

RESEARCH ARTICLE 

On off-line and on-line Bayesian filtering for uncertainty quantification of structural deterioration

Antonios Kamariotis^{1,2} , Luca Sardi¹, Iason Papaioannou¹, Eleni Chatzi^{2,3}  and Daniel Straub¹

¹Engineering Risk Analysis Group, Technical University of Munich, Munich, Germany

²Institute for Advanced Study, Technical University of Munich, Garching, Germany

³Institute of Structural Engineering, ETH Zurich, Zurich, Switzerland

Corresponding author: Antonios Kamariotis; Email: antonis.kamariotis@tum.de

Received: 16 August 2022; **Revised:** 02 June 2023; **Accepted:** 20 June 2023


Keywords: Bayesian filtering; Gaussian mixture; Markov Chain Monte Carlo; particle filter; structural deterioration; uncertainty quantification

Abstract

Data-informed predictive maintenance planning largely relies on stochastic deterioration models. Monitoring information can be utilized to update sequentially the knowledge on model parameters. In this context, on-line (recursive) Bayesian filtering algorithms typically fail to properly quantify the full posterior uncertainty of time-invariant model parameters. Off-line (batch) algorithms are—in principle—better suited for the uncertainty quantification task, yet they are computationally prohibitive in sequential settings. In this work, we adapt and investigate selected Bayesian filters for parameter estimation: an on-line particle filter, an on-line iterated batch importance sampling filter, which performs Markov Chain Monte Carlo (MCMC) move steps, and an off-line MCMC-based sequential Monte Carlo filter. A Gaussian mixture model approximates the posterior distribution within the resampling process in all three filters. Two numerical examples provide the basis for a comparative assessment. The first example considers a low-dimensional, nonlinear, non-Gaussian probabilistic fatigue crack growth model that is updated with sequential monitoring measurements. The second high-dimensional, linear, Gaussian example employs a random field to model corrosion deterioration across a beam, which is updated with sequential sensor measurements. The numerical investigations provide insights into the performance of off-line and on-line filters in terms of the accuracy of posterior estimates and the computational cost, when applied to problems of different nature, increasing dimensionality and varying sensor information amount. Importantly, they show that a tailored implementation of the on-line particle filter proves competitive with the computationally demanding MCMC-based filters. Suggestions on the choice of the appropriate method in function of problem characteristics are provided.

Impact Statement

Stochastic models describing time-evolving deterioration processes are widespread in engineering. In the modern data-rich engineering landscape, Bayesian methods can exploit monitoring data to sequentially update knowledge on underlying model parameters. The precise probabilistic characterization of these parameters is indispensable for several real-world tasks, where decisions need to be taken in view of the evaluated margins of risk and uncertainty. This work investigates and compares on-line and off-line Bayesian filters and adapts the

 This research article was awarded an Open Materials badge for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.



former for posterior uncertainty quantification of time-invariant parameters. We show that tailored on-line particle filters are competitive alternatives to off-line Bayesian filters, especially in certain high-dimensional settings.

1. Introduction

Structural deterioration of various forms is present in most mechanical and civil structures and infrastructure systems. Accurate and effective tracking of structural deterioration processes can help to effectively manage it and minimize the total life-cycle costs (Cadini et al., 2009a; Frangopol, 2011; Kim et al., 2017; Kamariotis et al., 2022). The deployment of sensors on structural components/systems can enable long-term monitoring of such processes. Monitoring data obtained sequentially at different points in time must be utilized in an efficient manner within a Bayesian framework to enable data-informed estimation and prediction of the deterioration process evolution.

Monitored structural deterioration processes are commonly modeled using Markovian state-space representations (Myötyri et al., 2006; Cadini et al., 2009b; Baraldi et al., 2013), whereby the deterioration state evolution is represented by a recursive Markov process equation, and is subject to stochastic process noise (Corbetta et al., 2018). Monitoring information is incorporated by means of the measurement equation. The deterioration models further contain time-invariant uncertain parameters. The state-space can be augmented to include these parameters, if one wishes to obtain updated estimates thereof conditional on the monitoring information (Saha et al., 2009; Straub, 2009; Sun et al., 2014; Corbetta et al., 2018; Yi and Song, 2018; Cristiani et al., 2021; Kamariotis et al., 2023); this is referred to as joint state-parameter estimation (Särkkä, 2013; Kantas et al., 2015).

A different approach entails defining the structural deterioration state evolution solely as a function of uncertain time-invariant model parameters (Ditlevsen and Madsen, 1996; Vu and Stewart, 2000; Elingwood, 2005; Stewart and Mullard, 2007), which can be updated in view of the monitoring data. This updating, referred to herein as Bayesian parameter estimation, is often the primary task of interest. In this approach, the deterioration state variables are obtained as outputs of the calibrated deterioration model with posterior parameter estimates (Kennedy and O'Hagan, 2001; Ramanacha et al., 2022). The parameter estimation problem can be cast into a Markovian state-space representation. Quantifying the full posterior uncertainty of the time-invariant model parameters is essential for performing monitoring-informed predictions on the deterioration process evolution, the subsequent monitoring-informed estimation of the time-variant structural reliability (Melchers and Beck, 2017; Straub et al., 2020) or the remaining useful life (Sun et al., 2014; Kim et al., 2017), and eventually for predictive maintenance planning.

Bayesian estimation of time-invariant deterioration model parameters is the main focus of this article. In long-term deterioration monitoring settings, where data is obtained sequentially at different points in time, Bayesian inference can be performed either in an on-line or an off-line framework (Storvik, 2002; Kantas et al., 2015; Azam et al., 2017). In literature, these are also referred to as recursive (on-line) and batch (off-line) estimation (Särkkä, 2013). Parameter estimation is cast into a state-space setup to render it suitable for application with on-line Bayesian filtering algorithms (Kantas et al., 2015), such as the Kalman filter (Kalman, 1960) and its nonlinear variants (Jazwinski, 1970; Julier and Uhlmann, 1997; Daum, 2005; Song et al., 2020), the ensemble Kalman filter (Evensen, 2006), and particle filters (Doucet et al., 2001; Chopin, 2004; Doucet and Johansen, 2008; Hu et al., 2008; Särkkä, 2013; Tatsis et al., 2022). We employ on-line particle filter methods for pure recursive estimation of time-invariant deterioration model parameters. This is not the typical use case for such methods, which often yield degenerate and impoverished posterior estimates, hence, failing to effectively characterize the posterior uncertainty (Del Moral et al., 2006; Särkkä, 2013). In this work, we tailor on-line particle filters for quantifying the full posterior uncertainty of time-invariant model parameters. Subsequently, we provide a formal investigation and discussion on their suitability with respect to this task.

In its most typical setting within engineering applications, Bayesian parameter estimation is commonly performed with the use of off-line Markov Chain Monte Carlo (MCMC) methods, which have been used extensively in statistics and engineering to sample from complex posterior distributions of model parameters (Hastings, 1970; Gilks et al., 1995; Beck and Au, 2002; Haario et al., 2006; Ching and Chen, 2007; Papaioannou et al., 2015; Wu et al., 2017; Lye et al., 2021). However, use of off-line methods for on-line estimation tasks is computationally prohibitive (Del Moral et al., 2006; Kantas et al., 2015). Additionally, when considering off-line inference, in settings when measurements are obtained sequentially at different points in time, off-line MCMC methods tend to induce a larger computational cost than on-line particle filter methods, which can be important, for example, when optimizing inspection and monitoring (Papakonstantinou and Shinozuka, 2014; Luque and Straub, 2019; Kamariotis et al., 2023). Questions that we investigate in this context include: Can one precisely quantify the posterior uncertainty of time-invariant parameters when employing on-line particle filter methods? How does this quantification compare against the posterior estimates obtained with off-line MCMC methods? How does the estimation accuracy depend on the nature of the problem, that is, dimensionality, nonlinearity, or non-Gaussianity? What is the computational cost induced by the different methods? Ideally, one would opt for the method which can provide sufficiently accurate posterior results at the expense of the least computational cost. To address these questions, this article selects and adapts algorithms in view of parameter estimation, and performs a comparative assessment of selected off-line and on-line Bayesian filters specifically tailored for posterior uncertainty quantification of time-invariant parameters. The innovative comparative assessment results in a set of suggestions on the choice of the appropriate algorithm in function of problem characteristics.

The article is structured as follows. **Section 2** provides a detailed description of on-line and off-line Bayesian inference in the context of parameter estimation. Three different selected and adapted Bayesian filters are presented in algorithmic detail, namely an on-line particle filter with Gaussian mixture-based resampling (PFGM) (van der Merwe and Wan, 2003; McLachlan and Krishnan, 2007), the on-line iterated batch importance sampling filter (IBIS) (Chopin, 2002), which performs off-line MCMC steps with a Gaussian mixture as a proposal distribution, and an off-line MCMC-based sequential Monte Carlo (SMC) filter (Del Moral et al., 2006), which enforces tempering of the likelihood function (known as simulated annealing) to sequentially arrive to the single final posterior density of interest (Neal, 2001; Jasra et al., 2011). The tPFGM and tIBIS variants, which adapt the PFGM and IBIS filters by employing tempering of the likelihood function of each new measurement, are further presented and proposed for problems with high sensor information amount. **Section 3** describes the two case studies that serve as the basis for numerical investigations, one nonlinear, non-Gaussian, and low-dimensional, and one linear, Gaussian, and high-dimensional. MATLAB codes implementing the different algorithms and applying them on the two case studies introduced in this article are made publicly available via a GitHub repository.¹ **Section 4** summarizes the findings of this comparative assessment, provides suggestions on choice of the appropriate method according to the nature of the problem, discusses cases which are not treated in our investigations, and concludes this work.

2. On-Line and Off-Line Bayesian Filtering for Time-Invariant Parameter Estimation

This work assumes the availability of a stochastic deterioration model D , parametrized by a vector $\theta \in \mathbb{R}^d$ containing the d uncertain time-invariant model parameters. We collect the uncertain parameters influencing the deterioration process in the vector θ . In the Bayesian framework, θ is modeled as a vector of random variables with a prior distribution $\pi_{\text{pr}}(\theta)$. We assume that the deterioration process is monitored via a permanently installed monitoring system. Long-term monitoring of a deterioration process leads to sets of noisy measurements $\{y_1, \dots, y_n\}$ obtained sequentially at different points in time $\{t_1, \dots, t_n\}$ throughout the lifetime of a structural component/system. Such measurements can be used to update

¹ https://github.com/antoniskam/Offline_online_Bayes.

the distribution of θ ; this task is referred to as Bayesian parameter estimation. Within a deterioration monitoring setting, Bayesian parameter estimation can be performed either in an on-line or an off-line framework (Kantas et al., 2015), depending on the task of interest.

In an on-line framework, one is interested in updating the distribution of θ in a sequential manner, that is, at every time step t_n when a new measurement y_n becomes available, conditional on all measurements available up to t_n . Thus, in an on-line framework, inference of the sequence of posterior densities $\{\pi_{\text{pos}}(\theta|\mathbf{y}_{1:n})\}_{n \geq 1}$ is the goal, where $\mathbf{y}_{1:n}$ denotes the components $\{y_1, \dots, y_n\}$. We point out that in this article the term on-line does not relate to “real-time” estimation, although on-line algorithms are also used in real-time estimation (Chatzi and Smyth, 2009; Russell and Norvig, 2021).

In contrast, in an off-line framework, inference of θ is performed at a fixed time step t_N using a fixed set of measurements $\{y_1, \dots, y_N\}$, and the single posterior density $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:N})$ is sought, which can be estimated via Bayes’ rule as:

$$\pi_{\text{pos}}(\theta|\mathbf{y}_{1:N}) \propto L(\mathbf{y}_{1:N}|\theta)\pi_{\text{pr}}(\theta), \quad (1)$$

where $L(\mathbf{y}_{1:N}|\theta)$ denotes the likelihood function of the whole measurement set $\mathbf{y}_{1:N}$ given the parameters θ . With the assumption that measurements are independent given the parameter state, $L(\mathbf{y}_{1:N}|\theta)$ can be expressed as a product of the likelihoods $L(y_n|\theta)$ as:

$$L(\mathbf{y}_{1:N}|\theta) = \prod_{n=1}^N L(y_n|\theta). \quad (2)$$

MCMC methods sample from $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:N})$ via simulation of a Markov chain with $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:N})$ as its stationary distribution, for example, by performing Metropolis Hastings (MH) steps (Hastings, 1970). MCMC methods do not require estimation of the normalization constant in equation (1). However, in the on-line framework, MCMC methods are impractical, since they require simulating anew a different Markov chain for each new posterior $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:n})$, and the previously generated Markov chain for the posterior estimation of $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:n-1})$ is not accounted for, except when choosing the seed for initializing the new Markov chain. This implies that MCMC methods quickly become computationally prohibitive in the on-line framework, already for a small n . An additional computational burden stems from the fact that each step within the MCMC sampling process requires evaluation of the full likelihood function $L(\mathbf{y}_{1:n}|\theta)$, that is, the whole set of measurements $\mathbf{y}_{1:n}$ needs to be processed. This leads to increasing computational complexity for increasing n , and can render use of MCMC methods computationally inefficient even for off-line inference, especially when N is large.

