



This is a repository copy of *BINDy: Bayesian identification of nonlinear dynamics with reversible-jump Markov-chain Monte Carlo*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/230272/>

Version: Published Version

Article:

Champneys, M.D. orcid.org/0000-0002-3037-7584 and Rogers, T.J. orcid.org/0000-0002-3433-3247 (2025) BINDy: Bayesian identification of nonlinear dynamics with reversible-jump Markov-chain Monte Carlo. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 481 (2319). 20240620. ISSN 1364-5021

<https://doi.org/10.1098/rspa.2024.0620>

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/>



Research



Cite this article: Champneys MD, Rogers TJ. 2025 BINDy: Bayesian identification of nonlinear dynamics with reversible-jump Markov-chain Monte Carlo. *Proc. R. Soc. A* **481**: 20240620.
<https://doi.org/10.1098/rspa.2024.0620>

Received: 15 August 2024

Accepted: 27 January 2025

Subject Areas:

applied mathematics, differential equations

Keywords:

Bayesian statistics, nonlinear dynamics, reversible-jump, Markov-chain Monte Carlo

Author for correspondence:

Max D. Champneys

e-mail: max.champneys@gmail.com

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.7724895>.

BINDy: Bayesian identification of nonlinear dynamics with reversible-jump Markov-chain Monte Carlo

Max D. Champneys and Timothy J. Rogers

Dynamics Research Group, University of Sheffield, Mappin St, Sheffield, UK

MDC, 0000-0002-3037-7584; TJR, 0000-0002-3433-3247

Model parsimony is an important *cognitive bias* in data-driven modelling that aids interpretability and helps to prevent over-fitting. Sparse identification of nonlinear dynamics (SINDy) methods are able to learn sparse representations of complex dynamics directly from data, given a basis of library functions. In this work, a novel Bayesian treatment of dictionary learning system identification, as an alternative to SINDy, is envisaged. The proposed method—Bayesian identification of nonlinear dynamics (BINDy)—is distinct from previous approaches in that it targets the full joint posterior distribution over both the terms in the library and their parametrization in the model. This formulation confers the advantage that an arbitrary prior may be placed over the model structure to produce models that are sparse in the model space rather than in parameter space. Because this posterior is defined over parameter vectors that can change in dimension, the inference cannot be performed by standard techniques. Instead, a Gibbs sampler based on reversible-jump Markov-chain Monte Carlo is proposed. BINDy is shown to compare favourably to ensemble SINDy in three benchmark case-studies. In particular, it is seen that the proposed method is better able to assign high probability to correct model terms.

1. Introduction

Describing the dynamics of nonlinear processes analytically is of fundamental interest to many branches of scientific modelling. In cases where governing differential equations are unavailable, the practitioner has little choice but to try and discern differential equations directly from data. This difficulty is compounded in situations whereby both the model structure (which terms should be included) and its parametrization are unknown. In high-noise and low-data regimes, model uncertainty quantification (UQ) becomes critical. Limited data, measurement noise and an unknown model structure all contribute significant uncertainty.

An important facet of data-driven modelling is model parsimony. Recently, there has been increasing interest in methods that are able to develop parsimonious models directly from measured data. Methods such as sparse identification of nonlinear dynamics (SINDy) [1] have been extremely effective at learning sparse representations of complex dynamics. However, a key feature of all SINDy approaches is the specification of heuristic sparsity-inducing methods and hyperparameters that control the inevitable trade-off between fidelity and parsimony in data-driven modelling.

In this paper, the authors take the view that developing parsimonious models from data is inherently an uncertain task. Parsimonious model selection in practice requires interpretable answers to questions in the vein of ‘what is the probability that a given term should be included given the data?’ and ‘if one were to include this term, how is its value distributed?’. In order to access principled answers to the above, authors propose an alternative to SINDy within a Bayesian framework.

Bayesian uncertainty quantification is a natural framework for conducting UQ in data-driven modelling. Indeed, much has been written on the topic, e.g. [2,3]. Physical knowledge and inductive biases (such as parsimony) can be incorporated as *prior knowledge*, making such modelling assumptions explicit. Indeed, many literature contributions propose Bayesian methods for UQ in both model parametrization [4–6] and model structure [7–10]. Further discussion of specific related works is presented in §3(b).

