

Using Approximate Bayesian Computation by Subset Simulation for Efficient Posterior Assessment of Dynamic State-Space Model Classes

Majid K. Vakilzadeh^{a,b}, James L. Beck^{a,*}, Thomas Abrahamsson^b

^a*Division of Engineering and Applied Science, California Institute of Technology, CA, USA*

^b*Department of Applied Mechanics, Chalmers University of Technology, Gothenburg, Sweden*

Abstract

Approximate Bayesian Computation (ABC) methods have gained in their popularity over the last decade because they expand the horizon of Bayesian parameter inference methods to the range of models for which only forward simulation is available. The majority of the ABC methods rely on the choice of a set of summary statistics to reduce the dimension of the data. However, as has been noted in the ABC literature, the lack of convergence guarantees that is induced by the absence of a vector of sufficient summary statistics that assures inter-model sufficiency over the set of competing models, hinders the use of the usual ABC methods when applied to Bayesian model selection or assessment. In this paper, we present a novel ABC model selection procedure for dynamical systems based on a newly appeared multi-level Markov chain Monte Carlo method, self-regulating ABC-SubSim, and a hierarchical state-space formulation of dynamic models. We show that this formulation makes it possible to independently approximate the model evidence required for assessing the posterior probability of each of the competing models. We also show that ABC-SubSim not only provides an estimate of the model evidence as a simple by-product but also it gives the posterior probability of each model as a function of the tolerance level, which allows the ABC model choices made in previous studies to be understood. We illustrate the performance of the proposed framework for ABC model updating and model class selection by applying it to two problems in Bayesian system identification: a single degree-of-freedom bilinear hysteretic oscillator and a three-story shear building with Masing hysteresis, both of which are subject to a seismic excitation.

Keywords: Approximate Bayesian Computation, Subset Simulation, Bayesian model selection, system identification, Bilinear and Masing hysteretic models

1. Introduction

In many areas of science such as biology, economics, social sciences and engineering, it is desired to make inference about the parameters of a mathematical model based on the experimental data from a real system in order to make more accurate predictions of the system behavior and to better understand it. Furthermore, there are invariably multiple candidate models with different mathematical forms to represent

*Corresponding author

Email address: jimbeck@caltech.edu (James L. Beck)

the system behavior and so there is a need to assess their plausibility based on the experiment data. The fully probabilistic Bayesian approach provides a rigorous framework to achieve these goals while also properly quantifying the uncertainty in the model parameters induced by uncertainty in the measurements and the accuracy of the mathematical model. In the Bayesian approach, a key idea is to construct a *stochastic model class* \mathcal{M} consisting of the following fundamental probability distributions [1]: a set of parameterized input-output probability models $p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{u}, \mathcal{M})$ for predicting the system behavior of interest \mathbf{y} for given input \mathbf{u} and a *prior* probability density function (PDF) $p(\boldsymbol{\theta}|\mathcal{M})$ over the parameter space $\Theta \in \mathbb{R}^{N_p}$ of \mathcal{M} that reflects the relative degree of plausibility of each input-output model in the set. When data \mathcal{D} consisting of the measured system input $\hat{\mathbf{u}}$ and output $\hat{\mathbf{z}}$ are available, the prior PDF $p(\boldsymbol{\theta}|\mathcal{M})$ can be updated through Bayes' Theorem to obtain the *posterior* PDF for the uncertain model parameters $\boldsymbol{\theta}$ as:

$$p(\boldsymbol{\theta}|\mathcal{D}, \mathcal{M}) = \frac{p(\hat{\mathbf{z}}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})}{p(\hat{\mathbf{z}}|\hat{\mathbf{u}}, \mathcal{M})} \propto p(\hat{\mathbf{z}}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M}) \quad (1)$$

where $p(\hat{\mathbf{z}}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M})$ denotes the *likelihood function* of $\boldsymbol{\theta}$ which gives the probability of getting the data based on the input-output probability model $p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{u}, \mathcal{M})$ and $p(\hat{\mathbf{z}}|\hat{\mathbf{u}}, \mathcal{M}) = \int_{\Theta} p(\hat{\mathbf{z}}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})d\boldsymbol{\theta}$ denotes the *evidence*, or marginal likelihood, for model class \mathcal{M} . Despite the fact that $p(\hat{\mathbf{z}}|\hat{\mathbf{u}}, \mathcal{M})$ is a constant and does not affect the shape of the posterior PDF, it is well known that it plays a crucial role in model class selection and averaging (e.g., [1–3]).

If $\mathbf{M} \equiv \{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_L\}$ is a set of competing candidate model classes for a real system, then the posterior probability of each model class is given by Bayes' Theorem at the model-class level:

$$p(\mathcal{M}_l|\mathcal{D}, \mathbf{M}) \propto p(\hat{\mathbf{z}}|\hat{\mathbf{u}}, \mathcal{M}_l)p(\mathcal{M}_l|\mathbf{M}), \quad l = 1, \dots, L \quad (2)$$

in which $p(\mathcal{M}_l|\mathbf{M})$ denotes the prior probability of the model class \mathcal{M}_l . This posterior distribution quantifies the plausibility of each \mathcal{M}_l to represent the uncertain behavior of the real system. If there are multiple model parameters treated as continuous stochastic variables for \mathcal{M}_l then calculation of its evidence, $p(\hat{\mathbf{z}}|\hat{\mathbf{u}}, \mathcal{M}_l)$, involves evaluation of a high dimensional integral over the parameter space that is computationally prohibitive. In addition, there are some model classes, e.g., hidden Markov models or dynamical state-space models, for which the likelihood function is difficult or even impossible to compute, but one might still be interested to perform Bayesian parameter inference or model selection. Approximate Bayesian Computation (ABC) methods were originally conceived to circumvent the need for computation of the likelihood by simulating samples from the corresponding input-output probability model $p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{u}, \mathcal{M})$.

The basic idea behind ABC is to avoid evaluation of the likelihood function in the posterior PDF $p(\boldsymbol{\theta}|\mathcal{D}, \mathcal{M}) \propto p(\hat{\mathbf{z}}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})$ over the parameter space $\boldsymbol{\theta}$ by using an augmented posterior PDF:

$$p(\boldsymbol{\theta}, \mathbf{y}|\mathcal{D}, \mathcal{M}) \propto P(\hat{\mathbf{z}}|\mathbf{y}, \boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M}) \quad (3)$$

over the joint space of the model parameters $\boldsymbol{\theta}$ and the model output \mathbf{y} that is simulated using the distribution $p(\mathbf{y}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M})$. The interesting point of this formulation is the degree of freedom brought by the choice of function $P(\hat{\mathbf{z}}|\mathbf{y}, \boldsymbol{\theta})$. The original ABC algorithm [4] defines $P(\hat{\mathbf{z}}|\mathbf{y}, \boldsymbol{\theta}) = \delta_{\hat{\mathbf{z}}}(\mathbf{y})$, where $\delta_{\hat{\mathbf{z}}}(\mathbf{y})$ is equal to 1

when $\hat{\mathbf{z}} = \mathbf{y}$ and equal to 0 otherwise, to retrieve the target posterior distribution when \mathbf{y} exactly matches $\hat{\mathbf{z}}$. However, the probability of generating exactly $\hat{\mathbf{z}} = \mathbf{y}$ is zero for continuous stochastic variables.

Pitchard et al. [5] broadened the realm of the applications for which ABC algorithm can be used by replacing the point mass at the observed output data $\hat{\mathbf{z}}$ with an indicator function $\mathbb{I}_{S(\epsilon)}(\mathbf{y})$, where $\mathbb{I}_{S(\epsilon)}(\mathbf{y})$ gives 1 over the set $S(\epsilon) = \{\mathbf{y} : \rho(\boldsymbol{\eta}(\hat{\mathbf{z}}) - \boldsymbol{\eta}(\mathbf{y})) \leq \epsilon\}$ and 0 elsewhere, for some chosen metric ρ and low-dimensional summary statistic $\boldsymbol{\eta}$. In this case, the approximate posterior PDF can be written as:

$$p(\boldsymbol{\theta}, \mathbf{y} | \mathcal{D}, \epsilon, \mathcal{M}) \propto \mathbb{I}_{S(\epsilon)}(\mathbf{y}) p(\mathbf{y} | \boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M}) p(\boldsymbol{\theta} | \mathcal{M}) \quad (4)$$

In this manner, Pitchard et al. [5] imposed two layers of approximation on the target posterior PDF $p(\boldsymbol{\theta}, \mathbf{y} | \mathcal{D}, \epsilon, \mathcal{M})$: (a) a tolerance parameter ϵ to assign a non-zero probability only for a region in $(\boldsymbol{\theta}, \mathbf{y})$ space where \mathbf{y} closely approximates $\hat{\mathbf{z}}$, i.e., $\hat{\mathbf{z}} \approx \mathbf{y}$; and (b) summary statistics $\boldsymbol{\eta}(\cdot)$ to summarize data in a low-dimensional space giving a weak form of the comparisons. If the summary statistics are sufficient for identification of $\boldsymbol{\theta}$, the approximation error vanishes, that is, $p(\boldsymbol{\theta}, \mathbf{y} | \mathcal{D}, \epsilon, \mathcal{M}) = p(\boldsymbol{\theta}, \mathbf{y} | \mathcal{D}, \mathcal{M})$ as $\epsilon \rightarrow 0$. Otherwise, summarizing the data induces a second level of approximation of the target posterior distribution. Algorithm 1 gives a pseudo-code to draw J samples from the approximate posterior distribution $p(\boldsymbol{\theta}, \mathbf{y} | \mathcal{D}, \epsilon, \mathcal{M})$.

Algorithm 1 Standard ABC rejection algorithm [4]

```

for  $j = 1$  to  $J$  do
  while  $\rho(\boldsymbol{\eta}(\mathbf{y}'), \boldsymbol{\eta}(\hat{\mathbf{z}})) > \epsilon$  do
    Draw a candidate sample  $\boldsymbol{\theta}' \sim p(\boldsymbol{\theta} | \mathcal{M})$ .
    Generate  $\mathbf{y}' \sim p(\mathbf{y} | \boldsymbol{\theta}', \mathcal{M})$ .
  end while
  Set  $(\boldsymbol{\theta}^{(j)}, \mathbf{y}^{(j)}) = (\boldsymbol{\theta}', \mathbf{y}')$ .
end for

```

Algorithm 1 gives samples from the true posterior distribution when the tolerance parameter ϵ is sufficiently small and the summary statistics $\boldsymbol{\eta}(\cdot)$ are sufficient. These conditions pose some difficulties for computer implementation of this algorithm which renders it far from a routine use for parameter inference and model selection. Firstly, a sufficiently small tolerance parameter ϵ means that only predicted model outputs \mathbf{y} lying in a small local neighborhood centered on the observed data vector $\hat{\mathbf{z}}$ are accepted. However, this leads to a problem of rare-event simulation and so if Algorithm 1 is used, the model output \mathbf{y} must be computed for a huge number of candidate samples in order to produce an acceptable sample size in the data-approximating region $S(\epsilon)$. Thus, many ABC algorithms have emerged to enhance the computational efficiency of the basic ABC rejection algorithm, e.g., ABC-MCMC [6–8], ABC-PRC [9, 10], ABC-SMC [11–13], ABC-PMC [14] and ABC-SubSim [15].

Secondly, the lack of a reasonable vector of summary statistics that works across models hinders the use of an ABC algorithm for model selection [16]. The ABC solution to the Bayesian model selection problem is to perform the inference at a model class level by incorporating a model index within the model parameters.

In this approach, a prior distribution $p(\mathcal{M}_l)$ for $l = 1, \dots, L$ is assigned to the competing models in the set $\mathbf{M} \equiv \{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_L\}$ along with a prior distribution $p(\boldsymbol{\theta}|\mathcal{M}_l)$ for the parameters conditional on the model index \mathcal{M}_l . Then, following the standard ABC rejection algorithm at the model class level, as given by the pseudo-code in Algorithm 2, the posterior probability of each of the candidate models can be readily estimated as follows:

$$p(\mathcal{M}_l|\hat{\mathbf{z}}, \epsilon) \approx \frac{1}{J} \sum_{j=1}^J \mathbb{I}_{m^{(j)}=\mathcal{M}_l}, \quad \text{for } l = 1, \dots, L \quad (5)$$

which is basically the frequency of acceptance from model \mathcal{M}_l . Let the vector $\boldsymbol{\eta}_{\mathbf{M}}(\mathbf{y}) = [\boldsymbol{\eta}_1^T(\mathbf{y}), \dots, \boldsymbol{\eta}_L^T(\mathbf{y})]^T$ denote the concatenation of the summary statistics used for all models. Grelaud et al. [17] reported that sufficiency of the summary statistics $\boldsymbol{\eta}_l(\mathbf{y})$ for models in the set \mathbf{M} does not guarantee sufficiency of $\boldsymbol{\eta}_{\mathbf{M}}(\mathbf{y})$ for comparison of those models. In fact, forming sufficient statistics for model comparison is not feasible in most problems for which ABC model selection has been implemented [18]. In this setting, the approximations made by Algorithm 2 does not always converge to the true model posterior probability [19].