On-line particle filters (Särkkä, 2013; Kantas et al., 2015) operate in a sequential fashion by making use of the Markovian property of the employed state-space representation, that is, they compute $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:n})$ solely based on $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:n-1})$ and the new measurement y_n . The typical use of particle filters targets the tracking of a system’s response (dynamic state) by means of a state-space representation (Gordon et al., 1993; Särkkä, 2013), while they are often also used also for joint state-parameter estimation tasks, wherein the state-space is augmented to include the model parameters to be estimated (Särkkä, 2013; Kantas et al., 2015). In addition, particle filters can also be applied for pure recursive estimation of time-invariant parameters, for which the noise in the dynamic model is formally zero (Del Moral et al., 2006; Särkkä, 2013), although this is not the typical setting for application of particle filters. A model of the Markovian discrete-time state-space representation for the case of time-invariant parameter estimation is given in equations (3a) and (3b).

$$\theta_n = \theta_{n-1}, \quad (3a)$$

$$y_n = D_n(\theta_n) \exp(\varepsilon_n), \quad (3b)$$

where ε_n models the error/noise of the measurement at time t_n , and θ_n denotes the time-invariant parameter vector at time step n . The dynamic equation for the time-invariant parameters equation (3a) is introduced for the sole purpose of casting the problem into a state-space representation. Since the measurements are

assumed independent given the parameter state, the errors ε_n in equation (3b) are independent. It should be noted that the measurement error, which is introduced in multiplicative form in equation (3b), is commonly expressed in an additive form (Corbetta et al., 2018). Indeed, equation (3b) can be reformulated in the logarithmic scale, whereby the measurement error is expressed in an additive form. The multiplicative form of the measurement error in equation (3b) is consistent with the fact that—in the context of structural deterioration—measurements y_n cannot be negative. All target distributions of interest in the sequence $\pi_{\text{pos}}(\theta_n | \mathbf{y}_{1:n})$ are defined on the same space of $\theta \in \mathbb{R}^d$. In the remainder of this article, the subscript n will therefore be dropped from θ_n . As previously discussed, particle filters are mainly used for on-line inference. However, these can also be used in exactly the same way for off-line inference, where only a single posterior density $\pi_{\text{pos}}(\theta | \mathbf{y}_{1:N})$ is of interest. In this case, particle filters use the sequence of measurements successively to sequentially arrive to the final single posterior density of interest via estimating all the intermediate distributions.

2.1. On-line particle filter

Particle filter (PF) methods, also referred to as sequential Monte Carlo (SMC) methods, are importance sampling-based techniques that use a set of weighted samples $\left\{ \left(\theta_n^{(i)}, w_n^{(i)} \right) : i = 1, \dots, N_{\text{par}} \right\}$, called particles, to represent the posterior distribution of interest at estimation time step n , $\pi_{\text{pos}}(\theta | \mathbf{y}_{1:n})$. PFs form the following approximation to the posterior distribution of interest:

$$\pi_{\text{pos}}(\theta | \mathbf{y}_{1:n}) \approx \sum_{i=1}^{N_{\text{par}}} w_n^{(i)} \delta(\theta - \theta_n^{(i)}), \tag{4}$$

where δ denotes the Dirac delta function.

When a new measurement y_n becomes available, PFs shift from $\pi_{\text{pos}}(\theta | \mathbf{y}_{1:n-1})$ to $\pi_{\text{pos}}(\theta | \mathbf{y}_{1:n})$ by importance sampling using an appropriately chosen importance distribution, which results in a reweighting procedure (updating of the weights). An important issue that arises from this weight-updating procedure is the sample degeneracy problem (Särkkä, 2013). This relates to the fact that the importance weights $w_n^{(i)}$ become more unevenly distributed with each updating step. In most cases, after a certain number of updating steps, the weights of almost all the particles assume values close to zero (see Figure 1).

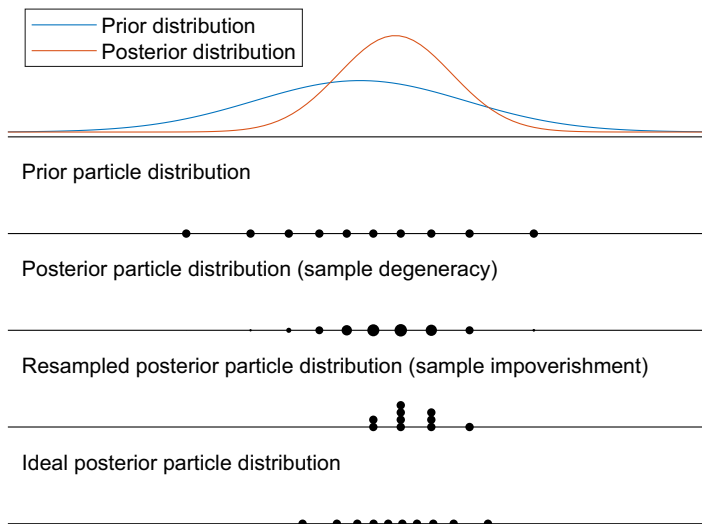


Figure 1. Sample degeneracy and impoverishment.

This problem is alleviated by the use of adaptive resampling procedures based on the effective sample size $N_{\text{eff}} = 1 / \sum_{i=1}^{N_{\text{par}}} (w_n^{(i)})^2$ (Liu and Chen, 1998). Most commonly, resampling is performed with replacement according to the particle weights whenever N_{eff} drops below a user-defined threshold $N_T = cN_{\text{par}}$, $c \in [0, 1]$. Resampling introduces additional variance to the parameter estimates (Särkkä, 2013). In the version of the PF algorithm presented in Algorithm 1, the dynamic model of equation (3) is used as the importance distribution, as originally proposed in the bootstrap filter by Gordon et al. (1993). Theoretical analysis of PF algorithms can be found in numerous seminal sources (e.g., Chopin, 2004; Del Moral et al., 2006; Hu et al., 2008).

Algorithm 1 Particle Filter (PF)

- 1: generate N_{par} initial particles $\theta^{(i)}$ from $\pi_{\text{pr}}(\theta)$, $i = 1, \dots, N_{\text{par}}$
 - 2: assign initial weights $w_0^{(i)} = 1/N_{\text{par}}$, $i = 1, \dots, N_{\text{par}}$
 - 3: **for** $n = 1, \dots, N$ **do**
 - 4: evaluate likelihood of the particles based on new measurement y_n , $L_n^{(i)} = L(y_n | \theta^{(i)})$
 - 5: update particle weights $w_n^{(i)} \propto L_n^{(i)} \cdot w_{n-1}^{(i)}$ and normalize s.t. $\sum_{i=1}^{N_{\text{par}}} w_n^{(i)} = 1$
 - 6: evaluate $N_{\text{eff}} = \frac{1}{\sum_{i=1}^{N_{\text{par}}} (w_n^{(i)})^2}$
 - 7: **if** $N_{\text{eff}} < N_T$ **then**
 - 8: resample particles $\theta^{(i)}$ with replacement according to $w_n^{(i)}$
 - 9: reset particle weights to $w_n^{(i)} = 1/N_{\text{par}}$
 - 10: **end if**
 - 11: **end for**
-

When using PFs to estimate time-invariant parameters, for which the process noise in the dynamic equation is zero, one runs into the issue of sample impoverishment (Särkkä, 2013). The origin of this issue is the resampling process. More specifically, after a few resampling steps, most (or in extreme cases all) of the particles in the sample set end up assuming the exact same value, that is, the particle set consists of only few (or one) distinct particles (see Figure 1). The sample impoverishment issue poses the greatest obstacle for time-invariant parameter estimation with PFs. A multitude of techniques has been suggested in literature to alleviate the sample impoverishment issue in joint state-parameter estimation setups (see, e.g., Gilks and Berzuini, 2001; Liu and West, 2001; Musso et al., 2001; Storvik, 2002; Andrieu et al., 2004; Andrieu et al., 2010; Carvalho et al., 2010; Chatzi and Smyth, 2013; Chopin et al., 2013). Fewer works have proposed solutions for resolving this issue in parameter estimation setups (see, e.g., Chopin, 2002; Del Moral et al., 2006). One of the simplest and most commonly used approaches consists in introducing artificial dynamics in the dynamic model of the parameter vector, that is, the dynamic model $\theta_n = \theta_{n-1} + \varepsilon_{n-1}$ is employed, where ε_{n-1} is a small artificial process noise (Kitagawa, 1998). In this way, the time-invariant parameter vector is transformed into a time-variant one, therefore, the parameter estimation problem deviates from the original one (Särkkä, 2013; Kantas et al., 2015). This approach can introduce a bias and an artificial variance inflation in the estimates (Kantas et al., 2015). For these reasons, this approach is not considered in this article.

To resolve the sample impoverishment issue encountered when using the PF Algorithm 1 for parameter estimation, this work employs the particle filter with Gaussian mixture resampling (PFGM), described in Algorithm 2. The PFGM algorithm relates to preexisting concepts (van der Merwe and Wan, 2003; Vehtil and Chakravorty, 2016), and is here specifically tailored for the parameter estimation task, with its main goal being, in contrast to previous works, the quantification of the full posterior parameter uncertainty. A comparison between Algorithms 1 and 2 shows that the only difference lies in the way that the resampling step is performed. PFGM replaces the standard resampling process of PF by first approximating the posterior distribution at estimation step n by a Gaussian mixture model (GMM), which is fitted via the Expectation–Maximization (EM) algorithm (McLachlan and Krishnan, 2007; Chen et al., 2010) on

the weighted particle set. The degenerating particle set is then rejuvenated by sampling N_{par} new particles from the GMM of equation (5),

$$p(\boldsymbol{\theta}|\mathbf{y}_{1:n}) \approx \sum_{i=1}^{N_{\text{GM}}} \phi_i \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i), \quad (5)$$

where ϕ_i represents the weight of the Gaussian component i , while $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ are the respective mean vector and covariance matrix. The number of Gaussians in the mixture N_{GM} , has to be chosen in advance, or can be estimated by use of appropriate algorithms (Schubert et al., 2017; Celeux et al., 2019; Geyer et al., 2019). In the numerical investigations of Section 3, we set $N_{\text{GM}} = 8$. We point out that the efficacy of PFGM strongly depends on the quality of the GMM posterior approximation. The reason for applying a GMM (and not a single Gaussian) is that the posterior distribution can deviate from the normal distribution, and can even be multimodal or heavy-tailed.

Algorithm 2 Particle Filter with Gaussian mixture resampling (PFGM)

- 1: generate N_{par} initial particles $\boldsymbol{\theta}^{(i)}$ from $\pi_{\text{pr}}(\boldsymbol{\theta})$, $i = 1, \dots, N_{\text{par}}$
 - 2: assign initial weights $w_0^{(i)} = 1/N_{\text{par}}$, $i = 1, \dots, N_{\text{par}}$
 - 3: **for** $n = 1, \dots, N$ **do**
 - 4: evaluate likelihood of the particles based on new measurement y_n , $L_n^{(i)} = L(y_n|\boldsymbol{\theta}^{(i)})$
 - 5: update particle weights $w_n^{(i)} \propto L_n^{(i)} \cdot w_{n-1}^{(i)}$ and normalize s.t. $\sum_{i=1}^{N_{\text{par}}} w_n^{(i)} = 1$
 - 6: evaluate $N_{\text{eff}} = \frac{1}{\sum_{i=1}^{N_{\text{par}}} (w_n^{(i)})^2}$
 - 7: **if** $N_{\text{eff}} < N_T$ **then**
 - 8: EM: fit a Gaussian mixture proposal distribution $g_{\text{GM}}(\boldsymbol{\theta})$ according to $\{\boldsymbol{\theta}^{(i)}, w_n^{(i)}\}$
 - 9: sample N_{par} new particles $\boldsymbol{\theta}^{(i)}$ from $g_{\text{GM}}(\boldsymbol{\theta})$
 - 10: reset particle weights to $w_n^{(i)} = 1/N_{\text{par}}$
 - 11: **end if**
 - 12: **end for**
-

The simple reweighting procedure used in the on-line PFs is based on the premise that $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n-1})$ and $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n})$ are likely to be similar, that is, that the new measurement y_n will not cause a very large change in the posterior. However, when that is not the case, this simple reweighting procedure is bound to perform poorly, leading to very fast degeneration of the particle set. In cases where already the first measurement set y_1 is strongly informative relative to the prior, the PF is bound to strongly degenerate already in the first weight updating step (e.g., we observe this in the second case study of subsec:RF in the case of 10 sensors). To counteract this issue, in this article, we incorporate the idea of simulated annealing (enforcing tempering of the likelihood function) (Neal, 2001) when needed within the on-line PFGM algorithm, which we term the tPFGM Algorithm 3. The tPFGM algorithm draws inspiration from previous works (Deutscher et al., 2000; Gall et al., 2007), but is here tailored for the parameter estimation task, opting for the quantification of the full posterior parameter uncertainty. The algorithm operates as follows: At estimation time step n , before performing the reweighting operation, the algorithm first checks the updated effective sample size for indication of sample degeneracy. If no degeneracy is detected, tPFGM operates exactly like PFGM. When sample degeneracy occurs, tPFGM employs adaptive tempering of the likelihood $L(y_n|\boldsymbol{\theta})$ of the new measurement y_n in order to “sequentially” sample from $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n-1})$ to $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n})$ by visiting a sequence of artificial intermediate posteriors, as defined by the tempered likelihood function $L^q(y_n|\boldsymbol{\theta})$. The tempering factor q takes values between 0 and 1. When $q = 0$, the new measurement y_n is neglected, while $q = 1$ entails considering the whole likelihood function of y_n , thus reaching to $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n})$. The intermediate values of q are adaptively selected via solution of the optimization problem in line 11 of Algorithm 3, which ensures that the effective sample size does not drop below the threshold N_T for the chosen q value. Naturally, use

of tPFGM can trigger more resampling events than PFGM, as resampling can occur more than once within a time step n .

