



This is a repository copy of *A latent restoring force approach to nonlinear system identification*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/230689/>

Version: Published Version

Article:

Rogers, T.J. orcid.org/0000-0002-3433-3247 and Friis, T. (2022) A latent restoring force approach to nonlinear system identification. *Mechanical Systems and Signal Processing*, 180. 109426. ISSN: 0888-3270

<https://doi.org/10.1016/j.ymssp.2022.109426>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>



A latent restoring force approach to nonlinear system identification

T.J. Rogers^{a,*}, T. Friis^b

^a Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield, S1 3JD, UK

^b Department of Civil Engineering, Technical University of Denmark, Brovej, 2800 Kgs., Lyngby, Denmark

ARTICLE INFO

Communicated by J. Noël

Keywords:

Bayesian
Nonlinear system identification
Gaussian process
Grey-box model

ABSTRACT

Identification of nonlinear dynamic systems remains a significant challenge across engineering. This work suggests an approach based on Bayesian filtering to extract and identify the contribution of an unknown nonlinear term in the system which can be seen as an alternative viewpoint on restoring force surface type approaches. To achieve this identification, the contribution which is the nonlinear restoring force is modelled, initially, as a Gaussian process in time. That Gaussian process is converted into a state-space model and combined with the linear dynamic component of the system. Then, by inference of the filtering and smoothing distributions, the internal states of the system and the nonlinear restoring force can be extracted. In possession of these states a nonlinear model can be constructed. The approach is demonstrated to be effective in both a simulated case study and on an experimental benchmark dataset.

1. Introduction

Identification of nonlinear dynamic systems remains an outstanding challenge across engineering. Although this paper will focus on mechanical applications, nonlinear systems emerge across a range of disciplines from electrical and control systems to bioengineering. In identifying these systems there are generally two aims for the engineer; first, to gain insight into the physical phenomena driving the behaviour of that system, and secondly to make predictions of the response of that system to new inputs. Owing to its broad applicability and challenge, nonlinear system identification has seen a great deal of research interest sustained over many years, for a good review see [1]. Although often presented, in a mechanical setting, as second-order nonlinear differential equations, a nonlinear dynamic system can be considered equivalently in a state-space framework. The distinction being that several internal states are modelled explicitly, as is a defined “measurement model” which relates these states to some observed quantity; this model also allows for the concepts of both process noise and measurement noise. Mathematically, the general continuous time state-space model is shown in Eq. (1).

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), \mathbf{v}(t)) \quad (1a)$$

$$\mathbf{y}(t) = \mathbf{g}(\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t)) \quad (1b)$$

where $\mathbf{x}(t)$ is a vector of internal states at time t whose time derivatives are contained in $\dot{\mathbf{x}}(t)$. These derivatives are given by some function $\mathbf{f}(\cdot)$ of the states $\mathbf{x}(t)$, some known external inputs $\mathbf{u}(t)$ and some process noise $\mathbf{v}(t)$. The internal states may not be directly observed, instead a quantity $\mathbf{y}(t)$ is the measured signal at time t which itself is a function $\mathbf{g}(\cdot)$ of the states $\mathbf{x}(t)$, some known external inputs $\mathbf{u}(t)$ and some measurement noise $\mathbf{w}(t)$.

The work shown in this paper will consider one specific task which a practitioner may be interested in when confronted with a nonlinear system. That is to identify from measured inputs and outputs a nonlinear system where the governing equation of motion is

* Corresponding author.

E-mail address: tim.rogers@sheffield.ac.uk (T.J. Rogers).

unknown *a priori*. In particular, a modern viewpoint on a classic nonparametric approach will be shown, that is a Bayesian treatment of the nonlinear restoring force surface method of Masri and Caughey [2] when limited measurement information is available. The methodology presented is of particular value in cases where measurement noise means that numerical estimates of the internal states of the system (displacement, velocity, etc.) are poor.

The key concept employed in this work is to hypothesise that the unknown *nonlinear restoring force* can be modelled as a *Gaussian process* (GP) in time [3,4] and therefore incorporated into a Gaussian process *latent force model* (LFM) [5,6]. In other words the contribution of the unknown nonlinear terms to the equation of motion is replaced with a GP as a function in time which recovers the unknown nonlinear restoring force. The GP provides a prior over a function [4], in this case the signal in time which is the nonlinear restoring force, the function is updated in a Bayesian manner as observations of the system response are made to recover a posterior distribution over the potential contribution of the nonlinear restoring force. The recovered nonlinear restoring force along with estimates of the displacement and velocity can then be used to infer the unknown terms in the nonlinear equation of motion in isolation. Therefore, the contribution of this paper is to show how the Gaussian process Latent Force Model (GPLFM) might be employed in a novel manner to aid system identification of nonlinear systems with nonlinear terms which are unknown *a priori*, through separating their contribution from that of the linear system components.

1.1. Related work

The scope of nonlinear system identification is far reaching and cross-disciplinary. Within a mechanical context, which is the focus of this work, good reviews of progress being made can be found in [7] and more recently, in the subsequent work of [1]. The range of applied technologies has been broad ranging from physically derived models to machine learning approaches. Examples of physically grounded methods include the control-based continuation approach of [8] or Approximate Bayesian Computation to learn models and parameters in [9] or applications of nonlinear filtering algorithms [10]. From a purely data-driven perspective, as early as [11,12] black-box tools such as neural networks were being considered as models of nonlinear dynamic systems and this continues with the current popularity of deep learning approaches; e.g. [13]. [11] in fact builds on the restoring force [2] approach by using a neural network to learn a model of the nonlinear function rather than the Chebyshev polynomials used in the original work. The general approach seen in [2] led to widespread use, for example [14,15], and in [16] is cited as one of the major families of approach in time-domain identification. The power of the restoring force approach in [2] is evidenced by its inclusion in the updated review paper of [1] which includes numerous examples of its ongoing application. The machine learning community has more generally considered a range of problems within the scope of nonlinear dynamic system identification. For example, methodologies such as the neural network approaches of [17,18] or GPs in the case of [19].

Of particular relevance to the work shown in this paper is the work of [5] with the introduction of the GPLFM. This is a mechanistically inspired machine learning model as discussed in the previous section. In [20] applications in gene expression and the diffusion of pollutants are discussed as the dynamic processes being considered. Within a mechanical engineering context the ideas of the GPLFM have been exploited for solving the joint input-state-parameter problem by the authors in [21,22]. [23] also discussed the use of such a model for input estimation with known system parameters. Here the task considered was to infer, from a set of measured accelerations, the unmeasured states (displacements and velocities), parameters (mass, stiffness and damping matrices) and inputs (excitation forces) in a joint framework based on the GPLFM. The approach was extended in [24] to an input-state estimation method for nonlinear systems where the parameters and form of the system were known but the inputs were not.

1.2. Overview of proposed approach

An overview of the proposed methodology presented in this work is shown in Fig. 1. The key steps of the method and main contributions of this work are those shown in blocks C, D and E. Specifically, the restoring force estimation problem is recast as a state-estimation problem where the unknown component (in time) is modelled as a GP. This gives rise to a linear SSM over which inference can be performed, since the SSM remains linear computation time is greatly reduced. In addition, the identification aims to decouple the dynamic model from the fitting of the nonlinear term (blocks G and H). By inferring the states and the missing contribution to the system along with the parameters (F.1 and F.2) it is possible to extract the *static* form of the nonlinearity, i.e. it is possible to plot the restoring force curve/surface as in block G. Fitting a static model to this data is far less demanding than fitting and comparing multiple nonlinear dynamic models. Having given this very general overview of the proposed approach, the specific construction of the model can now be presented.

2. From restoring force surfaces to latent restoring forces

Before discussing the particular approach that is adopted in this work, it is worth reviewing the classic restoring force surface approach of Masri and Caughey [2]. The ongoing applicability of this method is a testament to both its intuitive approach and the quality of that original work. The authors in that paper note that for a single-degree-of-freedom (SDOF) system the differential equation of motion can be written as,

$$m\ddot{z} + f(z, \dot{z}) = U(t) \quad (2)$$

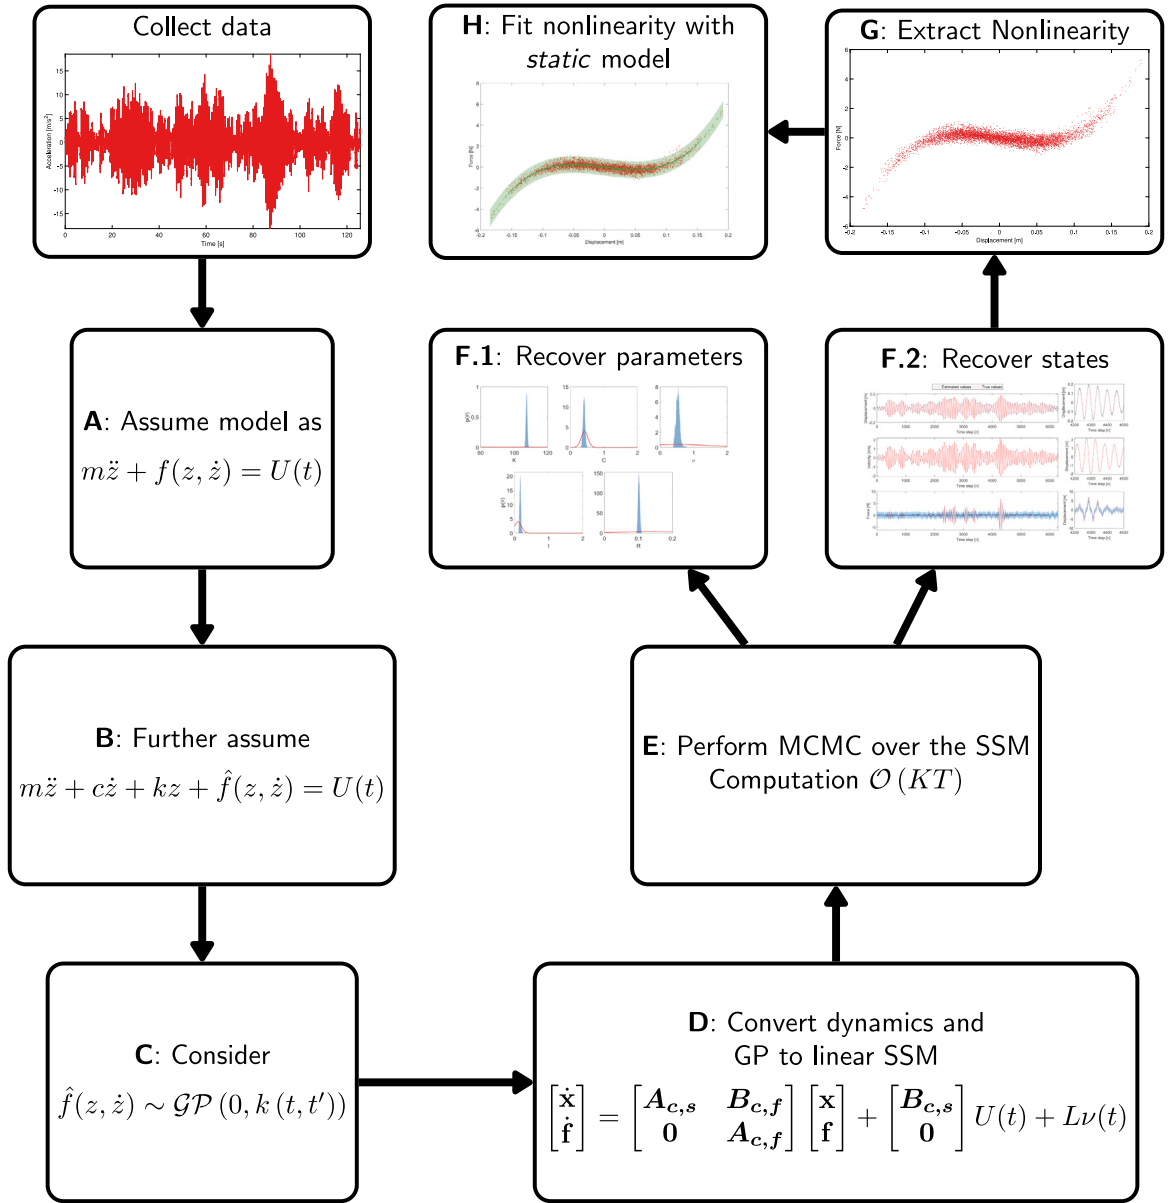


Fig. 1. Process for GPLFM based restoring force identification.