Note that one can resort to entire data in Algorithm 2 instead of using the summary statistics. This avoids a loss of information in the metric $\rho(\boldsymbol{\eta}_{\mathbf{M}}(\mathbf{y}'), \boldsymbol{\eta}_{\mathbf{M}}(\hat{\mathbf{z}}))$ and, in turn, leads to a consistent decision for model choice [18]. Toni et al. [13] and Toni and Stumpf [20] developed an ABC algorithm based on Sequential Monte Carlo (ABC-SMC) and modified it such that a model index is incorporated within the model parameters. The posterior probability for each model class is then estimated based on the distance between the entire measured and simulated output. However, the estimated model posterior probabilities by ABC-SMC are affected by the choice of the tolerance parameter ϵ and the variance of the proposal distribution of the Markov chain used in ABC-SMC, which imposes a type of dependency on the estimates of $p(\mathcal{M}_l|\hat{\mathbf{z}}, \epsilon)$ that is irrelevant to the statistical problem under investigation [21].

Algorithm 2 Standard ABC algorithm for model comparison [21]

```

for  $j = 1$  to  $J$  do
  repeat
    Draw  $\mathcal{M}'$  from the prior distribution  $p(\mathcal{M}_l)$  for  $l = 1, \dots, L$ .
    Draw a candidate sample  $\boldsymbol{\theta}' \sim p(\boldsymbol{\theta}|\mathcal{M}')$ .
    Generate  $\mathbf{y}' \sim p(\mathbf{y}|\boldsymbol{\theta}', \mathcal{M}')$ .
  until  $\rho(\boldsymbol{\eta}_{\mathbf{M}}(\mathbf{y}'), \boldsymbol{\eta}_{\mathbf{M}}(\hat{\mathbf{z}})) \leq \epsilon$ 
  Set  $m^{(j)} = \mathcal{M}'$  and  $(\boldsymbol{\theta}^{(j)}, \mathbf{y}^{(j)}) = (\boldsymbol{\theta}', \mathbf{y}')$ .
end for

```

Wilkinson [22, 23] proposed to replace the indicator function $\mathbb{I}_{S(\epsilon)}(\mathbf{y})$ over the data approximating region $S(\epsilon)$ with a probability distribution function $p(\hat{\mathbf{z}}|\mathbf{y}, \epsilon)$, centered at \mathbf{y} , to obtain the following approximate posterior distribution:

$$p(\boldsymbol{\theta}, \mathbf{y}|\mathcal{D}, \epsilon, \mathcal{M}) \propto p(\hat{\mathbf{z}}|\mathbf{y}, \epsilon) p(\mathbf{y}|\boldsymbol{\theta}, \hat{\mathbf{u}}, \mathcal{M}) p(\boldsymbol{\theta}|\mathcal{M}) \quad (6)$$

This suggests a departure from the previous perspective so that $p(\boldsymbol{\theta}, \mathbf{y}|\hat{\mathbf{z}}, \mathcal{M})$ is now interpreted as an exact

posterior PDF for a new model in which the summary statistics are corrupted by a uniform error. Didelot et al. [16] showed that if one formulates the posterior distribution in (6) using the entire data, its normalizing constant converges to the marginal likelihood $p(\hat{\mathbf{z}}|\mathcal{M})$ as $\epsilon \rightarrow 0$.

In this study, we show that formulating a dynamical system as a general hierarchical state-space model enables us to independently estimate the model evidence for each model class. The recently proposed multi-level MCMC algorithm called Approximate Bayesian Computation by Subset Simulation (ABC-SubSim), is applied to solve the Bayesian inference problem of the uncertain parameters of the stochastic state-space model. We show that not only can the model evidence be estimated as a by-product of the ABC-SubSim algorithm, but also using the MCMC samples one can estimate the probability that model output \mathbf{y} falls into the data-approximating region $S(\epsilon)$ as a function of ϵ . The inherent difficulty of the ABC method for estimation of the parameters of the uncertain prediction error for stochastic state-space model is addressed and a new solution based on Laplace’s method of asymptotic approximation is presented. The effectiveness of the ABC-SubSim algorithm for Bayesian model updating and model class selection with simulated data is illustrated using two Bayesian system identification examples selected from the literature: (i) a single degree-of-freedom bilinear hysteretic oscillator [2] for which the true system is not included in the set of competing models; and (ii) a three-story shear building with Masing hysteresis [24], both of which are subject to seismic excitation. These numerical examples demonstrate the performance of ABC-SubSim for solving the Bayesian model updating and model class selection problem for dynamic models with a relatively large parameter space.

2. Formulation

In this section, we review the formulation of a Bayesian hierarchical model class for dynamical systems and then we employ the recently-appeared algorithm for Bayesian updating, Approximate Bayesian Computation by Subset Simulation (ABC-SubSim) [15] to explore its posterior PDF. We finally address the Bayesian model selection approach for the hierarchical stochastic state-space models.

2.1. Formulation of hierarchical stochastic model class

We construct a hierarchical stochastic state-space model class $\mathcal{M}(\epsilon)$ to predict the uncertain input-output behavior of a system. The reason for the dependence on a parameter ϵ will become evident later in this section.

We start with the general case of a discrete-time finite-dimensional state-space model of a real dynamic system:

$$\begin{aligned} \forall n \in \mathbb{Z}^+, \mathbf{x}_n &= \mathbf{f}_n(\mathbf{x}_{n-1}, \mathbf{u}_{n-1}, \boldsymbol{\theta}_s) && \text{(State evolution)} \\ \mathbf{y}_n &= \mathbf{g}_n(\mathbf{x}_n, \mathbf{u}_n, \boldsymbol{\theta}_s) && \text{(Output)} \end{aligned} \tag{7}$$

where $\mathbf{u}_n \in \mathbb{R}^{N_I}$, $\mathbf{x}_n \in \mathbb{R}^{N_s}$ and $\mathbf{y}_n \in \mathbb{R}^{N_o}$ denote the (external) input, model state and output vector at time t_n , and $\boldsymbol{\theta}_s \in \mathbb{R}^{N_p}$ is a vector of uncertain-valued model parameters. For the general case of stochastic embedding, we introduce the uncertain state and output prediction errors \mathbf{w}_n and \mathbf{v}_n into this deterministic

model to account for the model being always an approximation of the real system behavior, regardless of the choice of $\boldsymbol{\theta}_s$ [1]:

$$\begin{aligned}\forall n \in \mathbb{Z}^+, \mathbf{x}_n &= \mathbf{f}_n(\mathbf{x}_{n-1}, \mathbf{u}_{n-1}, \boldsymbol{\theta}_s) + \mathbf{w}_n \\ \mathbf{y}_n &= \mathbf{g}_n(\mathbf{x}_n, \mathbf{u}_n, \boldsymbol{\theta}_s) + \mathbf{v}_n\end{aligned}\tag{8}$$

where we now re-define \mathbf{x}_n and \mathbf{y}_n to be the dynamic system state and output vectors at time t_n , not the model state and output. The prior distributions, $\mathcal{N}(\mathbf{w}_n | \mathbf{0}, \mathbf{Q}_n(\boldsymbol{\theta}_w))$ and $\mathcal{N}(\mathbf{v}_n | \mathbf{0}, \mathbf{R}_n(\boldsymbol{\theta}_v))$, $\forall n \in \mathbb{Z}^+$, are chosen for the \mathbf{w}_n and \mathbf{v}_n based on the Principle of Maximum (Information) Entropy [25] under first and second moment constraints, where $\{\mathbf{w}_n\}_{n=1}^N$ and $\{\mathbf{v}_n\}_{n=1}^N$ are sequences of independent stochastic variables [1]. We add the uncertain parameters that specify these priors to the model parameters $\boldsymbol{\theta}_s$ and use $\boldsymbol{\theta} = [\boldsymbol{\theta}_s^T \ \boldsymbol{\theta}_w^T \ \boldsymbol{\theta}_v^T]^T$ to denote the uncertain parameter vector for the stochastic state-space model. If the initial state \mathbf{x}_0 is uncertain, we also add it to $\boldsymbol{\theta}$ and then choose a prior $p(\boldsymbol{\theta} | \mathcal{M}(\epsilon))$ for all of the model class parameters.

The defined stochastic state-space model defines a “hidden” Markov chain for the state time history $\{\mathbf{x}\}_{n=1}^N$ (which will also be denoted by the vector $\mathbf{x}_{1:N} = [\mathbf{x}_1^T, \dots, \mathbf{x}_N^T]^T \in \mathbb{R}^{NN_s}$) by implying a state transition PDF:

$$\forall n \in \mathbb{Z}^+, p(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{u}_{n-1}, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) = \mathcal{N}(\mathbf{x}_n | \mathbf{f}_n(\mathbf{x}_{n-1}, \mathbf{u}_{n-1}, \boldsymbol{\theta}), \mathbf{Q}_n(\boldsymbol{\theta}))\tag{9}$$

along with a state-to-output PDF:

$$\forall n \in \mathbb{Z}^+, p(\mathbf{y}_n | \mathbf{x}_n, \mathbf{u}_n, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) = \mathcal{N}(\mathbf{y}_n | \mathbf{g}_n(\mathbf{x}_n, \mathbf{u}_n, \boldsymbol{\theta}), \mathbf{R}_n(\boldsymbol{\theta}))\tag{10}$$

These, in turn, imply the following two probability models connecting the input, state and output discrete-time histories (which are readily sampled because each factor is Gaussian):

$$p(\mathbf{x}_{1:N} | \mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) = \prod_{n=1}^N p(\mathbf{x}_n | \mathbf{x}_{n-1}, \mathbf{u}_{n-1}, \boldsymbol{\theta}, \mathcal{M}(\epsilon))\tag{11}$$

$$p(\mathbf{y}_{1:N} | \mathbf{x}_{1:N}, \mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) = \prod_{n=1}^N p(\mathbf{y}_n | \mathbf{x}_n, \mathbf{u}_n, \boldsymbol{\theta}, \mathcal{M}(\epsilon))\tag{12}$$

The stochastic input-output model (or forward model) for given parameter vector $\boldsymbol{\theta}$ is then:

$$p(\mathbf{y}_{1:N} | \mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) = \int p(\mathbf{y}_{1:N} | \mathbf{x}_{1:N}, \mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) p(\mathbf{x}_{1:N} | \mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) d\mathbf{x}_{1:N}\tag{13}$$

This high-dimensional integral usually cannot be done analytically. We will therefore structure the stochastic input-output model using a Bayesian hierarchical model to avoid the integration in (13), where $\mathbf{y}_{1:N}$ and $\mathbf{x}_{1:N}$ are treated in the same way as the model parameters $\boldsymbol{\theta}$.

This can be done by extending the stochastic model to predict the measured system output \mathbf{z}_n at time t_n :

$$\mathbf{z}_n = \mathbf{y}_n + \mathbf{e}_n = \mathbf{g}_n(\mathbf{x}_n, \mathbf{u}_n, \boldsymbol{\theta}) + \mathbf{v}_n + \mathbf{e}_n\tag{14}$$

where \mathbf{e}_n denotes the uncertain measurement error at time t_n . Unlike $\{\mathbf{w}_n\}_{n=1}^N$ and $\{\mathbf{v}_n\}_{n=1}^N$, the stochastic sequence $\{\mathbf{e}_n\}_{n=1}^N$ will not necessarily be modeled as a set of independent stochastic variables; rather, we allow for dependence by specifying a joint PDF for $\mathbf{e}_{1:n}$ for any $n \in \mathbb{Z}^+$. A simple choice for the probability model for the measurement error $\mathbf{e}_{1:n}$ is a uniform PDF:

$$p(\mathbf{e}_{1:n}|\mathcal{M}(\epsilon)) = \frac{\mathbb{I}_\epsilon(\mathbf{e}_{1:n})}{V_n(\epsilon)}, \quad \forall \mathbf{e}_{1:n} \in \mathbb{R}^{N_o n} \quad (15)$$

where $\mathbb{I}_\epsilon(\mathbf{e}_{1:n})$ is the indicator function for the set $S(\epsilon) = \{\mathbf{e}_{1:n} \in \mathbb{R}^{N_o n} : \|\mathbf{e}_{1:n}\| \leq \epsilon\}$ for some vector norm (e.g., $\|\cdot\|_\infty$ or $\|\cdot\|_2$) on $\mathbb{R}^{N_o n}$, and $V_n(\epsilon) = \int_{\mathbb{R}^{N_o n}} \mathbb{I}_\epsilon(\mathbf{e}_{1:n}) d\mathbf{e}_{1:n}$ is the volume of region $S(\epsilon)$. This finally reveals that ϵ in $\mathcal{M}(\epsilon)$ is a scalar upper bound on the measurement errors that parameterizes the chosen prior PDF for $\mathbf{e}_{1:n}$, $\forall n \in \mathbb{Z}^+$. Thus, the predictive PDF for the observed system output (sensor output) $\mathbf{z}_{1:n}$, given the actual system output $\mathbf{y}_{1:n}$, is then given by:

$$p(\mathbf{z}_{1:n}|\mathbf{y}_{1:n}, \mathcal{M}(\epsilon)) = p(\mathbf{e}_{1:n}|\mathcal{M}(\epsilon)) \Big|_{\mathbf{e}_{1:n}=\mathbf{z}_{1:n}-\mathbf{y}_{1:n}} = \begin{cases} V_n(\epsilon)^{-1} & \text{if } \|\mathbf{z}_{1:n} - \mathbf{y}_{1:n}\| \leq \epsilon \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

The specification of the hierarchical prior PDF:

$$p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\mathbf{u}_{0:N}, \mathcal{M}(\epsilon)) = p(\mathbf{y}_{1:N}|\mathbf{x}_{1:N}, \mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon))p(\mathbf{x}_{1:N}|\mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon))p(\boldsymbol{\theta}|\mathcal{M}(\epsilon)) \quad (17)$$

completes the definition of the stochastic model class $\mathcal{M}(\epsilon)$. Thus, the PDFs in (16) and (17) define a hierarchical stochastic model class $\mathcal{M}(\epsilon)$ for each value of ϵ [26]. Here, $p(\boldsymbol{\theta}|\mathcal{M}(\epsilon))$ is assumed to be independent of the system input history $\mathbf{u}_{0:N}$.

