covariance:

$$k(x, x') = \sigma_k^2 \exp\left(-\frac{\|x - x'\|^2}{2\lambda^2}\right). \quad (2.3)$$

Conditioning the Gaussian process on the available datapoints, we obtain the predictive process

$$y_* \sim \text{GP}(\bar{y}(x), \bar{k}(x, x')), \quad (2.4)$$

where the predictive mean, $\bar{y}(x)$ and predictive covariance $\bar{k}(x, x')$ are given in terms of the datapoints:

$$\left. \begin{aligned} \bar{y}(x) &= k(x, X)K(X, X)^{-1}Y \\ \text{and} \quad \bar{k}(x, x') &= k(x, x') - k(x, X)K(X, X)^{-1}k(X, x'), \end{aligned} \right\} \quad (2.5)$$

where $k(x, X) = [k(x, x_1), \dots, k(x, x_N)] \in \mathbb{R}^N$ and $K(X, X) = \{k(x_i, x_j)\}_{i,j=1}^N \in \mathbb{R}^{N \times N}$.

The native space for GPR schemes is the reproducing kernel Hilbert space corresponding to the kernel k , defined as follows:

Definition 2.1. A Hilbert space \mathcal{H}_k of functions $y : \mathbb{R}^n \rightarrow \mathbb{R}$, with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$, is called the reproducing kernel Hilbert space (RKHS) corresponding to a symmetric, positive-definite kernel k if

1. for all $x \in \mathbb{R}^n$, $k(x, x')$, as a function of its second argument, x' , belongs to \mathcal{H}_k ; and
2. for all $x \in \mathbb{R}^n$ and $y \in \mathcal{H}_k$, $\langle y, k(x, \cdot) \rangle_{\mathcal{H}_k} = y(x)$.

For special choices of kernels k , the RKHS can be characterized through its spectrum. Specifically, we have the following theorem for the characterization of the RKHS in the case of Gaussian kernels:

Theorem 2.2 (Wendland [21], theorem 10.12). *Let $k(x, x') = \sigma_k^2 \exp(-\|x - x'\|^2/2\lambda^2)$ be the squared-exponential kernel. The corresponding RKHS \mathcal{H}_k can be written as*

$$\mathcal{H}_k = \left\{ y \in L_2(\mathbb{R}^n) \cap C(\mathbb{R}^n) : \|y\|_{\mathcal{H}_k} = \frac{1}{c_0} \int |\mathcal{F}[y](\omega)|^2 \exp\left(\frac{\lambda^2 \|\omega\|^2}{2}\right) d\omega < \infty \right\}, \quad (2.6)$$

where c_0 is a constant that depends on n and λ , and \mathcal{F} is the Fourier transform.

The above results shows that for any $y \in \mathcal{H}_k$, the magnitude of its Fourier transform $|\mathcal{F}[y](\omega)|$ decays exponential fast as $|\omega| \rightarrow \infty$ and the speed of decay increases with λ . Analogous results exist for the case of Matérn kernels [21].

A fundamental property of GPR schemes is the fact that one can obtain *a priori* estimates for the accuracy of the surrogate model. In particular, for the case of GPR schemes, we have the following error estimate:

Proposition 2.3 (Stuart & Teckentrup [22], proposition 3.5). *Suppose that $\bar{y}(x)$ and $\bar{k}(x, x')$ are given by the GPR scheme (2.5). Then*

$$\sup_{\|y\|_{\mathcal{H}_k}=1} |y(x) - \bar{y}(x)| = \bar{k}(x, x)^{1/2} \triangleq \bar{\sigma}(x), \quad (2.7)$$

where the supremum occurs when the functions $y(\cdot)$ and $\bar{k}(\cdot, x)$ are linearly dependent.

Based on this result which involves the accuracy of the surrogate map, we will derive the corresponding results for the pdf of the output variable, $p_y(s|\mathcal{D})$.

(b) Optimality condition for data selection

We formulate an active-sampling criterion that aims directly for the convergence of the output pdf, $p_y(s|\mathcal{D})$. Let $h \in \mathbb{R}^n$ be the new candidate input point. As output, we employ the approximation by the surrogate model, $\bar{y}(h)$. In this way, we have the augmented dataset $\mathcal{D}' = \{[X, h], [Y, \bar{y}(h)]\}$, which results in the same predictive mean, $\bar{y}(x)$, and a new predictive covariance $\bar{k}(x, x'; h)$. Ideally, we would want the new input to be chosen by minimizing the distance

$$\mathbb{D}(p_y, p_{\bar{y}}) = \int_{S_y} |\log p_y(s) - \log p_{\bar{y}}(s|\mathcal{D}, h)| ds. \quad (2.8)$$