(a) Contribution

In this work, interpretable prior distributions are placed over terms in the model (independently of their parametrization) and over parameters (independently of their inclusion). This prior structure invokes a joint posterior distribution that can be marginalized to obtain useful distributions, that can robustly address modelling questions such as those introduced above.

The key contribution of this work is a novel Bayesian approach to the identification of nonlinear dynamics from a library of basis terms. In particular,

- Heuristic sparsity-inducing regression is replaced with interpretable prior distributions over models—all modelling assumptions are made *up-front*.
- An efficient sampler is proposed to produce samples from the target joint posterior distribution over model terms and parameters.
- The proposed approach is shown to compare favourably to ensemble-SINDy in three case studies including a popular population dynamics dataset consisting of only 21 data points.

The remainder of this paper is structured as follows: The following section introduces the necessary background, summarizing the SINDy methodology, holding some discussion on parsimony in data-driven modelling motivating the use of reversible-jump Monte Carlo. A third section introduces the proposed approach and its position in the literature, providing a framework for Bayesian identification of nonlinear dynamics. A fourth section compares the proposed approach to the ensemble SINDy method in three benchmark case studies. A final section presents some discussion and directions for future work.

2. Background

(a) SINDy methods

The SINDy method, proposed in 2016 by Brunton *et al.* [1], has received considerable attention in the scientific literature as a computationally efficient way to learn differential equations directly from data. Since it was first proposed, the method has been proven to be effective in correctly identifying governing equations from both simulated and real-world data.

At their core, SINDy methods make the key assumption that the observed dynamics admit a first-order formulation that is linear in the parameters of a library of user-specified basis functions Θ . Let f denote the first-order ordinary differential equation (ODE) that describes the evolution of some state-space dynamics $x(t)$,¹

$$\dot{x}(t) = f(x(t)), \quad (2.1)$$

where the overdot represents a time derivative. Now the foundational assumption of the SINDy method can be stated as

$$\dot{x} \approx \Theta(x), \quad (2.2)$$

where x is the measured state, $\Theta(x) \in \mathbb{R}^{N \times D}$ is a library of D user-selected basis functions computed from x . In the case that x is a vector, the library of terms may also include functions of all elements in x and cross terms. Thus, it is expected that D will grow much faster than the dimension of the ODE under investigation. Critically, the coefficient vector \mathcal{E} is assumed to be *sparse*, whereby many elements of \mathcal{E} are equal to 0. In practice, it is not often the case that measurements of both x and \dot{x} are available. It is usual in the SINDy literature to assume that only the states, x , are measured, and their derivatives \dot{x} must be computed numerically (with sufficient accuracy such that (2.2) holds) [1]. The SINDy framework (for identification of ODEs) can thus be summarized by three steps:

- (1) The selection of a library of candidate basis terms.
- (2) The selection of a numerical differentiation scheme to compute \dot{x} .
- (3) Sparse regression to the coefficient vector \mathcal{E} by heuristic means.

It is clear that a great deal of methods can be applied to each of the steps above. In this way, SINDy-type methods have come to encompass a family of approaches. Since the original SINDy algorithm was introduced in [1], extensions have been proposed to partial differential equation (PDE) systems [11], implicit dynamics [12], discrete dynamics and weak PDE solutions [13]. Furthermore, many methods have been applied to the sparse regression task including sequentially thresholded least-squares [1], sparse-relaxed regression [14] and forward orthogonal least-squares regression (FROLS) [15]. Several approaches also include uncertainty quantification such as the ensemble formulation in [16] and sparsity-inducing Bayesian methods [10,17,18]. Many of these methods have been made readily available to practitioners via an open-source python library [19].

Of the three components in SINDy methods described here, this paper is concerned only with the third—determination of appropriate model terms and their parameters. At this stage, it is useful to present some discussion as to the role of parsimony in data-driven modelling.