2.2. Bayesian model updating

If measured system input and system output data:

$$\mathcal{D}_N = \{\hat{\mathbf{u}}_{0:N}, \hat{\mathbf{z}}_{1:N}\}$$

are available from the dynamic system, then the predictive PDF in (16) with $n = N$ gives the likelihood function:

$$p(\hat{\mathbf{z}}_{1:N}|\mathbf{y}_{1:N}, \mathcal{M}(\epsilon)) = \frac{\mathbb{I}_{\mathcal{D}(\epsilon)}(\mathbf{y}_{1:N})}{V_N(\epsilon)} \quad (18)$$

with the indicator function defined over the set $\mathcal{D}(\epsilon) = \{\mathbf{y}_{1:N} \in \mathbb{R}^{N N_o} : \|\mathbf{y}_{1:N} - \hat{\mathbf{z}}_{1:N}\| < \epsilon\}$, where $\|\cdot\|$ is some vector norm on $\mathbb{R}^{N N_o}$.

The posterior PDF for stochastic model class $\mathcal{M}(\epsilon)$ is then given by Bayes' Theorem:

$$p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\mathcal{D}_N, \mathcal{M}(\epsilon)) = E(\epsilon)^{-1} \frac{\mathbb{I}_{\mathcal{D}(\epsilon)}(\mathbf{y}_{1:N})}{V_N(\epsilon)} p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}(\epsilon)) \quad (19)$$

where the evidence for $\mathcal{M}(\epsilon)$ is then defined as:

$$\begin{aligned} E(\epsilon) &= p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}(\epsilon)) \\ &= \int p(\hat{\mathbf{z}}_{1:N}|\mathbf{y}_{1:N}, \epsilon) p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}(\epsilon)) d\mathbf{y}_{1:N} d\mathbf{x}_{1:N} d\boldsymbol{\theta} \\ &= \int \frac{\mathbb{I}_{\mathcal{D}(\epsilon)}(\mathbf{y}_{1:N})}{V_N(\epsilon)} p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}(\epsilon)) d\mathbf{y}_{1:N} d\mathbf{x}_{1:N} d\boldsymbol{\theta} \end{aligned} \quad (20)$$

The theory for the hierarchical model and its updating presented so far in Section 2 is general and valid for any $\epsilon > 0$ deemed appropriate. For the application of ABC, we suppose that $\mathcal{M}(0) \equiv \mathcal{M}(\epsilon \rightarrow 0)$ is actually the stochastic model class of interest. In this case, the output prediction errors \mathbf{v}_n in (8) represent a combination of measurement errors and modeling errors with respect to the system output. For ϵ sufficiently small, the set $\mathcal{D}(\epsilon)$ of outputs $\mathbf{y}_{1:N}$ will converge to the observed output vector $\hat{\mathbf{z}}_{1:N}$ and the posterior PDF in (19) for stochastic model class $\mathcal{M}(\epsilon)$ will converge to the desired posterior distribution of the model parameters $p(\boldsymbol{\theta}|\mathcal{D}_N, \mathcal{M}(0))$ after marginalization. This can be shown in the following way:

$$\begin{aligned}
p(\boldsymbol{\theta}|\mathcal{D}_N, \mathcal{M}(\epsilon)) &= \int p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\mathcal{D}_N, \mathcal{M}(\epsilon)) d\mathbf{y}_{1:N} d\mathbf{x}_{1:N} = \\
&\frac{1}{E(\epsilon)} \int \frac{\mathbb{I}_{\mathcal{D}(\epsilon)}(\mathbf{y}_{1:N})}{V_N(\epsilon)} p(\mathbf{y}_{1:N}|\mathbf{x}_{1:N}, \boldsymbol{\theta}, \hat{\mathbf{u}}_{0:N}, \mathcal{M}(\epsilon)) p(\mathbf{x}_{1:N}|\boldsymbol{\theta}, \hat{\mathbf{u}}_{0:N}, \mathcal{M}(\epsilon)) p(\boldsymbol{\theta}|\mathcal{M}(\epsilon)) d\mathbf{y}_{1:N} d\mathbf{x}_{1:N}
\end{aligned} \tag{21}$$

Utilizing (13) to simplify the integrand in (21) and allowing $\epsilon \rightarrow 0$ gives:

$$\begin{aligned}
p(\boldsymbol{\theta}|\mathcal{D}_N, \mathcal{M}(0)) &= \frac{1}{E(0)} \int \delta_{\hat{\mathbf{z}}_{1:N}}(\mathbf{y}_{1:N}) p(\mathbf{y}_{1:N}|\boldsymbol{\theta}, \hat{\mathbf{u}}_{0:N}, \mathcal{M}(0)) p(\boldsymbol{\theta}|\mathcal{M}(0)) d\mathbf{y}_{1:N} \\
&= \frac{p(\hat{\mathbf{z}}_{1:N}|\boldsymbol{\theta}, \hat{\mathbf{u}}_{0:N}, \mathcal{M}(0)) p(\boldsymbol{\theta}|\mathcal{M}(0))}{E(0)}
\end{aligned} \tag{22}$$

in which $E(0) \equiv E(\epsilon \rightarrow 0) = \int p(\hat{\mathbf{z}}_{1:N}|\boldsymbol{\theta}, \hat{\mathbf{u}}_{0:N}, \mathcal{M}(0)) p(\boldsymbol{\theta}|\mathcal{M}(0)) d\boldsymbol{\theta}$, the evidence for $\mathcal{M}(0)$. However, if the tolerance ϵ is small, the acceptance rate is small so that the posterior distribution is estimated by only a few points unless the Algorithm 1 is run for a very long time. On the contrary, if the tolerance is too large, ($\epsilon \rightarrow \infty$) then the samples come from the hierarchical prior $p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}(\epsilon))$. Thus, a rational choice for ϵ should strike a balance between computability and accuracy.

Remark 1. Vakilzadeh et al. [26] showed that for the hierarchical stochastic model class $\mathcal{M}(\epsilon)$, the exact posterior PDF (19) using a uniformly-distributed uncertain measurement error in the output space is identical to the ABC posterior PDF given for no measurement error, see (16) and (20) in Vakilzadeh et al. [26]. Thus, ABC-SubSim that was originally developed by Chiachio et al. [15] to draw samples from an ABC posterior PDF can be used to solve the exact Bayesian problem for the hierarchical stochastic model class $\mathcal{M}(\epsilon)$ given by (16) and (17).

2.2.1. Approximate Bayesian Computation by Subset Simulation (ABC-SubSim)

For a good approximation of the posterior distribution for stochastic model class $\mathcal{M}(0)$, we want $\mathcal{D}(\epsilon)$ to be small neighborhood centered on the data vector $\hat{\mathbf{z}}_{1:N}$ in \mathbb{R}^{NN_0} . The probability $P(\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon)|\epsilon)$ will then be small and so, on average, many candidate samples will be required to generate an acceptable sample having $\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon)$ (on average, $1/P(\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon)|\epsilon)$ candidate samples will be required). Henceforth, we use $P(\mathcal{D}(\epsilon)) \equiv P(\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon)|\epsilon)$ for a simpler notation. ABC-SubSim [15] was originally developed to address this problem by exploiting the Subset Simulation method for efficient rare-event simulation [27]. The reader is referred to [27] and [15] for a detailed explanation of how Subset Simulation and ABC-SubSim work.

The basic idea behind ABC-SubSim is to define the data-approximating region $\mathcal{D}(\epsilon)$ as the intersection of a set of nested decreasing data-approximating regions, $\mathcal{D}(\epsilon_j)$ of “radius” ϵ_j , as defined above after (18),

where $\epsilon_{j+1} < \epsilon_j$. The probability $P(\mathcal{D}(\epsilon))$ can be then estimated as a product of conditional probabilities:

$$P(\mathcal{D}(\epsilon)) = P(\mathcal{D}(\epsilon_1)) \prod_{j=2}^m P(\mathcal{D}(\epsilon_j) | \mathcal{D}(\epsilon_{j-1})) \quad (23)$$

The intermediate data-approximating regions are adaptively selected such that all conditional probabilities can be made large. Thus, ABC-SubSim replaces a problem involving rare-event simulation by a sequence of problems involving simulation of more frequent events.

The simulation algorithm starts by drawing N_t independent and identically distributed samples $(\mathbf{y}_{1:N}^{i,(1)}, \mathbf{x}_{1:N}^{i,(1)}, \boldsymbol{\theta}^{i,(1)})$ from the hierarchical prior $p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta} | \mathbf{u}_{0:N}, \mathcal{M})$. The corresponding metric value $\epsilon^{i,(1)} = \|\mathbf{y}_{1:N}^{i,(1)} - \hat{\mathbf{z}}_{1:N}\|$ is then evaluated and samples are sorted in decreasing order of magnitude of their metric value so that $\epsilon^{1,(1)} \geq \epsilon^{2,(1)} \geq \dots \geq \epsilon^{N_t,(1)}$. Thus, probability $P(\mathcal{D}(\epsilon^{i,(1)}))$ corresponding to tolerance level $\epsilon^{i,(1)}$ can be approximated based on the samples by:

$$P(\mathcal{D}(\epsilon^{i,(1)})) \approx \frac{1}{N_t} \sum_{i=1}^{N_t} \mathbb{I}_{\mathcal{D}(\epsilon^{i,(1)})}(\mathbf{y}_{1:N}^{i,(1)}) = \frac{N_t - i}{N_t} = \hat{P}_{\epsilon^{i,(1)}} \quad (24)$$

In ABC-SubSim, the initial tolerance level ϵ_1 is chosen using (24) so that $\hat{P}_{\epsilon_1} = P_0$, an assigned probability whose value is best selected from the range $[0.1, 0.3]$ (see Remark 2).

For higher simulation levels $j \geq 2$, sampling from the conditional PDF $(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta} | \mathbf{y}_{1:N} \in \mathcal{D}(\epsilon_{j-1}), \hat{\mathbf{u}}_{0:N})$ can be achieved by means of a component-wise MCMC algorithm, called Modified Metropolis Algorithm (MMA) in [27], at the expense of generating dependent samples. Using the MCMC samples $(\mathbf{y}_{1:N}^{i,(j)}, \mathbf{x}_{1:N}^{i,(j)}, \boldsymbol{\theta}^{i,(j)})$, the conditional probability can be estimated as:

$$P(\mathcal{D}(\epsilon_j) | \mathcal{D}(\epsilon_{j-1})) \approx \frac{1}{N_t} \sum_{i=1}^{N_t} \mathbb{I}_{\mathcal{D}(\epsilon_j)}(\mathbf{y}_{1:N}^{i,(j)}) = \hat{P}_{\epsilon_j | \epsilon_{j-1}} \quad (25)$$

where $P(\mathcal{D}(\epsilon_j) | \mathcal{D}(\epsilon_{j-1})) \equiv P(\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon_j) | \mathbf{y}_{1:N} \in \mathcal{D}(\epsilon_{j-1}))$ is the conditional probability at the j th simulation level. In ABC-SubSim, the intermediate tolerance levels ϵ_j ($j \geq 2$) are adaptively determined as in Subset Simulation [27], such that the sample estimate $\hat{P}_{\epsilon_j | \epsilon_{j-1}}$ of the conditional probabilities $P(\mathcal{D}(\epsilon_j) | \mathcal{D}(\epsilon_{j-1}))$ is equal to an assigned value P_0 . To this end, we rearrange the samples generated by MCMC at the j th simulation level in decreasing order of the magnitude for their associated metric values $\{\epsilon^{i,(j)}, i = 1, \dots, N_t\}$. Then, the tolerance ϵ_j can be determined as the $100P_0$ percentile of the set of metric values $\epsilon^{i,(j)}, i = 1, \dots, N_t$. For instance, we define $\epsilon_j = (\epsilon^{N_t(1-P_0),(j)} + \epsilon^{N_t(1-P_0)+1,(j)})/2$.