However, this is not possible as p_y is *a priori* unknown. To this end, we will use as a selection criterion the supremum of the above distance over the unit sphere of the functional space $y \in \mathcal{H}_k$ (we fix the norm of the unknown function without loss of generality). Therefore, the selection criterion takes the form of minimizing the acquisition function

$$Q(h|\mathcal{D}) \triangleq \sup_{\|y\|_{\mathcal{H}_k}=1} \mathbb{D}(p_y, p_{\bar{y}}) = \sup_{\|y\|_{\mathcal{H}_k}=1} \int_{S_y} |\log p_y(s) - \log p_{\bar{y}}(s|\mathcal{D}, h)| ds. \quad (2.9)$$

While the criterion is targeting the output pdf, it is not easily computable, especially in high dimensions. The rest of the paper aims to derive a computable version appropriate for high-dimensional input spaces.

3. Asymptotically optimal criterion for data selection

Our efforts focus on obtaining a computable version of the criterion (2.9). We plan to achieve this by assuming small variance $\bar{\sigma}^2(x) = \bar{k}(x, x)$. We first recall an asymptotic result that connects the error for a map and the error between the induced pdfs defined by the corresponding maps.

Theorem 3.1 (Mohamad & Sapsis [23], theorem 2). Let $\hat{y}(x)$ and $y(x) : \mathbb{R}^n \rightarrow \mathbb{R}$ be two continuous functions with difference $\Delta y(x) = \hat{y}(x) - y(x)$, which is assumed to be small. Let also $p_x(x)$ be the probability density function of the random vector $x \in \mathbb{R}^n$. The difference between the induced pdfs for \hat{y} and y has the following asymptotic behaviour:

$$p_y(s) - p_{\hat{y}}(s) = -\frac{d}{ds} \int_{\hat{y}(x)=s} p_x(x) \Delta y(x) dx + \mathcal{O}(|\Delta y|^2).$$

Building on this result, we have the main theorem that characterizes the supremum in the RKHS between the pdf induced by the surrogate approximation and the pdf induced by the exact map:

Theorem 3.2. Let $y(x) \in \mathcal{H}_k$ be an arbitrary function and its GPR approximation with kernel k given by $\bar{y}(x)$ with corresponding variance $\bar{\sigma}^2(x)$. Then the following property holds for small $\bar{\sigma}$:

$$\sup_{\|y\|_{\mathcal{H}_k}=1} \int_{S_y} |\log p_{\bar{y}}(s) - \log p_y(s)| ds = \int_{\bar{y}^{-1}(S_y)} \frac{p_x(x) |p'_y(\bar{y}(x))|}{p_y(\bar{y}(x))^2} \bar{\sigma}(x) dx + \mathcal{O}(\bar{\sigma}^2). \quad (3.1)$$

Proof. Using theorem 3.1, we have for the case where \bar{y} is close to y (i.e. $\Delta y/y = (\hat{y} - y)/y \ll 1$)

$$\begin{aligned} \log p_{\bar{y}}(s) - \log p_y(s) &= \frac{p_{\bar{y}}(s) - p_y(s)}{p_y(s)} + \mathcal{O}(|\Delta p_y|^2) \\ &= -\frac{(d/ds) \int_{\bar{y}(x)=s} p_x(x) \Delta y(x) dx}{p_y(s)} + \mathcal{O}(|\Delta y|^2). \end{aligned}$$

Expressing the right-hand side as a volume integral with a delta function and using properties of generalized derivatives, we have

$$\begin{aligned} \frac{(d/ds) \int_{\bar{y}(x)=s} p_x(x) \Delta y(x) dx}{p_y(s)} &= \frac{(d/ds) \int p_x(x) \Delta y(x) \delta(s - \bar{y}(x)) dx}{p_y(s)} \\ &= \int \frac{p_x(x) \Delta y(x) \delta'(s - \bar{y}(x))}{p_y(s)} dx \\ &= \int \frac{p_x(x) \Delta y(x) p'_y(s) \delta(s - \bar{y}(x))}{p_y^2(s)} dx \\ &= p'_y(s) \int \frac{p_x(x) \Delta y(x) \delta(s - \bar{y}(x))}{p_y^2(s)} dx. \end{aligned}$$

We note that $p_x(x) \delta(s - \bar{y}(x)) / p_y^2(s) \geq 0$. By employing proposition 2.3, we obtain the tight bound

$$\int \frac{p_x(x) \Delta y(x) \delta(s - \bar{y}(x))}{p_y^2(s)} dx \leq \int \frac{p_x(x) \bar{\sigma}(x) \delta(s - \bar{y}(x))}{p_y^2(s)} dx,$$

where equality holds for the case where the functions $y(\cdot)$ and $\bar{k}(\cdot, x)$ are linearly dependent [22].