Algorithm 3 Particle Filter with Gaussian mixture resampling and likelihood tempering (tPFGM)

```

1: generate  $N_{\text{par}}$  initial particles  $\theta^{(i)}$  from  $\pi_{\text{pr}}(\theta)$ ,  $i = 1, \dots, N_{\text{par}}$ 
2: assign initial weights  $w_0^{(i)} = 1/N_{\text{par}}$ ,  $i = 1, \dots, N_{\text{par}}$ 
3: for  $n = 1, \dots, N$  do
4:   evaluate likelihood of the particles based on new measurement  $y_n$ ,  $L_n^{(i)} = L(y_n | \theta^{(i)})$ 
5:   set  $q = 0$  and create auxiliary particle weights  $w_a^{(i)} = w_{n-1}^{(i)}$ 
6:   while  $q \neq 1$  do
7:     if  $N_{\text{eff}} = \left( \sum_{i=1}^{N_{\text{par}}} w_a^{(i)} \cdot L_n^{(i)1-q} \right)^2 / \sum_{i=1}^{N_{\text{par}}} \left( w_a^{(i)} \cdot L_n^{(i)1-q} \right)^2 > N_T$  then
8:       update auxiliary particle weights  $w_a^{(i)} \propto w_a^{(i)} \cdot L_n^{(i)1-q}$  and normalize s.t.  $\sum_{i=1}^{N_{\text{par}}} w_a^{(i)} = 1$ 
9:       set  $q = 1$ 
10:    else
11:      solve  $\left( \sum_{i=1}^{N_{\text{par}}} w_a^{(i)} \cdot L_n^{(i)dq} \right)^2 / \sum_{i=1}^{N_{\text{par}}} \left( w_a^{(i)} \cdot L_n^{(i)dq} \right)^2 - N_T = 0$  for  $dq$ 
12:      set  $q_{\text{new}} = \min[q + dq, 1]$ 
13:      set  $dq = q_{\text{new}} - q$  and  $q = q_{\text{new}}$ 
14:      update auxiliary particle weights  $w_a^{(i)} \propto w_a^{(i)} \cdot L_n^{(i)dq}$  and normalize s.t.  $\sum_{i=1}^{N_{\text{par}}} w_a^{(i)} = 1$ 
15:      EM: fit a Gaussian mixture proposal distribution  $g_{\text{GM}}(\theta)$  according to  $\left\{ \theta^{(i)}, w_a^{(i)} \right\}$ 
16:      sample  $N_{\text{par}}$  new particles  $\theta^{(i)}$  from  $g_{\text{GM}}(\theta)$ 
17:      reset auxiliary particle weights to  $w_a^{(i)} = 1/N_{\text{par}}$ 
18:    end if
19:  end while
20:  set  $w_n^{(i)} = w_a^{(i)}$ 
21: end for

```

The PFGM and tPFGM filters rely entirely on the posterior approximation via a GMM for sampling N_{par} new particles during the resampling process. However, there is no guarantee that these new particles follow the true posterior distribution of interest. The IBIS filter of the following Section 2.2 aims at addressing this issue.

2.2. Iterated batch importance sampling

Implementing MCMC steps within PF methods to move the particles after a resampling step was originally proposed by Gilks and Berzuini (2001), in the so-called resample-move algorithm. Chopin (2002) introduced a special case of the resample-move algorithm, specifically tailored for application to static parameter estimation purposes, namely the iterated batch importance sampling (IBIS) filter. IBIS was originally introduced as an iterative method for solving off-line estimation tasks by incorporating the sequence of measurements one at a time. In doing this, the algorithm visits the sequence of intermediate posteriors within its process, and can therefore also be used to perform on-line estimation tasks. An on-line version of the IBIS filter is presented in Algorithm 5, used in conjunction with the MCMC routine of Algorithm 4.

The core idea of the IBIS filter is the following: At estimation step n , if sample degeneracy is identified, first the particles are resampled with replacement, and subsequently the resampled particles are moved with a Markov chain transition kernel whose stationary distribution is $\pi_{\text{pos}}(\theta | \mathbf{y}_{1:n})$. More specifically, each of the N_{par} resampled particles is used as the seed to perform a single MCMC step. This approach is inherently different to standard applications of MCMC, where a transition kernel is applied multiple times on one particle.

A question that arises is how to choose the Markov chain transition kernel. Chopin (2002) argues for choosing a transition kernel that ensures that the proposed particle only weakly depends on the seed particle value. It is therefore recommended to use an independent Metropolis-Hastings (IMH) kernel, wherein the proposed particle is sampled from a proposal distribution g , which has to be as close as possible to the target distribution $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n})$. In obtaining such a proposal distribution, along the lines of what is described in Section 2.1, in this work we employ a GMM approximation (see equation (5)) of the target distribution as the proposal density $g_{\text{GM}}(\boldsymbol{\theta})$ within the IMH kernel (Papaioannou et al., 2016; South et al., 2019). The IMH kernel with a GMM proposal distribution is denoted IMH-GM herein. The acceptance probability (line 6 of Algorithm 4) of the IMH-GM kernel is a function of both the initial seed particle and the GMM proposed particle. The acceptance rate can indicate how efficient the IMH-GM kernel is in performing the MCMC move step within the IBIS algorithm. It is important to note that when computing the acceptance probability, a call of the full likelihood function is invoked, which requires the whole set of measurements $\mathbf{y}_{1:n}$ to be processed; this leads to a significant additional computational demand, which pure on-line methods are not supposed to accommodate (Doucet et al., 2001).

The performance of the IBIS sampler depends highly on the mixing properties of the IMH-GM kernel. If the kernel leads to slowly decreasing chain auto-correlation, the moved particles are bound to remain in regions close to the particles obtained by the resampling step. This can lead to an under-representation of the parameter space of the intermediate posterior distribution. It might therefore be beneficial to add a burn-in period within the IMH-GM kernel (Del Moral et al., 2006). Implementing that is straightforward and is shown in Algorithm 4, where n_{B} is the user-defined number of burn-in steps. Naturally, the computational cost of the IMH-GM routine increases linearly with the number of burn-in steps.

Algorithm 4 Independent Metropolis Hastings with GM proposal (IMH-GM)

```

1: IMH-GM Input:  $\{\boldsymbol{\theta}^{(i)}, L^{(i)} \cdot \pi_{\text{pr}}(\boldsymbol{\theta}^{(i)})\}$ ,  $\pi_{\text{pr}}(\boldsymbol{\theta})$ ,  $L(\mathbf{y}_{1:n}|\boldsymbol{\theta})$  and  $g_{\text{GM}}(\boldsymbol{\theta})$ 
2: for  $i = 1, \dots, N_{\text{par}}$  do
3:   for  $j = 1, \dots, n_{\text{B}} + 1$  do
4:     sample candidate particle  $\boldsymbol{\theta}_{c,j}^{(i)}$  from  $g_{\text{GM}}(\boldsymbol{\theta})$ 
5:     evaluate  $L_{c,j}^{(i)} = L(\mathbf{y}_{1:n}|\boldsymbol{\theta}_{c,j}^{(i)})$  for candidate particle
6:     evaluate acceptance ratio  $\alpha = \min \left[ 1, \frac{L_{c,j}^{(i)} \cdot \pi_{\text{pr}}(\boldsymbol{\theta}_{c,j}^{(i)}) \cdot g_{\text{GM}}(\boldsymbol{\theta}^{(i)})}{L^{(i)} \cdot \pi_{\text{pr}}(\boldsymbol{\theta}^{(i)}) \cdot g_{\text{GM}}(\boldsymbol{\theta}_{c,j}^{(i)})} \right]$ 
7:     generate uniform random number  $u \in [0, 1]$ 
8:     if  $u < \alpha$  then
9:       replace  $\{\boldsymbol{\theta}^{(i)}, L^{(i)} \cdot \pi_{\text{pr}}(\boldsymbol{\theta}^{(i)})\}$  with  $\{\boldsymbol{\theta}_{c,j}^{(i)}, L_{c,j}^{(i)} \cdot \pi_{\text{pr}}(\boldsymbol{\theta}_{c,j}^{(i)})\}$ 
10:    end if
11:  end for
12: end for
13: IMH-GM Output:  $\{\boldsymbol{\theta}^{(i)}, L^{(i)} \cdot \pi_{\text{pr}}(\boldsymbol{\theta}^{(i)})\}$ 

```

Algorithm 5 details the workings of the IMH-GM-based IBIS filter used in this work. In line 11 of this algorithm, the IMH-GM routine of Algorithm 4 is called, which implements the IMH-GM kernel for the MCMC move step. Comparing Algorithms 2 and 5, it is clear that both filters can be used for on-line inference within a single run, but the IBIS filter has significantly larger computational cost, as will also be demonstrated in the numerical investigations of Section 3. In the same spirit as the proposed tPFGM Algorithm 3, which enforces simulated annealing (tempering of the likelihood function) in cases when $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n-1})$ and $\pi_{\text{pos}}(\boldsymbol{\theta}|\mathbf{y}_{1:n})$ are likely to be quite different, the same idea can be implemented also within the IBIS algorithm. That leads to what we refer to as the tIBIS algorithm in this article.

Algorithm 5 IMH-GM-based Iterated Batch Importance Sampling (IBIS)

- 1: generate N_{par} initial particles $\theta^{(i)}$ from $\pi_{\text{pr}}(\theta)$, $i = 1, \dots, N_{\text{par}}$
 - 2: assign initial weights $w_0^{(i)} = 1/N_{\text{par}}$, $i = 1, \dots, N_{\text{par}}$
 - 3: **for** $n = 1, \dots, N$ **do**
 - 4: evaluate likelihood of the particles based on new measurement y_n , $L_n^{(i)} = L(y_n|\theta^{(i)})$
 - 5: evaluate the new target distribution, $L(\mathbf{y}_{1:n}|\theta^{(i)}) \cdot \pi_{\text{pr}}(\theta^{(i)}) = L_n^{(i)} \cdot L(\mathbf{y}_{1:n-1}|\theta^{(i)}) \cdot \pi_{\text{pr}}(\theta^{(i)})$
 - 6: update particle weights $w_n^{(i)} \propto L_n^{(i)} \cdot w_{n-1}^{(i)}$ and normalize s.t. $\sum_{i=1}^{N_{\text{par}}} w_n^{(i)} = 1$
 - 7: evaluate $N_{\text{eff}} = \frac{1}{\sum_{i=1}^{N_{\text{par}}} (w_n^{(i)})^2}$
 - 8: **if** $N_{\text{eff}} < N_T$ **then**
 - 9: EM: fit a Gaussian mixture proposal distribution $g_{\text{GM}}(\theta)$ according to $\{\theta^{(i)}, w_n^{(i)}\}$
 - 10: resample N_{par} new particles $\{\theta^{(i)}, L(\mathbf{y}_{1:n}|\theta^{(i)}) \cdot \pi_{\text{pr}}(\theta^{(i)})\}$ with replacement according to $w_n^{(i)}$
 - 11: IMH-GM step with inputs $\{\theta^{(i)}, L(\mathbf{y}_{1:n}|\theta^{(i)}) \cdot \pi_{\text{pr}}(\theta^{(i)})\}$, $\pi_{\text{pr}}(\theta)$, $L(\mathbf{y}_{1:n}|\theta)$, and $g_{\text{GM}}(\theta)$
 - 12: reset particle weights to $w_n^{(i)} = 1/N_{\text{par}}$
 - 13: **end if**
 - 14: **end for**
-

2.3. Off-line sequential Monte Carlo sampler

In Section 4 of Del Moral et al. (2006), the authors presented a generic approach to convert an off-line MCMC sampler into a sequential Monte Carlo (SMC) sampler tailored for performing off-line estimation tasks, that is, for estimating the single posterior density of interest $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:N})$. The off-line SMC sampler used in this work is presented in Algorithm 6 based on Del Moral et al. (2006) and Jasra et al. (2011). The key idea of this sampler is to adaptively construct the following artificial sequence of densities,

$$\pi_j(\theta|\mathbf{y}_{1:N}) \propto L^{q_j}(\mathbf{y}_{1:N}|\theta)\pi_{\text{pr}}(\theta), \tag{6}$$

where q_j is a tempering parameter which obtains values between 0 and 1, in order to “sequentially” sample in a smooth manner from the prior to the final single posterior density of interest. Once $q_j = 1$, $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:N})$ is reached. Similar to what was described in tPFGM, the intermediate values of q_j are adaptively found via solution of the optimization problem in line 5 of Algorithm 6. The GMM approximation of the intermediate posteriors and the IMH-GM kernel of Algorithm 4 in order to move the particles after resampling are also key ingredients of this SMC sampler. Unlike PFGM and IBIS, this SMC algorithm cannot provide the on-line solution within a single run, and has to be rerun from scratch for every new target posterior of interest. In this regard, use of Algorithm 6 for on-line inference is impractical. We choose to include this algorithm for the purpose of the comparative assessment in this article, as off-line algorithms are generally considered to be better suited for the full posterior uncertainty quantification of time-invariant parameters.