Observe that the MCMC samples generated at the j th simulation level that fell in the data-approximation region $\mathcal{D}(\epsilon_j)$, i.e., samples corresponding to $\{\epsilon^{N_t(1-P_0)+i,(j)}, i = 1, \dots, N_t P_0\}$, are distributed as $p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta} | \mathbf{y}_{1:N} \in \mathcal{D}(\epsilon_j), \hat{\mathbf{u}}_{0:N})$, and thus they provide $N_t P_0$ seeds in $\mathcal{D}(\epsilon_j)$. Thus, a Markov chain of length $(1 - 1/P_0)$ can be initiated from each of the seeds to populate $\mathcal{D}(\epsilon_j)$ with N_t samples and it will be in its stationary state from the start, giving *perfect sampling* (e.g., [28]).

Each of the sorted metric values $\epsilon^{i,(j)}$ gives a corresponding probability $P(\mathcal{D}(\epsilon^{i,(j)}))$, which can be ap-

proximated based on samples by:

$$\begin{aligned}
 P(\mathcal{D}(\epsilon^{i,(j)})) &= P(\mathcal{D}(\epsilon_1)) \left[\prod_{k=2}^{j-1} P(\mathcal{D}(\epsilon_k) | \mathcal{D}(\epsilon_{k-1})) \right] P(\mathcal{D}(\epsilon^{i,(j)}) | \mathcal{D}(\epsilon_{j-1})) \\
 &\approx P_0^{j-1} \frac{N_t - i}{N_t} = \hat{P}_{\epsilon^{i,(j)}}
 \end{aligned}
 \tag{26}$$

The algorithm proceeds in this way until ϵ_m becomes smaller than an appropriate final tolerance level ϵ .

Remark 2. To run the ABC-SubSim Algorithm, one needs to specify appropriate values for: (i) the number of samples per simulation level N_t , and (ii) the conditional probability P_0 of each stage of ABC-SubSim. Choosing a large value for P_0 increases the number of simulation levels required to achieve a specified tolerance level ϵ for a fixed N_t . Thus, the higher P_0 is, the higher is the computational burden of the algorithm. On the other hand, a choice of a small value for P_0 decreases the quality of the posterior approximation. Recently, Zuev et al. [29] implemented a rigorous sensitivity analysis of Subset Simulation and reported that an optimal choice for P_0 lies in the range $[0.1, 0.3]$. Furthermore, for convenience $P_0 N_t$ and $1/P_0$ are selected as positive integers.

Remark 3. (24) and (26) produce an estimate of the probability $P(\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon))$ as a function of tolerance level ϵ , covering the large probabilities to small probability regimes (e.g., see Figure 5). This in turn means that the calculation of the evidence of a candidate model class for different tolerance levels is a simple by-product of the ABC-SubSim algorithm. In the next section, we will show that this property makes ABC-SubSim an effective algorithm for Bayesian model selection.

Remark 4. In [26], a modification of ABC-SubSim, called self-regulating ABC-SubSim, has been proposed. The key idea behind this method is to enhance efficient exploration of the posterior distribution over the parameter space. A way to achieve this goal is to learn the proposal variance for the MMA algorithm in each simulation level on-the-fly in order to coerce the mean acceptance probability for a candidate sample to be close to a desired target value. Another key benefit of incorporating the self-regulating algorithm in ABC-SubSim is that it gives a heuristic rule to automatically determine the number of simulation levels m as follows: *stop the algorithm when the average acceptance rate drops significantly*. This recently proposed variant of ABC-SubSim is used in this study to draw samples from the posterior PDF $p(\mathbf{y}_{1:N}, \boldsymbol{\theta} | \mathcal{D}_N, \mathcal{M}(\epsilon))$ of stochastic model class $\mathcal{M}(\epsilon)$.

2.3. Bayesian model class assessment

Bayesian model class selection provides a rigorous framework to compare the performance of a set of candidate model classes in describing the experimental data. As exposed by Marin et al. [21], there are well-known limitations of the ABC approach to the Bayesian model selection problem, mainly due to lack of a sufficient summary statistics that work across models. However, in Section 2.2 we showed that formulating the standard ABC posterior distribution as the exact posterior PDF (19) for the hierarchical state-space

model class allows us to independently estimate the evidence $E(\epsilon)$ for each alternative candidate model $\mathcal{M}(\epsilon)$ as the normalizing constant associated with the exact posterior PDF (19). Didelot et al. [16] demonstrated that under mild continuity conditions, this normalizing constant converges to the true model evidence $E(0)$ as $\epsilon \rightarrow 0$. In this section, our objective is to show that calculation of the evidence $E(\epsilon)$ is a simple by-product of the ABC-SubSim algorithm.

Consider a set $\mathbf{M} \equiv \{\mathcal{M}_1(\epsilon_{\mathcal{M}_1}), \mathcal{M}_2(\epsilon_{\mathcal{M}_2}), \dots, \mathcal{M}_L(\epsilon_{\mathcal{M}_L})\}$ of L Bayesian hierarchical model classes for representing a system. In Bayesian model selection, models in \mathbf{M} are ranked based on their probabilities conditioned on the data \mathcal{D}_N that is given by Bayes' Theorem:

$$P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j})|\mathcal{D}_N) = \frac{p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}_j(\epsilon_{\mathcal{M}_j}))P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j})|\mathbf{M})}{\sum_{i=1}^L p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}_i(\epsilon_{\mathcal{M}_i}))P(\mathcal{M}_i(\epsilon_{\mathcal{M}_i})|\mathbf{M})} \quad (27)$$

where $P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j})|\mathbf{M})$ denotes the prior probability of $\mathcal{M}_j(\epsilon_{\mathcal{M}_j})$ that indicates the modeler's belief about the initial relative plausibility of $\mathcal{M}_j(\epsilon_{\mathcal{M}_j})$ within the set \mathbf{M} . The factor $p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \mathcal{M}_j(\epsilon_{\mathcal{M}_j}))$, which is the evidence (or marginal likelihood) for $\mathcal{M}_j(\epsilon_{\mathcal{M}_j})$, indicates the probability of data \mathcal{D}_N according to $\mathcal{M}_j(\epsilon_{\mathcal{M}_j})$.

For the specific choice of Bayesian hierarchical model class, the evidence can be estimated by (20). However, its calculation requires the evaluation of a high-dimensional integral which is the computationally challenging step in Bayesian model selection, especially as $\epsilon \rightarrow 0$. ABC-SubSim provides a straightforward approximation for it via the conditional probabilities involved in the Subset Simulation. Indeed, the last integral in (20) is the probability $P(\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon_{\mathcal{M}_j})|\mathcal{M}_j)$ that $\mathbf{y}_{1:N}$ belongs to $\mathcal{D}(\epsilon_{\mathcal{M}_j}) = \{\mathbf{y}_{1:N} \in \mathbb{R}^{NN_o} : \|\mathbf{y}_{1:N} - \hat{\mathbf{z}}_{1:N}\| \leq \epsilon_{\mathcal{M}_j}\}$. This probability can be readily estimated as a by-product of ABC-SubSim by using (24) and (26). Thus, for a particular tolerance level $\epsilon_{\mathcal{M}_j}$ and model class $\mathcal{M}_j(\epsilon_{\mathcal{M}_j})$, the evidence is estimated by:

$$\hat{E}_{\mathcal{M}_j} = \frac{P(\mathbf{y}_{1:N} \in \mathcal{D}(\epsilon_{\mathcal{M}_j})|\mathcal{M}_j)}{V_N(\epsilon_{\mathcal{M}_j})} = \frac{1}{V_N(\epsilon_{\mathcal{M}_j})} P_0^{i-1} P_i \quad (28)$$

where i would be such that $\epsilon_i \leq \epsilon_{\mathcal{M}_j} < \epsilon_{i-1}$, in which the intermediate "radii" ϵ_i 's are automatically chosen by ABC-SubSim, and P_i is the fraction of samples generated in $\mathcal{D}(\epsilon_{i-1})$ that lie in $\mathcal{D}(\epsilon_{\mathcal{M}_j})$. Here, $V_N(\epsilon)$ is the volume of the ball centered at $\hat{\mathbf{z}}_{1:N}$, with radius ϵ and norm $\|\cdot\|$. If $\mathcal{D}(\epsilon)$ is defined by using $\|\cdot\|_\infty$ on \mathbb{R}^{NN_o} , then $V_N(\epsilon) = (2\epsilon)^{NN_o}$ and if is defined by the Euclidean norm, then $V_N(\epsilon) = \pi^{NN_o/2}/\Gamma(NN_o/2 + 1)\epsilon^{NN_o}$. It is worth noting that $V_N(\epsilon_{\mathcal{M}_j})$ is not needed for posterior model class assessment if we choose the same tolerance level in ABC for each of the L candidate model classes. Wilkinson [22, 23] showed that a standard ABC posterior gives an exact posterior distribution for a new model under the assumption that the summary statistics are corrupted with a uniform additive error term. However, formulating standard ABC based on summary statistics hinders the independent approximation of evidence for each candidate model [30]. Here, the estimate of the model evidence in (28) is a result of formulating a dynamic problem in terms of a general hierarchical stochastic state-space model where the likelihood function $p(\hat{\mathbf{z}}_{1:N}|\mathbf{y}_{1:N}, \mathcal{M}(\epsilon))$ is expressed using the entire data $\hat{\mathbf{z}}_{1:N}$ and ABC-SubSim readily produces an unbiased approximation of the evidence.

2.4. Treatment of uncertain prediction error model parameters

In reality the parameters specifying the covariance matrix of the uncertain prediction errors are unknown and they should be learned from the data \mathcal{D}_N . However, Vakilzadeh et al. [26] found that shrinking the data-approximation region, $\mathcal{D}(\epsilon \rightarrow 0)$, drives the covariance matrix of the prediction errors to zero. This is due to the fact that the formulation of the stochastic state-space model has built-in prediction errors such that a sample of the model prediction of the system output includes a realization of these error signals. For high-dimensional data vectors, the probability of drawing sequences of random numbers $\{\mathbf{w}_n\}_{n=1}^N$ and $\{\mathbf{v}_n\}_{n=1}^N$ such that the measured data exactly matches the simulated system output is essentially zero, even if the measured data is synthetic and generated by adding realizations of the prediction-error signals when calculating the output from the chosen model for system identification. Therefore, one can conclude that not only are the parameters specifying the uncertain prediction errors unidentifiable based on the likelihood function $p(\hat{\mathbf{z}}_{1:N}|\mathbf{y}_{1:N}, \mathcal{M}(\epsilon))$ constructed from the entire data, but also that when $\epsilon \rightarrow 0$, the stochastic input-to-output model (13) reduces to the underlying deterministic model.

To get around this problem, we treat the parameters $\boldsymbol{\theta}_w$ and $\boldsymbol{\theta}_v$ specifying the state and output uncertain errors as nuisance parameters and eliminate them from the analysis to achieve a posterior distribution for the structural model parameters $\boldsymbol{\theta}_s$ alone. Vakilzadeh et al. [26] defined appropriate conjugate priors for the $\boldsymbol{\theta}_w$ and $\boldsymbol{\theta}_v$ and then integrated them out from the posterior distribution (19) and worked with the resulting marginal posterior distribution for $\boldsymbol{\theta}_s$. Although the marginalization technique automatically incorporates the posterior uncertainty induced by the nuisance parameters [31], it relies on choosing appropriate prior distributions for the nuisance parameters. Here, we present an alternative approach that avoids having to choose these prior distributions by using Laplace's method to replace the marginal posterior distribution of $\boldsymbol{\theta}_s$ by its asymptotic approximation [32], which turns out to be insensitive to the prior distribution adopted for the nuisance parameters. However, the use of this approximate method to solve the Bayesian problem for the hierarchical model class $\mathcal{M}(\epsilon)$ given by (16) and (17) is restricted to the cases where the state prediction error \mathbf{w}_n is suppressed in (8) so that the hierarchical prior PDF in (17) is replaced by:

$$p(\mathbf{y}_{1:N}, \boldsymbol{\theta}|\mathbf{u}_{0:N}, \mathcal{M}(\epsilon)) = p(\mathbf{y}_{1:N}|\mathbf{u}_{0:N}, \boldsymbol{\theta}, \mathcal{M}(\epsilon)) p(\boldsymbol{\theta}|\mathcal{M}(\epsilon)) \quad (29)$$

Actually, it is quite common in Bayesian inference for dynamic systems that the uncertainty in the system output is modeled by only an output prediction error \mathbf{v}_n . Furthermore, we model the covariance matrix for \mathbf{v}_n as a time-invariant isotropic diagonal matrix $\mathbf{R}_n(\boldsymbol{\theta}_v) = \sigma^2 \mathbf{I}_{N_o}$ giving the scalar nuisance parameter $\theta_v = \sigma^2$. Throughout this section, we drop the conditioning on the model class $\mathcal{M}(\epsilon)$ since the formulation is valid for any choice of model class.