for an object of mass m with displacement z , velocity¹ \dot{z} and acceleration \ddot{z} ; forced by a known signal $U(t)$. Then the function $f(z, \dot{z})$ is termed the *restoring force* of the system. In the special case where $f(z, \dot{z}) = c\dot{z} + kz$, for scalar coefficients c and k , the equation of motion of a linear SDOF mass-spring-damper system is recovered. However, the case of interest is when this restoring force is generated by some unknown, possibly nonlinear, function. Masri and Caughey [2] observe that if the mass m , acceleration \ddot{z} and excitation $U(t)$ are known, for example from an experiment study, the equation of motion can be rearranged to find this unknown function $f(z, \dot{z})$.

$$f(z, \dot{z}) = U(t) - m\ddot{z} \quad (3)$$

Having made this rearrangement, a functional form for the restoring force is estimated. While there are many possible options for this modelling, the one highlighted in the original paper is the use of Chebyshev orthogonal polynomials. Through fitting this

¹ Throughout this paper the notation $\dot{(\cdot)}$ is adopted to indicate the first derivative with respect to time and likewise $\ddot{(\cdot)}$ for the second derivative.

function $f(z, \dot{z})$, given the observations of the excitation $U(t)$ and the inertial term $m\ddot{z}$, it is possible to then predict the response of the nonlinear system to some new excitation signal or to inspect the form of this function for physical insight.

In practice, this methodology can perform very well, however, one drawback that may be encountered is the need to measure or estimate the displacement, velocity and acceleration of the system in order to fit the function $f(z, \dot{z})$. This may increase the cost of testing to introduce more sensor modalities or require the user to resort to some estimation method for the missing quantities. By far the most common measurement, of a mechanical dynamic system, is the acceleration of the mass \ddot{z} . From this, the use of numerical integration techniques could be employed to estimate the displacement and velocity of the mass. The issue with this approach is that error is often introduced during this procedure, which may distort results or require further post-processing. This difficulty increases significantly with noise on the measurements which are collected. Worden [25] discusses in further detail some of the challenges associated with employing these numerical methods within the restoring force surface methodology.

Inspired by the restoring force surface approach, this paper looks to overcome the issues associated with obtaining measurements or estimates of the displacement, velocity and acceleration by estimating them within a Bayesian state-space framework. The aim then being to recover simultaneously, the parameters of the underlying linear model, the unmeasured quantities and the missing nonlinear restoring force.

2.1. Gaussian process latent force models

Alvarez et al. [5] introduced a model where a data-driven regression model, a Gaussian process, is linked to a mechanistic representation of the data — an ordinary differential equation (ODE). For example to model data generated by a second-order ODE in the form,

$$m\ddot{z} + c\dot{z} + kz = U(t), \quad U(t) \sim \mathcal{GP}(0, k(t, t')) \quad (4)$$

where the unknown forcing on the differential equation $U(t)$ is modelled as a GP with zero mean and a covariance function $k(t, t')$ which is stationary and depends on time.

The original purpose of the GPLFM was to increase the flexibility and expressive power of the GP in modelling data. This is shown to work well on a number of examples. One limiting factor to the approach shown in Alvarez et al. [5] is the computational burden of computing the latent force model, this is in some way addressed in [26] through the introduction of sparse inference over the GP.² An alternative approach to increasing the computational efficiency of the latent force model is made available through work of [28], who show that a GP in time with zero mean and a stationary covariance function can be rewritten as a stochastic differential equation (SDE) which has an equivalent form as a linear Gaussian state-space model (LGSSM). Once in possession of this LGSSM, inference can be performed using the well known Kalman filter [29] and Rauch–Tung–Streibel (RTS) smoother [30]. Importantly, [28] show that the smoothing solution to the obtained model is equivalent to the posterior solution of the full GP under the usual Gaussian likelihood. Once able to model the forcing of the system as a GP in time, which can be written as linear state-space model, it is a sensible step to link this with a state-space representation of the dynamic system on the left-hand side of Eq. (4) and perform inference over the resulting joint state-space model (SSM). This combination of the two ideas was shown in [6] which allows $\mathcal{O}(T)$ inference of the latent force model to take place via the Kalman filtering, RTS smoothing algorithms. For the full details of these approaches, the reader is referred to the original papers; [5] as an introduction followed by [28] then [6] for the implementation as an LGSSM.

2.2. A latent restoring force model

The approach of the GPLFM assumes a linear ODE driven by a GP which is unmeasured and must be inferred. This lends itself naturally to problems in mechanical engineering such as the joint input-state-parameter estimation problem where it has been previously applied, see Section 1.1. However, in this work a different problem is considered. It is assumed that measurements of the input to the system $U(t)$ are available with the exact form of the ODE describing the physical behaviour of the system being the unknown.

It remains to show that a similar approach to that of the GPLFM can be adopted to infer the unmeasured states and the nonlinear restoring force of the system simultaneously in a state-space framework. Recalling Eq. (2), the model being considered is one of a (potentially nonlinear) second order differential equation with an unknown restoring force function $f(z, \dot{z})$. Assuming that there is some linear component, this can be rewritten as,

$$m\ddot{z} + c\dot{z} + kz + \hat{f}(z, \dot{z}) = U(t) \quad (5)$$

where the nonlinear component of the restoring force is represented by a new but still unknown function $\hat{f}(z, \dot{z})$. It will also be considered that only measurement of the acceleration is available from testing, although an equivalent method could be written down trivially for different sensor modalities, e.g. velocity measurements from a laser Doppler vibrometer. It is necessary to infer the unknown displacement z , velocity \dot{z} and the unknown restoring force function $f(z, \dot{z}) = kz + c\dot{z} + \hat{f}(z, \dot{z})$ simultaneously as all three quantities are unmeasured.

² There are a number of approaches for sparse computationally efficient inference in GPs, one of the most common frameworks is that which employs a number of pseudo-points or inducing points, a good overview is provided in [27] which also links two of the most popular methods.

Inference over the coupled quantities (states, parameters and the unknown nonlinear function) directly is difficult. For example, it is possible to consider multiple possible forms for $f(z, \dot{z})$ and calibrate each of those models, eventually selecting the one which is optimal in some sense. This approach has been seen in the literature, an example already discussed is Abdesslem et al. [9] using ABC to perform joint model and parameter selection. While effective, it — along with many other direct approaches — requires multiple nonlinear simulations which impose a high computational cost.

Therefore, a strategy is proposed in the same spirit as the GPLFM where it is assumed that $\hat{f}(z, \dot{z})$ may be represented as a GP in time with zero mean and a stationary kernel,

$$m\ddot{z} + c\dot{z} + kz + \hat{f}(z, \dot{z}) = U(t), \quad \hat{f}(z, \dot{z}) \sim \mathcal{GP}(0, k(t, t')) \quad (6)$$

This construction of the model gives rise to two related benefits; firstly inference may now be performed in a joint manner over z, \dot{z} and $\hat{f}(z, \dot{z})$; secondly, the nonlinear system is approximated by a linear SSM leading to significant computational advantages. These computational gains mean that Bayesian inference over the parameters of the model: the mass, stiffness and damping; as well as the hyperparameters of the GP is feasible. Having conducted Bayesian inference over the linear dynamic model which is far less computationally intensive, the nonlinear (and linear if desired) components of the model can be fit in a *static* manner. In other words, extracting z, \dot{z} and $\hat{f}(z, \dot{z})$ allows the ODE describing the system to be learnt without the need for further costly nonlinear simulations of the full dynamic model.

The formulation of such a model can now be shown; although the finer details of representing the GP in a state-space form will not be covered. These details can be found in the original papers of [6,28] or in the context of structural dynamics in [22]. The process of inferring the model begins by converting the system into a state space form in continuous time given a state vector $\mathbf{x} = [z \quad \dot{z}]^T$, as in Eq. (7).

$$\begin{aligned} \dot{\mathbf{x}} &= \underbrace{\begin{bmatrix} 0 & 1 \\ -\frac{k}{m} & -\frac{c}{m} \end{bmatrix}}_{\mathbf{A}_{c,s}} \mathbf{x} + \underbrace{\begin{bmatrix} 0 \\ \frac{1}{m} \end{bmatrix}}_{\mathbf{B}_{c,s}} (U(t) - \hat{f}(z, \dot{z})) \\ \mathbf{y} &= \mathbf{C}\mathbf{x} + \mathbf{D}(U(t) - \hat{f}(z, \dot{z})) \end{aligned} \quad (7)$$

It can be observed that this has the form of a standard continuous-time state-space model, with the dynamics governed by $\mathbf{A}_{c,s}$ and $\mathbf{B}_{c,s}$, and the observations made via \mathbf{C} and \mathbf{D} . These matrices \mathbf{C} and \mathbf{D} will be dependent on the particular sensing modality used; for example, if observing acceleration $\mathbf{C} = [-k/m \quad -c/m]$ and $\mathbf{D} = 1/m$.

Since the nonlinear restoring force $\hat{f}(z, \dot{z})$ is not a known or measured quantity, it must be modelled approximately using some suitable prior assumptions which can be updated through the observation data. As discussed, this is achieved through modelling $\hat{f}(z, \dot{z})$ as a Gaussian process in time with a zero mean and a stationary kernel $k(t, t')$. Notationally, the model of the nonlinear restoring force is comprised of a state vector \mathbf{f} , a state transition matrix $\mathbf{A}_{c,f}$ and a white noise process $\mathbf{v}(t)$ with a spectral density q entering on one of the states of \mathbf{f} :

$$\dot{\mathbf{f}} = \mathbf{A}_{c,f}\mathbf{f} + \mathbf{L}\mathbf{v}(t)$$

The size of \mathbf{f} , structure of the $\mathbf{A}_{c,f}$ matrix and spectral density q are fully defined by the form of the kernel $k(t, t')$ via its spectral density $S(\omega)$ [28]. For certain classes of kernel, including the widely used Matérn type [31], these quantities are available in a closed form which gives rise to a finite dimensional linear state-space model. The derivation of the matrix $\mathbf{A}_{c,f}$ from a particular kernel is not included in this work for brevity, it can be found in full in the literature, for example see [6,22,23,28,32].