Combining the above, we have

$$|\log p_{\bar{y}}(s) - \log p_y(s)| \leq |p'_y(s)| \int \frac{p_x(x) \bar{\sigma}(x) \delta(s - \bar{y}(x))}{p_y^2(s)} dx + \mathcal{O}(\bar{\sigma}^2).$$

Integrating over s , we have the final result:

$$\begin{aligned} \sup_{\|y\|_{\mathcal{H}_k}=1} \int_{S_y} |\log p_{\bar{y}}(s) - \log p_y(s)| ds &= \int_{S_y} |p'_y(s)| \int \frac{p_x(x) \bar{\sigma}(x) \delta(s - \bar{y}(x))}{p_y^2(s)} dx ds + \mathcal{O}(\bar{\sigma}^2) \\ &= \int_{\bar{y}^{-1}(S_y)} \frac{p_x(x) |p'_y(\bar{y}(x))|}{p_y^2(\bar{y}(x))} \bar{\sigma}(x) dx + \mathcal{O}(\bar{\sigma}^2). \end{aligned}$$

This completes the proof. ■

Next, we use this asymptotic form to reformulate the data selection criterion (2.9). Specifically, we apply the above theorem for the augmented dataset \mathcal{D}' and obtain the asymptotic reformulation of the selection criterion:

$$\sup_{\|y\|\tau_k=1} \int_{S_y} |\log p_y(s) - \log p_{\bar{y}}(s|\mathcal{D}, h)| ds = \int_{\bar{y}^{-1}(S_y)} \frac{p_x(x)|p'_{\bar{y}}(\bar{y}(x))|}{p_y(\bar{y}(x))^2} \bar{\sigma}(x; h) dx + \mathcal{O}(\sigma^2), \quad (3.2)$$

where $\bar{\sigma}(x; h) = \bar{k}(x, x; h)$ is the predictive variance based on the augmented dataset \mathcal{D}' . We note that the right-hand side involves the pdf p_y , which is unknown but can always be approximated by $p_{\bar{y}}$ with negligible error, given the assumptions of theorem 3.2. This gives us the final asymptotic approximation:

$$Q(h|\mathcal{D}) = \sup_{\|y\|\tau_k=1} \int_{S_y} |\log p_y(s) - \log p_{\bar{y}}(s|\mathcal{D}, h)| ds \simeq \int_{\bar{y}^{-1}(S_y)} \frac{p_x(x)|p'_{\bar{y}}(\bar{y}(x))|}{p_{\bar{y}}(\bar{y}(x))^2} \bar{\sigma}(x; h) dx. \quad (3.3)$$

It is worth emphasizing the term in the denominator, which promotes sampling of regions associated with low probability, i.e. large deviations of the output of the reduced-order model.

We can simplify the right-hand side further by using the Cauchy–Schwarz inequality (with weight $p_x(x)/p_{\bar{y}}(\bar{y}(x))$) to obtain the less conservative upper bound:

$$Q(h|\mathcal{D}) \simeq \int_{\bar{y}^{-1}(S_y)} \frac{p_x(x)|p'_{\bar{y}}(\bar{y}(x))|}{p_{\bar{y}}^2(\bar{y}(x))} \bar{\sigma}(x; h) dx \leq c \left[\int_{\bar{y}^{-1}(S_y)} \frac{p_x(x)}{p_{\bar{y}}(\bar{y}(x))} \bar{\sigma}^2(x; h) dx \right]^{1/2}, \quad (3.4)$$

where

$$c = \left[\int_{\bar{y}^{-1}(S_y)} \frac{p_x(x)p_{\bar{y}}^2(\bar{y}(x))}{p_{\bar{y}}^3(\bar{y}(x))} dx \right]^{1/2} \quad (3.5)$$

is a constant that depends only on $p_x(x)$, $p_{\bar{y}}(y)$, and $\bar{y}(x)$, i.e. not on h . This form, also referred to as an output-weighted (or likelihood-weighted) criterion, is appropriate for computations involving even high-dimensional input spaces, since it allows for the analytical computation of $\sigma^2(x; h)$ in terms of simpler integrals [16]. It has been studied numerically in recent papers [15,16,18], showing significantly favourable convergence properties compared with existing active-learning criteria.