Algorithm 6 IMH-GM-based Sequential Monte Carlo (SMC)

- 1: generate N_{par} initial particles $\theta^{(i)}$ from $\pi_{\text{pr}}(\theta)$, $i = 1, \dots, N_{\text{par}}$
- 2: evaluate for every particle the full likelihood $L^{(i)} = L(\mathbf{y}_{1:N}|\theta^{(i)})$ and the prior $\pi_{\text{pr}}(\theta^{(i)})$
- 3: set $q = 0$
- 4: **while** $q \neq 1$ **do**
- 5: solve $\left(\sum_{i=1}^{N_{\text{par}}} L^{(i)dq}\right)^2 / \sum_{i=1}^{N_{\text{par}}} L^{(i)2dq} - N_T = 0$ for dq
- 6: set $q_{\text{new}} = \min[q + dq, 1]$
- 7: set $dq = q_{\text{new}} - q$ and $q = q_{\text{new}}$

- 8: evaluate particle weights $w^{(i)} \propto L^{(i)dq}$ and normalize s.t. $\sum_{i=1}^{N_{\text{par}}} w^{(i)} = 1$
- 9: EM: fit a Gaussian mixture proposal distribution $g_{\text{GM}}(\theta)$ according to $\{\theta^{(i)}, w^{(i)}\}$
- 10: resample N_{par} new particles $\{\theta^{(i)}, L^{(i)q} \cdot \pi_{\text{pr}}(\theta^{(i)})\}$ with replacement according to $w^{(i)}$
- 11: IMH-GM step with inputs $\{\theta^{(i)}, L^{(i)q} \cdot \pi_{\text{pr}}(\theta^{(i)})\}$, $\pi_{\text{pr}}(\theta)$, $L^q(\mathbf{y}_{1:N}|\theta)$, and $g_{\text{GM}}(\theta)$
- 12: reset particle weights to $w^{(i)} = 1/N_{\text{par}}$
- 13: **end while**

2.4. Computational remarks

The algebraic operations in all presented algorithms are implemented in the logarithmic scale, which employs evaluations of the logarithm of the likelihood function and, hence, ensures computational stability. Furthermore, the EM step for fitting the GMM is performed after initially transforming the prior joint probability density function of θ to an underlying vector \mathbf{u} of independent standard normal random variables (Kiureghian and Liu, 1986). In standard normal space, the parameters are decorrelated, which enhances the performance of the EM algorithm.

3. Numerical Investigations

3.1. Low-dimensional case study: Paris–Erdogan fatigue crack growth model

A fracture mechanics-based model serves as the first case study. This describes the fatigue crack growth evolution under increasing stress cycles (Paris and Erdogan, 1963; Ditlevsen and Madsen, 1996). The crack growth follows the following first-order differential equation (7), known as Paris–Erdogan law,

$$\frac{da(n, \theta)}{dn} = \exp(C_{\text{ln}}) \left[\Delta S \sqrt{\pi a(n)} \right]^m, \tag{7}$$

where $\theta = [a, \Delta S, C_{\text{ln}}, m]$ is a vector containing the uncertain time-invariant model parameters. Specifically, a is the crack length, n is the number of stress cycles, ΔS is the stress range per cycle when assuming constant stress amplitudes, C and m represent empirically determined model parameters; C_{ln} corresponds to the natural logarithm of C .

The solution to this differential equation, with boundary condition $a(n=0) = a_0$, can be written as a function of the number of stress cycles n and the vector $\theta = [a_0, \Delta S, C_{\text{ln}}, m]$ as (for the derivation see, e.g., Ditlevsen and Madsen (1996)):

$$a(n, \theta) = \left[\left(1 - \frac{m}{2}\right) \exp(C_{\text{ln}}) \Delta S^m \pi^{m/2} n + a_0^{(1-m/2)} \right]^{(1-m/2)^{-1}}. \tag{8}$$

We assume that noisy measurements of the crack y_n are obtained sequentially at different values of n . The measurement equation (9) assumes a multiplicative lognormal measurement error, $\exp(\varepsilon_n)$.

$$y_n = a(n, \theta) \exp(\varepsilon_n). \tag{9}$$

In numerical investigations that follow, the measurement equation (9) is used for generating synthetic measurements of the deterioration state. In this context, a multiplicative lognormal measurement error ensures that nonnegative generated measurements of the deterioration state are not feasible.

Under this assumption in equation (9), the likelihood function for a measurement at a given n is shown in equation (10).

$$L(y_n; a(n, \theta)) = \frac{1}{\sigma_{\varepsilon_n} \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{\ln(y_n) - \mu_{\varepsilon_n} - \ln(a(n, \theta))}{\sigma_{\varepsilon_n}} \right)^2 \right]. \tag{10}$$

Table 1 shows the prior probability distribution model for each random variable in the vector θ (Ditlevsen and Madsen, 1996; Straub, 2009), as well as the assumed probabilistic model of the measurement error. In this case study we are dealing with a nonlinear model and a parameter vector with non-Gaussian prior distribution.

Table 1. Prior distribution model for the fatigue crack growth model parameters and the measurement error.

Parameter	Distribution	Mean	Standard Deviation	Correlation
a_0	Exponential	1	1	—
ΔS	Normal	60	10	—
C_{ln}, m	Binormal	(-33;3.5)	(0.47;0.3)	$\rho_{C_{ln},m} = -0.9$
$\exp(\varepsilon_n)$	Lognormal	1.0	0.1508	

3.1.1. Markovian state-space representation for application of on-line filters

A Markovian state-space representation of the deterioration process is required for application of on-line filters. The number of stress cycles is discretized as $n = k\Delta n$, with $k = 1, \dots, 100$ denoting the estimation time step and $\Delta n = 1 \times 10^5$ the number of stress cycles per time step. The dynamic and measurement equations of the discrete-time state-space representation of the fatigue crack growth model with unknown time-invariant parameters $\theta = [a_0, \Delta S, C_{ln}, m]$ are shown below:

$$\theta_k = \theta_{k-1}$$

$$y_k = a(k\Delta n, \theta_k) \exp(\varepsilon_k) = \left[\left(1 - \frac{m_k}{2} \right) \exp(C_{ln,k}) \Delta S_k^{m_k} \pi^{m_k/2} k \Delta n + a_0^{(1-m_k/2)} \right]^{(1-m_k/2)^{-1}} \exp(\varepsilon_k). \tag{11}$$

The state-space model of equation (11) is nonlinear and the prior is non-Gaussian. For reasons explained in Section 2, the subscript k in θ_k is dropped in the remainder of this section.

3.1.2. Reference posterior solution

For the purpose of performing a comparative assessment of the different filters, an underlying “true” realization of the fatigue crack growth process $a^*(n, \theta)$ is generated for $n = k\Delta n$, with $k = 1, \dots, 100$ and $\Delta n = 1 \times 10^5$. This realization corresponds to the randomly generated “true” vector of time-invariant parameters $\theta^* = [a_0^* = 2.0, \Delta S^* = 50.0, C_{ln}^* = -33.5, m^* = 3.7]$. Sequential synthetic crack monitoring measurements y_k are sampled from the measurement equation (9) for $a(k\Delta n, \theta^*)$, and for randomly generated measurement noise samples $\exp(\varepsilon_k)$. These measurements are scattered in green in Figure 3.

Based on the generated measurements, the sequence of reference posterior distributions $\pi_{pos}(\theta | y_{1:k})$ is obtained using the prior distribution as an envelope distribution for rejection sampling (Smith and Gelfand, 1992; Rubinstein and Kroese, 2016). More specifically, for each of the 100 posterior distributions of interest $\pi_{pos}(\theta | y_{1:k})$, 10^5 independent samples are generated. The results of this reference posterior estimation of the four time-invariant model parameters are plotted in Figure 2.




Figure 2. Reference posterior solution: mean and credible intervals for the sequence of posterior distributions $\pi_{pos}(\theta | y_{1:k})$.

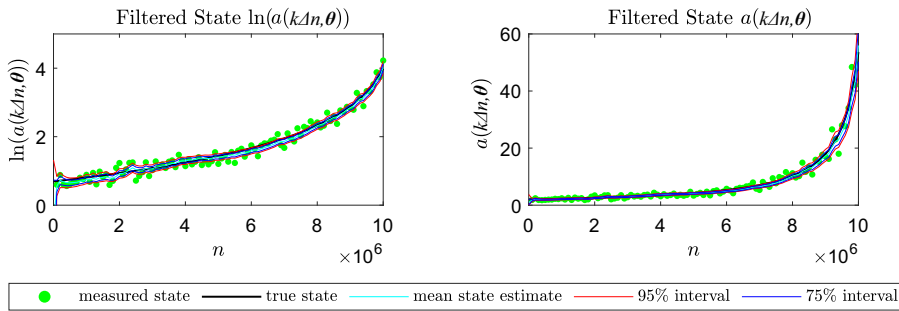


Figure 3. Reference mean and credible intervals for the filtered crack growth state $a(k\Delta n, \theta)$.




Figure 4. Reference final posterior: prior and single posterior distribution of interest $\pi_{\text{pos}}(\theta | \mathbf{y}_{1:100})$.

With posterior samples, the reference filtered estimate of the crack length $a(k\Delta n, \theta)$ at each estimation step is also obtained via the model of equation (8) and plotted in Figure 3. In the left panel of this figure, the filtered state is plotted in logarithmic scale. In an off-line estimation, a single posterior density is of interest. One such reference posterior estimation result for the last estimation step, $\pi_{\text{pos}}(\theta | \mathbf{y}_{1:100})$, is plotted for illustration in Figure 4.

3.1.3. Comparative assessment of the investigated on-line and off-line filters

We apply the PFGM filter with 5,000 and 50,000 particles, the IBIS filter with 5,000 particles, and the SMC filter with 5,000 particles for performing on-line and off-line time-invariant parameter estimation tasks. We evaluate the performance of each filter by taking the relative error of the estimated mean and standard deviation of each of the four parameters with respect to the reference posterior solution. For example, the relative error in the estimation of the mean of parameter a_0 at a certain estimation step k is computed as $|\frac{\mu_{a_0,k} - \hat{\mu}_{a_0,k}}{\mu_{a_0,k}}|$, where $\mu_{a_0,k}$ is the reference posterior mean from rejection sampling (Section 3.1.2), and $\hat{\mu}_{a_0,k}$ is the posterior mean estimated with each filter. Each filter is run 50 times, and the mean relative error of the mean and the standard deviation of each parameter, together with the 90% credible intervals (CI), are obtained. These are plotted in Figure 5.

Figure 6 plots the L^2 relative error norm of the mean and the standard deviation of all four parameters, that is, the quantity of equation (12) (here formulated for the mean at estimation step k).

$$\sqrt{\frac{\sum_{i=1}^d (\mu_{i,k} - \hat{\mu}_{i,k})^2}{\sum_{i=1}^d (\mu_{i,k})^2}}, \tag{12}$$

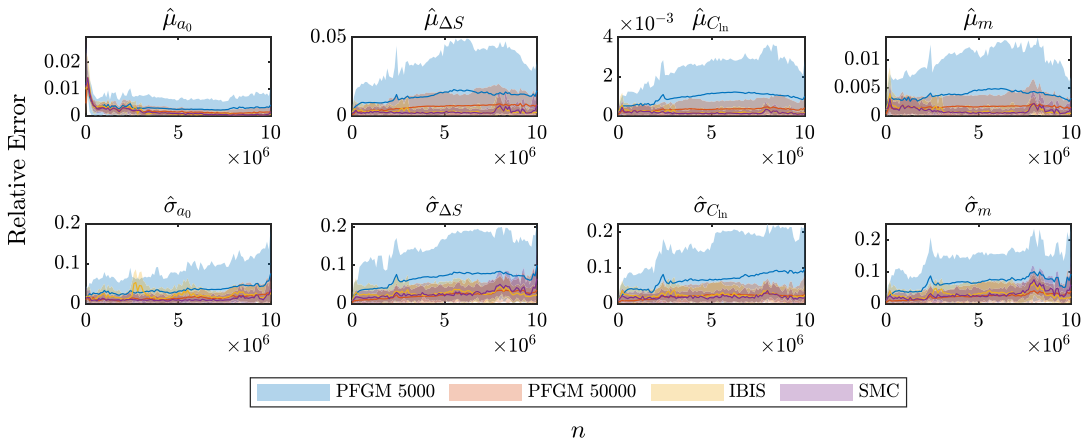


Figure 5. Comparison of the relative error of the mean and standard deviation of the parameters evaluated for each filter. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 50 repeated runs of each filter. In the horizontal axis, n is the number of stress cycles.