What remains is the combination of these two models into a single SSM, such that inference can be performed jointly over the unknown states of the system (the displacement and velocity) and the unknown nonlinear restoring force. This involves one modification to Eq. (7), that being to move the nonlinear restoring force $\hat{f}(z, \dot{z})$ into the state vector \mathbf{x} , creating a new extended state-transition matrix \mathbf{A}_c and control matrix \mathbf{B}_c . The observation matrices \mathbf{C} and \mathbf{D} will also be updated to account for the relation of the new extended state vector to the measured quantity or quantities. Appending the state vector \mathbf{x} with the additional states related to the GP representation of the nonlinear restoring force gives,

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{f}} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{c,s} & \mathbf{B}_{c,f} \\ \mathbf{0} & \mathbf{A}_{c,f} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{f} \end{bmatrix} + \begin{bmatrix} \mathbf{B}_{c,s} \\ \mathbf{0} \end{bmatrix} U(t) + \mathbf{L}\mathbf{v}(t) \quad (8)$$

where $\mathbf{B}_{c,f}$ models how the force will act on the physical dynamic system. In the SDOF cases considered in this work, this is a matrix of zeros except in the lower left-hand corner where it equals $1/m$, this can be seen from the rearrangement of Eq. (5). The matrix \mathbf{L} controls how the white noise process $\mathbf{v}(t)$ enters the system, in this case it is a column vector where all elements are zero except the final one which is unity. The formulation in Eq. (8) provides a linear system approximation of the nonlinear system, where the nonlinear restoring force is represented by a Gaussian process in time. The object of interest is the smoothing distribution of the LGSSM which can be efficiently computed via the Kalman filter and RTS smoother, this distribution is the posterior over the internal states of the model which include the “physical states” and the additional states related to the GP which represents the nonlinear restoring force. In order to practically employ these methods, the system must first be discretised which is a standard procedure and can be found in many good textbooks, for example see [33,34].

Once the posterior distributions over the states have been recovered, these may be used to build a model of the nonlinear component of the system in a manner similar to [2]. Samples of the states can be drawn which capture the uncertainty in the estimated displacement and velocities as well as in the unknown nonlinear restoring force. In possession of these samples, the task




Fig. 2. Acceleration measurement from the simulated Duffing oscillator.

is then to fit a model of the nonlinearity, this step is not covered in detail in this work since many options exist. For instance, one could adopt a polynomial method (or orthogonal polynomial) as in the original work of [2]. Alternatively, a more physically-informed approach could be taken by inspection of the system itself or the nonlinear restoring force now isolated from the rest of the dynamics. If not in possession of strong prior knowledge, then a purely “black-box” approach may be taken, for example a neural network or GP could be used to model the nonlinear function. After a model is developed of the nonlinearity, on the basis of the state estimates, this can be used in future prediction tasks.

At this point it is worth noting that the authors would recommend the use of numerically stable square-root forms of the Kalman filter and RTS smoother. This can overcome issues when the measurement noise in the system is very low. Specifically, the formulations shown in [35] were used in the case studies shown in this paper, including the procedure for sampling from the posterior distributions over the states. Without the use of these more stable methods, it was found that the variance in the samples of the states could be artificially inflated or in extreme cases lead to over/underflow issues in the computation.

3. Case studies

This section will present two examples of how the proposed methodology could be applied to the identification of a nonlinear system. The first, a simulated study with a Duffing oscillator; second, an experimental dataset known as the Silverbox produced by [36] which electrically exhibits response similar to a Duffing oscillator. This will demonstrate the effectiveness of the method in recovering a model in the case where it can be compared to a ground truth and also when attempting to identify a physical test object which is more representative of an industrial system identification task.

3.1. Simulated Duffing system

The Duffing oscillator is arguably one of the simplest interesting nonlinear systems that an engineer may wish to study. Its equation of motion is identical to that of a linear mass–spring–damper system except the addition of a term which is cubic in the displacement, i.e.,

$$m\ddot{z} + c\dot{z} + kz + k_3z^3 = U(t) \quad (9)$$

A system following Eq. (9) is simulated using a Newmark integration scheme [37] with a mass $m = 1$ kg, linear stiffness $k = 100$ N m⁻¹, viscous damping $c = 0.4$ N s m⁻¹ and cubic stiffness $k_3 = 1000$ N m⁻³. The system is simulated for 12566 time steps with a sample frequency of 100 Hz. The excitation used is a random phase multisine based on 1000 sampled frequencies and amplitudes from a JONSWAP wave spectrum [38] with $H_s = 2.5$ and $T_p = 1$.

It is assumed that only measurements of the acceleration of the oscillator are available and these are polluted with an artificial noise through addition of i.i.d. samples from a Gaussian distribution with zero mean and a standard deviation close to 5% of the standard deviation of the noise free signal, this gives a noise variance $R = 0.05$. This noise addition replicates the sort of behaviour that may be expected when testing a physical system. Fig. 2 shows the time domain signal which is used as the “measurement” in this simulated Duffing study.

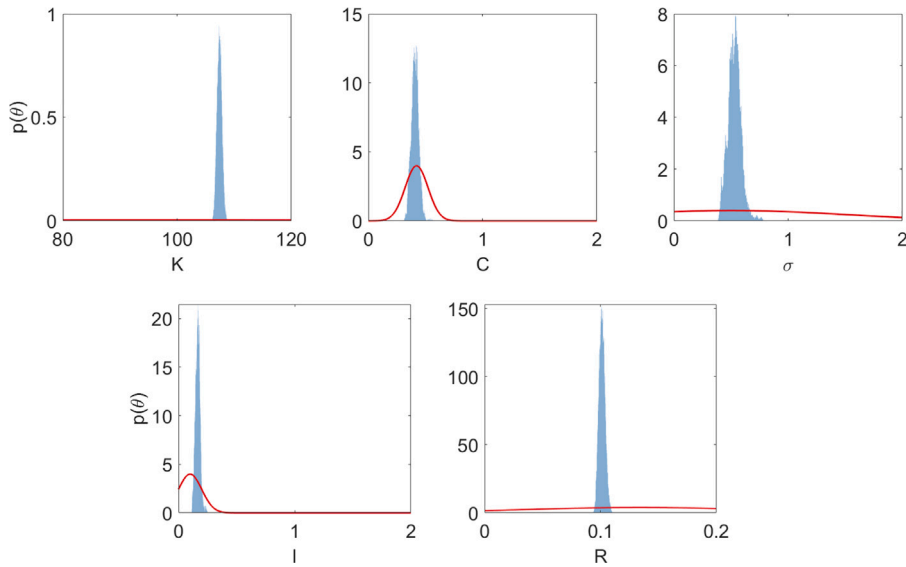


Fig. 3. Parameter posteriors from the MCMC inference over the state-space model. Solid (red) lines indicate the probability density function of the prior and the histograms (blue) are built from samples of the posterior.

Table 1
Prior distributions used in Case Study 1.

Prior	Distribution
$p(k)$	$\mathcal{N}(96.68, 100)$
$p(c)$	$\mathcal{N}(0.422, 0.1)$
$p(\sigma_f^2)$	$\mathcal{N}(0.5, 1)$
$p(\ell)$	$\mathcal{N}(0.1, 0.1)$
$p(R)$	$\mathcal{N}(0.1337, 1)$

It can now be shown how the framework proposed will allow identification of this nonlinear dynamic system. From Eq. (6) and Eq. (9) it can be seen that this system is well suited to identification in the manner described in this paper. Therefore, it is trivial to form Eq. (8) for this SDOF nonlinear system based on the known state-space form of a linear SDOF system (for example see [33], chapter 4) which gives the matrices $\mathbf{A}_{c,s}$ and $\mathbf{B}_{c,f}$. What remains is to specify a form for the GP which is chosen to model the nonlinear restoring force in the system. In this case a Matérn kernel is chosen with a roughness of $1/2$. In possession of this kernel, it is possible to write down the form of $\mathbf{A}_{c,f}$ and the spectral density of the white noise process $v(t)$, as functions of the two kernel hyperparameters: the signal variance σ_f^2 and length scale ℓ , see [32],

Given this construction of the model there are a number of unknown quantities which must be inferred. It is assumed that the mass of the oscillator is known but the remaining system parameters, the damping coefficient c and linear stiffness k are not. The nonlinear stiffness k_3 is not modelled at this stage but is also assumed unknown. There are two hyperparameters of the GP which are not known, as discussed there are σ_f^2 and ℓ . Finally, it is also assumed that the measurement noise variance R is unknown. This leaves five free (hyper)parameters which must be identified. The posterior distributions of these unknowns are estimated by means of a Markov Chain Monte Carlo (MCMC) inference where the (unnormalised log) likelihood is available in closed form from the Kalman filter and the priors are defined independently as follows in Table 1 after perturbing the true values, noting that the ground truth values for the hyperparameters of the GP are not available.

Using the standard Metropolis–Hastings accept reject kernel with a random walk proposal [39], samples are drawn from the Markov Chain until 20,000 are accepted and the first 2000 of these are discarded as a “burn-in”. Tuning the random walk proposals gave an acceptance rate of 17.9% which is within reasonable bounds. From this set of samples it is possible to infer the posterior distributions in a Monte Carlo manner and histograms of these posteriors are shown in Fig. 3. A certain degree of bias can be seen in the parameters of the linear system, both of which are overestimated when compared to the ground truth values known from the simulation. This is an interesting phenomenon which will be discussed fully later in the context of estimating the nonlinear component of the system from the estimated restoring forces.

Once the parameter posteriors have been estimated, it is also possible to estimate the distributions over the states in the model. It is worth commenting that these states represent the displacement and velocity of the oscillator as well as the unknown nonlinear restoring force. Fig. 4 shows the estimation of these states obtained via Kalman filtering and RTS smoothing for the batch of data used in the identification alongside a reconstructed time series of the accelerations from those states. The uncertainty in the reconstructed acceleration is also shown through linear combination of the Gaussian uncertainty on all of the states. Very high quality recovery of

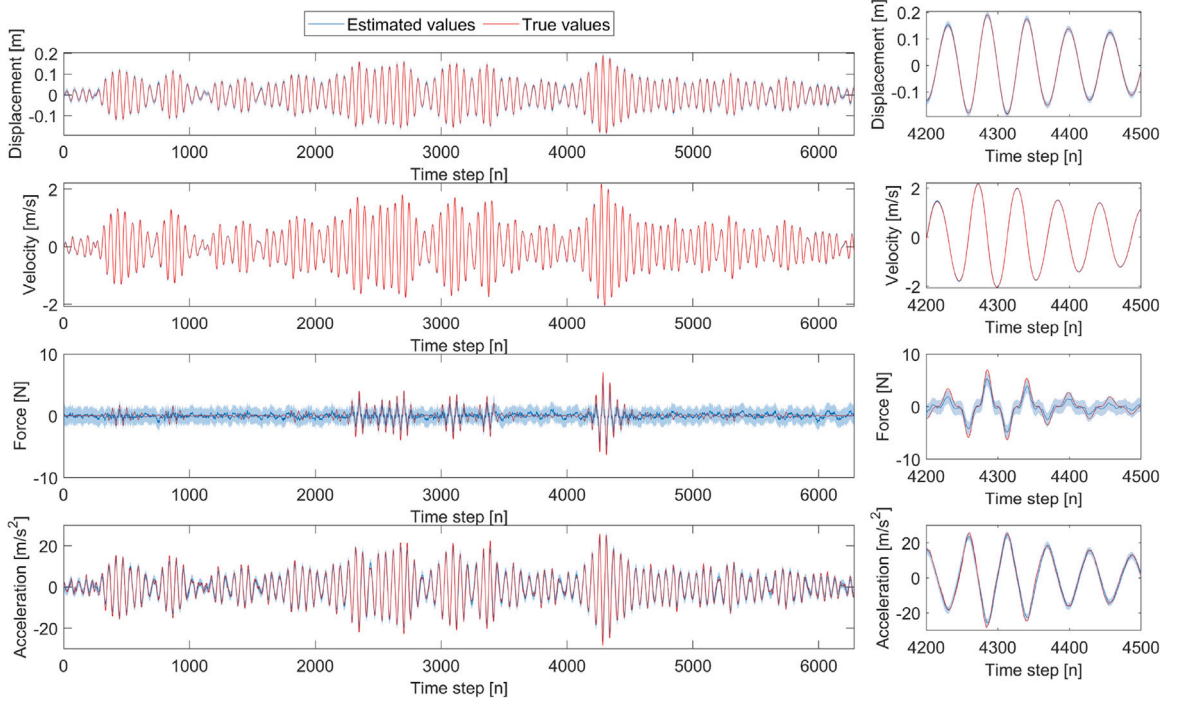


Fig. 4. Top three frames GPLFM state estimates for the simulated Duffing model, on the right hand side sections are enlarged to see the quality of the estimates. Bottom frame shows reconstructed acceleration from the state estimates.