(a) The case of extreme-event quantiles

For a wide range of applications, the focus is on characterizing the probability of exceeding a certain level, $P_y(s_*) = P[y \leq s_*]$, rather than characterizing the full pdf $p_y(s)$ (the case of non-exceeding probability can be addressed in a similar fashion). For this type of problem, the appropriate selection criterion is minimizing the acquisition function

$$R(h|\mathcal{D}, s_*) \triangleq \sup_{\|y\|\tau_k=1} |P_{\bar{y}}(s_*|\mathcal{D}, h) - P_y(s_*)|. \quad (3.6)$$

For this case, we have the following asymptotic result for small σ :

Theorem 3.3. *Let $y(x) \in \mathcal{H}_k$ and its GPR approximation with kernel k given by $\bar{y}(x)$ with variance $\bar{\sigma}^2(x)$. Then, for any given s_* , the following property holds for small $\bar{\sigma}$:*

$$\sup_{\|y\|\tau_k=1} |P_{\bar{y}}(s_*) - P_y(s_*)| = \int_{\bar{y}(x)=s_*} \bar{\sigma}(x)p_x(x) dx + \mathcal{O}(\bar{\sigma}^2).$$

Proof. The starting point is theorem 3.1. We integrate from $-\infty$ to s_* to obtain the cumulative distribution function on the left-hand side, when \bar{y} is close to y (i.e. $\Delta y/y = (\hat{y} - y)/y \ll 1$):

$$P_{\bar{y}}(s) - P_y(s) = - \int_{\bar{y}(x)=s} p_x(x) \Delta y(x) dx + \mathcal{O}(|\Delta y|^2).$$

We then employ proposition 2.3 and bound the difference Δy on the right-hand side. This completes the proof. ■

Applying the above theorem to the augmented dataset \mathcal{D}' , we obtain the asymptotic form for the optimal selection criterion involving quantiles:

$$R(h|\mathcal{D}, s_*) \triangleq \sup_{\|y\|_{\tau_{\epsilon_k}}=1} |P_{\bar{y}}(s_*|\mathcal{D}, h) - P_y(s_*)| = \int_{\bar{y}(x)=s_*} \bar{\sigma}(x; h) p_x(x) dx + \mathcal{O}(\bar{\sigma}^2). \quad (3.7)$$

We observe that for this case, the optimal data selection criterion takes a different form, focusing primarily on reducing the error around the contour $\bar{y}(x) = s_*$. This is not surprising given that the extreme-event quantile is essentially equivalent to a classification problem, so what is most important is to have low error around the *a priori* unknown contour of interest.

(b) Convergence of spatial approximation error

Given an active-learning criterion, it is important to characterize the convergence of the spatial error of the GPR approximation. This can be obtained with the help of proposition 2.3 and standard results of measurable functions. Here we give a proof for the convergence properties using the sampling criterion appearing on the right-hand side in (3.4). Specifically, we have the following result:

Theorem 3.4. *Suppose that \bar{y}_N and $\bar{\sigma}_N$ are given by a GPR with kernel k , using N samples. Moreover, assume that sampling is performed using the optimal criterion so that*

$$\lim_{N \rightarrow \infty} \int \frac{p_x(x)}{p_{\bar{y}_N}(\bar{y}_N(x))} \bar{\sigma}_N^2(x) dx = 0.$$

Then we have convergence in measure, i.e.

$$\lim_{N \rightarrow \infty} P \left[x : \sup_{\|y\|_{\tau_{\epsilon_k}}=1} (\bar{y}_N(x) - y(x))^2 \leq q \frac{p_{\bar{y}_N}(\bar{y}_N(x))}{p_x(x)} \right] = 1,$$

for every $q > 0$.

Proof. From standard results of measurable functions [24], for every sequence of functions $\phi_k(x) \in L^p$ ($1 \leq p \leq \infty$) for which $\|\phi_k(x)\|_p \rightarrow 0$, we have convergence in measure:

$$\lim_{k \rightarrow \infty} P[x : |\phi_k(x)| \geq q] = 0, \quad (3.8)$$

for every $q > 0$. Applying this result to the sequence $p_x(x) \bar{\sigma}_N^2(x) / p_{\bar{y}_N}(\bar{y}_N(x))$, we obtain

$$\lim_{N \rightarrow \infty} P \left[x : \frac{p_x(x)}{p_{\bar{y}_N}(\bar{y}_N(x))} \bar{\sigma}_N^2(x) \leq q \right] = 1.$$

We then utilize proposition 2.3, which immediately leads to the desired result. ■

This result provides a description of the spatial convergence properties for the approximation error of the reduced-order model. Specifically, it shows that the convergence is accelerated in

regions of the input space that are (i) most probable according to the pdf $p_x(x)$, and (ii) associated with small probability of the output pdf $p_{\tilde{y}_N}(\tilde{y}_N(x))$, i.e. large deviations of the output y . This type of convergence guarantees a balanced distribution of resources between the most probable inputs and those that result in large deviations for the output, resulting in an effective way of sampling towards data-driven reduced-order modelling.