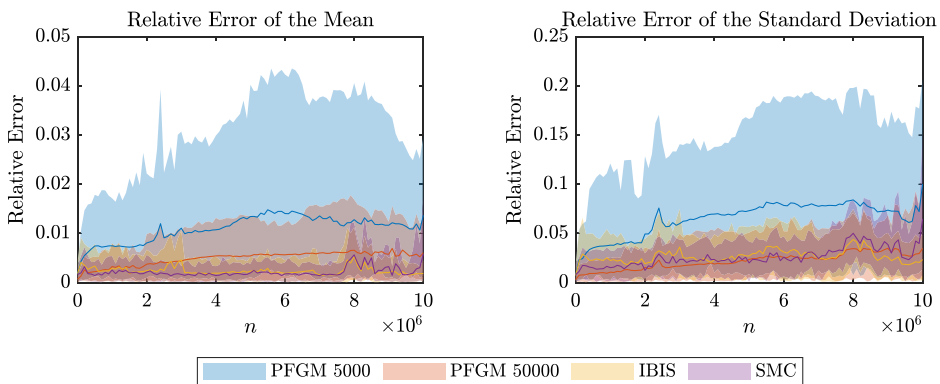


Figure 6. Comparison of the L^2 relative error norm of the mean and the standard deviation of the parameters evaluated for each filter. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 50 repeated runs of each filter. In the horizontal axis, n is the number of stress cycles.

where d is the dimensionality of the time-invariant parameter vector θ (in this example $d=4$). More specifically, Figure 6 plots the mean and credible intervals of the L^2 relative error norm of the estimated mean and standard deviation, as obtained from 50 runs of each filter.

Figures 5 and 6 reveal that, when all three filters are run with the same number of particles, the IBIS and SMC filters yield superior performance over PFGM. When the number of particles in the PFGM filter is increased to 50,000, the PFGM filter performance is comparable to the one of the IBIS and SMC filters. In estimating the mean, the mean L^2 relative error norm obtained from the PFGM filter with 50,000 particles is slightly larger than the corresponding error obtained from IBIS and SMC with 5,000 particles, while the 90% credible intervals of the PFGM filter estimation are still wider. In estimating the standard deviation, the PFGM filter with 50,000 particles proves competitive.

Figures 5 and 6 show the estimation accuracy of each filter when used for on-line inference, that is, for estimating the whole sequence of 100 posterior distributions $\pi_{\text{pos}}(\theta|\mathbf{y}_{1:k})$, $k=1, \dots, 100$. The PFGM and IBIS filters, being intrinsically on-line filters, provide the whole posterior sequence with one run. On the

Table 2. Average number of model evaluations (equation (8)) for the fatigue crack growth model parameter estimation.

Method	PFGM 5,000	PFGM 50,000	IBIS	SMC (final posterior)	SMC (all posteriors)
Model evaluations	5×10^5	5×10^6	3.4×10^6	4.5×10^6	1.9×10^8

other hand, the off-line SMC filter is run anew for each of the 100 required posterior estimations. Hence, Figures 5 and 6 enclose the results of both the on-line and the off-line inference. If one is interested in the off-line estimation accuracy at a specific stress cycle n , one can simply consider a vertical “cut” at n .

Table 2 documents the computational cost associated with each filter, expressed in the form of required model evaluations induced by calls of the likelihood function. When using the term model, we refer to the model defined in equation (8), which corresponds to an analytical expression with negligible associated runtime. However, unlike the simple measurement equation that we have assumed in this example, in many realistic deterioration monitoring settings, the deterioration state cannot be measured directly (e.g., in vibration-based structural health monitoring Kamarotis et al., 2022). In such cases, each deterioration model evaluation often entails evaluation of a finite element (FE) model, which has substantial runtime. It, therefore, appears appropriate to evaluate the filters’ computational cost in terms of required model evaluations. The on-line PFGM filter with 5,000 particles requires 5×10^5 model evaluations, and yields by far the smallest computational cost, while at the same time providing the solution to both on-line and off-line estimation tasks. However, it also yields the worst performance in terms of accuracy of the posterior estimates. Running the IBIS filter with 5,000 particles, which performs MCMC move steps, leads to 3.4×10^6 model evaluations. Comparing this value against the 5×10^5 model evaluations required by the PFGM filter with 5,000 particles for performing the same task distinctly shows the computational burden associated with MCMC move steps, which require a complete browsing of the whole measurement data set in estimating the acceptance probability. However, the IBIS filter also leads to enhanced estimation accuracy, which might prove significant when the subsequent tasks entail prognosis of the deterioration evolution, the structural reliability or the remaining useful lifetime, and eventually the predictive maintenance planning. Using 50,000 particles, the PFGM filter performance increases significantly with a computational cost that is comparable to the IBIS filter with 5,000 particles. For the off-line SMC algorithm, 4.5×10^6 model evaluations are required only for the task of estimating the final posterior density. The 1.9×10^8 model evaluations required by the SMC for obtaining the whole sequence of posteriors $\pi_{\text{pos}}(\boldsymbol{\theta} | \mathbf{y}_{1:k})$, $k = 1, \dots, 100$, clearly demonstrate that off-line MCMC techniques are unsuited to on-line estimation tasks.

3.2. High-dimensional case study: Corrosion deterioration spatially distributed across beam

As a second case study, we employ the deterioration model of equation (13), which describes the spatially and temporally varying corrosion deterioration across the structural beam shown in Figure 7.

$$D(t, x) = A(x)t^{B(x)}, \quad t = 0, \dots, 50. \quad (13)$$

$A(x), B(x), x \in \Omega$ are random fields defined on $\Omega = [0, L]$, with L denoting the length of the beam taken as $L = 4\text{m}$. $A(x)$ models the deterioration rate, while $B(x)$ models the nonlinearity effect of the deterioration process in terms of a power law in time. The corrosion deterioration $D(t, x)$ is therefore also a spatial random field.

A random field, by definition, contains an infinite number of random variables, and must therefore be discretized (Vanmarcke, 2010). One of the most common methods for discretization of random fields is the midpoint method (Der Kiureghian and Ke, 1988). Thereby the spatial domain Ω is discretized into m elements, and the random field is approximated within each element through the random variable that corresponds to midpoint of the element. In that case, the uncertain time-invariant deterioration model




Figure 7. Structural beam subjected to spatially and temporally varying corrosion deterioration. The deterioration process is monitored from sensors deployed at specific sensor locations (in green).

parameter vector is $\theta = [A_1, \dots, A_m, B_1, \dots, B_m]$, where $A_i := A(x_i), B_i := B(x_i), i = 1, \dots, m$, are the random variables corresponding to the element midpoints. With the midpoint discretization, the spatial deterioration $D(t, x)$ is parametrized by θ .

We assume that noisy measurements of the corrosion deterioration state $D_{t,l}(\theta) := D(t, x_l, \theta)$ at time t and at certain locations l of the beam are obtained sequentially (summarized in one measurement per year) from n_l sensors deployed at these locations ($n_l = 10$ sensor locations are shown in Figure 7). The measurement equation (14), describing the corrosion measurement at time t and sensor location l , assumes a multiplicative measurement error, $\exp(\varepsilon_{t,l})$.

$$y_{t,l} = D_{t,l}(\theta) \exp(\varepsilon_{t,l}) = A_{i_l} t^{B_{i_l}} \exp(\varepsilon_{t,l}), \tag{14}$$

where i_l returns the discrete element number of the midpoint discretization within which the measurement location l lies. Table 3 shows the prior distribution model for the two random fields of the deterioration model of equation (13) and the assumed probabilistic model of the multiplicative measurement error. Since $A(x)$ models a lognormal random field, $\ln(A(x))$ follows the normal distribution. For both random fields $\ln(A(x))$ and $B(x)$, the exponential correlation model with correlation length of 2 m is applied (Sudret and Der Kiureghian, 2000).

The goal is to update the time-invariant deterioration model parameters $\theta = [A_1, \dots, A_m, B_1, \dots, B_m]$ given sequential noisy corrosion measurements $y_{t,l}$ from n_l deployed sensors. The dimensionality of the problem is $d = 2 \times m$. Hence, the more elements in the midpoint discretization, the higher the dimensionality of the parameter vector.

The main goal of this second case study is to investigate the effect of the problem dimensionality and the amount of sensor information on the posterior results obtained with each filter. We choose the following three midpoint discretization schemes:

1. $m = 25$ elements: $d = 50$ time-invariant parameters to estimate.
2. $m = 50$ elements: $d = 100$ time-invariant parameters to estimate.
3. $m = 100$ elements: $d = 200$ time-invariant parameters to estimate.

Furthermore, we choose the following three potential sensor arrangements:

1. $n_l = 2$ sensors (the 4th and 7th sensors of Figure 7).
2. $n_l = 4$ sensors (the 1st, 4th, 7th, and 10th sensors of Figure 7).
3. $n_l = 10$ sensors of Figure 7.

We, therefore, study nine different cases of varying problem dimensionality and the number of sensors.

Table 3. Prior distribution model for the corrosion deterioration model parameters and the measurement error.

Parameter	Distribution	Mean	Standard Deviation	Corr. length (m)
$A(x)$	Lognormal	0.8	0.24	2
$B(x)$	Normal	0.8	0.12	2
$\exp(\varepsilon_{t,l})$	Lognormal	1.0	0.101	–

3.2.1. Markovian state-space representation for application of on-line filters

A Markovian state-space representation of the deterioration process is required for application of on-line filters. The dynamic and measurement equations are shown in equation (15). The measurement equation is written in the logarithmic scale. Time t is discretized in yearly estimation time steps k , that is, $k = 1, \dots, 50$, and the subscript $l = 1, \dots, n_l$ corresponds to the sensor location.

$$\begin{aligned} \theta_k &= \theta_{k-1} \\ \ln(y_{k,l}) &= \ln(D_{k,l}(\theta_k)) + \varepsilon_{k,l} \Rightarrow \ln(y_{k,l}) = \ln(A_{k,i_l}) + B_{k,i_l} \ln(k) + \varepsilon_{k,l}, \end{aligned} \tag{15}$$

where θ_k denotes the time-invariant parameter vector at time step k . In the logarithmic scale, both the dynamic and measurement equations are linear functions of Gaussian random variables. For reasons explained in Section 2, the subscript k in θ_k is dropped in the following.

3.2.2. Underlying “true” realization

To generate a high-resolution underlying “true” realization of the two random fields $A(x)$ and $B(x)$, and the corresponding synthetic monitoring data set, we employ the Karhunen–Loève (KL) expansion using the first 400 KL modes. The KL expansion is an alternative random field discretization scheme to the midpoint method that represents the random field in terms of the eigenfunctions of its autocovariance function (Ghanem and Spanos, 1991; Sudret and Der Kiureghian, 2000). The implementation of the KL expansion can be found in the Matlab codes accompanying this article.² We remark that a fine resolution KL expansion is chosen to represent the “true” realization in order to avoid the inverse crime (Wirgin, 2004). These realizations are shown in the left panel of Figure 8. Given these $A(x)$ and $B(x)$ realizations, the underlying “true” realizations of the deterioration process at 10 specific beam locations are generated, which correspond to the 10 potential sensor placement locations shown in Figure 7. Subsequently, a synthetic corrosion sensor measurement data set (one measurement per year) at these 10 locations is generated from the measurement equation (14). These are shown in the right panel of Figure 8.

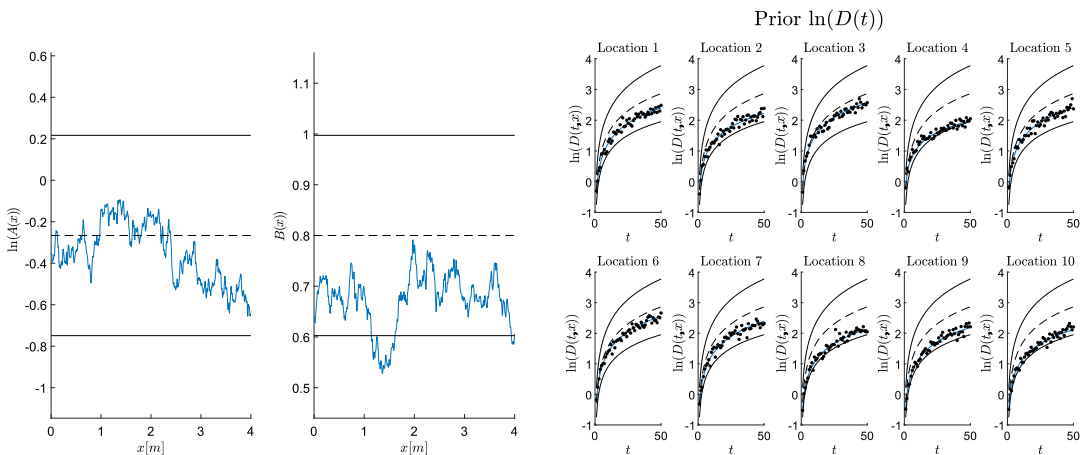


Figure 8. Left: the blue solid line plots the underlying “true” realization of $\ln(A(x))$ and $B(x)$ created using the KL expansion. Right: the blue solid line plots the underlying “true” realization of $\ln(D(t,x))$ at 10 specific sensor locations and the corresponding synthetic sensor monitoring data are scattered in black. In both figures, the black dashed lines plot the prior mean and the black solid lines the prior 90% credible intervals.

² https://github.com/antoniskam/Offline_online_Bayes.