(b) Parsimony in data-driven modelling

Inherently, model parsimony is a *cognitive bias*, injected by modellers to promote interpretability and to prevent overfitting. In many sparsity-promoting methods (such as SINDy), parsimony is enforced by proxy; the importance of each term in the library is related to the corresponding

¹The method can be extended to systems of ODEs, PDEs, etc., but for notational simplicity, the exposition in this paper will consider scalar-valued ODEs. The authors note that the extension to systems of ODEs (as seen in the case-studies of this manuscript) can be handled by treating each of the regression targets $\dot{x}_i(t)$ independently, from a shared state space comprised all x_i .

size of the parameter in \mathcal{E} . Although this is a convenient proxy for parsimony, it is not without limitations. A critical limitation in practice is that small parameters in \mathcal{E} do not necessarily correspond to small effects in the dynamics. Consider the effect of neglecting a small but negative damping term in the model leading to instability in extrapolation. Practitioners cannot know in advance how small parameters will behave in regions of the state that are not observed in the training data. Normalization of the columns of Θ can help to address this shortcoming, but can also introduce ambiguity between correlated library functions. For a motivating example, consider the problem of selecting between the terms x^2 and x^4 , $x \in [0, 1]$ in the presence of noise. If all terms in $\Theta(x)$ are normalized to unit standard deviation, the difference between the two terms might appear trivial in the regime of the training data; the difference in extrapolation is evident. Introducing correlations of this form can ultimately harm model parsimony.

A particular issue with parameter size as a proxy for term importance is the selection of a threshold value, below which model terms should be excluded. In practice, this level can be difficult to select in a principled manner, leading to a spectrum of models at different levels of sparsity. In [20], an approach based on the Akaike information criterion is used to establish a Pareto front of models trading off error against the number of identified terms. This approach is successful in that the user is offered an interpretable choice from a small subset of possible models. However, the theoretical interpretation of this Pareto front is limited and distributional estimates are not explicitly available. The method also relies on parameter size as a proxy for term importance which can lead to the issues described above.

Alternative proxies for term inclusion have been considered based on greedy reduction of some error metric. A key example is forward regression least-squares optimization (FROLS) [15,19] that attempts to greedily add (or remove) model terms that maximize an error-reduction ratio. Such approaches circumvent problems of small parameter values described above. In practice, however, these methods require the specification of a convergence threshold (or directly the number of terms in the model, as in the implementation in [19]) which can also be difficult to set in advance in a principled manner.

It is the position of the authors that data alone cannot inform model parsimony. In order to select models that are parsimonious, cognitive biases must be applied. In this work, a wider probabilistic view of parsimony in SINDy-type methods is taken. Rather than use parameter magnitude as a proxy for sparsity, the authors propose to target the *parameter inclusion probability* directly. Taking a Bayesian view of equation (2.2), and assuming an additive Gaussian noise model, one has,

$$\dot{x} = \Theta(x)\mathcal{E} + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2), \quad (2.3)$$

where σ^2 is the noise variance in the observation model. The model identification task at hand is to identify a subset of the columns of $\Theta(x)$. Intuitively, this model set includes all subsets of library functions (the powerset) and thus all possible ways that the parameter vector may be sparsified. Let

$$\mathcal{M} = \{m_i\}_{i=1}^{2^n} \quad (2.4)$$

be the set of all such possible models for $\Theta(x)$ consisting of n library functions². Key to the formulation here is that each model in $m \in \mathcal{M}$ is parametrized by a corresponding parameter vector \mathcal{E} , the dimension of which depends on the model m .

The objective of this work is to infer the joint posterior distribution $p(\mathcal{E}, m | \dot{x}, \Theta(x))$. The observation model in equation (2.3) gives rise to a Gaussian likelihood $p(\mathcal{E}, m | \dot{x}, \Theta(x))$ for each

²The method exposed in this work does not require n to be finite and can be generalized to the case of infinite libraries (for example, all polynomials), however, the exposition and results herein consider finite n only. Some additional discussion is held in §5.

model m and corresponding parameter vector \mathcal{E} . With an appropriately defined prior $p(\mathcal{E}, m)$, and applying Bayes rule, one may write,

$$p(\mathcal{E}, m|\dot{x}, \Theta(x)) = \frac{p(\mathcal{E}, m)p(\dot{x}|\mathcal{E}, m, \Theta(x))}{p(\dot{x})}. \quad (2.5)$$