Fig. 5. Frequency domain comparison of the third state related to the GP estimate.

the first two states can be observed from the estimation, all data points lie within a 3σ confidence interval and if taking the expected values as point estimates, the error compared to the ground truth is 0.35% for the displacement and 0.04% for the velocity.³ The associated error in the reconstructed acceleration is 2.81% using the NMSE metric, which remains a very good score. The slight increase in relation to the displacement and velocity states is associated with the increased noise on the measured signal (as opposed to the noise free ground truth that the states are compared to) which is captured in the predictive uncertainty.

³ To compute this error a normalised mean square error (NMSE) metric is used such that $J = 100/N/\mathbb{V}[y] \sum_{i=1}^N (y_i - \hat{y}_i)^2$ where y is the ground truth, \hat{y} the estimated value and $\mathbb{V}[y]$ the variance of y for N data points.




Fig. 6. Estimation of the true restoring force in the model and its residuals. True values shown in red, estimates in blue and the residual in green. Panel (a) shows the time series; (b) a zoomed in section of that time series; (c) the frequency domain view of the restoring force and its residuals; (d) the distribution of residuals.

In view of the goal of inferring the missing nonlinear component of the equation of motion, the second stage of the analysis is to extract samples of the unknown nonlinear restoring force. These can be obtained by sampling from the smoothing distribution of the constructed model, $p(\mathbf{x}_{1:T}; \mathbf{y}_{1:T}, \theta)$. These samples can be obtained in a backwards recursive manner based on the filtered system, see [35]. The third estimated state in Fig. 4 is that related to the GP which is being used to model the unknown nonlinear restoring force. Since this is a simulated system it is possible to compare this state with a ground truth as well. It can be seen that the fit of the nonlinear restoring force to the known ground truth is not as close to the ground truth as may be desired. The ground truth values of the nonlinear restoring force remain within the confidence bounds (three standard deviations) of the model, however, a significant increase in variance is observed and the expected value of this state does not lie as close to the ground truth as would be expected (60.3% NMSE). A frequency domain view of the prediction and residual when compared to the ground truth $k_3 z^3$ is shown in Fig. 5. Here the content of the difference between the GP estimate and the $k_3 z^3$ term can be explored. It is seen that there is significant low frequency deviation in the GP close to 1Hz and that there is increase high(er) frequency content from 7–25Hz. The low frequency content is a result of the biased estimate of the linear stiffness which will be discussed shortly. The higher frequency deviation is a result of the relative roughness of the GP kernel prior as opposed to the polynomial nonlinearity, see Appendix A where the choice of a kernel which is more times differentiable reduces this high frequency content. It is important to remember that the estimation of this state is not the final goal of the model, the aim of this procedure is to isolate the contribution of the nonlinear component such that it can be more easily identified and modelled, this will be expanded on shortly.

Before showing how the estimated states may be used to learn the nonlinearity in the system, it is worth discussing further the bias in the linear system parameters and the difference between the estimated nonlinear restoring force (the GP state) and the known contribution. During the first stage of the estimation procedure, MCMC is employed to jointly infer the parameters of the model (including the linear stiffness and damping coefficients) and the states of the SSM. Considering the results shown in Fig. 3, the damping coefficient c is estimated accurately by the MCMC procedure, however, significant positive bias is observed in the linear stiffness k . The maximum a posteriori (MAP) estimate of the stiffness is 107.38 N m^{-1} which is significantly higher than the ground truth of 100 N m^{-1} and counter to the perturbed prior which pulls that value lower. It will be seen that the total estimate of the restoring force is very accurate but this linear component is overestimated and the contribution of the GP is underestimated. This effect is not isolated to the current model under consideration, in fact it has been well established when combining GPs with physical models [40]. The prior imposes a weak constraint on the values of the physical parameters which is informed by engineering insight into the system and in this case study the mass is assumed known which alleviates the most serious nonidentifiability problems. However, there is still a notable positive bias in the linear stiffness of approximately 7 N m^{-1} , this increases the restoring force contribution of the linear stiffness term and conversely a lower contribution is made by the GP, leading to its difference from the ground truth. This phenomenon of bias in the physical parameters of a model when coupling it with a flexible machine learner is well established, for example see [40], where even in simple examples the flexibility of the GP leads to a bias in the learnt physical




Fig. 7. Plot of the nonlinear component of the model extracted from the GPLFM using estimated z and $\hat{f}(z, \dot{z})$.

parameters. In the case of a hardening stiffness nonlinearity in a dynamic system, it is intuitive that the linear stiffness may be overestimated to compensate for the contribution of the nonlinear terms, this is what is observed here.

It is reasonable to question why, given the construction of the model, the third state in the system fails to properly capture the expected nonlinear restoring force $\hat{f}(z, \dot{z})$ which, for the Duffing system, should be simply z^3 . To investigate, the estimate of the total restoring force is plotted in Fig. 6. This estimate is constructed by summing the estimated displacement multiplied by the MAP estimate of k , the estimated velocity multiplied by the MAP estimate of c and the contribution of the GP which is the third state in the model. This can be compared to the known true restoring force in the model, $kz + c\dot{z} + k_3z^3$. Considering this full restoring force as shown in Fig. 6 (frame (a) and in detail in frame (b)), it can be seen that the estimate of the total restoring force is very accurate (0.34% NMSE). The residuals of the estimate also appear close to white Gaussian noise indicating that there is no dynamic structure which has not been captured. This is further confirmed considering the frequency content of the signal (frame (c)) where the spectrum of the residuals is very close to flat and lies well below the spectrum of the true forcing close to its dominant frequencies. Additionally, considering the distribution of the residuals (frame (d)) a Gaussian shape is observed which is confirmed through a one-sample Kolmogorov–Smirnov test which fails to reject the null hypothesis at 0.1% significance indicating that the samples are in fact Gaussian. Given this evidence, it is possible to conclude that the model has captured well the total restoring force present in the model, however, this reveals a point for further discussion regarding the quantity estimated by the Gaussian process.

Since it has been seen that the total restoring force in the system is well estimated in the GPLFM, the question remains of how to fit the nonlinear component of the system. As mentioned previously, there are a number of options for how to approach this task as it is close to the task in the classical restoring force surface method [2]. For example, one could fit the extracted value from the GP component with another black-box model, e.g. a neural network (as in [11]) or another GP! However, one strength of the approach presented here as opposed to methods which simulate multiple nonlinear systems, e.g. the ABC approach of [9], is that the GPLFM has estimated the latent components of the model, which are the unmeasured displacement, velocity and the (nonlinear) restoring force. Since the equation of motion is assumed in the form $m\ddot{z} + c\dot{z} + kz + \hat{f}(z, \dot{z}) = U(t)$ and the quantities for z , \dot{z} and $\hat{f}(z, \dot{z})$ have now been estimated, the form of $\hat{f}(z, \dot{z})$ can now be plotted and inspected visually, which is not possible in other approaches. Using the estimated values for z and $\hat{f}(z, \dot{z})$, the plot in Fig. 7 is produced. To learn the nonlinear system, one must fit this function which is a far easier task because it does not require running multiple nonlinear dynamic simulations. Importantly, the plot produced in Fig. 7 has been produced without any prior knowledge of the nonlinearity in the system, it is a result of the construction of the GPLFM.

At this point it is necessary to determine which model to fit to the data in Fig. 7. There are numerous approaches available, as discussed a purely black-box approach may be adopted, however, it may be at this point that engineering insight is able to add value. Having extracted the nonlinearity and in combination with insight into the physical structure an engineer may be able to make a reasonable judgement about the form of the nonlinearity, e.g. “there is only a nonlinear stiffness component”. Of course, in the absence of certainty one could perform checks to confirm if the model suggested is in fact correct. This task is also greatly simplified since the function being learnt is now static. To demonstrate this a Bayesian Information Criterion (BIC) check over multiple polynomial model orders is performed on this data to confirm the visually suggested cubic shape. The BIC is calculated as:

$$BIC = P \log N - 2 \log \mathcal{L}$$

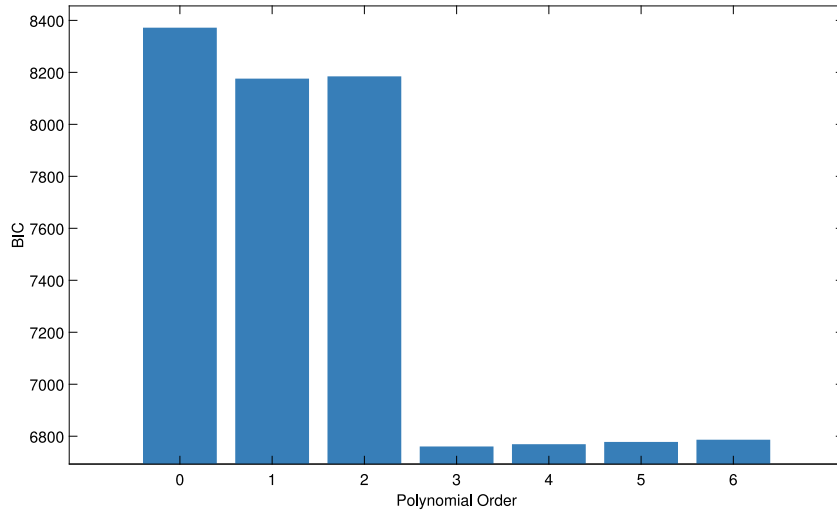


Fig. 8. BIC for increasing polynomial model orders fitting the restoring force data.