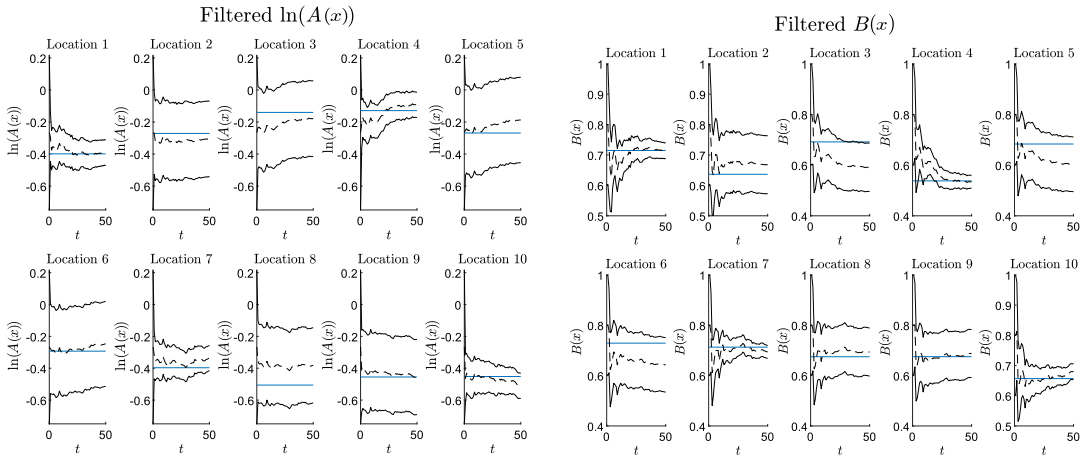


Figure 9. Case with $m=25, n_l=4$: reference on-line posterior solution at 10 locations across the beam obtained by applying the Kalman filter for solving equation (15). The solid blue horizontal line represents the underlying “true” values of $\ln(A(x))$ and $B(x)$ at these locations. The black dashed lines plot the posterior mean and the black solid lines the posterior 90% credible intervals. Locations 1, 4, 7, and 10 correspond to the four assumed sensor placement locations.

3.2.3. Reference posterior solution

For the investigated linear Gaussian state space representation of equation (15), we create reference on-line posterior solutions for each of the nine considered cases by applying the Kalman filter (KF) (Kalman, 1960), which is the closed-form solution to the Bayesian filtering equations. The process noise covariance matrix in the KF equations is set equal to zero. The linear Gaussian nature of the chosen problem ensures existence of an analytical reference posterior solution obtained with the KF. One such reference on-line posterior solution for the case described by $m=25$ elements ($d=50$) and $n_l=4$ sensors is shown in Figure 9.

3.2.4. Comparative assessment of the investigated on-line and off-line filters

The goal of this section is to offer a comparative assessment among the three Bayesian filtering algorithms presented in this article when applied on a high-dimensional problem. To be able to derive a reference solution, as described above, a linear Gaussian state space representation of a structural deterioration problem has been defined (equation (15)). We apply the tPFGM filter, the tIBIS filter, and the SMC filter, all with $N_{par}=2000$ particles, for estimating the time-invariant parameter vector θ . For each of the nine cases of varying problem dimensionality and number of sensors described above, we compute the L^2 relative error norm of the estimated means, correlation coefficients, and standard deviations of the parameters with respect to the corresponding KF reference posterior solution, that is, we estimate a quantity as in equation (12) for all estimation steps $k=1, \dots, 50$. In Figures 10–12, we plot the mean and credible intervals of these relative errors as obtained from 50 different runs. The off-line SMC filter, which does not provide the on-line solution within a single run, is run anew for estimating the single posterior density of interest at years 10, 20, 30, 40, and 50, and in between, the relative error is linearly interpolated. Although each of the nine panels in the figures corresponds to a different case with a different underlying KF reference solution, their y-axes have the same scaling. Table 4 documents the computational cost of each filter in each considered case, measured by average number of evaluations of the model of equation (13).

Figures 10 and 11 show that the off-line IMH-GM-based SMC filter yields the best performance in estimating the KF reference posterior mean and correlation, for all nine considered cases, while at the same time producing the narrowest credible intervals. Comparison of the relative errors obtained with the SMC and tIBIS filters reveals that, although they are both reliant on the IMH-GM MCMC move step, the

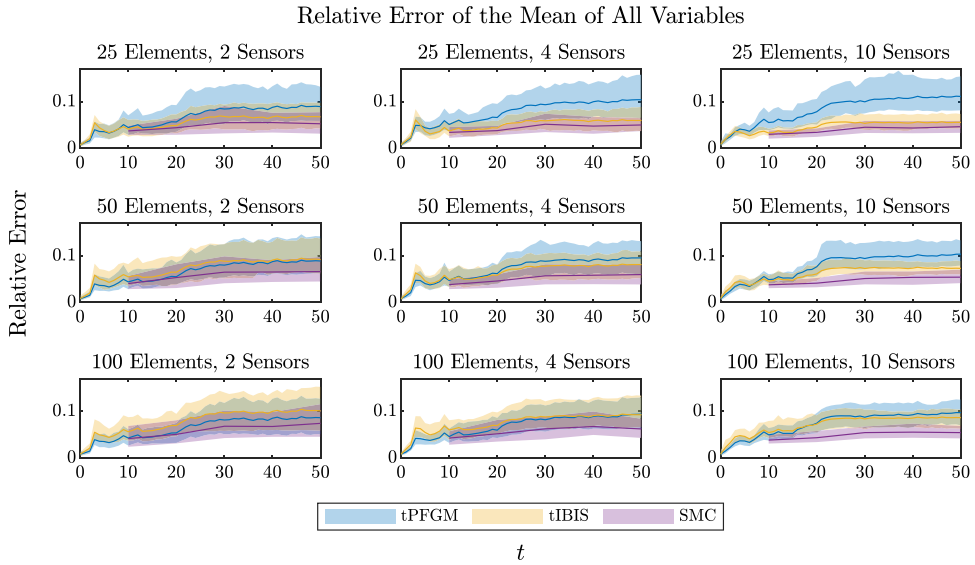


Figure 10. Comparison of the L^2 relative error norm of the means of the parameters evaluated for each filter. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 50 repeated runs of each filter.

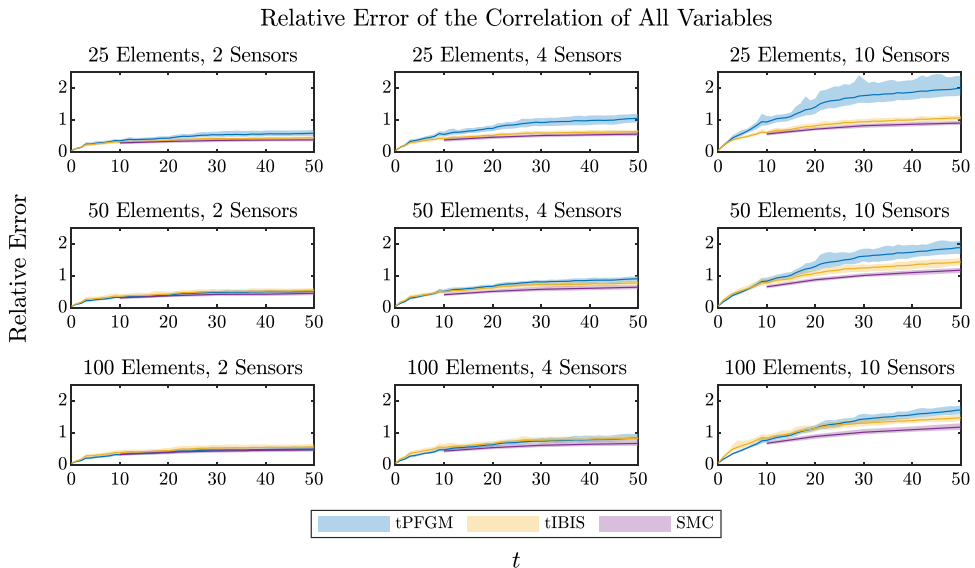


Figure 11. Comparison of the L^2 relative error norm of the correlation coefficients of the parameters evaluated for each filter. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 50 repeated runs of each filter.

on-line tIBIS filter leads to larger estimation errors. The on-line tPFGM and tIBIS filters generate quite similar results in estimating the reference posterior mean and correlation, thus rendering the benefit of the MCMC move step in tIBIS unclear, except in cases with more sensors and lower parameter dimension. Figures 10 and 11 reveal a slight trend, indicating that for fixed dimensionality, the availability of more sensors (i.e., stronger information content in the likelihood function) leads to a slight decrease in the

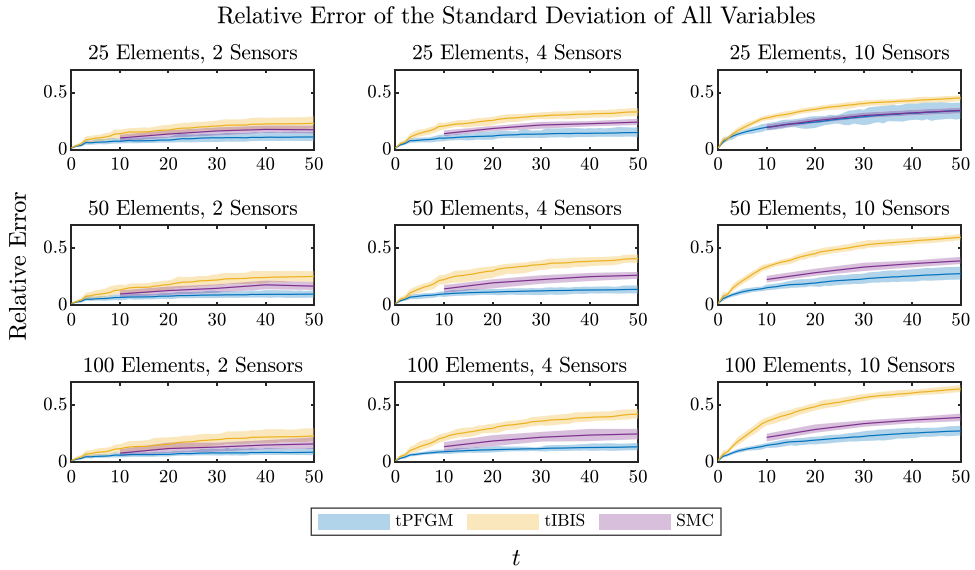


Figure 12. Comparison of the L^2 relative error norm of the standard deviations of the parameters evaluated for each filter. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 50 repeated runs of each filter.

relative errors when using the SMC and tIBIS filters, whereas the opposite trend can be identified for the tPFGM filter. Increasing problem dimensionality (for fixed number of sensors) does not appear to have strong influence on the posterior results in any of the columns of Figures 10–12, a result that initially appears puzzling.

Figure 12 conveys that the tPFGM filter, which entirely depends on the GMM posterior approximation, induces the smallest relative errors for the estimation of the standard deviation of the parameters in all considered cases. This result reveals a potential inadequacy of the single application of the IHM-GM kernel for the move step within the tIBIS and SMC filters in properly exploring the space of θ . In all 50 runs of the tIBIS and SMC filters, the standard deviation of the parameters is consistently underestimated compared to the reference, unlike when applying the tPFGM filter.

Based on the discussion of Section 2.2, we introduce a burn-in period of $n_B = 5$ in the IMH-GM kernel of Algorithm 4 and perform 50 new runs of the tIBIS and SMC filters. One can expect that inclusion of a burn-in is more likely to ensure sufficient exploration of the intermediate posterior distributions. However, at the same time, the computational cost of tIBIS and SMC increases significantly, with a much larger number of required model evaluations than in Table 4. In Figures 13 and 14, we plot the mean and credible intervals for the relative errors in the estimation of the mean and standard deviation of the parameters. Comparing Figures 10 and 13, inclusion of burn-in is shown to lead to an improved performance of tIBIS and SMC in estimating the mean of the parameters in all cases. This improvement is more evident in the lower-dimensional case with 25 elements, and lessens as the problem dimension increases. Hence, it is only after the inclusion of burn-in, which leads to an enhanced posterior solution, that one starts observing the anticipated deterioration of the tIBIS and SMC filters’ performance with increasing dimensionality. This effect was masked in the results of Figure 10 without burn-in. This point becomes more evident when looking at the relative errors of the estimated standard deviation in Figure 14. With burn-in, the tIBIS and SMC filters provide better results than the tPFGM filter in estimating the standard deviation in the case of 25 elements, but perform progressively worse as the dimensionality increases, where they underestimate the KF reference standard deviation. This underestimation is clearly illustrated in Figure 15. The reason for this behavior is the poor performance of the IMH-GM algorithm in high dimensions, which is numerically demonstrated in Papaioannou et al. (2016). We suspect that this

Table 4. Average number of model evaluations (equation (13)) for the high-dimensional case study.

Elements	25			50			100		
Sensors	2	4	10	2	4	10	2	4	10
tPFGM	129,480	154,000	194,440	129,440	155,560	195,760	130,480	157,120	199,040
tIBIS	602,400	1,038,440	1,878,000	603,240	1,049,400	1,909,880	567,720	1,017,280	1,876,240
SMC	1,130,000	1,596,000	2,298,000	1,108,000	1,582,000	2,250,000	1,100,000	1,504,000	2,150,000

Note. For the SMC, the required model evaluations for obtaining the single final posterior density are reported.