For notational simplicity in the exposition that follows, the dependence of the likelihood and posterior on $\Theta(x)$ will be dropped, and the observed (or computed) state derivatives are denoted by $\mathcal{D} = \dot{x}$. The above can thus be written,

$$p(\mathcal{E}, m|\dot{x}, \mathcal{D}) = \frac{p(\mathcal{E}, m)p(\mathcal{D}|\mathcal{E}, m)}{p(\mathcal{D})}. \quad (2.6)$$

Access to the posterior distribution over both the model terms and its parametrization confers a number of advantages for the data-driven discovery of nonlinear dynamics. If samples from the posterior are available, a major advantage is that they can be marginalized (in the Monte Carlo sense) to give direct access to a posterior distribution over models, independently of their parametrization, i.e. $p(m|\mathcal{D})$. This allows the practitioner to evaluate the *probability* that a term should be included in a model (given the observed data). This permits more robust ways to introduce model parsimony, for example by excluding terms with an inclusion probability below a certain threshold. Alternatively, the full posterior over model terms and parameters can be propagated to further analyses.

The joint posterior distribution $p(\mathcal{E}, m|\dot{x}, \mathcal{D})$ over both models and their parameters is a very challenging object to approach. As is usual in complex Bayesian inference tasks, the model evidence term $p(\mathcal{D})$ is unavailable analytically. However, this is not the only obstacle. The major inferential challenge is that the *dimension* of the parameter vector \mathcal{E} has come to depend on the particular model it parametrizes. In order to overcome these difficulties, a sampling scheme based on reversible-jump Markov-chain Monte Carlo (RJMCMC) [21] will be employed.

(c) Reversible-jump Markov-chain Monte Carlo

It will be useful here to briefly review both Metropolis–Hastings (MH) and reversible-jump Markov-chain Monte Carlo theory in a general setting in order to motivate the proposed approach. Much of the exposition here is available in additional detail in references [22] for MH and [21] for reversible-jump MCMC. Consider the well-studied problem of sampling from a target distribution $\pi(\theta)$ (that is known up to some constant), by constructing a Markov chain with π as its stationary distribution. A sufficient condition for convergence is that of *detailed balance*, whereby,

$$\pi(\theta')\kappa(\theta|\theta') = \pi(\theta)\kappa(\theta'|\theta). \quad (2.7)$$

Where $\kappa(\theta'|\theta)$ is the kernel of the Markov chain moving from state θ to new location θ' . The MH algorithm further divides this kernel into a transition density $p(\theta'|\theta)$ and an acceptance probability $\alpha(\theta \rightarrow \theta')$,

$$k(\theta'|\theta) = p(\theta'|\theta)\alpha(\theta \rightarrow \theta'), \quad (2.8)$$

the familiar MH acceptance probability can be found by maximizing the acceptance probability while retaining the condition of detailed balance. The optimal choice is found to be

$$\alpha(\theta \rightarrow \theta') = \min \left\{ 1, \underbrace{\frac{\pi(\theta')}{\pi(\theta)}}_{\text{Target ratio}} \underbrace{\frac{p(\theta|\theta')}{p(\theta'|\theta)}}_{\text{Proposal ratio}} \right\}, \quad (2.9)$$

whereby an unknown normalization constant in π can be cancelled from the numerator and denominator.

In the case that θ' and θ have different dimensions, detailed balance will not hold for standard choices of the transition density. To overcome this issue, in 1995 Green introduced the reversible-jump MCMC (RJMCMC) algorithm [21] as a method to sample from distributions defined over

parameters of different dimension. The RJMCMC method overcomes this difficulty via *dimension matching*. Let k be the dimension of the state of a Markov chain at θ . In order to move to a new state with dimension k' and position θ' , an auxiliary variable u' with dimension j' is sampled. To compensate for the mismatch between dimensions, it is required that $j + k = j' + k'$, where j is the current dimension of the auxiliary variable u . Then, a bijective map $g_{k \rightarrow k'} : \{u, \theta\} \rightarrow \{u', \theta'\}$ between each pair of dimensions is defined, such that the dimension of $\{u, \theta\}$ is unchanged by the bijection. If the probability of ‘jumping’ from one dimension to another (via the appropriate corresponding bijection) is given by $J(k', \theta' | k, \theta) q(u' | \theta')$, then the RJMCMC acceptance probability can be reformulated as