4. Numerical illustration

We demonstrate the optimal sampling criteria in a mechanical oscillator subject to stochastic forcing and in the reduced-order modelling of a beam under axial and transverse stochastic loads. We consider the optimal sampling criteria appearing on each side of the inequality in (3.4). The left-hand side will be referred to as the ‘B’ criterion, and the right-hand side as ‘IVR-LW’ as it is strictly equivalent to the eponym criterion introduced in [15,16]. We also consider two criteria commonly used in the literature which do not account for the importance of the output relative to the input; namely, uncertainty sampling (US) and input-weighted integrated variance reduction (IVR-IW), whose definitions can be found in [15,16]. For each example below, we run 100 Bayesian experiments, each differing in the choice of the $n + 1$ points making up the initial dataset. Observations are assumed to be corrupted by Gaussian noise with zero mean and unknown (i.e. to be learned) variance σ_n^2 . Performance at each iteration is evaluated using the median of $\mathbb{D}(p_y, p_{\tilde{y}_N})$ across the 100 randomized experiments. Note that we employ this particular quantity to measure performance so we can better emphasize the accuracy in the tails of the resulting probability distributions.

(a) Forced nonlinear oscillator exhibiting extreme events

We begin with the stochastic oscillator of Mohamad & Sapsis [23],

$$\ddot{u} + \delta \dot{u} + F(u) = \xi(t), \quad t \in [0, T], \quad (4.1)$$

where $u(t) \in \mathbb{R}$ is the state variable, F a nonlinear restoring force, and $\xi(t)$ a stationary stochastic process which we parametrize using a Karhunen–Loève expansion with n modes:

$$\xi(t) \simeq x\Phi(t),$$

where $\{\Lambda, \Phi(t)\}$ contains the first n eigenpairs of the correlation matrix, and $x \in \mathbb{R}^n$ is a vector of random coefficients having mean zero and diagonal covariance matrix Λ . The quantity of interest is taken to be the mean value of $u(t)$ over the interval $[0, T]$. (Further details about system parameters can be found in [16,23].)

We consider the case $n = 2$ as it allows visual comparison of the decisions made by the sampling criteria as more points are being acquired. Even with $n = 2$, the output pdf has heavy tails (figure 1), a consequence of the strong nonlinearity in (4.1). For $\sigma_n^2 = 10^{-3}$, figure 1 shows that the derived optimal criteria accelerate convergence of the output pdf quite dramatically. The error for B remains close to, but always slightly below, that for IVR-LW, which is consistent with the mathematical derivation laid out in the previous section.

To explain the success of the proposed optimal criteria, we investigate the decisions made by US, B and IVR-LW in the case where observations are noiseless and σ_n^2 is set to zero in the GPR model (i.e. it is not learned from data). Consistent with [15,16], figure 2 shows that US attempts to reduce uncertainty somewhat evenly across the space as it has no mechanism to discriminate between relevant and irrelevant regions. With IVR-IW, the algorithm does not explore beyond the centre region where p_x is large, and consequently the interesting regions are not visited. On the other hand, both B and IVR-LW decide to focus on a diagonal band, with the former being even more surgical and localized than the latter. Specifically, IVR-LW focuses on input regions with important probability as well as those input regions associated with large outputs. In this way the resulting surrogates predict the output statistics much better than with US.

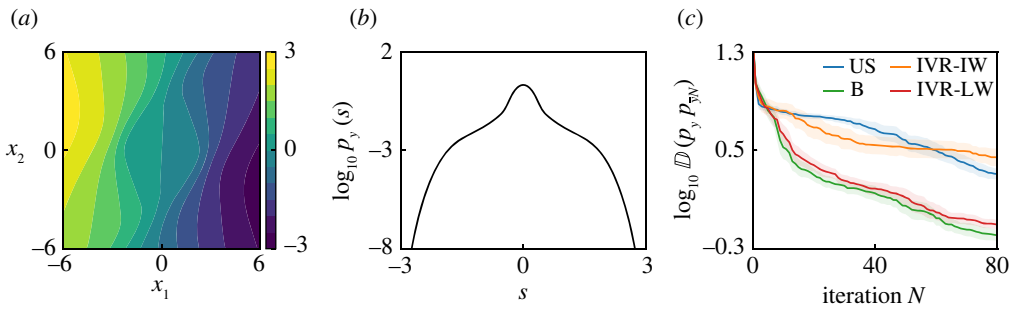


Figure 1. For the stochastic oscillator (4.1) with $n = 2$, contour plot of the output y (a) and its pdf (b), and performance of several sampling criteria for $\sigma_n^2 = 10^{-3}$ (c). The error bands indicate one half of the median absolute deviation. (Online version in colour.)