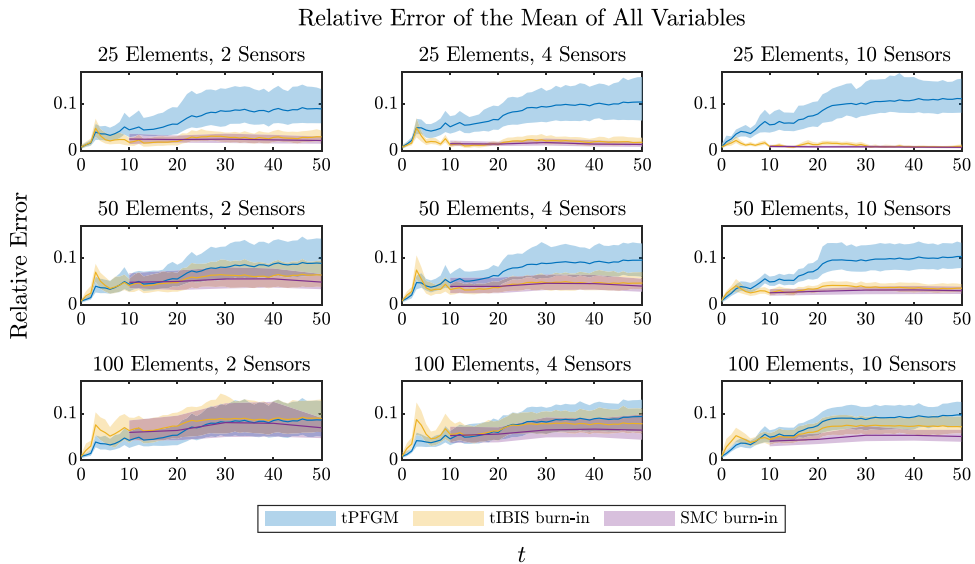


Figure 13. Comparison of the L^2 relative error norm of the mean of the parameters evaluated for each filter. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 50 repeated runs of each filter. Burn-in $n_B = 5$.

behavior is related to the degeneracy of the acceptance probability of MH samplers in high dimensions, which has been extensively discussed in the literature for random walk samplers, for example, in Gelman et al. (1997), Au and Beck (2001), Katafygiotis and Zuev (2008), Beskos and Stuart (2009), Cotter et al. (2013), and Papaioannou et al. (2015). Single application of the IHM-GM kernel without burn-in yielded acceptance rates of around 50% for all cases. With inclusion of burn-in, in higher dimensions, the acceptance rate in IMH-GM drops significantly in the later burn-in steps, leading to rejection of most proposed particles. To alleviate this issue, one could consider using the preconditioned Crank Nicolson (pCN) sampler to perform the move step within the IBIS and SMC filters, whose performance is shown to be independent of the dimension of the parameter space when the prior is Gaussian (Cotter et al., 2013).

Increase of dimensionality does not seem to have any influence on the results obtained with the tPFGM filter. The illustrated efficacy of the tPFGM filter in estimating the time-invariant parameters in all considered cases of increasing dimensionality is related to the nature of the studied problem. The tPFGM filter relies entirely on the GMM approximation of the posterior distribution within its resampling process, in that it simply “accepts” all the N_{par} GMM-proposed particles, unlike the tIBIS and SMC filters, which contain the degenerating acceptance-rejection step within the IMH-GM move step. Clearly, the worse the GMM fit, the worse the expected performance of the tPFGM filter. The particular case investigated here has a Gaussian reference posterior solution, hence the GMM fitted by EM proves effective in approximating the posterior with a relatively small number of particles, even when going up to $d = 200$




Figure 14. Comparison of the L^2 relative error norm of the standard deviation of the parameters evaluated for each filter. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 50 repeated runs of each filter. Burn-in $n_B = 5$.

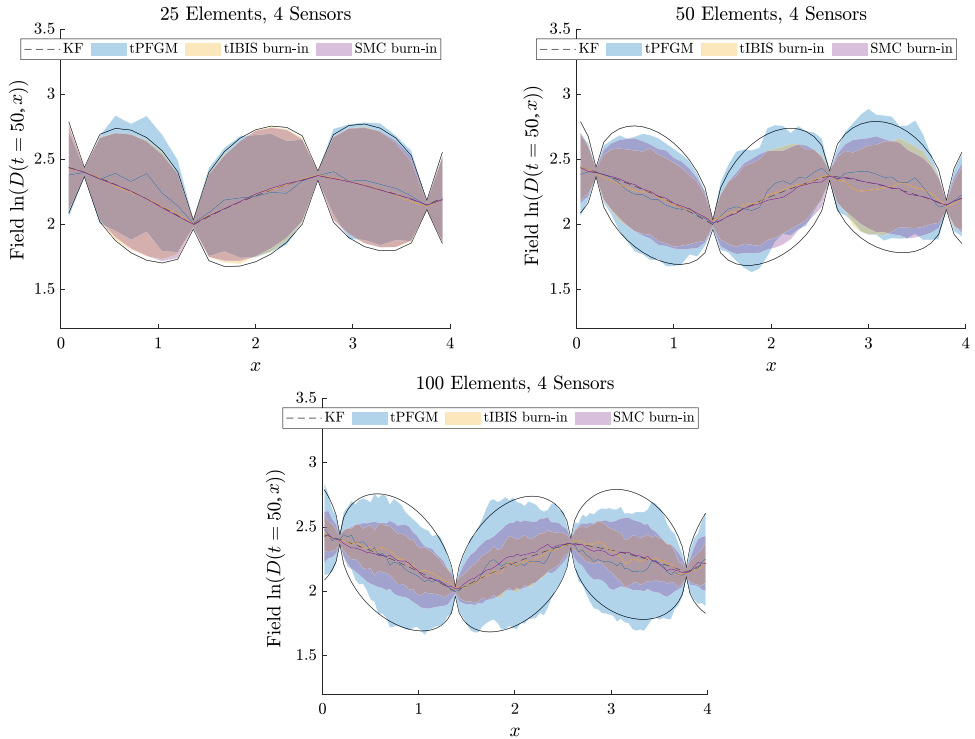


Figure 15. Updating of the random field $\ln(D(t=50, x))$ in three different cases of varying problem dimensionality. The solid lines show the mean and the shaded areas the 90% credible intervals inferred from 10 repeated runs of each filter. The black dashed line represented the posterior mean obtained via the KF, and the black solid lines the KF 90% credible intervals.

dimensions, thus leading to a good proposal distribution for sampling N_{par} new particles in tPFGM. As reported in Table 4, the tPFGM filter is associated with a significantly lower computational cost than its MCMC-based counterparts.

4. Concluding Remarks

In this article, we present in algorithmic detail three different on-line and off-line Bayesian filters. The on-line filters are specifically tailored for quantifying the full posterior uncertainty of time-invariant deterioration model parameters in long-term monitoring settings. Specifically, the three presented methods are an on-line particle filter with Gaussian mixture resampling (PFGM), an on-line iterated batch importance sampling (IBIS) filter, and an off-line sequential Monte Carlo (SMC) filter, which applies simulated annealing to sequentially arrive to a single posterior density of interest. The IBIS and SMC filters perform Markov Chain Monte Carlo (MCMC) move steps via application of an independent Metropolis Hastings kernel with a Gaussian mixture proposal distribution (IMH-GM) whenever degeneracy is identified. A simulated annealing process (tempering of the likelihood function) is further incorporated within the update step of the on-line PFGM and IBIS filters for cases when each new measurement is expected to have a strong information content; this leads to the presented tPFGM and tIBIS filters. The SMC filter can be employed only for off-line inference, while the PFGM, tPFGM, IBIS, and tIBIS filters can perform both on-line and off-line inference tasks.

With the aid of two numerical examples, a comparative assessment of these algorithms for off-line and on-line Bayesian filtering of time-invariant deterioration model parameters is performed. In contrast to other works, the main focus here lies on the efficacy of the investigated Bayesian filters in quantifying the full posterior uncertainty of deterioration model parameters, as well as on the induced computational cost.

For the first nonlinear, non-Gaussian, and low-dimensional case study, the IBIS and SMC filters, which both contain IMH-GM-based MCMC move steps, are shown to consistently outperform the purely on-line PFGM filter in estimating the parameters' reference posterior distributions. However, they induce a computational cost of at least an order of magnitude larger than the PFGM filter, when the same initial number of particles is used in all three filters. With similar computational cost, that is, when increasing the number of particles in PFGM, it achieves enhanced posterior accuracy, comparable to the IBIS and SMC filters.

The second case study involves a linear, Gaussian, and high-dimensional model. The focus here lies on evaluating the performance of the investigated filters in increasing dimensions. The linear Gaussian nature of the problem allows access to an exact reference posterior solution with the Kalman filter. For this case study, the results vary with increasing problem dimensionality and number of sensors. The on-line tPFGM filter achieves a consistently satisfactory quality with increasing dimensionality, a behavior explained by the linear Gaussian nature of the problem, while a slight drop in the posterior quality is observed for increasing amount of sensor information. The tIBIS and SMC filters are shown to consistently outperform the tPFGM filter in lower dimensions, they however perform progressively worse in higher dimensions, a behavior likely explained by the degeneracy of the acceptance probability of MH samplers in high dimensions. The computational cost of the tIBIS and SMC filters is an order of magnitude larger than the tPFGM filter.

Some general conclusions drawn from the delivered comparative assessment are listed below.

1. The purely on-line PFGM (and its tPFGM variant) filter is competitive with MCMC-based filters, especially for higher-dimensional well-behaved problems.
2. The IBIS (and its tIBIS variant) and SMC filters, which contain MCMC move steps, offer better approximations of the posterior mean of the model parameters than the purely on-line PFGM (and its tPFGM variant) filter with the same number of samples, as shown in both studied examples.
3. The independent Metropolis Hastings (IMH)-based MCMC move step performed within the IBIS, tIBIS, and SMC filters proves inadequate in properly exploring the posterior parameter space in high-dimensional problems.

Table 5. Set of suggestions on choice of the appropriate method in function of problem characteristics.

Criterion		PFGM (tPFGM)	IMH-GM-based IBIS (tIBIS)	IMH-GM-based SMC
On-line inference		✓	◦	×
Computational cost		C_1	C_2	C_3
Applicability to different problems	Mean estimation in low-dimensional, nonlinear, non-Gaussian problems	Q_3	Q_4	Q_4
	Uncertainty quantification in low-dimensional, nonlinear, non-Gaussian problems	Q_3	Q_4	Q_4
	Mean estimation in high-dimensional, well-behaved problems	Q_2	Q_3	Q_4
	Uncertainty quantification in high-dimensional, well-behaved problems	Q_3	Q_1	Q_1
Increasing sensor information amount		× (✓)	× (✓)	✓
Q_1 : low quality, Q_2 : moderate quality, Q_3 : moderate to high quality, Q_4 : high quality		✓: applicable, ◦: partly applicable, ×: not applicable	C_1 : moderate, C_2 : moderate to high, C_3 : high	

Finally, to support the reader with the selection of the appropriate algorithm for a designated scenario, we provide Table 5, which contains an assessment of the methods presented in this article in function of problem characteristics.

This article does not investigate the performance of these filters when applied to high-dimensional and highly non-Gaussian problems. Such problems are bottlenecks for most existing filters and we expect the investigated filters to be confronted with difficulties in approximating the posterior distributions.

Author contribution. Conceptualization: A.K., L.S., I.P., E.C., D.S.; Formal analysis: A.K., L.S.; Funding acquisition: E.C., D.S.; Methodology: A.K., L.S., I.P., E.C., D.S.; Software: A.K., L.S.; Visualization: A.K., L.S.; Writing—original draft: A.K., Writing—review and editing: A.K., L.S., I.P., E.C., D.S. All authors approved the final submitted draft.

Competing interest. The authors declare no competing interests exist.

Data availability statement. The data and code that support the findings of this study are openly available at https://github.com/antoniskam/Offline_online_Bayes.

Funding statement. The work of A.K. and E.C. has been carried out with the support of the Technical University of Munich—Institute for Advanced Study, Germany, funded by the German Excellence Initiative and the TÜV SÜD Foundation.

References

Andrieu C, Doucet A and Holenstein R (2010) Particle Markov Chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(3), 269–342.

Andrieu C, Doucet A, Singh S and Tadić V (2004) Particle methods for change detection, system identification, and control. *Proceedings of the IEEE* 92 (3), 423–438.

Au S-K and Beck J.L. (2001) Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics* 16(4), 263–277.