$$\alpha(\theta \rightarrow \theta') = \min \left\{ 1, \underbrace{\frac{\pi(\theta')}{\pi(\theta)}}_{\text{Target ratio}} \underbrace{\frac{J(k', \theta' | k, \theta)}{J(k, \theta | k', \theta')}}_{\text{Jump ratio}} \underbrace{\frac{q(u' | \theta')}{q(u | \theta)}}_{\text{Auxiliary ratio}} \underbrace{\frac{|\partial g_{k \rightarrow k'}|}{|\partial(u, \theta)|}}_{\text{Jacobian determinant}} \right\}. \quad (2.10)$$

In the general setting, working with RJMCMC can be cumbersome for the practitioner. The specification of the bijective maps is non-trivial and has a strong effect on the sampling efficiency of the scheme. Fortunately, a significant simplification is sometimes available and will be employed here. The trick is to sample model parameters independently between models. This is achieved by setting $u \triangleq \theta'$ and $u' \triangleq \theta$ in each transition. Thus, each bijection can be defined as ‘exchanging’ the auxiliary parameters with new ones in the model move,

$$g_{k \rightarrow k'}(u, \theta) = \{u', \theta'\} = \{\theta, \theta'\}. \quad (2.11)$$

Considering the above, it is clear that the dimension matching constraint is satisfied and that the Jacobian is a row-wise re-ordering of the identity matrix. Therefore, the Jacobian determinant term in the acceptance ratio must be identically equal to one. The acceptance probability may now be simplified to

$$\alpha(\theta \rightarrow \theta') = \min \left\{ 1, \underbrace{\frac{\pi(\theta')}{\pi(\theta)}}_{\text{Target ratio}} \underbrace{\frac{J(k', \theta' | k, \theta)}{J(k, \theta | k', \theta')}}_{\text{Jump ratio}} \underbrace{\frac{q(\theta | \theta')}{q(\theta' | \theta)}}_{\text{Proposal ratio}} \right\}. \quad (2.12)$$

3. Bayesian identification of nonlinear dynamics

With the relevant RJMCMC theory established, and the inference problem formalized, attention can be returned to the problem of sampling from the posterior density $p(\mathcal{E}, m | \mathcal{D})$. Development of this posterior over both the model terms and their parametrization requires the specification of both prior and proposal densities. It is clear that there are many appropriate choices for these objects and, in practice, it can be expected that prior selection will be guided by the problem at hand. However, it can be shown that certain choices of the parameter proposal can lead to drastic simplification of the inference scheme.

Considering for now the case whereby the noise variance σ^2 in equation (2.3) is known and by substituting the required posterior distribution into the acceptance probability (equation (2.12)) above, one has

$$\alpha(\mathcal{E} \rightarrow \mathcal{E}') = \min \left\{ 1, \underbrace{\frac{p(\mathcal{E}', m' | \mathcal{D}, \sigma^2)}{p(\mathcal{E}, m | \mathcal{D}, \sigma^2)}}_{\text{Posterior ratio}} \underbrace{\frac{J(m' | m)}{J(m | m')}}_{\text{Jump ratio}} \underbrace{\frac{q(\mathcal{E} | \mathcal{E}')}{q(\mathcal{E}' | \mathcal{E})}}_{\text{Proposal ratio}} \right\}. \quad (3.1)$$

The proposal density q in equation (3.1) is freely chosen by the user. A natural choice (after Troughton & Godsill [23] and similar to the approach of Brooks *et al.* [24]) is to employ the full

conditional density,

$$q(\mathcal{E}'|\mathcal{E}) = p(\mathcal{E}'|m', \mathcal{D}, \sigma^2). \quad (3.2)$$