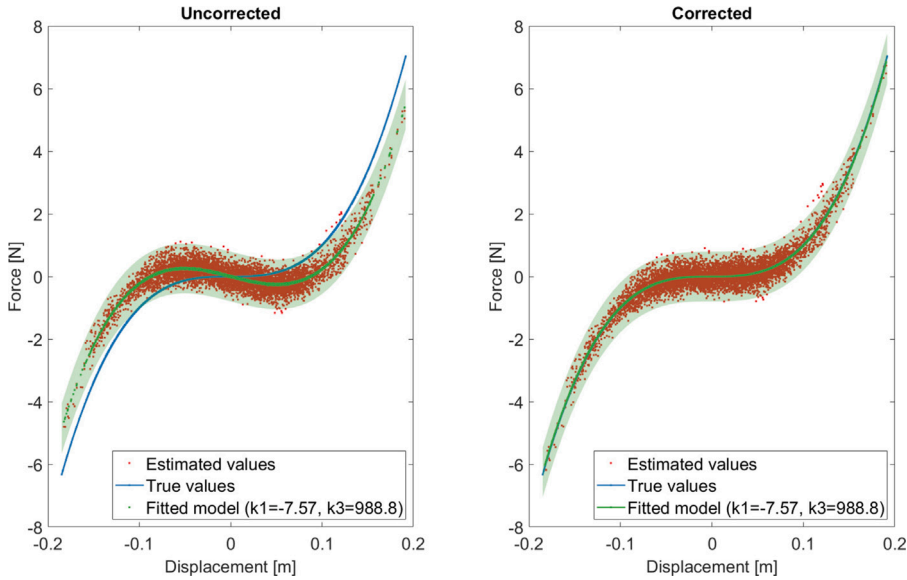


Fig. 9. Estimating the cubic nonlinear term from the estimated states, the bias in the linear parameter from the GP is removed in the right-hand frame. The shaded area indicates the 3σ confidence intervals of the Bayesian linear regression.

for P parameters fitting N data points and where the log likelihood of the fitted model is given by $\log \mathcal{L}$. Using a Bayesian linear regression the log likelihood of the model is available in closed form and the whole procedure can be conducted in under one second on a laptop computer. The lowest value of the BIC indicates the model with the “optimal” trade-off of model fit against model complexity. Taking the data extracted from the GPLFM and testing multiple polynomial models gives the results shown in Fig. 8. It can be clearly seen that the BIC check confirms that this is indeed a cubic nonlinearity present in the system and identification can proceed. One additional factor which should be considered is the choice of the kernel used for the GP. Since this choice only defines the prior over the restoring force the model should be relatively insensitive. In Appendix A the same system is identified with a Matérn 3/2 kernel and the results shown very similar performance to those with the Matérn 1/2 kernel used in the results of this section.

This procedure can now be seen in the context of the Duffing oscillator, where the nonlinear restoring force has been extracted as previously shown. Selecting the cubic model for the nonlinearity one can fit the data in Fig. 7 to obtain a nonlinear ODE which can be used in future simulations. This is done in Fig. 9.

Remembering that the contribution of the GP does not perfectly capture the contribution of only the nonlinear term, the cubic function which is fit also contains the linear coefficient. Since the total restoring force $f(z, \dot{z}) = c\dot{z} + kz + \hat{f}(z, \dot{z})$ is estimated correctly




Fig. 10. Comparison of simulation performance of the nonlinear model to that of an estimated linear model of the same system.

(see Fig. 6) one can investigate the combination of the learnt $\hat{f}(z, \dot{z})$ with the previously identified linear parameters. Since the model being fit to the restoring force data is $\hat{f}(z, \dot{z}) = \alpha z + \beta z^3$, the learnt model may be substituted into $f(z, \dot{z})$ and the effect on the estimate of the linear coefficient is seen.

$$\begin{aligned}
 f(z, \dot{z}) &= c\dot{z} + kz + \hat{f}(z, \dot{z}) \\
 &= c\dot{z} + kz + \alpha z + \beta z^3 \\
 &= c\dot{z} + (k + \alpha)z + \beta z^3
 \end{aligned} \tag{10}$$

Therefore, given that the estimate of $f(z, \dot{z})$ is correct, the fit of the cubic to $\hat{f}(z, \dot{z})$ with a linear term in allows for the correction of the bias in the linear stiffness parameter identified in the MCMC procedure. This is shown visually in Fig. 9 where the true $f(z, \dot{z}) = z^3$ term is shown in blue and the effect of correcting for the linear term is demonstrated. It is worth noting that this approach will only work if the model fit to $\hat{f}(z, \dot{z})$ contains linear components which can be combined as in Eq. (10).

It is also possible to see the benefit of the proposed approach by considering the performance of the learnt model when used to simulate the system that was identified. Given the identified models, using the expected values of the parameters, with the nonlinear component as learnt in Fig. 9, the response of the system can be simulated. This is shown in Fig. 10 for both the linear model and the nonlinear model which have been fitted. The performance of the linear model, i.e. only the dynamic model learnt in the MCMC when the GP is removed, is seen to be significantly lower than the estimated nonlinear model. When using the identified linear model, the NMSE is 74.9% a level which for almost all applications would be considered unsatisfactory. However, once the nonlinearity has been identified and included in the model used for simulation this NMSE falls to 0.2%. This represents a significant increase in quality of fit, which is to be expected. Having seen that the proposed methodology is effective on a simulated dataset, it can now be shown on an experimental benchmark.

3.2. Silverbox benchmark

The Silverbox benchmark [36] is an experimental dataset produced by an electrical system which implements a nonlinearity, similar in behaviour to the theoretical Duffing oscillator. It has been investigated previously in the literature as a benchmark dataset for nonlinear system identification, for example see [41–44]. The benchmark itself aims to replicate, electronically, the behaviour of a Duffing oscillator seen previously in Eq. (9). Although it may be the case that it does not perfectly produce such a response it is a very good approximation and serves as a challenging benchmark. Two datasets are available from the benchmark but one, referred to as the ‘arrowhead’ is the most commonly used in identification and testing. Two excitation signals are applied, first is a filtered Gaussian white noise with an increasing amplitude (40,000 samples) which forms the ‘head’ of the arrow then ten realisations of an odd random phase multisine of length 8192 samples at a fixed amplitude (see Fig. 11). See [36] for the complete details.

In keeping with most approaches seen in the literature, this work uses the ‘tail’ of the data to identify a nonlinear model of the system which is then tested on the ‘head’ of the arrow (points 1 to 40,500). The identification is based on samples from within one realisation of the odd random phase multi-sine in the tail of the arrowhead, specifically data points 49,278 to 52,350 are used in the training phase (a total of 3072 observations). In training and testing these data were upsampled by a factor of four using a cubic spline interpolation. When reporting testing errors, the simulated data are downsampled and compared against the original collected corresponding time points. A GP with a Matérn 1/2 kernel is used to estimate the unknown nonlinear restoring force as




Fig. 11. The measured ‘arrowhead’ response of the Silverbox benchmark data.




Fig. 12. Parameter posteriors from the MCMC inference over the state-space model for the Silverbox. Solid (red) lines indicate the probability density function of the prior and the histograms (blue) are built from samples of the posterior.

a function in time. Its hyperparameters σ_f^2 and ℓ are estimated along with the parameters of the linear dynamic system by means of an MCMC inference over the linear SSM.

Priors over the parameters were set on the basis of the results reported in [44]. The priors were defined as in Table 2, these priors are established on the basis of the best engineering knowledge of the system which would be available in application of the method. To explore the sensitivity of the posteriors to the choice of prior, further results are shown in Appendix B. Given this specification of the priors and using the likelihood of the linear state-space model a Metropolis–Hastings based MCMC inference was run. Ten thousand samples are drawn from the Markov Chain of which two thousand are discarded. From these samples the posterior distributions over the system parameters and the hyperparameters of the GP are recovered and shown in Fig. 12. Since this is a physical system, there are no ground truth values to be compared to but it can be seen that the posterior distributions have not shifted too far from the specified priors and that there has been a decrease in the variance of the estimates. The real test of this identification will be in the simulation of the test data shown later.

The estimates of the system parameters also allow samples of the state trajectories to be generated as before by sampling from the smoothing distribution of the LGSSM. Since only the output voltage of the system, corresponding to the displacement in a mechanical oscillator, is collected, the ground truth is only available for that state. It should be noted that this also affected the

Table 2
Prior distributions used in the Silverbox case study.

Prior	Distribution
$p(m)$	$\mathcal{N}(5.3732 \times 10^{-6}, 1 \times 10^{-5})$
$p(c)$	$\mathcal{N}(2.2653 \times 10^{-4}, 1 \times 10^{-4})$
$p(k)$	$\mathcal{N}(0.99, 0.05)$
$p(\sigma_f^2)$	$\mathcal{N}(0.005, 0.05)$
$p(\mathcal{E})$	$\mathcal{N}(0.4, 0.05)$
$p(R)$	$\mathcal{N}(1.252 \times 10^{-6}, 5 \times 10^{-6})$




Fig. 13. GPLFM state estimates for the Silverbox. Note that ground truth values are not available for dV/dt or the nonlinear restoring force state.

observation model used in the system, leading to $C = [1 \ 0 \ 0]$ and $D = 0$ when constructing the state space model. It can be seen in Fig. 13 that there is very good agreement in this state. The estimates of the voltage (output) state and the first derivate (dV/dt) states show very low levels of uncertainty in the estimates and the accuracy of the first state would lead to the reasonable assumption that dV/dt has also been well captured. This assumption is additionally supported by the results seen in the previous numerical case study. With respect to the forcing state which relates to the input voltage, the degree of uncertainty is significantly higher as was seen in the previous case. There are also periods as seen before where in regions of high displacement clear spikes in the nonlinear restoring force can be seen which may be related to increased effects of the nonlinear component in the system.

Using these samples it is possible to visualise and construct a regression which models the nonlinearity in the system. Since some knowledge regarding the nonlinearity is known *a priori* in the case of the Silverbox dataset it was possible to specify a parametric form of nonlinearity to be fit based on these samples of the states. That parametric form is one of a cubic polynomial with respect to the displacement state. It is assumed that none of the damping behaviour of the system has been transferred into the third state which relates to the GP model of the nonlinear restoring force. The samples can be drawn from the GP in time which is now the estimate of the missing contribution of the nonlinear restoring force, given the assumption of a parametric cubic nonlinearity the regression in Fig. 14 can be learnt. As before, a Bayesian linear regression is used to additionally estimate the uncertainty in the estimate. Of interest in this result is that the residuals of the model are clearly not Gaussian i.i.d., this would suggest that there remains, at least in the samples, some dynamic behaviour not captured by the chosen parametric model of the nonlinearity and this could be another potential source of bias in the parameters. A subject of further investigation is determining whether this is an artefact of the methodology proposed or a result of the approximation of the Duffing behaviour through its realisation in an electronic form.

Despite the discrepancies seen between the fitted model and the samples of the states, for a large proportion of the data the fitted model is a good representation of the behaviour. In possession of this nonlinear model it is also now possible to consider the performance of the overall identification approach through simulation of the test data. The linear parameters of the model can be taken as the MAP estimates of the posterior distributions in Fig. 12 and the parameters of the cubic model learnt from samples of the states are added to this to establish the nonlinear model, the values used in simulation are shown in Table 3. Given the description of the cubic ODE, the inputs for the ramp up at the beginning of the arrowhead are used to simulate the response.

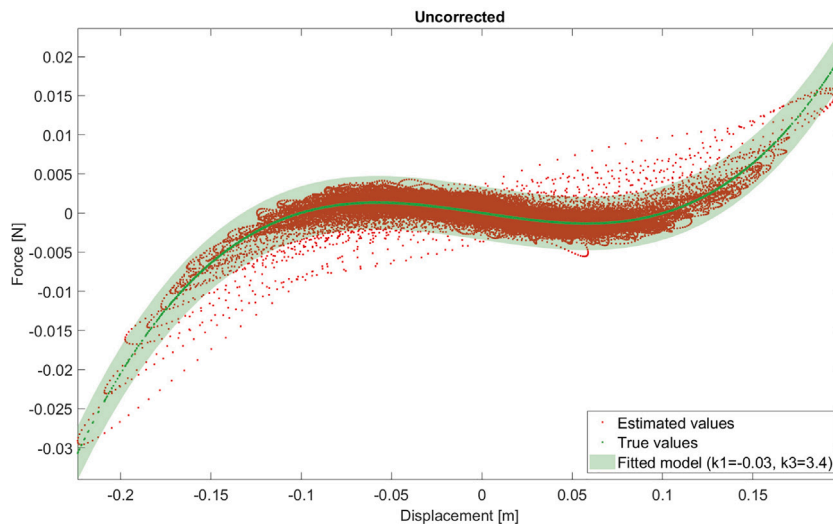


Fig. 14. Fitting of the cubic parametric nonlinearity to the samples of the nonlinear restoring force.