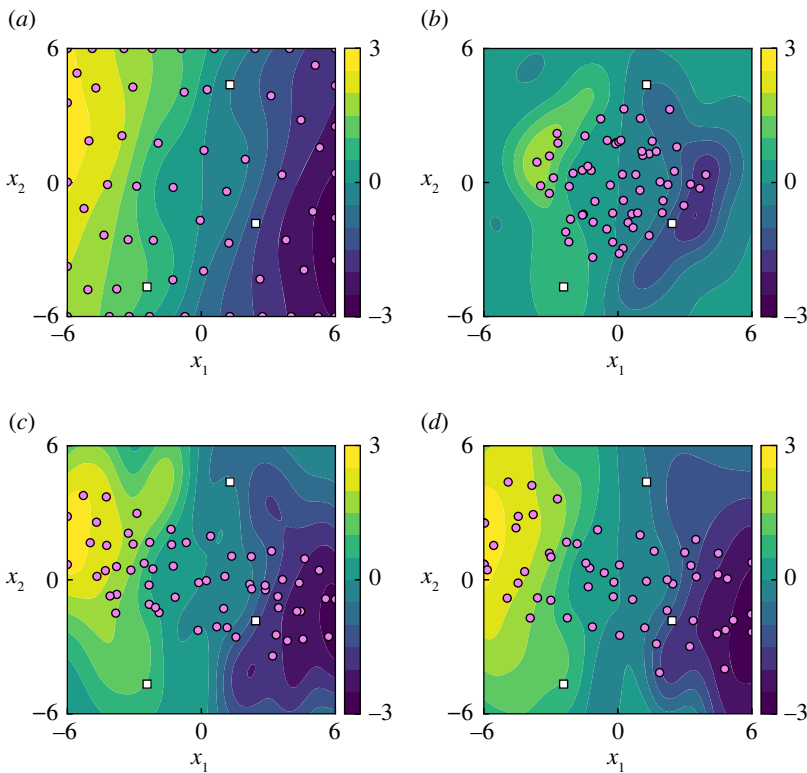


Figure 2. For the stochastic oscillator (4.1) with $n = 2$ and $\sigma_n^2 = 0$, progression of the sampling algorithm for several criteria after 60 iterations; the contours denote the posterior mean of the GPR model, the open squares the initial dataset, and the filled circles the optimized samples. (a) US, (b) IVR-IW, (c) B, (d) IVR-LW. (Online version in colour.)

(b) Buckling of a beam under stochastic axial and transverse excitation

Next, we consider the case of a beam of length l subject to both axial and transverse stochastic loads. The linearized equation of motion takes the form

$$\frac{\partial^2 w}{\partial t^2} + 2\zeta\omega_0 \frac{\partial w}{\partial t} + \omega_0^2 \frac{\partial^4 w}{\partial x^4} + P(t) \frac{\partial^2 w}{\partial x^2} = R(x, t), \quad (4.2)$$

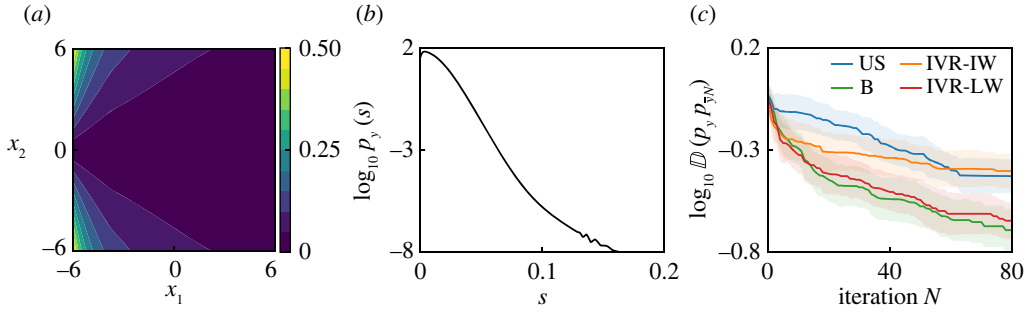


Figure 3. For the beam under random load with $n = 2$ ($J = 1, n_{\text{KL}} = 1$), contour plot of the output y (a) and its pdf (b), and performance of several sampling criteria for $\sigma_n^2 = 10^{-3}$ (c). The error bands indicate one half of the median absolute deviation. (Online version in colour.)