Substituting the proposal into [equation \(3.1\)](#), one finds that since

$$p(m|\mathcal{D}, \sigma^2) = \frac{p(\mathcal{E}, m|\mathcal{D}, \sigma^2)}{p(\mathcal{E}|m, \mathcal{D}, \sigma^2)}, \quad (3.3)$$

the acceptance ratio can be further simplified to

$$\alpha(\mathcal{E} \rightarrow \mathcal{E}') = \min \left\{ 1, \underbrace{\frac{p(m'|\mathcal{D}, \sigma^2)}{p(m|\mathcal{D}, \sigma^2)}}_{\text{Ratio of model posteriors}} \underbrace{\frac{J(m'|m)}{J(m|m')}}_{\text{Jump ratio}} \right\}. \quad (3.4)$$

This simplification is only available for particular choices of the prior $p(\mathcal{E}, m)$ that are conjugate, e.g. the Gaussian likelihood defined above (and defined independently to the prior over models). Moreover, it gives direct access to the marginal posteriors $p(m|\mathcal{D})$ (see [25] (Chapter 6) or [26] (Chapter 3) for example). Considering here a conjugate Gaussian prior at each model order (conditioned on the model choice), one has

$$p(\mathcal{E}|m) \triangleq \mathcal{N}(\mu_m^{(0)}, \Sigma_m^{(0)}). \quad (3.5)$$

Because the prior is conjugate, for a given model and noise variance, one may write,

$$p(\mathcal{E}|m, \mathcal{D}, \sigma^2) = \mathcal{N}(\mu_m, \Sigma_m), \quad (3.6)$$

where

$$\Sigma_m = \left[\frac{1}{\sigma^2} \Theta_m^T \Theta_m + \Sigma_m^{(0)-1} \right]^{-1} \quad (3.7)$$

and

$$\mu_m = \Sigma_m \Sigma_m^{(0)-1} \mu_m^{(0)} + \frac{1}{\sigma^2} \Sigma_m^T \Theta^T \dot{x} \quad (3.8)$$

are the posterior mean and covariance of the parameters for a given model m (e.g. [25,26]). The marginal posteriors are thus available as

$$\begin{aligned} p(m|\mathcal{D}, \sigma^2) &\propto p(m)p(\mathcal{D}|m, \sigma^2), \\ &= p(m)(2\pi\sigma^2)^{\frac{N}{2}} |\Sigma_m^{(0)}|^{-\frac{1}{2}} |\Sigma_m|^{\frac{1}{2}} \exp \left(-\frac{1}{2} [\sigma^{-2} \mathcal{D}^T \mathcal{D} - \mu_m^T \Sigma_m^{-1} \mu_m + [\mu_m^{(0)}]^T [\Sigma_m^{(0)}]^{-1} \mu_m^{(0)}] \right). \end{aligned} \quad (3.9)$$

Now assuming a zero-mean prior over the model parameters \mathcal{E} ,³ the overall acceptance probability of a move from \mathcal{E} to \mathcal{E}' can be written,

$$\alpha(\mathcal{E} \rightarrow \mathcal{E}') = \min \left\{ 1, \underbrace{\frac{p(m')}{p(m)} \frac{J(m'|m)}{J(m|m')}}_{\text{Model prior Jump ratio}} \underbrace{\frac{|\Sigma_{m'}^{(0)}|^{-1/2} |\Sigma_{m'}|^{1/2} \exp \left(\frac{1}{2} \mu_{m'}^T \Sigma_{m'}^{-1} \mu_{m'} \right)}{|\Sigma_m^{(0)}|^{-1/2} |\Sigma_m|^{1/2} \exp \left(\frac{1}{2} \mu_m^T \Sigma_m^{-1} \mu_m \right)}}_{\text{Model posterior ratio}} \right\}, \quad (3.10)$$

where several terms that do not depend on m have cancelled out in the model posterior ratio. The acceptance probability in [equation \(3.10\)](#) can now be readily used to sample from the required posterior. In practice, the noise variance in [equation \(2.3\)](#) will be unknown. This problem can be simply addressed by including a Gibbs move for the noise variance given the model and

³Note that in the case of non-zero mean parameter priors, an additional exponentiated quadratic form will appear in the numerator and denominator of the acceptance ratio.