Table 3

MAP estimates of parameters used for test simulation.

Parameter	Value
m	5.3722×10^{-6}
c	2.1905×10^{-4}
k	0.9932
k_3	3.4239

For comparison, the simulation performance [33] of the model is compared for the learnt linear and nonlinear models in Fig. 15.⁴ Starting from assumed zero initial conditions, the models attempt to recreate the previously observed behaviour of the Silverbox with access only to the inputs to the system. Using only the learnt linear model, the results in the top frame of Fig. 15 are obtained. With the ground truth in red and simulation in blue it can be seen that the linear model does quite well especially in the initial portion of the data, roughly the first 20 s. In this lower amplitude excitation region the effect of the nonlinearity in the system is not as pronounced and the methodology adopted appears to have effectively captured the underlying linear subsystem. However, as the input amplitude has increased the error in this linear prediction grows considerably as can be seen by the large increase in the residual of the model, shown in green. The nonlinear model on the other hand (lower frame of Fig. 15) preserves very good performance at low excitation levels, in fact showing even lower residuals in the initial portion of the data, as well as retaining excellent performance as the input magnitude increases. There is still an increase in the residual magnitude as the ramp continues, especially when the level of input exceeds that seen in the training phase of the methodology. This is likely due to some bias in the identified parameters of the nonlinear model or as a result of some unmodelled nonlinearity in the system. It is reassuring, however, that very good performance is maintained by the identified nonlinear model even in ‘extrapolation’, in the sense of the input magnitude level. For some quantitative comparison the normalised mean squared errors are computed for both the linear and nonlinear model as before. The resulting errors are 18.47% for the linear model and 0.17% for the nonlinear model, the normalised error of 0.17% corresponds to a root-mean-squared-error of 0.0021 in the testing data which is competitive with results seen in the literature, e.g. Schoukens and Tóth [42].

In addition to the time series comparison of the linear and nonlinear simulation of the system it is also possible to explore the frequency content of the fits and their relative quality. This comparison is shown in Fig. 16. While both models appear to have relatively well captured the frequency content of the arrowhead, the plots of the residuals show the significant improvement which is observed when moving to the identified nonlinear model. This is consistent with the results seen in the time series and also indicates that the most significant benefit from the nonlinear model is seen near the resonance of the system which is to be expected.

In summary, the proposed approach for inference over the nonlinear restoring force in the system and the linear system parameters has proven to be very effective in identification of the Silverbox. The use of MCMC to learn the underlying linear system and generate samples of the potential nonlinear restoring force as well as the unobserved states of the system has allowed a nonlinear model of the dynamics to be learnt which was proven to perform well on an independent test set in simulation.

⁴ Note, that owing to the good performance in large regions for both models the ground truth (red) is overlapped by the simulated signal (blue), the residuals (green) make clear the performance increase from the linear to the nonlinear model.




Fig. 15. Testing portion of the 'arrowhead' signal. Ground truth is shown in red, simulated response in dashed blue and the residual in green. The final 750 points of test data are enlarged to the right hand side of the figure for more detailed comparison.




Fig. 16. Frequency domain comparison of the arrowhead fit.

4. Discussion

Having seen the performance of the method, it is now possible to make some observations and comments on the proposed approach. It will be possible to highlight the potential benefits of the GPLFM viewpoint but also present some of the potential challenges in applying this method as it currently stands. Those challenges should be seen to inform avenues for further investigation. The authors are also mindful that, while not formally proved, it is reasonable to assume that there is *no free lunch* in system identification, as is famously the case in optimisation [45]. In view of this, the authors will not (and should not) claim that the GPLFM approach presented in this work supersedes other efforts in the system identification field for all models but do argue its utility for certain classes of problem.

At the core of the method is the idea to break down the complex nonlinear system identification problem into two (hopefully) simpler tasks:

1. To identify the displacement and velocity of a system alongside some additional unknown forcing represented by the GP
2. To use these identified quantities to fit a nonlinear model in a *static* manner

It has been seen that the GPLFM, in its state-space form, provides a model with sufficient flexibility to achieve the first task. The estimates of the displacement and velocity, which can be traditionally challenging in the case of noisy data, appear well estimated in the first case study where they can be compared to the ground truth. Of note, as discussed, is the bias in the estimated linear parameters of the system in this initial stage of fitting. However, the value of this initial stage, despite the linear parameter bias, seen is that the identification problem is translated from one which is inherently dynamic to one which is now static, i.e. all the component terms in the ODE are now known and $f(z, \dot{z})$ can be fitted directly and, by including linear terms, the linear parameter bias can be corrected.

Contrasting the proposed approach to a more common nonlinear system identification approach, e.g. a deterministic optimisation of a nonlinear model, e.g. [43], or a Bayesian technique such as ABC-SMC [9], there are two advantages to the methodology in this work. Firstly, during the first stage of the identification (recovering the restoring force) the modeller is not required to suggest any candidate nonlinear system models, instead the restriction in the current work is that the system should be a single-degree-of-freedom, be representable by Eq. (5) and that the unknown contribution of the nonlinearity must be capable of being modelled by a GP in time with a stationary kernel. Addressing the first limitation is the most pressing within this framework, although the authors are hopeful that this will be the subject of future work in the near future. The main challenge is to address identifiability problems with how the latent forces enter the multiple-degree-of-freedom system so as to maintain physical interpretability. The second limitation on the method is found to be very small when considering the class of nonlinear problems encountered in mechanical engineering where it is generally assumed that systems exhibit second order behaviour (in the single-degree-of-freedom case). Finally, the requirement for a stationary kernel in the GP is also a relatively weak constraint on the problem, one point of interest for the modeller is to consider the ability of the GP to model the force if there is a discontinuity in the restoring force *time series*, bearing in mind that GP will still attempt to model the signal but will smooth the “sharp corners” in the signal.

The second advantage of the proposed GPLFM approach is in its significant computational efficiency as compared to optimisation of a nonlinear dynamic model [43], a procedure such as ABC-SMC [9] or even MCMC over a nonlinear state-space model. The source of this advantage is in the retention of a linear state-space model despite the nonlinearity in the system, where the mismatch between the linear dynamic prior (GP) and the nonlinear contribution is corrected through the posterior estimation of the states from the smoother. This means that computationally, the estimation can be carried out in *linear time* ($\mathcal{O}(T)$) as opposed to in a simulation based approach (optimisation or ABC, including its more efficient variants [46]) which requires multiple time series simulations (i.e. Runge-Kutta or similar integration approaches) or a method which required estimation of the smoothing distribution of a nonlinear state-space model. Once the states and missing contribution have been recovered in this computationally efficient manner the remaining fitting of the nonlinearity happens statically, that is without the need to simulate time series responses. For example, in the cases shown here (with the choice of a polynomial basis for the nonlinearity) the fitting of the nonlinearity reduces to linear regression. This fitting of the nonlinearity will be much faster in the static setting than the dynamic one and allows the modeller to interrogate this static relationship graphically, providing physical insight into the form of the nonlinearity before fitting begins which can guide the class of models to be tested. Even if additional engineering insight cannot be embedded at this stage, the faster fitting of this static function still enables more efficient model selection, demonstrated here was the use of BIC although more sophisticated approaches are also possible.

It is worth considering then when the approach proposed in this work may not be recommended by the authors. In its current form the framework has not yet been extended to multiple-degree-of-freedom systems, therefore, unless each degree-of-freedom can be isolated in testing the method is not applicable. Secondly, the modeller will encounter difficulty when the nonlinear term in the equation of motion is not readily fit in the second stage. This situation will occur, for example, if there is a *dynamic nonlinearity* in the system, e.g. hysteresis, as opposed to a static functional form for $f(z, \dot{z})$, e.g. polynomial in z and/or \dot{z} . If in a situation where this form of nonlinearity is expected, it may still be useful to complete the first stage in the proposed identification framework, extracting the restoring force, to inspect it and confirm that a static function is not appropriate. In future it would be an interesting avenue of investigation to consider if the extracted restoring force could be fit in the second stage of the GPLFM process by means of fitting another ODE to that signal. At that point, however, the difficulty in identification may not be significantly reduced and a one-step approach such as the ones discussed may be more appropriate, especially if a range of candidate models for the system are readily available. Despite these potential limitations of the method, the authors would propose that the presented GPLFM approach should be adopted as another available tool for the modeller which would excel when there is little *a priori* knowledge regarding the nonlinearity (only that it is static if wanting to fit the data directly in the second step) and which provides a comparatively fast, interpretable approach to nonlinear system identification.

5. Conclusion

This paper has proposed a methodology for identification of nonlinear dynamic systems when the form of the nonlinearity is not known *a priori* and when the measurements may be noisy. The identification procedure requires two steps. The first step is concerned with extracting the contribution of the nonlinear component of the system, modelling the system as a linear oscillator with an additional force acting on it, in the form of a GP in time. The second step, which has not been the focus of this paper, is to

build a model based on that contribution, although it is shown in the first case study how given the extracted data model selection techniques can be used to fit the nonlinearity.

In order to extract the action of the nonlinear component of the model, the nonlinear restoring force, the approach proposed in this work models that force as a Gaussian process in time. Through Bayesian filtering and smoothing, this GP prior is updated to give an estimate over the possible contributions from the nonlinear term. Given this estimates of the internal states of the system and of the ‘missing’ nonlinear component can be sampled and used to construct the nonlinear model. This nonlinear model need not be known *a priori* since it is fit to the samples of the nonlinear restoring force which are extracted nonparametrically from the system since they are modelled by the Gaussian process. Once in possession of those samples, the practitioner may employ a wide range of modelling techniques, from expert intuition regarding the parametric form of the nonlinearity to fully *black-box* modelling of those samples. It was discussed how the inclusion of the linear terms in the initial identification of the model can introduce bias into the true linear parameters of the ODE. It was additionally shown that the estimation of the total restoring force, linear and nonlinear, remains very accurate despite this bias. The model for the nonlinearity could then be learnt either solely from the GP component as was shown in this work or using this estimated total restoring force. When the learnt nonlinearity from the GP included a linear term related to the displacement, this could be used to correct the bias in the estimated linear parameters.

The case studies in this paper have then shown that it is possible to realise such an identification procedure. Very good results have been shown both on simulated and experimental datasets. However, a number of future avenues of research remain open. The first of these is to address more fully the second step of the identification procedure when a parametric form for the nonlinearity is not easily described or available. Some example has been given in this work where a parametric identification is performed with an number of polynomial models as candidates. BIC confirmed that a cubic stiffness model was an appropriate model in this case, however, in the future a fully nonparametric approach could also be explore or a more sophisticated model selection regime. The second future challenge which should be addressed is the extension of this methodology to a multi-degree-of-freedom case where the nonlinear restoring forces may act on more than one degree of freedom in the model. One example of this may be a two degree-of-freedom lumped mass system where the masses are connected by a nonlinear spring (e.g. cubic).

In conclusion, this work has introduced a methodology for identification of nonlinear dynamic systems. The core of the approach has been to separate the contribution of the nonlinear terms from the underlying nonlinear system by application of a Gaussian process latent force approach. The success of this method in the case studies shown would suggest that further investigation into approaches of this type could be a valuable pursuit in the future.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors are grateful for the assistance of E.J. Cross in proof reading this manuscript and for many interesting conversations related to the work. This research was supported by funds from the Centre for Oil and Gas - DTU/Danish Hydrocarbon Research and Technology Centre (DHRTC), in support of T. Friis, as a visiting student in Sheffield. T.J. Rogers is supported by the University of Sheffield, UK and the Engineering and Physical Sciences Research Council (EPSRC), UK through grant number EP/W002140/1.