- Azam SE, Chatzi E, Papadimitriou C and Smyth A** (2017) Experimental validation of the Kalman-type filters for online and real-time state and input estimation. *Journal of Vibration and Control* 23(15), 2494–2519.
- Baraldi P, Cadini F, Mangili F and Zio E** (2013) Model-based and data-driven prognostics under different available information. *Probabilistic Engineering Mechanics* 32, 66–79.
- Beck JL and Au S-K** (2002) Bayesian updating of structural models and reliability using Markov Chain Monte Carlo simulation. *Journal of Engineering Mechanics* 128(4), 380–391.
- Beskos A and Stuart A** (2009) Computational complexity of metropolis-hastings methods in high dimensions. In L' Ecuyer P and Owen AB (eds), *Monte Carlo and Quasi-Monte Carlo Methods 2008*, Berlin: Springer, pp. 61–71.
- Cadini F, Zio E and Avram D** (2009a) Model-based Monte Carlo state estimation for condition-based component replacement. *Reliability Engineering & System Safety* 94(3), 752–758.
- Cadini F, Zio E and Avram D** (2009b) Monte Carlo-based filtering for fatigue crack growth estimation. *Probabilistic Engineering Mechanics* 24(3), 367–373.
- Carvalho CM, Johannes MS, Lopes HF and Polson NG** (2010) Particle learning and smoothing. *Statistical Science* 25(1), 88–106.
- Celeux G, Fruewirth-Schnatter S and Robert CP** (2019) Model selection for mixture models - Perspectives and strategies. In *Handbook of Mixture Analysis*. London: Chapman and Hall.
- Chatzi EN and Smyth AW** (2009) The unscented Kalman filter and particle filter methods for nonlinear structural system identification with non-collocated heterogeneous sensing. *Structural Control and Health Monitoring* 16(1), 99–123.
- Chatzi E and Smyth A** (2013) Particle filter scheme with mutation for the estimation of time-invariant parameters in structural health monitoring applications. *Structural Control and Health Monitoring* 20, 1081–1095.
- Chen Y, Gupta MR, Chen Y and Gupta MR** (2010) *EM Demystified: An Expectation-Maximization Tutorial*. Seattle, WA: University of Washington, Dept. of EE.
- Ching J and Chen Y-C** (2007) Transitional markov chain Monte Carlo method for Bayesian model updating, model class selection, and model averaging. *Journal of Engineering Mechanics* 133(7), 816–832.
- Chopin N** (2002) A sequential particle filter method for static models. *Biometrika* 89(3), 539–552.
- Chopin N** (2004) Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *The Annals of Statistics* 32(6), 2385–2411.
- Chopin N, Jacob PE and Papaspiliopoulos O** (2013) Smc2: An efficient algorithm for sequential analysis of state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 75(3), 397–426.
- Corbetta M, Sbarufatti C, Giglio M and Todd MD** (2018) Optimization of nonlinear, non-gaussian bayesian filtering for diagnosis and prognosis of monotonic degradation processes. *Mechanical Systems and Signal Processing* 104, 305–322.
- Cotter SL, Roberts GO, Stuart AM and White D** (2013) MCMC methods for functions: Modifying old algorithms to make them faster. *Statistical Science* 28(3), 424–446.
- Cristiani D, Sbarufatti C and Giglio M** (2021) Damage diagnosis and prognosis in composite double cantilever beam coupons by particle filtering and surrogate modelling. *Structural Health Monitoring* 20(3), 1030–1050.
- Daum F** (2005) Nonlinear filters: Beyond the Kalman filter. *IEEE Aerospace and Electronic Systems Magazine* 20(8), 57–69.
- Del Moral P, Doucet A and Jasra A** (2006) Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 68(3), 411–436.
- Der Kiureghian A and Ke J-B** (1988) The stochastic finite element method in structural reliability. *Probabilistic Engineering Mechanics* 3(2), 83–91.
- Deutscher J, Blake A and Reid I** (2000) Articulated body motion capture by annealed particle filtering. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000* (Cat. No. PR00662), vol. 2. Hilton Head, SC: IEEE, pp. 126–133.
- Ditlevsen O and Madsen H** (1996) *Structural Reliability Methods*. New York: Wiley.
- Doucet A, de Freitas N and Gordon N** (2001) *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag.
- Doucet A and Johansen AM** (2008) A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later.
- Elingwood B** (2005) Risk-informed condition assessment of civil infrastructure: State of practice and research issues. *Structure and Infrastructure Engineering* 1(1), 7–18.
- Evensen G** (2006) *Data Assimilation: The Ensemble Kalman Filter*. Berlin: Springer-Verlag.
- Frangopol DM** (2011) Life-cycle performance, management, and optimisation of structural systems under uncertainty: Accomplishments and challenges. *Structure and Infrastructure Engineering* 7(6), 389–413.
- Gall J, Potthoff J, Schnörr C, Rosenhahn B and Seidel H-P** (2007) Interacting and annealing particle filters: Mathematics and a recipe for applications. *Journal of Mathematical Imaging and Vision* 28, 1–18.
- Gelman A, Gilks WR and Roberts GO** (1997) Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability* 7(1), 110–120.
- Geyer S, Papaioannou I and Straub D** (2019) Cross entropy-based importance sampling using Gaussian densities revisited. *Structural Safety* 76, 15–27.
- Ghanem RG and Spanos PD** (1991) *Stochastic Finite Elements: A Spectral Approach*. New York: Springer.
- Gilks WR and Berzuini C** (2001) Following a moving target-Monte Carlo inference for dynamic bayesian models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 63(1), 127–146.

- Gilks W, Richardson S and Spiegelhalter D** (1995) *Markov Chain Monte Carlo in Practice*. London: Chapman & Hall/CRC Interdisciplinary Statistics. Taylor & Francis.
- Gordon N, Salmond D and Smith A** (1993) Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings Part F: Radar and Signal Processing* 140(2), 107.
- Haario H, Laine M, Mira A and Saksman E** (2006) DRAM: Efficient adaptive MCMC. *Statistics and Computing* 16, 339–354.
- Hastings WK** (1970) Monte Carlo sampling methods using markov chains and their applications. *Biometrika* 57(1), 97–109.
- Hu X-L, Schon TB and Ljung L** (2008) A basic convergence result for particle filtering. *IEEE Transactions on Signal Processing* 56(4), 1337–1348.
- Jasra A, Stephens DA, Doucet A and Tsagaris T** (2011) Inference for lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scandinavian Journal of Statistics* 38(1), 1–22.
- Jazwinski A** (1970) *Stochastic Processes and Filtering Theory. Mathematics in Science and Engineering*. Amsterdam, The Netherlands: Elsevier Science.
- Julier S and Uhlmann J** (1997) New extension of the Kalman filter to nonlinear systems. In Kadar, I., editor, *Signal Processing, Sensor Fusion, and Target Recognition VI*, 3068, 182–193. International Society for Optics and Photonics, SPIE, Orlando, FL, United States.
- Kalman R** (1960) A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering* 82, 35–45.
- Kamariotis A, Chatzi E and Straub D** (2022) Value of information from vibration-based structural health monitoring extracted via bayesian model updating. *Mechanical Systems and Signal Processing* 166, 108465.
- Kamariotis A, Chatzi E and Straub D** (2023) A framework for quantifying the value of vibration-based structural health monitoring. *Mechanical Systems and Signal Processing* 184, 109708.
- Kantas N, Doucet A, Singh S and Chopin N** (2015) On particle methods for parameter estimation in state-space models. *Statistical Science* 30(3), 328–351.
- Katafygiotis L and Zuev K** (2008) Geometric insight into the challenges of solving high-dimensional reliability problems. *Probabilistic Engineering Mechanics* 23(2), 208–218.
- Kennedy MC and O'Hagan A** (2001) Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(3), 425–464.
- Kim N-H, An D and Choi J-H** (2017) *Prognostics and Health Management of Engineering Systems. An introduction*. Cham: Springer.
- Kitagawa G** (1998) A self-organizing state-space model. *Journal of the American Statistical Association* 93(443), 1203–1215.
- Kiureghian AD and Liu P** (1986) Structural reliability under incomplete probability information. *Journal of Engineering Mechanics* 112(1), 85–104.
- Liu JS and Chen R** (1998) Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association* 93(443), 1032–1044.
- Liu J and West M** (2001) *Combined Parameter and State Estimation in Simulation-Based Filtering*. New York: Springer.
- Luque J and Straub D** (2019) Risk-based optimal inspection strategies for structural systems using dynamic bayesian networks. *Structural Safety* 76, 68–80.
- Lye A, Cicirello A and Patelli E** (2021) Sampling methods for solving Bayesian model updating problems: A tutorial. *Mechanical Systems and Signal Processing* 159, 107760.
- McLachlan GJ and Krishnan T** (2007) *The EM Algorithm and Extensions, Second Edition*. Hoboken, NJ: John Wiley & Sons.
- Melchers R and Beck A** (2017) *Structural Reliability Analysis and Prediction*, 3rd Edn. Hoboken, NJ: John Wiley & Sons.
- Musso C, Oudjane N and Le Gland F** (2001) *Improving Regularised Particle Filters*. New York: Springer.
- Myötyri E, Pulkkinen U and Simola K** (2006) Application of stochastic filtering for lifetime prediction. *Reliability Engineering & System Safety* 91(2), 200–208.
- Neal RM** (2001) Annealed importance sampling. *Statistics and Computing* 11, 125–139.
- Papaioannou I, Betz W, Zwirgmaier K and Straub D** (2015) MCMC algorithms for subset simulation. *Probabilistic Engineering Mechanics* 41, 89–103.
- Papaioannou I, Papadimitriou C and Straub D** (2016) Sequential importance sampling for structural reliability analysis. *Structural Safety* 62, 66–75.
- Papakonstantinou K and Shinozuka M** (2014) Planning structural inspection and maintenance policies via dynamic programming and markov processes. Part II: POMDP implementation. *Reliability Engineering & System Safety* 130, 214–224.
- Paris P and Erdogan F** (1963) A critical analysis of crack propagation laws. *Journal of Basic Engineering* 85(4), 528–533.
- Ramanacha MK, Astroza R, Madarshahian R and Conte JP** (2022) Bayesian updating and identifiability assessment of nonlinear finite element models. *Mechanical Systems and Signal Processing* 167, 108517.
- Rubinstein RY and Kroese DP** (2016) *Simulation and the Monte Carlo Method*, 3rd Edn. Hoboken, NJ: Wiley Publishing.
- Russell SJ and Norvig P** (2021) *Artificial Intelligence: A Modern Approach*, 4th Edn. London: Pearson Education, Inc.
- Saha B, Goebel K, Poll S and Christophersen J** (2009) Prognostics methods for battery health monitoring using a bayesian framework. *IEEE Transactions on Instrumentation and Measurement* 58(2), 291–296.
- Särkkä S** (2013) *Bayesian Filtering and Smoothing*. Cambridge: Cambridge University Press.
- Schubert E, Sander J, Ester M, Kriegel HP and Xu X** (2017) DbSCAN revisited, revisited: Why and how you should (still) use dbSCAN. *ACM Transactions on Database Systems* 42(3), 1–21.

- Smith AFM and Gelfand AE** (1992) Bayesian statistics without tears: A sampling-resampling perspective. *The American Statistician* 46(2), 84–88.
- Song M, Astroza R, Ebrahimi H, Moaveni B and Papadimitriou C** (2020) Adaptive Kalman filters for nonlinear finite element model updating. *Mechanical Systems and Signal Processing* 143, 106837.
- South LF, Pettitt AN and Drovandi CC** (2019) Sequential Monte Carlo samplers with independent Markov chain Monte Carlo proposals. *Bayesian Analysis* 14(3), 753–776.
- Stewart MG and Mullard JA** (2007) Spatial time-dependent reliability analysis of corrosion damage and the timing of first repair for RC structures. *Engineering Structures* 29(7), 1457–1464.
- Storvik G** (2002) Particle filters for state-space models with the presence of unknown static parameters. *IEEE Transactions on Signal Processing* 50(2), 281–289.
- Straub D** (2009) Stochastic modeling of deterioration processes through dynamic Bayesian networks. *Journal of Engineering Mechanics* 135(10), 1089–1099.
- Straub D, Schneider R, Bismut E and Kim H-J** (2020) Reliability analysis of deteriorating structural systems. *Structural Safety* 82, 101877.
- Sudret B and Der Kiureghian A** (2000) Stochastic Finite Element Methods and Reliability a State-of-the-Art Report.
- Sun J, Zuo H, Wang W and Pecht MG** (2014) Prognostics uncertainty reduction by fusing on-line monitoring data based on a state-space-based degradation model. *Mechanical Systems and Signal Processing* 45(2), 396–407.
- Tatsis KE, Dertimanis VK and Chatzi EN** (2022) Sequential Bayesian inference for uncertain nonlinear dynamic systems: A tutorial. *Journal of Structural Dynamics* 1, 236–262.
- van der Merwe R and Wan E** (2003) Gaussian mixture sigma-point particle filters for sequential probabilistic inference in dynamic state-space models. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03)*, vol. 6, IEEE, Hong Kong, China.
- Vanmarcke E** (2010) *Random Fields*. Singapore: World Scientific.
- Veettil DRA and Chakravorty S** (2016) *Particle Gaussian Mixture (PGM) Filters*. *arXiv:1603.04510*.
- Vu KAT and Stewart MG** (2000) Structural reliability of concrete bridges including improved chloride-induced corrosion models. *Structural Safety* 22(4), 313–333.
- Wirgin A** (2004) The Inverse Crime. *arXiv E-prints*, p. math-ph/0401050.
- Wu S, Angelikopoulos P, Papadimitriou C and Koumoutsakos P** (2017) Bayesian annealed sequential importance sampling: An unbiased version of transitional Markov Chain Monte Carlo. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering* 4(1), 011008.
- Yi S-R and Song J** (2018) Particle filter based monitoring and prediction of spatiotemporal corrosion using successive measurements of structural responses. *Sensors* 18(11), 3909.