the parameters, which are available following a reversible-jump move. For example, an inverse-gamma prior, conjugate to the Gaussian likelihood can be employed. The Gibbs step in this case is

$$p(\sigma^2) \triangleq \mathcal{IG}(a^{(0)}, b^{(0)}). \quad (3.11)$$

Then by the conjugacy of the prior and likelihood, the posterior over the noise variance is available exactly as an inverse-gamma distribution,

$$p(\sigma^2 | \mathcal{E}, m, \mathcal{D}) = \mathcal{IG}(a, b), \quad (3.12)$$

where

$$a = a^{(0)} + \frac{n}{2} \quad (3.13)$$

and

$$b = b^{(0)} + \frac{\sum_i^n (\dot{x} - \Theta(x)\mathcal{E})^2}{2}, \quad (3.14)$$

where n is the number of data in \dot{x} .

(a) A framework for Bayesian identification of nonlinear dynamics

It is useful to consider here how the above may be used by the practitioner in order to arrive at a useful estimate of the posterior distribution $p(\mathcal{E}, m, \sigma^2 | \dot{x}, \Theta(x))$ in the context of nonlinear system identification. In practice, there are several quantities that must be selected by the user. These are the parameter and noise variance priors, the model prior and the jump kernel. The choices of these objects will affect the cognitive biases in the inference (including the sparsity of the posterior) as well as affecting the convergence properties of the sampling scheme. Some discussion on how these objects might be selected in practice is held here.

(i) Parameter and noise variance priors, $p(\mathcal{E})$ and $p(\sigma^2)$

A constraint of the method introduced here is that these priors must be conditionally conjugate to the Gaussian likelihood model in equation (2.3). Although this is a strong constraint it remains possible to reflect many types of prior belief with these objects. For the parameter priors, vague belief can be introduced by consideration of very wide Gaussian distributions centred around zero. Stronger belief can be imposed by the consideration of more concentrated variances. For the noise variance prior, lack of belief can be imposed by the improper parametrization of the inverse gamma distribution, $\mathcal{IG}(0, 0)$ (although there are potentially limitations if samples are required to be drawn from the prior). Stronger belief can be enforced by more concentrated parametrizations.

(ii) Model space priors, $p(m)$

One of the major advantages of the formulation considered here is that prior belief can be applied directly in the model space without using parameter values as a proxy. The RJMCMC scheme employed here places no restriction on the form of the prior distribution $p(m)$. Thus, very many cognitive biases can be injected into the inference scheme by the careful consideration of this parameter. Ignorance as to the model structure and level of sparsity can be trivially introduced by the consideration of a flat model prior. For example,

$$p(m) \propto 1, \quad \forall m \in \mathcal{M}. \quad (3.15)$$

The cognitive bias of model parsimony can also be introduced. For example, one could naively choose to prefer models with fewer terms (thus promoting sparsity) by the consideration of a monotonically decreasing discrete density defined over the number of terms in the model. Appropriate choices might include the geometric distribution,

$$p(m) = (1 - \theta)^d \theta, \quad (3.16)$$

where θ is a hyperparameter that controls the extent to which more parsimonious models are preferred. Other types of belief could be applied by assigning probabilities to model term

inclusion directly. An instructive example might be in the identification of structural dynamics in a near linear regime. One could envisage a prior structure that placed a high probability of selecting expected linear terms (corresponding to inertia, viscous damping and linear elasticity), while assigning less prior probability to nonlinear terms. The effect would be that the nonlinear terms would only appear in the posterior if there was significant evidence in the likelihood. As with all prior selection in Bayesian approaches, the choice of the model-space prior must ultimately be guided by the problem at hand and domain knowledge (or lack thereof) [27]. However, the authors note that in most SINDy-type applications, practitioners expect the observed dynamics to be sparse in the library of functions. This assumption motivates the use of regularizing priors in general.