Appendix A. Alternative kernel results

The same approach as seen applied to the Duffing oscillator in the main body of the text is applied to the same system with the same forcing, only varying the use of a Matérn 3/2 kernel as opposed to the 1/2 kernel used in the first study.

The results of the MCMC run over the system shown in Fig. A.1 show that the posteriors lie very close to those observed in the main body. Since the change of kernel does not affect the number of hyperparameters in the GP, the number of parameters to be learnt via the MCMC procedure remains the same.

The estimation of the states with the Matérn 3/2 kernel is shown in Fig. A.2 for the physically meaningful states of the model. The fourth state corresponding to the derivative of the force with respect to time is not shown as it is not simple to interpret its contribution to the model. As can be seen, the issues with overestimation of the linear stiffness remain and the estimate of the forcing state shows some continued underestimation. The quality of the fit in the displacement and velocity states is comparable to that seen when using the Matérn 1/2 kernel. As may be expected, in the forcing state estimated by the kernel there is a less pronounced high frequency component in the estimate. This can be attributed to the less rough nature of the Matérn 3/2 kernel.

Turning attention again to the estimation of the full restoring force, this is shown in Fig. A.3. The quality of the fit on this quantity is again very similar to that observed in the main body. If the NMSE of the estimation is considered it is slightly larger at 0.3% as opposed to 0.2% with the Matérn 1/2 kernel. However, both of these are excellent fits and it would be hard to attribute this change in NMSE to a drop in performance when changing the GP kernel.

Finally, given this identification with the Matérn 3/2 kernel, the most rigorous test of the performance is to compare the simulation on a new forcing signal as was done with the Matérn 1/2. This is shown in Fig. A.4 and a very similar pattern is observed. The nonlinear model vastly outperforms the linear model and its error is very low $< 1\%$ NMSE. Again the observed error

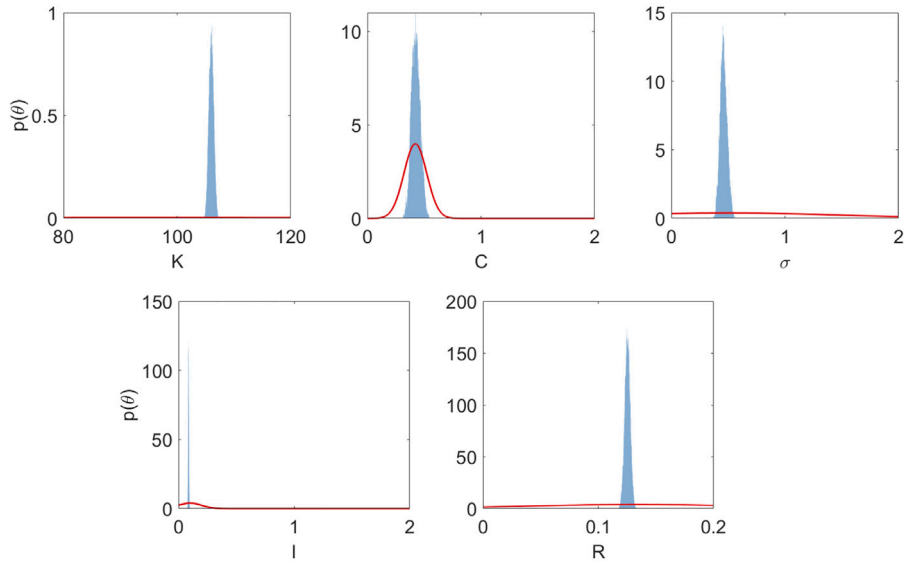


Fig. A.1. Parameter posteriors from the MCMC inference over the state-space model with a Matérn 3/2 kernel.

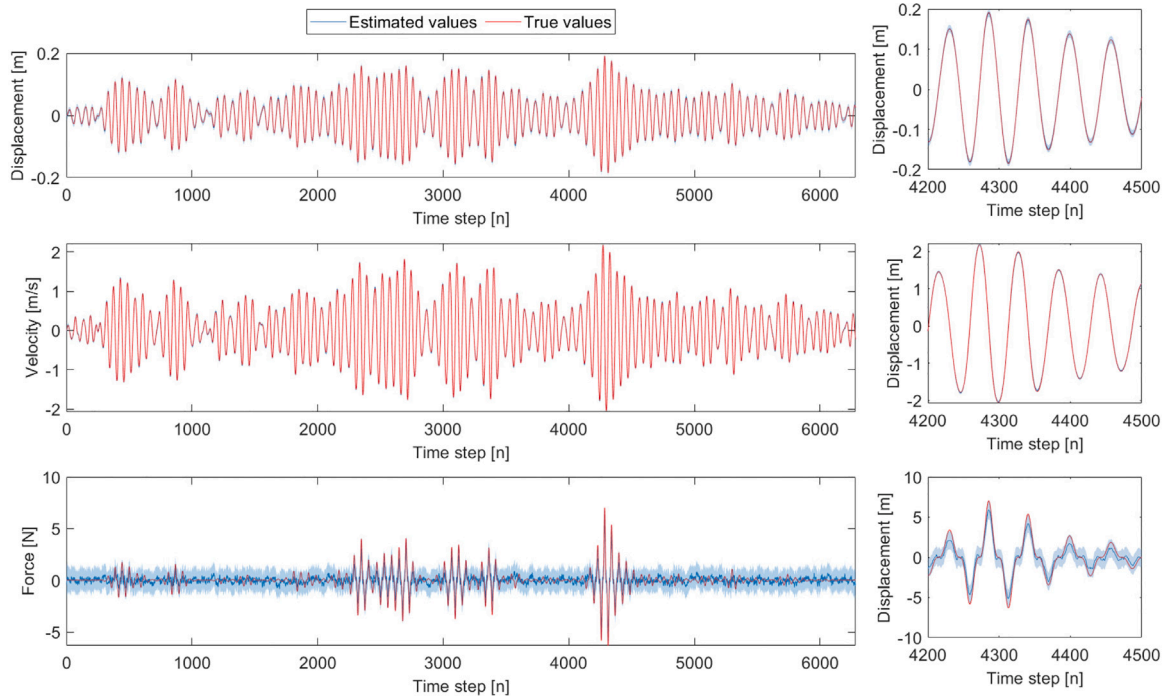


Fig. A.2. GPLFM state estimates for the simulated Duffing model, with a Matérn 3/2 kernel, on the right hand side sections are enlarged to see the quality of the estimates.

of 0.6% NMSE is slightly over that seen in the case with the Matérn 1/2 kernel but it is comparable and not sufficiently different to suggest that the kernel is playing a significant role.

In this short [Appendix](#) it has been discussed how alternative kernels may be used in the model. This was shown to have minimal effect on the performance of the model and the results were very closely aligned with those seen in the main body where a Matérn 1/2 kernel is used. This outcome follows from what would be expected given the use of a GP, that the GP imposes only a prior on the function being identified and that the inclusion of the data through filtering and smoothing diminishes the importance of the choice of kernel. This lack of sensitivity is reassuring to those wishing to apply this technique as it suggests that a full exploration

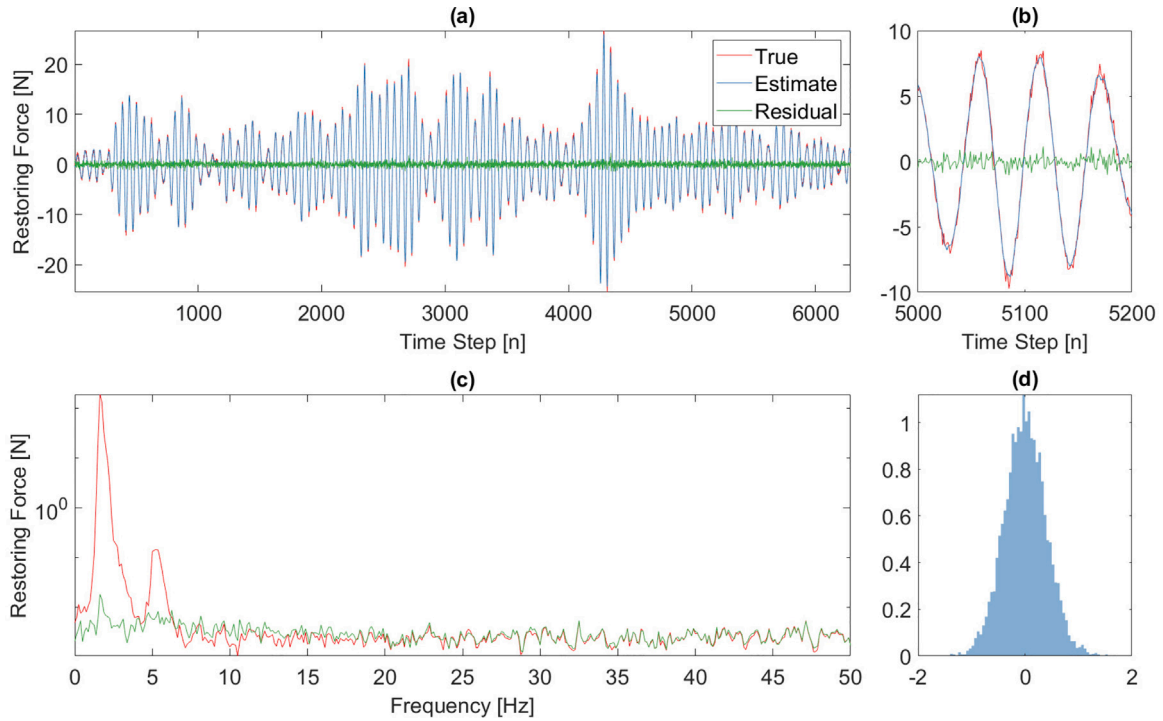


Fig. A.3. Estimation of the true restoring force in the model and its residuals when using the Matérn 3/2 kernel. True values shown in red, estimates in blue and the residual in green. Panel (a) shows the time series; (b) a zoomed in section of that time series; (c) the frequency domain view of the restoring force and its residuals; (d) the distribution of residuals.

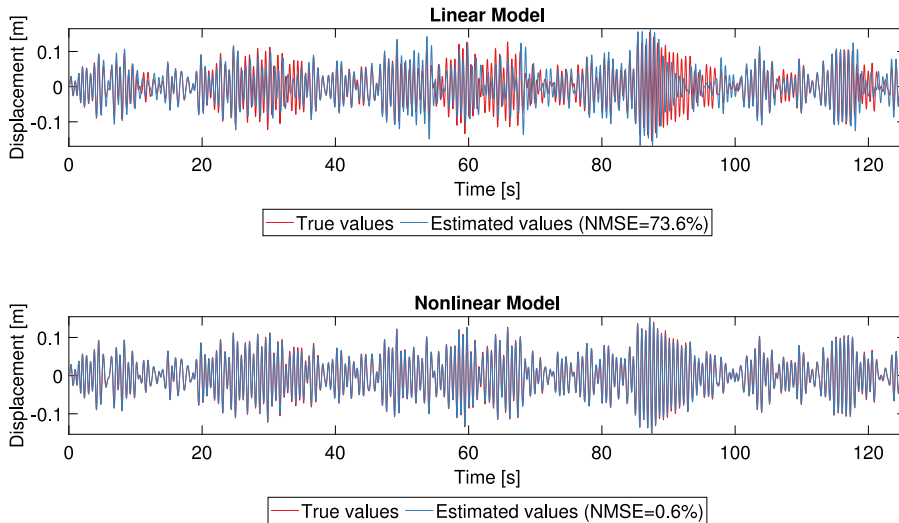


Fig. A.4. Comparison of simulation performance of the nonlinear model to that of an estimated linear model of the same system.

of all possible kernels may not be required. It has been the experience of the authors that the Matérn family, either 1/2 or 3/2, provide a good choice for a general purpose kernel in the application of LFM models.