where $P(t)$ is random axial load and $R(x, t)$ is a random distributed transverse load [25]. We assume that both functions have zero-mean Gaussian statistics and prescribed spectra. We also assume that the beam has pin boundary conditions at both ends. In this case and assuming weak damping, the solution can be expressed as

$$w(x, t) = \sum_{j=1}^{\infty} f_j(t) \sin\left(\frac{j\pi x}{l}\right), \quad (4.3)$$

which results in a set of modal equations:

$$\ddot{f}_j + 2\zeta\omega_j\dot{f}_j + \omega_j^2 \left[1 - \frac{P(t)}{c_j}\right] f_j = R_j(t), \quad j = 1, 2, \dots \quad (4.4)$$

where

$$\omega_j^2 = \omega_0^2 \left(\frac{j\pi}{l}\right)^4, \quad c_j = \left(\frac{j\pi}{l}\right)^4 \quad \text{and} \quad R_j(t) = \frac{2}{l} \int_0^l \sin\left(\frac{j\pi x}{l}\right) R(x, t) dx. \quad (4.5)$$

We define as a quantity of interest the maximum absolute displacement at $x = l/4$ over a prescribed time interval $[0, T]$:

$$y = \max_{t \in [0, T]} \left| w\left(\frac{l}{4}, t\right) \right| = \max_{t \in [0, T]} \left| \sum_{j=1}^J \sin\left(\frac{j\pi}{4}\right) f_j(t) \right|. \quad (4.6)$$

We use $\sigma_\xi^2 \exp[-t^2/(2\ell_\xi^2)]$ for the correlation function of $P(t)$ and each $R_j(t)$, which are expanded using Karhunen–Loève expansions with n_{KL} modes. The search space, therefore, has dimension $n = J(n_{\text{KL}} + 1)$. We use parameters $\sigma_\xi = 20$, $\ell_\xi = 0.1$, and $T = 5$.

For simplicity, we consider a modal truncation of $J = 1$ and $n_{\text{KL}} = 1$, leading to a two-dimensional search space. For the parameters considered, figure 3 shows that the pdf of the output has a heavy right tail, to which the B and IVR-LW criteria converge more quickly than US and IVR-IW. The reason is that the extreme displacement values are found in low-probability areas of the search space (i.e. small p_x), which US and IVR-IW have no mechanism to discover.

5. Conclusion

We have derived optimal acquisition functions for active-learning schemes utilized for reduced-order modelling based on Gaussian process regression. The key feature of these optimal acquisition functions is a mechanism that targets the output pdf of the surrogate model. The derivation begins by selecting each sample so that the distance between the exact output pdf and the approximated output pdf obtained from the reduced-order model is minimized. Given that

the exact pdf is *a priori* unknown, a supremum of this distance over the native Hilbert space for the Gaussian process regression is considered. The resulting selection criterion, although optimal, is generally computationally intractable. We show that this difficulty can be overcome by deriving successive asymptotic upper bounds, resulting in a sampling criterion which can be evaluated analytically along with its gradients. In addition to the data-selection criterion, we derive the corresponding bound for the approximation error of the resulting reduced-order model. This result shows that our approach enables the reduced-order model to find the optimal trade-off between most likely and most interesting (i.e. large-deviation) output for the process at hand, thereby greatly enhancing the effectiveness of the algorithm. Numerical results confirm the derived analytical findings.

Note that in this work and for the sake of simplicity we have chosen to represent the reduced-order model in the form of a function from the parameter space to the output space. The present framework can be adapted to represent the reduced-order model in the form of a low-dimensional dynamical system, focusing e.g. on capturing a specific mechanism of the full dynamics. We leave this topic as a possible direction for future work.

Data accessibility. This article has no additional data.

Authors' contributions. T.S.: conceptualization, formal analysis, funding acquisition, methodology, project administration, resources, supervision, writing—original draft, writing—review and editing; A.B.: data curation, investigation, software, validation, visualization, writing—original draft, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. The authors acknowledge support from the Air Force Office of Scientific Research (MURI grant no. FA9550-21-1-0058), the Defense Advanced Research Projects Agency (grant no. HR00112110002), and the Office of Naval Research (grant no. N00014-21-1-2357).