For reasons which will become clearer later, we temporarily neglect the measurement error in (14) and assume that the predictive PDF for the sensor output can be expressed by (13). For data \mathcal{D}_N , this gives the more common likelihood function $p(\hat{\mathbf{z}}_{1:N}|\mathbf{u}_{0:N}, \boldsymbol{\theta})$ where \mathbf{v}_n in (8) is often taken as a combination of

measurement and modeling errors. The marginal likelihood function of the model parameters $\boldsymbol{\theta}_s$ is now:

$$p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s) = \int p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \theta_v) p(\theta_v) d\theta_v \quad (30)$$

Our objective here is to develop an asymptotic approximation for this integral. To this end, suppose that $\hat{\theta}_v(\boldsymbol{\theta}_s)$ maximizes the likelihood function $p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \theta_v)$ for a fixed value of the model parameters $\boldsymbol{\theta}_s$. By expanding $\ln p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \theta_v)$ in a second-order Taylor series about $\hat{\theta}_v(\boldsymbol{\theta}_s)$, one obtains the following local approximation [33]:

$$p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \theta_v) = p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\theta}_v(\boldsymbol{\theta}_s)) \exp\left(-\frac{1}{2}H(\hat{\theta}_v(\boldsymbol{\theta}_s))[\theta_v - \hat{\theta}_v(\boldsymbol{\theta}_s)]^2\right) \quad (31)$$

where the Hessian $H(\hat{\theta}_v(\boldsymbol{\theta}_s))$ is given by:

$$H(\hat{\theta}_v(\boldsymbol{\theta}_s)) = -\left.\frac{\partial^2 \ln p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta})}{\partial \theta_v^2}\right|_{\theta_v=\hat{\theta}_v(\boldsymbol{\theta}_s)} \quad (32)$$

and the MLE (maximum likelihood estimate) of the nuisance parameter θ_v for a fixed value of the model parameters $\boldsymbol{\theta}_s$ is equal to the mean prediction error value [33]:

$$\hat{\theta}_v(\boldsymbol{\theta}_s) = \frac{1}{NN_o} \|\hat{\mathbf{z}}_{1:N} - \mathbf{g}_{1:N}(\boldsymbol{\theta}_s)\|_2^2 \quad (33)$$

where $\mathbf{g}_{1:N} = [\mathbf{g}_1^T, \mathbf{g}_2^T, \dots, \mathbf{g}_N^T]^T$. Here $\mathbf{g}_n(\boldsymbol{\theta}_s)$ denotes $\mathbf{g}_n(\hat{\mathbf{x}}_n, \mathbf{u}_n, \boldsymbol{\theta}_s)$ from (7) where $\hat{\mathbf{x}}_n$ is the deterministic solution of $\mathbf{x}_n = \mathbf{f}_n(\mathbf{x}_{n-1}, \hat{\mathbf{u}}_{n-1}, \boldsymbol{\theta}_s)$ from (7). For a large number N of sampling times, and given structural model parameters $\boldsymbol{\theta}_s$, $p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \theta_v)$ will be very peaked at the optimal parameter $\hat{\theta}_v(\boldsymbol{\theta}_s)$. Therefore, Laplace's method for asymptotic approximation can be applied to the integral in (30) to obtain the following approximation for the marginal likelihood distribution, which is essentially derived by substituting (31) into (30) [33, 34]:

$$p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s) = \sqrt{2\pi}H(\hat{\theta}_v(\boldsymbol{\theta}_s))^{-\frac{1}{2}} p(\hat{\theta}_v(\boldsymbol{\theta}_s)) p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\theta}_v(\boldsymbol{\theta}_s)) [1 + \mathcal{O}(N^{-1})] \quad (34)$$

Remark 5. The structure of the Hessian matrix of $\ln p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta})$ regarding the uncertain parameter vector $\boldsymbol{\theta} = [\boldsymbol{\theta}_s^T \theta_v]^T$ is block diagonal with one block being an $N_p \times N_p$ matrix corresponding to the structural model parameters $\boldsymbol{\theta}_s$ and the other block being simply the scalar $H(\hat{\theta}_v(\boldsymbol{\theta}_s)) = \frac{1}{2}NN_o\hat{\theta}_v(\boldsymbol{\theta}_s)^{-2}$ corresponding to the scalar nuisance parameter θ_v [33]. In this setting, Cox and Reid [35] showed that the MLE $\hat{\theta}_v(\boldsymbol{\theta}_s)$ of the nuisance parameter can be treated approximately as a constant $\hat{\theta}_v$ in $p(\hat{\theta}_v(\boldsymbol{\theta}_s))$ and $H(\hat{\theta}_v(\boldsymbol{\theta}_s))$. We will show next that using this feature, the posterior distribution of the model parameters $\boldsymbol{\theta}_s$ is approximately free from the prior adopted for the nuisance parameter θ_v [32].

By substituting the asymptotic approximation of the likelihood function (34) into Bayes' Theorem, the posterior PDF can be approximated by:

$$p(\boldsymbol{\theta}_s|\mathcal{D}_N) \cong \frac{\sqrt{2\pi}H(\hat{\theta}_v)^{-\frac{1}{2}} p(\hat{\theta}_v) p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\theta}_v(\boldsymbol{\theta}_s)) p(\boldsymbol{\theta}_s)}{\int \sqrt{2\pi}H(\hat{\theta}_v)^{-\frac{1}{2}} p(\hat{\theta}_v) p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\theta}_v(\boldsymbol{\theta}_s)) p(\boldsymbol{\theta}_s) d\boldsymbol{\theta}_s} \quad (35)$$

where \cong denotes an approximation to $\mathcal{O}(N^{-1})$. Using Remark 5, the asymptotic approximation for the marginal posterior distribution of the structural model parameters $\boldsymbol{\theta}_s$ can be rewritten as:

$$p(\boldsymbol{\theta}_s|\mathcal{D}_N) \cong \frac{p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\boldsymbol{\theta}}_v(\boldsymbol{\theta}_s)) p(\boldsymbol{\theta}_s)}{\int p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\boldsymbol{\theta}}_v(\boldsymbol{\theta}_s)) p(\boldsymbol{\theta}_s) d\boldsymbol{\theta}_s} \propto p(\hat{\mathbf{z}}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\boldsymbol{\theta}}_v(\boldsymbol{\theta}_s)) p(\boldsymbol{\theta}_s) \quad (36)$$

which shows that the prior distribution of the nuisance parameter cancels out. Following the same line of thought as for ABC methods (see Introduction), the approximation of the marginal posterior distribution (36) can be augmented to:

$$p(\boldsymbol{\theta}_s|\mathcal{D}_N) \propto p(\hat{\mathbf{z}}_{1:N}|\mathbf{y}_{1:N}, \epsilon) p(\mathbf{y}_{1:N}|\hat{\mathbf{u}}_{0:N}, \boldsymbol{\theta}_s, \hat{\boldsymbol{\theta}}_v(\boldsymbol{\theta}_s)) p(\boldsymbol{\theta}_s) \quad (37)$$

By defining $p(\hat{\mathbf{z}}_{1:N}|\mathbf{y}_{1:N}, \epsilon) = p(\mathbf{e}_{1:N}|\epsilon)$ as the uniform probability distribution given in (18), the ABC approximate marginal posterior distribution (37) gives the exact posterior PDF for the hierarchical state-space model class with a uniform measurement error for predictions of the sensor output and no state prediction errors \mathbf{w}_n [26].

3. Illustrative examples for Bayesian system identification

The utilization of Bayesian framework for system identification has received increasing attention in recent years (e.g., [1–3, 24, 33, 36–52]). Here, we study two numerical examples from Bayesian system identification literature to demonstrate the application of the ABC-SubSim algorithm to model class updating and selection, with a special focus on the estimation of the model class evidence for different values of the tolerance level ϵ . The first example, which is a single degree-of-freedom bilinear oscillator, compares the performance of ABC-SubSim for Bayesian model class selection with the method presented in Beck and Yuen [2], which provides an asymptotic approximation for the model evidence in the presence of large amounts of dynamic data. The second example, which is a three degree-of-freedom nonlinear structure, demonstrates the applicability of ABC-SubSim to perform Bayesian model class selection for a class of Masing hysteretic models and compares its behavior with the TMCMC algorithm [36]. Both examples use input and output data that is artificially generated by subjecting the dynamic model to recorded seismic excitations. For ABC-SubSim, the self-regulating algorithm presented in [26] is used for both examples with the number of samples in each level fixed to $N_t = 2000$ and with the adaptation probability $P_a = 0.1$ and the optimal acceptance rate $\alpha^* = 0.5$. However, the conditional probability for each level of ABC-SubSim is set differently for the first and second examples as $P_0 = 0.2$ and $P_0 = 0.1$, respectively. For both examples, our objective is to draw samples from the posterior $p(\mathbf{y}_{1:N}, \mathbf{x}_{1:N}, \boldsymbol{\theta}|\mathcal{D}_N, \epsilon)$ where $\mathbf{y}_{1:N}$ is constrained to lie in a small neighborhood, $\mathcal{D}(\epsilon)$, of the data vector defined by:

$$\mathcal{D}(\epsilon) = \{\mathbf{y}_{1:N} \in \mathbb{R}^{NN_o} : \|\mathbf{y}_{1:N} - \hat{\mathbf{z}}_{1:N}\|_2 \leq \epsilon\} \quad (38)$$

The model evidence is then estimated as a by-product of ABC-SubSim.

Example 1: Single degree-of-freedom bilinear hysteretic oscillator under seismic excitation

This example follows the spirit of the first example used in Beck and Yuen [2] where response measurements are used to select the most plausible model class for different levels of excitation. In their paper, the model evidence is approximated by using Laplace’s method for asymptotic approximation while we use ABC-SubSim here. The system is a single degree-of-freedom bilinear hysteretic oscillator with linear viscous damping. The equation of motion for this oscillator subject to ground acceleration can be represented by:

$$m\ddot{z}(t) + c\dot{z}(t) + f_h(z; k_1, k_2, z_y) = -mu(t) \quad (39)$$

where $z(t) \in \mathbb{R}$ is the horizontal displacement vector relative to ground; $u(t) \in \mathbb{R}$ is the horizontal seismic ground acceleration; $m, c \in \mathbb{R}$ denote the mass and linear viscous damping, respectively; $f_h \in \mathbb{R}$ denotes the hysteretic restoring force; $k_1 \in \mathbb{R}$ is the elastic stiffness; $k_2 \in \mathbb{R}$ is the post-yield stiffness; and $z_y \in \mathbb{R}$ is the yield displacement. The relationship between the restoring force and the displacement is shown in Figure 1.

The oscillator is assumed to have known mass $m = 1$ kg. In general, $\theta_s = [k_1 \ k_2 \ z_y \ c]^T$ forms the uncertain structural model parameter vector for the bilinear hysteresis oscillator where the actual values are defined as $k_1 = 1.0$ N/m, $k_2 = 0.1$ N/m, $c = 0.02$ N.s/m and $z_y = 2.0$ cm. Samples of the response time history $\mathbf{y}_{1:N}$ for given values of the uncertain parameters and ground acceleration time history are simulated using the function ‘ODE45’ in Matlab.




Fig. 1 *The hysteresis loop for the bilinear oscillator (Example 1).*

Three sets of data are studied here with 10, 15, and 20% scaling of the 1940 El Centro earthquake record (north-south component) as the excitation corresponding to each data set. The earthquake excitation and displacement of the oscillator are measured for 40 s with sampling frequency of 60 Hz to give $N = 2400$ data points. These synthetic data are contaminated with zero-mean Gaussian discrete white noise in which the variance σ_v^2 for each data set is selected so that the output prediction error gives a 5% RMS noise-to-signal ratio over the associated displacement data from the actual system. This gives σ_v equal to 4.6×10^{-4} cm, 5.9×10^{-4} cm, and 7.1×10^{-4} cm for data sets generated using 10, 15, and 20% scaling of the 1940 El Centro




Fig. 2 a) Oscillator response measurements for three levels of excitation; b) Oscillator hysteresis loops for the three levels of excitation (Example 1).

Table 1 Posterior sample mean of parameter values for each model class representing the oscillator and for different levels of excitations (Example 1), UN denotes Unidentifiable parameter.