Appendix B. Silverbox prior sensitivity

To explore the sensitivity of the solution for the silverbox system (Section 3.2) to the choice of prior, a number of different priors were run. In the main body, the prior distributions over the parameters are set using best practice, some estimation based on available engineering knowledge of the system. It is unknown if the priors are in fact a good representation of the system as the

Table B.1
Specification of different priors, all $\mathcal{N}(\cdot, \cdot)$.

Parameter		Prior mean					Prior variance
		1	2	3	4	5	All
m	($\times 10^{-6}$)	5.373	2.521	3.258	3.901	5.611	10
c	($\times 10^{-4}$)	2.190	9.719	6.571	5.630	2.477	1
k	($\times 10^0$)	0.993	1.117	0.160	4.407	9.032	0.5
σ	($\times 10^{-3}$)	5.000	4.467	7.034	2.513	10.5220	50
ℓ	($\times 10^0$)	0.400	0.435	0.427	0.105	0.787	50
R	($\times 10^{-6}$)	1.251	1.134	1.303	2.856	0.762	5




Fig. B.1. Parameter posteriors from the MCMC inference over the state-space model for the Silverbox under 5 different priors.

ground truth is unknown and the priors chosen could be biased. In this sense, they represent a realistic scenario for a Bayesian model in engineering, where partial information is available but this information may be biased or incorrect. Given this situation, there is a possibility that the estimated priors may be biased, it is worth exploring the effect this may have on the inference carried out. Deliberately, the priors were chosen to reflect a relatively large degree of uncertainty in the physical parameters of the model, see the diffuse nature of the priors shown in Fig. 12 relative to the posteriors. To test the sensitivity of the model to this prior selection the inference task was repeated with the means of those priors randomly perturbed, multiplying the mean value of each prior by the exponential of an independent standard Gaussian random variable. Table B.1 shows the values used for each prior, note that the variance (being large) is left constant throughout. For comparison, the first of these five priors is set to be the same as the priors used in the main body of this work.

Proceeding with inference as in the main-body text Section 3.2, the posteriors over the estimated parameters are shown in Fig. B.1, with each prior shown using a different colour. In this figure, the posteriors over the parameters are shown for five different randomly perturbed priors. It can be seen that the posteriors have very little variation owing to the change in the priors. The priors are not shown in these figures (since they are not clearly visible owing to their high variance compared to the posterior), for details of the priors see Table B.1. All of the priors have significantly larger variance than the posteriors showing that good information is recovered from the observed data. It is expected that not all of the posterior distributions will be identical, but it is reassuring that in the examples shown here the variation between them is not significantly affected by deliberately introducing large bias into the priors.

References

- [1] J.-P. Noël, G. Kerschen, Nonlinear system identification in structural dynamics: 10 more years of progress, *Mech. Syst. Signal Process.* 83 (2017) 2–35.
- [2] S.F. Masri, T.K. Caughey, A nonparametric identification technique for nonlinear dynamic problems, *J. Appl. Mech.* 46 (2) (1979) 433–447.
- [3] A. O'Hagan, J.F.C. Kingman, Curve fitting and optimal design for prediction, *J. R. Stat. Soc. Ser. B Stat. Methodol.* (1978) 1–42.
- [4] C.E. Rasmussen, C.K.I. Williams, *Gaussian Processes for Machine Learning*, Citeseer, 2006.
- [5] M. Alvarez, D. Luengo, N. Lawrence, Latent force models, in: *Artificial Intelligence and Statistics*, 2009, pp. 9–16.
- [6] J. Hartikainen, S. Sarkka, Sequential inference for latent force models, 2012, arXiv preprint [arXiv:1202.3730](https://arxiv.org/abs/1202.3730).

- [7] G. Kerschen, K. Worden, A.F. Vakakis, J.-C. Golinval, Nonlinear system identification in structural dynamics: current status and future directions, in: 25th International Modal Analysis Conference, Orlando, 2007, 2007.
- [8] D.A.W. Barton, Control-based continuation: bifurcation and stability analysis for physical experiments, *Mech. Syst. Signal Process.* 84 (2017) 54–64.
- [9] A.B. Abdessalem, N. Dervilis, D. Wagg, K. Worden, Model selection and parameter estimation in structural dynamics using approximate Bayesian computation, *Mech. Syst. Signal Process.* 99 (2018) 306–325.
- [10] E.N. Chatzi, A.W. Smyth, The unscented Kalman filter and particle filter methods for nonlinear structural system identification with non-collocated heterogeneous sensing, *Struct. Control Health Monit.: Off. J. Int. Assoc. Struct. Control Monit. Eur. Assoc. Control Struct.* 16 (1) (2009) 99–123.
- [11] S.F. Masri, A.G. Chassiakos, T.K. Caughey, Identification of nonlinear dynamic systems using neural networks, *J. Appl. Mech.* 60 (1) (1993) 123.
- [12] J.E. Chance, K. Worden, G.R. Tomlinson, Higher-order frequency response functions for nonlinear systems from neural networks, in: Proceedings of the 14th International Modal Analysis Conference, Vol. 2768, 1996, p. 1027.
- [13] L. Ljung, C. Andersson, K. Tiels, T. Schön, Deep learning and system identification, in: Proc IFAC Congress, Berlin, 2020.
- [14] M.A. Al-Hadid, J.R. Wright, Application of the force-state mapping approach to the identification of non-linear systems, *Mech. Syst. Signal Process.* 4 (6) (1990) 463–482.
- [15] K.S. Mohammad, K. Worden, G.R. Tomlinson, Direct parameter estimation for linear and non-linear structures, *J. Sound Vib.* 152 (3) (1992) 471–499.
- [16] G. Kerschen, K. Worden, A.F. Vakakis, J.-C. Golinval, Past, present and future of nonlinear system identification in structural dynamics, *Mech. Syst. Signal Process.* 20 (3) (2006) 505–592.
- [17] R.T. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems, 2018, pp. 6572–6583.
- [18] M. Raissi, G.E. Karniadakis, Hidden physics models: Machine learning of nonlinear partial differential equations, *J. Comput. Phys.* 357 (2018) 125–141.
- [19] A. Damianou, N.D. Lawrence, Deep Gaussian processes, in: Artificial Intelligence and Statistics, PMLR, 2013, pp. 207–215.
- [20] M.A. Alvarez, D. Luengo, N.D. Lawrence, Linear latent force models using Gaussian processes, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (11) (2013) 2693–2705.
- [21] T.J. Rogers, K. Worden, G. Manson, U.T. Tygesen, E.J. Cross, A Bayesian filtering approach to operational modal analysis with recovery of forcing signals, in: Proceedings of ISMA 2018 - International Conference on Noise and Vibration Engineering and USD 2018 - International Conference on Uncertainty in Structural Dynamics, 2018.
- [22] T.J. Rogers, K. Worden, E.J. Cross, On the application of Gaussian process latent force models for joint input-state-parameter estimation: With a view to Bayesian operational identification, *Mech. Syst. Signal Process.* 140 (2020) 106580.
- [23] R. Nayek, S. Chakraborty, S. Narasimhan, A Gaussian process latent force model for joint input-state estimation in linear structural systems, *Mech. Syst. Signal Process.* 128 (2019) 497–530.
- [24] T.J. Rogers, K. Worden, E.J. Cross, Bayesian joint input-state estimation for nonlinear systems, *Vibration* 3 (3) (2020) 281–303.
- [25] K. Worden, Data processing and experiment design for the restoring force surface method, part I: integration and differentiation of measured time data, *Mech. Syst. Signal Process.* 4 (4) (1990) 295–319.
- [26] M. Alvarez, N.D. Lawrence, Sparse convolved Gaussian processes for multi-output regression, in: Advances in Neural Information Processing Systems, 2009, pp. 57–64.
- [27] T.D. Bui, J. Yan, R.E. Turner, A unifying framework for sparse Gaussian process approximation using power expectation propagation, 2016, arXiv preprint arXiv:1605.07066.
- [28] J. Hartikainen, S. Särkkä, Kalman filtering and smoothing solutions to temporal Gaussian process regression models, in: Machine Learning for Signal Processing (MLSP), 2010 IEEE International Workshop on, IEEE, 2010, pp. 379–384.
- [29] R.E. Kalman, A new approach to linear filtering and prediction problems, *J. Basic Eng.* 82 (1) (1960) 35–45.
- [30] H.E. Rauch, C.T. Striebel, F. Tung, Maximum likelihood estimates of linear dynamic systems, *AIAA J.* 3 (8) (1965) 1445–1450.
- [31] M.L. Stein, Interpolation of Spatial Data: Some Theory for Kriging, Springer Science & Business Media, 2012.
- [32] S. Särkkä, A. Solin, Applied Stochastic Differential Equations, Vol. 10, Cambridge University Press, 2019.
- [33] L. Ljung, System Identification, Springer, 1998.
- [34] S. Särkkä, Bayesian Filtering and Smoothing, Vol. 3, Cambridge University Press, 2013.
- [35] A. Wills, T.B. Schön, F. Lindsten, B. Ninness, Estimation of linear systems using a Gibbs sampler, *IFAC Proc. Vol.* 45 (16) (2012) 203–208.
- [36] T. Wigren, J. Schoukens, Three free data sets for development and benchmarking in nonlinear system identification, in: 2013 European Control Conference (ECC), IEEE, 2013, pp. 2933–2938.
- [37] N.M. Newmark, A method of computation for structural dynamics, *J. Eng. Mech. Div.* 85 (3) (1959) 67–94.
- [38] M.K. Ochi, Ocean Waves: The Stochastic Approach, Number 6, Cambridge University Press, 2005.
- [39] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, D.B. Rubin, Bayesian Data Analysis, CRC Press, 2013.
- [40] J. Brynjarsdóttir, A. O'Hagan, Learning about physical parameters: The importance of model discrepancy, *Inverse Problems* 30 (11) (2014) 114007.
- [41] A. Marconato, J. Sjöberg, J. Suykens, J. Schoukens, Identification of the silverbox benchmark using nonlinear state-space models, *IFAC Proc. Vol.* 45 (16) (2012) 632–637.
- [42] M. Schoukens, R. Tóth, From nonlinear identification to linear parameter varying models: Benchmark examples, *IFAC-PapersOnLine* 51 (15) (2018) 419–424.
- [43] K. Worden, R.J. Barthorpe, E.J. Cross, N. Dervilis, G.R. Holmes, G. Manson, T.J. Rogers, On evolutionary system identification with applications to nonlinear benchmarks, *Mech. Syst. Signal Process.* 112 (2018) 194–232.
- [44] J. Kocijan, Parameter estimation of a nonlinear benchmark system, *Sci. Eng. Educ.* 1 (2018) 3–10.
- [45] D.H. Wolpert, W.G. Macready, No free lunch theorems for optimization, *IEEE Trans. Evol. Comput.* 1 (1) (1997) 67–82.
- [46] A.B. Abdessalem, N. Dervilis, D. Wagg, K. Worden, Model selection and parameter estimation of dynamical systems using a novel variant of approximate Bayesian computation, *Mech. Syst. Signal Process.* 122 (2019) 364–386.