References

1. Brunton S, Noack B, Koumoutsakos P. 2020 Machine learning for fluid mechanics. *Ann. Rev. Fluid Mech.* **52**, 477–508. (doi:10.1146/annurev-fluid-010719-060214)
2. Fernex D, Semaan R, Albers M, Meysonnat PS, Schröder W, Noack BR. 2020 Actuation response model from sparse data for wall turbulence drag reduction. *Phys. Rev. Fluids* **5**, 073901. (doi:10.1103/PhysRevFluids.5.073901)
3. Ghattas O, Willcox K. 2021 Learning physics-based models from data: perspectives from inverse problems and model reduction. *Acta Numer.* **30**, 445–554. (doi:10.1017/S0962492921000064)
4. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. 2021 Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440. (doi:10.1038/s42254-021-00314-5)
5. Ma Z, Pan W. 2021 Data-driven nonintrusive reduced order modeling for dynamical systems with moving boundaries using Gaussian process regression. *Comput. Methods Appl. Mech. Eng.* **373**, 113495. (doi:10.1016/j.cma.2020.113495)
6. Bertalan T, Dietrich F, Mezič I, Kevrekidis IG. 2019 On learning Hamiltonian systems from data. *Chaos* **29**, 1–9. (doi:10.1063/1.5128231)
7. Jain S, Haller G. 2022 How to compute invariant manifolds and their reduced dynamics in high-dimensional finite-element models? *Nonlinear Dyn.* **107**, 1417–1450.
8. Kerschen G, Worden K, Vakakis AF, Golinval JC. 2006 Past, present and future of nonlinear system identification in structural dynamics. *Mech. Syst. Signal Process.* **20**, 505–592. (doi:10.1016/j.ymsp.2005.04.008)
9. Moore KJ, Mojahed A, Kurt M, Eriten M, McFarland DM, Bergman LA, Vakakis AF. 2019 *Advanced nonlinear system identification for modal interactions in nonlinear structures: a review*, vol. 94. New York, NY: Springer International Publishing.
10. Sapsis TP. 2021 Statistics of extreme events in fluid flows and waves. *Annu. Rev. Fluid Mech.* **53**, 85–111. (doi:10.1146/annurev-fluid-030420-032810)

11. Chaloner K, Verdinelli I. 1995 Bayesian experimental design: a review. *Stat. Sci.* **10**, 273–304. (doi:10.1214/ss/1177009939)
12. Gramacy RB, Lee HK. 2009 Adaptive design and analysis of supercomputer experiments. *Technometrics* **51**, 130–145. (doi:10.1198/TECH.2009.0015)
13. Sacks J, Welch WJ, Mitchell TJ, Wynn HP. 1989 Design and analysis of computer experiments. *Stat. Sci.* **4**, 409–423.
14. Shahriari B, Swersky K, Wang Z, Adams RP, De Freitas N. 2016 Taking the human out of the loop: a review of Bayesian optimization. *Proc. IEEE* **104**, 148–175. (doi:10.1109/JPROC.2015.2494218)
15. Sapsis T. 2020 Output-weighted optimal sampling for Bayesian regression and rare event statistics using few samples. *Proc. R. Soc. A* **476**, 20190834. (doi:10.1098/rspa.2019.0834)
16. Blanchard A, Sapsis T. 2021 Output-weighted optimal sampling for Bayesian experimental design and uncertainty quantification. *SIAM-ASA J. Uncertain. Quantif.* **9**, 564–592. (doi:10.1137/20M1347486)
17. Blanchard A, Maceda GC, Fan D, Li Y, Zhou Y, Noack B, Sapsis TP. 2021 Bayesian optimization for active flow control. *Acta Mech. Sin.*
18. Blanchard A, Sapsis T. 2021 Bayesian optimization with output-weighted importance sampling. *J. Comput. Phys.* **425**, 109901. (doi:10.1016/j.jcp.2020.109901)
19. Blanchard A, Sapsis T. 2021 Informative path planning for extreme anomaly detection in environment exploration and monitoring. *Ocean Eng.* **43**, 110242. (doi:10.1016/j.oceaneng.2021.110242)
20. Yang Y, Blanchard A, Sapsis T, Perdikaris P. 2021 Output-weighted sampling for multi-armed bandits with extreme payoffs. (<http://arxiv.org/abs/2102.10085v2>)
21. Wendland H. 2004 *Scattered data approximation*. Cambridge, UK: Cambridge University Press.
22. Stuart AM, Teckentrup AL. 2018 Posterior consistency for Gaussian process approximations of Bayesian posterior distributions. *Math. Comput.* **87**, 721–753. (doi:10.1090/mcom/3244)
23. Mohamad MA, Sapsis TP. 2018 Sequential sampling strategy for extreme event statistics in nonlinear dynamical systems. *Proc. Natl Acad. Sci. USA* **115**, 11138–11143. (doi:10.1073/pnas.1813263115)
24. Vulikh BZ. 1976 *A brief course in the theory of functions of a real variable*. Moscow, Russia: Mir Publishers.
25. Nayfeh AH, Mook DT. 1984 *Nonlinear oscillations*. New York, NY: Wiley-Interscience.