Excitation level	Model class	c (N.s/m)	k_1 (N/m)	k_2 (N/m)	z_y (cm)	σ_v (cm)
10% El Centro	\mathcal{M}_1	0.0199	1.0008	—	—	0.0005
	\mathcal{M}_2	—	1.0036	—	1.86	0.0010
	\mathcal{M}_3	—	1.0034	UN	1.84	0.0010
15% El Centro	\mathcal{M}_1	0.0944	0.9661	—	—	0.0021
	\mathcal{M}_2	—	0.9978	—	2.10	0.0012
	\mathcal{M}_3	—	1.0141	0.1485	1.98	0.0009
20% El Centro	\mathcal{M}_1	0.1978	0.9029	—	—	0.0061
	\mathcal{M}_2	—	0.9517	—	2.26	0.0039
	\mathcal{M}_3	—	1.0141	0.0950	1.97	0.0008

earthquake record, respectively. The left panel of Figure 2 demonstrates the synthetic measurements for the three levels of excitation and the right panel of Figure 2 indicates the corresponding hysteresis loops. This figure shows that the dynamics of the structure is linear for 10% scaling, mildly nonlinear for 15% scaling, and strongly nonlinear for 20% scaling of the El Centro earthquake record.

Here, we consider three different model classes: Model class \mathcal{M}_1 is a linear oscillator in which the uncertain model parameter vector $\boldsymbol{\theta}_s = [k_1, c]$, consisting of the elastic stiffness k_1 and viscous damping c . Model class \mathcal{M}_2 is a bilinear hysteretic oscillator with post-yield stiffness $k_2 = 0$ (elastoplastic oscillator) and no viscous damping $c = 0$. Thus, the model parameter vector $\boldsymbol{\theta}_s = [k_1, z_y]$, consisting of the elastic stiffness k_1 and yield displacement z_y . Model class \mathcal{M}_3 is a bilinear hysteretic oscillator with no viscous damping $c = 0$. In this model class, the uncertain model parameter vector $\boldsymbol{\theta}_s = [k_1, k_2, z_y]$, consisting of the elastic stiffness k_1 , post-yield stiffness k_2 , and yield displacement z_y . Note that none of these model classes match the exact model class used to generate the data. The prior distributions for the model parameters k_1 , k_2 , c , z_y are chosen as independent uniform distributions over the intervals $(0, 2)$ N/m, $(0, 0.5)$ N/m, $(0, 0.5)$ N.s/m, $(0, 0.1)$ m, respectively.

The mean estimates of the model parameters for the three model classes and three different excitation levels are reported in Table 1. The parameter estimates are the mean of 2000 posterior samples drawn by self-regulating ABC-SubSim. For model class \mathcal{M}_1 and for higher levels of excitation, Table 1 shows that the mean estimates have lower values of the linear stiffness and higher values of the damping coefficient to represent the hysteretic dissipated energy. Figure 3 shows the 2000 samples drawn in different levels of ABC-SubSim in the $\{k_1, z_y\}$ space when updating model class \mathcal{M}_2 using data from the 10% El Centro earthquake record. As shown, the posterior samples are tightly clustered around the value $z_y = 1.86$ cm. This might seem a counter-intuitive result since when the oscillator is excited with the 10% El Centro earthquake, it behaves perfectly linear and so one expects z_y to be unidentifiable. However, it seems that this value for z_y , which is slightly less than the maximum amplitude of the oscillations (1.875 cm), is an attempt to yield enough hysteretic dissipative energy to compensate for the lack of viscous damping in the model class \mathcal{M}_2 . A similar behavior can be observed in Figure 4 for model class \mathcal{M}_3 where z_y is pinned down around 1.84 cm while k_2 is unidentifiable (because there is little post-yield response), which seems to be again an attempt to generate enough dissipative energy due to yielding to compensate for the lack of viscous damping. It can be observed in Figure 4 that for the higher levels of excitation, the structure experiences stronger nonlinear behavior and so the post-yielding stiffness parameter k_2 is pinned down more tightly.

Table 2 presents the number of simulation levels m and the final tolerance levels $\epsilon_{\mathcal{M}_j}$ which are adaptively selected by the algorithm to explore the posterior distribution of the parameters of different models \mathcal{M}_j , $j = 1, 2, 3$, for the three levels of excitation. This table also shows the posterior probability of models $P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j})|\mathcal{D}_N, \mathbf{M})$ obtained from (26) by evaluation of evidence (27) at the final tolerance levels $\epsilon_{\mathcal{M}_j}$ and equal prior probabilities $P(\mathcal{M}_j|\mathbf{M}) = 1/3$. According to this result, in the case of 10% scaling of the 1940 El Centro earthquake record where the structure behaves linearly (see Figure 2b), \mathcal{M}_1 gives the largest posterior




Fig. 3 Scatter plot of 2000 samples in the $\{k_1, z_y\}$ space that are generated at levels 2, 5 and 11 (in blue) and their previous intermediate levels (in gray) of ABC-SubSim when updating model class \mathcal{M}_2 with data from the 10% El Centro earthquake (Example 1).




Fig. 4 Scatter plot of 2000 posterior samples in the $\{k_2, z_y\}$ space when updating model class \mathcal{M}_3 for the intermediate levels (in gray) and the final level (in blue) and for different excitation levels (Example 1).

Table 2 Posterior probability of different model classes together with final tolerance level and number of simulation levels for three-story Masing building (Example 2).

Model class	10% El Centro			15% El Centro			20% El Centro		
	\mathcal{M}_1	\mathcal{M}_2	\mathcal{M}_3	\mathcal{M}_1	\mathcal{M}_2	\mathcal{M}_3	\mathcal{M}_1	\mathcal{M}_2	\mathcal{M}_3
Sim. levels (m)	9	9	9	7	9	10	5	7	12
Tol. level ($\epsilon_{\mathcal{M}_j}$)	0.0007	0.0014	0.0014	0.0029	0.0016	0.0012	0.0084	0.0054	0.0011
$P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j}) \mathcal{D}_N, \mathbf{M})$	1	0	0	0	0	1	0	0	1

probability. For the higher levels of excitations, when the dissipated energy is dominated by the hysteretic dissipative energy, \mathcal{M}_3 is the most plausible model. This example reflects the important point made by the famous statement that “all models are wrong, but some are useful” [53] and, furthermore, the “best model” for the system analysis and response predictions depends on the data that is used for system identification.

One of the difficulties for any ABC algorithm is to select the final tolerance level for which the self-regulating ABC-SubSim algorithm brings a straightforward solution, as explained in Remark 4. The agreement between the approximate posterior probabilities $P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j})|\mathcal{D}_N, \mathbf{M})$ presented in Table 2 with those reported by Beck and Yuen [2] shows the validity of the stopping criterion used in self-regulating ABC-SubSim. To further emphasize the importance of choosing an appropriate final tolerance level, the probability that $\mathbf{y}_{1:N}$ lies in the ball of radius ϵ around the data vector $\hat{\mathbf{z}}_{1:N}$ and the posterior probability $P(\mathcal{M}_j(\epsilon)|\mathcal{D}_N, \mathbf{M})$ for different model classes are depicted versus the tolerance level ϵ in Figures 5 and 6, respectively. As ϵ goes down from 0.1 to $\epsilon_{\mathcal{M}_j}$, $P(\mathcal{M}_j(\epsilon)|\mathcal{D}_N, \mathbf{M})$ varies between the model prior probabilities at $\epsilon = 0.1$ and the true model posterior probabilities at $\epsilon_{\mathcal{M}_j}$.




Fig. 5 The probability of entering the data-approximating region $\mathcal{D}(\epsilon)$ against tolerance level ϵ (Example 1).




Fig. 6 The posterior probability of different model classes \mathcal{M}_j against tolerance level ϵ (Example 1).

Example 2: Three-story Masing shear-building under seismic excitation

The second example is taken from Muto and Beck [24] where the Transitional Markov Chain Monte Carlo (TMCMC) algorithm has been used for Bayesian updating and model selection of the class of Masing hysteretic structural models. This example considers a three-story shear building with the following equation of motion:

$$\mathbf{M}\ddot{\mathbf{z}}(t) + \mathbf{C}\dot{\mathbf{z}}(t) + \mathbf{f}_h = -\mathbf{M}\mathbf{1}u(t) \quad (40)$$

where $\mathbf{z}(t) \in \mathbb{R}^3$ is the horizontal displacement vector relative to the ground; $\mathbf{M}, \mathbf{C} \in \mathbb{R}^{3 \times 3}$ are the mass and damping matrices; $u(t)$ is the horizontal ground acceleration; and $\mathbf{1} = [1 \ 1 \ 1]^T$. The restoring force for the i th story is given by:

$$f_{h,i} = r_i - r_{i+1} \quad (41)$$

where the inter-story shear force-deflection relation is given by the differential equation:

$$\dot{r}_i = k_i(\dot{z}_i - \dot{z}_{i-1}) \left[1 - \left| \frac{r_i}{r_{u,i}} \right|^{\alpha_i} \right] \quad (42)$$

Here, k_i is the small-amplitude inter-story stiffness, $r_{u,i}$ is the story ultimate strength and α_i is the elastic-to-plastic transition parameter. The force-deflection relationship given in (42) defines the initial loading curve OA in Figure 7. Any other loading curve can be selected according to two *extended Masing rules* [54–56]:

1. The force-deflection relation for any loading curve other than the initial loading (42) is described by the differential equation:

$$\dot{r}_i = k_i(\dot{z}_i - \dot{z}_{i-1}) \left[1 - \left| \frac{r_i - r^*}{2r_{u,i}} \right|^{\alpha_i} \right] \quad (43)$$

where r^* is the restoring force at the latest load reversal point. For instance, (43) gives the loading curve AC in Figure 7 if $r^* = r_a$.

2. Once an interior loading curve crosses a curve from a previous load cycle, the load deformation continues that of the previous cycle. For instance, if the curve DE is continued to point C, it follows the force-deflection relation of curve ABC.

It should be noted that a wide variety of hysteretic models can be described by using the two extended Masing rules through the choice of the initial load curve. Thus, the class of Masing hysteretic model with restoring force-deflection relation (42) used here is only a special class of Masing models.




Fig. 7 *Hysteresis loops for transient loading of the extended Masing model [54].*

In this example, the structure has a known story mass of 1.25×10^5 kg and a known earthquake ground acceleration defining the excitation $u(t)$. Similar to [24], the east-west component of the Sylmar ground-motion record from the County Hospital Parking Lot during 1994 Northridge earthquake in California is used here as the excitation. The actual values for the model parameters for each story $i = 1, 2, 3$ are: small-amplitude inter-story stiffnesses $k_i = 2.5 \times 10^8$ N/m, ultimate strengths $r_{u,i} = 1.75 \times 10^6$ N, and elastic-to-plastic transition parameters $\alpha_i = 4$. The viscous damping matrix \mathbf{C} is modeled using Rayleigh damping $\mathbf{C} = c_M \mathbf{M} + c_K \mathbf{K}$ with coefficients $c_M = 0.293$ and $c_K = 2.64 \times 10^{-4}$. This set of structural parameters gives the three small-amplitude natural frequencies as $\omega_1 = 3.17$ Hz, $\omega_2 = 8.88$ Hz and $\omega_3 = 12.83$ Hz, and the modal damping ratios as $\zeta_1 = \zeta_2 = 0.01$ and $\zeta_3 = 0.012$.

In this study, we use the fixed-point exponential integrator [57] together with the particular numerical implementation of the two extended Masing rules that was proposed in [58] to simulate samples of the response time history $\mathbf{y}_{1:N}$ for given values of the uncertain parameters and the Sylmar ground-motion record.

The synthetic response data for system identification is the inter-story drift time histories when the uncertain parameters are set to their actual values and the standard deviation of the uncertain output error is set to 0.03 cm to give a 5% RMS noise-to-signal ratio. The simulated data, which are shown in Figure 8, correspond to 500 data points from each story with a sampling rate of 0.02 s. Figure 9 shows the hysteresis loops for each story generated by the Sylmar ground motion. This figure indicates that the simulated structure experiences strong yielding in the first story and moderate yielding in the second story whereas it exhibits almost linear behavior in the top story.

Four model classes are studied for system identification. For all model classes, the story masses are taken as known and set to their actual values. Generally, the two parameters, c_M and c_K , specifying the viscous damping matrix and the nine parameters of the hysteresis model k_i , $r_{u,i}$ and α_i , $i = 1, 2, 3$ in (42), make the vector of uncertain structural model parameters θ_s . For model classes \mathcal{M}_1 and \mathcal{M}_2 , the elastic-to-plastic transition parameters are constrained to be equal for all three stories whereas they are allowed to vary for




Fig. 8 Inter-story drift time histories and the Sylmar ground-motion record (Example 2).




Fig. 9 Simulated inter-story restoring forces against inter-story drifts (Example 2).

model classes \mathcal{M}_3 and \mathcal{M}_4 . The model classes \mathcal{M}_1 and \mathcal{M}_3 contain no viscous damping, but model classes \mathcal{M}_2 and \mathcal{M}_4 do and the Rayleigh damping coefficients c_M and c_K are estimated for these model classes. Therefore, in contrast to Example 1, a subset of the candidate model classes, i.e., \mathcal{M}_2 and \mathcal{M}_4 , contains the model used to generate the data.

The prior distribution over the nine-dimensional parameter space of the hysteresis model is selected to be the product of nine lognormal PDFs with logarithmic mean value of $\log(2.5 \times 10^8)$ for k_i , $i = 1, 2, 3$, $\log(2.5 \times 10^6)$ for $r_{u,i}$, $i = 1, 2, 3$ and $\log(4)$ for α_i , $i = 1, 2, 3$ and a logarithmic standard deviation of 0.5 for all of them. The prior distributions for the parameters of the viscous damping matrix are defined as independent uniform PDFs over the interval $[0, 1.5]$ for c_M and $[0, 1.5 \times 10^{-3}]$ for c_K .

Table 3 shows the MAP (maximum a posteriori) values and the standard deviations of the uncertain parameters obtained for all model classes. The MAP value for each parameter is estimated by fitting a lognormal PDF to the samples drawn from their posterior distribution meaning that this is different from the posterior sample mean. Figures 10-13 show 2000 samples obtained from the self-regulating ABC-SubSim algorithm for some of the uncertain parameters of the four model classes.

For model classes \mathcal{M}_1 and \mathcal{M}_2 , the posterior samples of the small-amplitude stiffnesses are tightly clustered (see Figures 10 and 11). The posterior samples of the ultimate strength for the first story is well-constrained, for the second story they show slightly higher level of uncertainty and for the third story they exhibit a high level of uncertainty (see Figures 10 and 11). This phenomenon can be understood by looking at Figure 9, which demonstrates that a noticeable yielding occurred in the first story. This means that there is enough information in the response data to estimate the ultimate strength for the first story. On the other hand, this figure shows a nearly linear behavior for the third story, so the response data only impose a lower bound on the ultimate strength. These results are very similar to those reported by Muto and Beck [24]. For model class \mathcal{M}_2 , the posterior distribution for the parameters of the viscous damping matrix reveals a high level of uncertainty. Presumably, this can be attributed to the fact that the response is less sensitive to the variation of the parameters of the hysteresis model.

For model classes \mathcal{M}_3 and \mathcal{M}_4 , the posterior distribution for the small-amplitude stiffnesses are also compactly clustered but they are not graphically shown here (see Table 3). Both model classes exhibit an almost similar behavior in the parameter space $\{r_{u,i}, \alpha_i\}$ (see Figures 12 and 13). As expected, the parameters for the first story, $r_{u,1}$ and α_1 , are globally identifiable for both models. For the third story, the posterior distribution of the parameters $r_{u,3}$ and α_3 shows a large spread over the parameter space with a clear lower bound. This can be attributed to the fact that the third story does not experience yielding and the lower bound is the only information that can be extracted from the data. However, the joint posterior distribution of $r_{u,2}$ and α_2 in model class \mathcal{M}_3 differs from its counterpart in model class \mathcal{M}_4 (see Figures 12 and 13). The lack of viscous damping in model class \mathcal{M}_3 apparently forces the posterior samples in $\{\alpha_2, r_{u,2}\}$ space to be clustered in a region around a lower value of α_2 and a higher value for $r_{u,2}$, so that the dissipated hysteretic energy can compensate for the lack of viscous damping. In model class \mathcal{M}_4 , the

Table 3 The maximum a posteriori parameter values and the standard deviations (in parentheses) obtained from fitting a lognormal distribution to the posterior samples (Example 2).

Model class	\mathcal{M}_1	\mathcal{M}_2	\mathcal{M}_3	\mathcal{M}_4
$k_1 (10^8 N/m)$	2.586 (0.026)	2.497 (0.018)	2.545 (0.029)	2.490 (0.015)
$k_2 (10^8 N/m)$	2.455 (0.044)	2.499 (0.025)	2.539 (0.042)	2.509 (0.034)
$k_3 (10^8 N/m)$	2.566 (0.054)	2.490 (0.023)	2.545 (0.061)	2.504 (0.025)
$r_{u,1} (10^6 N)$	1.737 (0.004)	1.749 (0.003)	1.746 (0.006)	1.751 (0.003)
$r_{u,2} (10^6 N)$	1.779 (0.064)	1.750 (0.037)	1.924 (0.152)	1.757 (0.054)
$r_{u,3} (10^6 N)$	2.056 (1.014)	2.140 (0.772)	2.358 (0.934)	2.154 (1.083)
α_1	3.430 (0.090)	3.981 (0.075)	3.447 (0.145)	4.041 (0.094)
α_2	$= \alpha_1$	$= \alpha_1$	2.626 (0.300)	3.863 (0.411)
α_3	$= \alpha_1$	$= \alpha_1$	2.552 (2.607)	3.332 (2.047)
$c_M (s^{-1})$	—	0.259 (0.071)	—	0.283 (0.069)
$c_K (10^{-4}s)$	—	2.295 (1.322)	—	2.116 (0.909)
$\sigma_v (10^{-4}m)$	5.293 (0.063)	3.197 (0.064)	5.163 (0.084)	3.176 (0.048)

estimated Rayleigh damping parameters are rather close to their actual values and so the need for a higher hysteretic dissipation energy is mitigated, explaining why the posterior samples in $\{\alpha_2, r_{u,2}\}$ are clustered around their actual values. We note that the results presented for models \mathcal{M}_3 and \mathcal{M}_4 are to some extent different from their counterparts reported by Muto and Beck [24]. This difference can be explained by the fact that simulating the response time history $\mathbf{y}_{1:N}$ from a structure with hysteretic restoring forces is very dependent on the numerical schemes, e.g., type of time integrator used for the numerical implementations, and since the schemes used in this study are different from those of their study, results can be expected to be different to some extent.

Table 4 shows the number of simulation levels m and the final tolerance levels $\epsilon_{\mathcal{M}_j}$ for different model classes. This table also presents the posterior probability of model classes $P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j})|\mathcal{D}_N, \mathbf{M})$, $j = 1, 2, 3, 4$ calculated from (26) by evaluation of evidence (27) at the final tolerance levels $\epsilon_{\mathcal{M}_j}$ and equal prior probabilities $P(\mathcal{M}_j|\mathbf{M}) = 1/4$ for the models. It is not surprising that the posterior probability for the model classes favors model class \mathcal{M}_2 since it contains the model used to generate the synthetic data and has two parameters less than model class \mathcal{M}_4 , which also contains the data-generating model. As shown by the information-theoretic expression for the log evidence in [24], the posterior probability of a model class is controlled by a trade-off between the posterior average data fit (the posterior mean of the log-likelihood) and the amount of information extracted from data (the relative entropy of the posterior with respect to the prior). \mathcal{M}_2 and \mathcal{M}_4 give essentially the same average data fit but \mathcal{M}_2 extracts less information about its parameters.




Fig. 10 Scatter plot of 2000 posterior samples plotted in $\{r_{u,1}, \alpha_1\}$ (left) and $\{r_{u,3}, \alpha_3\}$ (right) spaces when updating model class \mathcal{M}_1 for some intermediate levels (in gray) and the final level (in blue) (Example 2).




Fig. 11 Scatter plot of 2000 posterior samples plotted in $\{r_{u,2}, \alpha_2\}$ (left) and $\{c_M, c_K\}$ (right) spaces when updating model class \mathcal{M}_2 for some intermediate levels (in gray) and the final level (in blue) (Example 2).




Fig. 12 Scatter plot of 2000 posterior samples plotted in $\{r_{u,1}, \alpha_1\}$ (left) and $\{r_{u,2}, \alpha_2\}$ (middle) and $\{r_{u,3}, \alpha_3\}$ (right) spaces when updating model class \mathcal{M}_3 for some intermediate levels (in gray) and the final level (in blue) (Example 2).




Fig. 13 Scatter plot of 2000 posterior samples plotted in $\{r_{u,1}, \alpha_1\}$ (left) and $\{r_{u,2}, \alpha_2\}$ (middle) and $\{r_{u,3}, \alpha_3\}$ (right) spaces when updating model class \mathcal{M}_4 for some intermediate levels (in gray) and the final level (in blue) (Example 2).

Table 4 Posterior probability of different model classes together with final tolerance level and number of simulation levels for three-story Masing building (Example 2).

Model class	\mathcal{M}_1	\mathcal{M}_2	\mathcal{M}_3	\mathcal{M}_4
Sim. levels (m)	10	10	11	12
Tol. level ($\epsilon_{\mathcal{M}_j}$)	6.80×10^{-4}	4.25×10^{-4}	7.10×10^{-4}	4.25×10^{-4}
$P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j}) \mathcal{D}_N, \mathbf{M})$	0	0.982	0	0.018

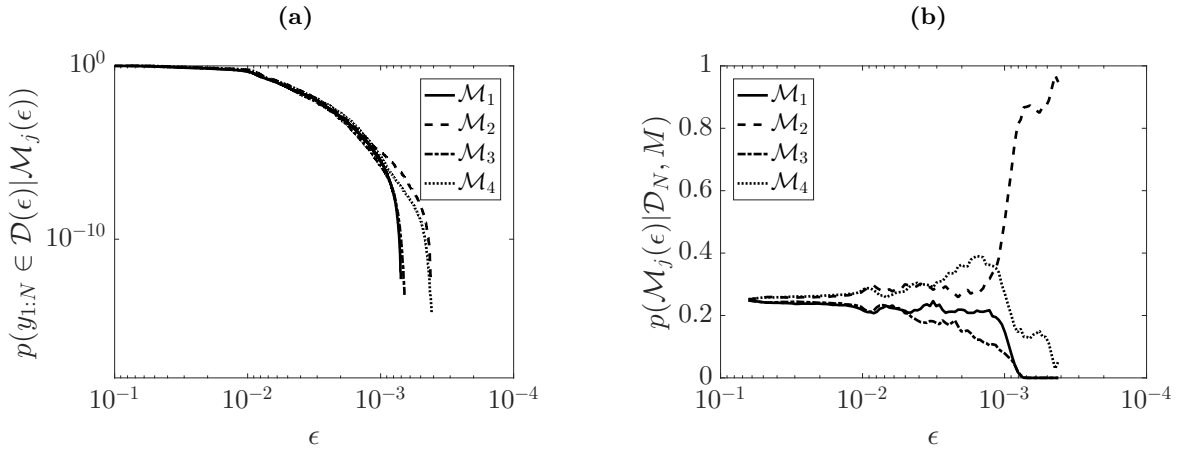


Fig. 14 a) The probability of entering the data-approximating region $\mathcal{D}(\epsilon)$ against tolerance level ϵ ; b) The posterior probability of different model classes \mathcal{M}_j against tolerance level ϵ (Example 2).

Similar to the previous example, the approximate posterior probabilities $P(\mathcal{M}_j(\epsilon_{\mathcal{M}_j})|\mathcal{D}_N, \mathbf{M})$ presented in Table 4 are in agreement with those reported by Muto and Beck [24] which shows that the self-regulating ABC-SubSim algorithm selected proper values for the final tolerance levels $\epsilon_{\mathcal{M}_j}$. Figures 14(a and b), respectively show the probability that $\mathbf{y}_{1:N}$ falls in the data-approximating region $\mathcal{D}(\epsilon)$ and the posterior probability $P(\mathcal{M}_j(\epsilon)|\mathcal{D}_N, \mathbf{M})$ for different model classes versus the tolerance level ϵ . As ϵ goes down from 0.1 to $\epsilon_{\mathcal{M}_j}$, $P(\mathcal{M}_j(\epsilon)|\mathcal{D}_N, \mathbf{M})$ varies between the model prior probabilities at $\epsilon = 0.1$ and the true model posterior probabilities at $\epsilon_{\mathcal{M}_j}$.

It is worth noting that both the parameter vector and the tolerance level are taken to be the same across models in the traditional ABC approach to model comparison explained in Algorithm 2 [59]. This makes the estimates of the posterior probability of model classes sensitive to (i) the proposal PDF of the Markov chain used within the sampling algorithm [21], and (ii) the choice of a final tolerance level ϵ . The former dependency should not occur since is not related to the inference problem under study. The choice of a unique tolerance level ϵ that works across all models is very delicate, since, as illustrated in Figure 6, a wrong choice of ϵ can result in a significant bias in the ABC approximation of the model posterior probabilities. However, the proposed model selection procedure which is based on the hierarchical state-space formulation of dynamic models and the self-regulating ABC-SubSim algorithm alleviates these type of difficulties by independently estimating the model evidence for each of the models under comparison.

4. Concluding remarks

In the current state of the art, ABC methods can only be used for model class selection in a very limited range of models for which a set of sufficient summary statistics can be found so that it also guarantees sufficiency across the set of models under study. In this paper, a new ABC model selection procedure has been presented which broadens the realm of ABC-based model comparison to be able to assess dynamic models. In the proposed procedure, a dynamic problem is formulated in terms of a general hierarchical state-space model such that the normalizing constant associated to its exact posterior distribution using the entire data provides an unbiased estimator of the model evidence as an error tolerance level $\epsilon \rightarrow 0$.

The self-regulating ABC-SubSim provides a straightforward way to estimate the model evidence and, as a result, the posterior probability of models as a function of the error tolerance level ϵ . This enables us to better understand the model choices made in the earlier applications of the ABC-based model comparison methods. Furthermore, a new solution based on the Laplace’s method of asymptotic approximations is presented to mitigate the fundamental difficulty of the ABC algorithms to learn the parameters specifying the uncertain state and output prediction errors in a stochastic state-space model. It has the key advantage that the approximated marginal distribution of the model parameters is insensitive to the prior adopted for the uncertain prediction error variance.

Two illustrative examples with synthetic data are selected from the Bayesian system identification literature to show the estimation of the model class evidences and posterior probabilities obtained by the self-

regulating ABC-SubSim algorithm. The first example shows the successful application of the self-regulating ABC-SubSim for Bayesian model class selection when the true system is not among the competing model classes. The second example shows the capability of the self-regulating ABC-SubSim algorithm to efficiently explore a posterior distribution with a relatively high-dimensional parameter space.

Acknowledgment

The first author of this paper wants to express his gratitude to the California Institute of Technology (Caltech) for kindly hosting him during the course of this work.

References

References

- [1] J. L. Beck, Bayesian system identification based on probability logic, *Structural Control and Health Monitoring* 17 (7) (2010) 825–847.
- [2] J. L. Beck, K.-V. Yuen, Model selection using response measurements: Bayesian probabilistic approach, *Journal of Engineering Mechanics* 130 (2) (2004) 192–203.
- [3] S. H. Cheung, J. L. Beck, Calculation of posterior probabilities for Bayesian model class assessment and averaging from posterior samples based on dynamic system data, *Computer-Aided Civil and Infrastructure Engineering* 25 (5) (2010) 304–321.
- [4] S. Tavaré, D. J. Balding, R. C. Griffiths, P. Donnelly, Inferring coalescence times from DNA sequence data, *Genetics* 145 (2) (1997) 505–518.
- [5] J. K. Pritchard, M. T. Seielstad, A. Perez-Lezaun, M. W. Feldman, Population growth of human Y chromosomes: a study of Y chromosome microsatellites., *Molecular Biology and Evolution* 16 (12) (1999) 1791–1798.
- [6] P. Bortot, S. G. Coles, S. A. Sisson, Inference for stereological extremes, *Journal of the American Statistical Association* 102 (477) (2007) 84–92.
- [7] P. Marjoram, J. Molitor, V. Plagnol, S. Tavaré, Markov chain Monte Carlo without likelihoods, *Proceedings of the National Academy of Sciences* 100 (26) (2003) 15324–15328.
- [8] S. A. Sisson, Y. Fan, Likelihood-free MCMC, *Handbook of Markov Chain Monte Carlo* (2011) 313–335.
- [9] S. Sisson, Y. Fan, M. Tanaka, A note on backward kernel choice for sequential Monte Carlo without likelihoods, *Tech. Rep.*, Technical report, University of New South Wales, 2008.
- [10] S. A. Sisson, Y. Fan, M. M. Tanaka, Sequential Monte Carlo without likelihoods, *Proceedings of the National Academy of Sciences* 104 (6) (2007) 1760–1765.

- [11] P. Del Moral, A. Doucet, A. Jasra, An adaptive sequential Monte Carlo method for approximate Bayesian computation, *Statistics and Computing* 22 (5) (2012) 1009–1020.
- [12] C. C. Drovandi, A. N. Pettitt, Estimation of parameters for macroparasite population evolution using approximate Bayesian computation, *Biometrics* 67 (1) (2011) 225–233.
- [13] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, M. P. Stumpf, Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems, *Journal of the Royal Society Interface* 6 (31) (2009) 187–202.
- [14] M. A. Beaumont, J.-M. Cornuet, J.-M. Marin, C. P. Robert, et al., Adaptive approximate Bayesian computation, *Biometrika* 96 (4) (2009) 983–990.
- [15] M. Chiachio, J. L. Beck, J. Chiachio, G. Rus, Approximate Bayesian computation by Subset Simulation, *SIAM Journal on Scientific Computing* 36 (3) (2014) A1339–A1358.
- [16] X. Didelot, R. G. Everitt, A. M. Johansen, D. J. Lawson, et al., Likelihood-free estimation of model evidence, *Bayesian Analysis* 6 (1) (2011) 49–76.
- [17] A. Grelaud, C. P. Robert, J.-M. Marin, F. Rodolphe, J.-F. Taly, et al., ABC likelihood-free methods for model choice in Gibbs random fields, *Bayesian Analysis* 4 (2) (2009) 317–335.
- [18] C. P. Robert, J.-M. Cornuet, J.-M. Marin, N. S. Pillai, Lack of confidence in approximate Bayesian computation model choice, *Proceedings of the National Academy of Sciences* 108 (37) (2011) 15112–15117.
- [19] G. O. Roberts, J. S. Rosenthal, Examples of adaptive MCMC, *Journal of Computational and Graphical Statistics* 18 (2) (2009) 349–367.
- [20] T. Toni, M. P. Stumpf, Simulation-based model selection for dynamical systems in systems and population biology, *Bioinformatics* 26 (1) (2010) 104–110.
- [21] J.-M. Marin, P. Pudlo, C. P. Robert, R. J. Ryder, Approximate Bayesian computational methods, *Statistics and Computing* 22 (6) (2012) 1167–1180.
- [22] R. D. Wilkinson, Bayesian inference of primate divergence times, Ph.D. thesis, University of Cambridge, 2008.
- [23] R. D. Wilkinson, Approximate Bayesian Computation (ABC) gives exact results under the assumption of model error, *Statistical Applications in Genetics and Molecular Biology* 12 (2) (2013) 129–141.
- [24] M. Muto, J. L. Beck, Bayesian updating and model class selection for hysteretic structural models using stochastic simulation, *Journal of Vibration and Control* 14 (1-2) (2008) 7–34.

- [25] E. T. Jaynes, Information theory and statistical mechanics, *Physical Review* 106 (4) (1957) 620.
- [26] M. K. Vakilzadeh, Y. Huang, J. L. Beck, T. Abrahamsson, Approximate Bayesian Computation by Subset Simulation using hierarchical state-space models, *Mechanical Systems and Signal Processing* (2016) in press.
- [27] S.-K. Au, J. L. Beck, Estimation of small failure probabilities in high dimensions by Subset Simulation, *Probabilistic Engineering Mechanics* 16 (4) (2001) 263–277.
- [28] C. Robert, G. Casella, *Monte Carlo statistical methods*, Springer Science & Business Media, 2013.
- [29] K. M. Zuev, J. L. Beck, S.-K. Au, L. S. Katafygiotis, Bayesian post-processor and other enhancements of Subset Simulation for estimating failure probabilities in high dimensions, *Computers and Structures* 92 (2012) 283–296.
- [30] S. Barthelmé, N. Chopin, Expectation propagation for likelihood-free inference, *Journal of the American Statistical Association* 109 (505) (2014) 315–333.
- [31] J. O. Berger, B. Liseo, R. L. Wolpert, et al., Integrated likelihood methods for eliminating nuisance parameters, *Statistical Science* 14 (1) (1999) 1–28.
- [32] L. Tierney, J. B. Kadane, Accurate approximations for posterior moments and marginal densities, *Journal of the American Statistical Association* 81 (393) (1986) 82–86.
- [33] J. L. Beck, L. S. Katafygiotis, Updating models and their uncertainties. I: Bayesian statistical framework, *Journal of Engineering Mechanics* 124 (4) (1998) 455–461.
- [34] T. Sweeting, Discussion on ‘Parameter orthogonality and approximate conditional inference’ (by D. R. Cox and N. Reid), *Journal of the Royal Statistical Society. Series B (Methodological)* 49 (1987) 20–21.
- [35] D. R. Cox, N. Reid, Parameter orthogonality and approximate conditional inference, *Journal of the Royal Statistical Society. Series B (Methodological)* (1987) 1–39.
- [36] J. Ching, Y.-C. Chen, Transitional Markov chain Monte Carlo method for Bayesian model updating, model class selection, and model averaging, *Journal of Engineering Mechanics* 133 (7) (2007) 816–832.
- [37] S. H. Cheung, J. L. Beck, Bayesian model updating using hybrid Monte Carlo simulation with application to structural dynamic models with many uncertain parameters, *Journal of Engineering Mechanics* 135 (4) (2009) 243–255.
- [38] P. Angelikopoulos, C. Papadimitriou, P. Koumoutsakos, X-TMCMC: Adaptive kriging for Bayesian inverse modeling, *Computer Methods in Applied Mechanics and Engineering* 289 (2015) 409–428.
- [39] S.-K. Au, F.-L. Zhang, Fundamental two-stage formulation for Bayesian system identification, Part I: General theory, *Mechanical Systems and Signal Processing* 66 (2016) 31–42.

- [40] J. L. Beck, S.-K. Au, Bayesian updating of structural models and reliability using Markov chain Monte Carlo simulation, *Journal of Engineering Mechanics* 128 (4) (2002) 380–391.
- [41] B. Goller, H. Pradlwarter, G. Schueller, Robust model updating with insufficient data, *Computer Methods in Applied Mechanics and Engineering* 198 (37) (2009) 3096–3104.
- [42] P. Green, K. Worden, Bayesian and Markov chain Monte Carlo methods for identifying nonlinear systems in the presence of uncertainty, *Phil. Trans. R. Soc. A* 373 (2051) (2015) 20140405.
- [43] H. Jensen, C. Vergara, C. Papadimitriou, E. Millas, The use of updated robust reliability measures in stochastic dynamical systems, *Computer Methods in Applied Mechanics and Engineering* 267 (2013) 293–317.
- [44] S. Madireddy, B. Sista, K. Vemaganti, A Bayesian approach to selecting hyperelastic constitutive models of soft tissue, *Computer Methods in Applied Mechanics and Engineering* 291 (2015) 102–122.
- [45] D. Straub, I. Papaioannou, Bayesian updating with structural reliability methods, *Journal of Engineering Mechanics* 141 (3) (2014) 04014134.
- [46] K. Worden, J. Hensman, Parameter estimation and model selection for a class of hysteretic systems using Bayesian inference, *Mechanical Systems and Signal Processing* 32 (2012) 153–169.
- [47] J. Ching, J. L. Beck, K. A. Porter, Bayesian state and parameter estimation of uncertain dynamical systems, *Probabilistic Engineering Mechanics* 21 (1) (2006) 81–96.
- [48] J. Ching, J. L. Beck, K. A. Porter, R. Shaikhutdinov, Bayesian state estimation method for nonlinear systems and its application to recorded seismic response, *Journal of Engineering Mechanics* 132 (4) (2006) 396–410.
- [49] J. Ching, M. Muto, J. L. Beck, Structural model updating and health monitoring with incomplete modal data using Gibbs sampler, *Computer-Aided Civil and Infrastructure Engineering* 21 (4) (2006) 242–257.
- [50] K.-V. Yuen, J. L. Beck, Updating properties of nonlinear dynamical systems with uncertain input, *Journal of Engineering Mechanics* 129 (1) (2003) 9–20.
- [51] K.-V. Yuen, Recent developments of Bayesian model class selection and applications in civil engineering, *Structural Safety* 32 (5) (2010) 338–346.
- [52] C. Papadimitriou, J. L. Beck, L. S. Katafygiotis, Updating robust reliability using structural test data, *Probabilistic Engineering Mechanics* 16 (2) (2001) 103–113.
- [53] G. E. Box, N. R. Draper, *Empirical model-building and response surfaces*, vol. 424, Wiley New York, 1987.

- [54] P. Jayakumar, J. L. Beck, System identification using nonlinear structural models, in: Structural safety evaluation based on system identification approaches, Springer, 82–102, 1988.
- [55] P. Jayakumar, Modeling and identification in structural dynamics, Tech. Rep. EERL 87-01, California Institute of Technology, 1987.
- [56] S. A. Ashrafi, A. W. Smyth, Generalized masing approach to modeling hysteretic deteriorating behavior, *Journal of Engineering Mechanics* 133 (5) (2007) 495–505.
- [57] Y. Chen, A. Linderholt, T. Abrahamsson, An efficient simulation method for large-scale systems with local nonlinearities, in: *Special Topics in Structural Dynamics, Volume 6*, Springer, 259–267, 2016.
- [58] R. S. Thyagarajan, Modeling and analysis of hysteretic structural behavior, Tech. Rep. EERL 89-03, California Institute of Technology, 1989.
- [59] O. Ratmann, C. Andrieu, C. Wiuf, S. Richardson, Model criticism based on likelihood-free inference, with an application to protein network evolution, *Proceedings of the National Academy of Sciences* 106 (26) (2009) 10576–10581.