(iii) Jump kernel, $J(m'|m)$

The selection of the jump kernel is one of the key components of the RJMCMC and may strongly affect the convergence of the sampler. Here, we propose the following ‘bit-flipping’ scheme. In each proposal step a random index in $[1, n]$ (where n is the number of rows in Θ) is selected with equal probability. If that term is present in the model at the current state of the chain then it is removed, if not then it is added. If a move is rejected then a new index is drawn, again with even probability. While this approach is computationally convenient, it is likely that for certain posterior geometries there would be significant limitations. Firstly, the proposed scheme always proposes to leave the current model. For tightly concentrated posteriors (wherein most of the probability mass is concentrated in a single model), this approach might result in low acceptance rates and inefficient exploration. A second limitation may occur in situations whereby the posterior has several modes, each corresponding to highly different model structures. In this second scenario, the proposed approach may struggle to mix from one mode to another. A potential solution would be to increase the number of indices drawn, although this could exacerbate the issues with tightly concentrated posteriors.

An alternative to the ‘bit-flipping’ approach employed here could be to incorporate the idea of neighbourhood proposals [28,29]. These methods generalize the bit-flipping scheme above to balance exploration of the model space with exploitation near to high-probability models. These schemes typically consider moves from the current model to one in a nearby ‘neighbourhood’ that is adaptively varied based on data or posterior probabilities. These ideas are related to the more general idea of adaptive MCMC schemes [30] that seek to improve the efficiency of MCMC samplers by including adaptive proposal mechanisms.

(iv) BINDy algorithm

The overall sampling approach is summarized in algorithm 1. In the algorithm, it is assumed that there is a Gaussian prior over the parameters \mathcal{E} and an inverse gamma prior over the noise variance. Note that because the parameter vector does not appear in equation (3.10), there is no need to specify an initial condition for the parameter vector.

(v) Posterior interpretation

Once the practitioner has collected a sufficient set of samples (convergence of the chain can be verified by observing the trace of the parameter values in \mathcal{E} , for example), the posterior density can be examined in the usual Monte Carlo fashion by counting terms and by forming histograms over model parameters. In the context of nonlinear system identification, it is often of interest to compare the performance of the identified models in *simulation*. The simulation performance of a model is a better measure of the quality of the fit to data and as such represents a more robust challenge than prediction of state derivative from known states and library. Because the proposed approach here generates a probability distribution and not a single model, the simulation performance of the posterior can be evaluated sample-wise, by integrating samples from $p(\mathcal{E}, m, \sigma^2 | \mathcal{D})$ forward in time (with or without added noise depending on the application).

Algorithm 1. Bayesian identification of nonlinear dynamics.

Require: Parameter prior mean and variance $\mu_m^{(0)}, \Sigma_m^{(0)}$, noise variance prior parameters $a^{(0)}, b^{(0)}$, prior over models $p(m)$, jump kernel $J(m'|m)$, term library $\Theta(x)$, state derivatives \dot{x} , initial chain state (m_0, σ_0^2) .
Initialize the chain at $m = m_0$ and $\sigma^2 = \sigma_0^2$.
for Number of samples N **do**
 Propose a model move $m \rightarrow m'$ by sampling from $J(m'|m)$.
 Sample a new parameter vector \mathcal{E}' for the model m' by sampling from $p(\mathcal{E}|\mathcal{D}, \sigma^2)$ in equation (3.6).
 Accept the model and parameter move jointly with probability $\alpha(\mathcal{E} \rightarrow \mathcal{E}')$ given by equation (3.10) else remain in place.
 Propose a new noise variance σ^2 using the Gibbs update given by equation (3.12).
end for
return Samples from $p(\mathcal{E}, m, \sigma^2|\mathcal{D})$.

The overall simulation performance can then be established by considering the distribution over sample trajectories. Access to robust uncertainty quantification in simulation has important applications of system identification including active learning [16], model predictive control [31] and structural health monitoring [32].

(b) Related work

Several authors have approached the joint model structure and parametrization problem from a Bayesian perspective. A key contribution is the work of Abdessalem *et al.* [8,9] wherein an approximate Bayesian computation (ABC) method was applied to the identification of a nonlinear system from within a small number of candidate models. Although the results of the method are compelling (approximate posteriors, jointly over models and parameters), the utility of the method is limited by the extreme computational cost of ABC.

The proposed method is also not the first time that the SINDy method has been formulated in a Bayesian context. Previous attempts have made use of sparsity inducing priors [5,17,33–36]. These priors are placed over parameters in \mathcal{E} such that much of the posterior mass is concentrated at or around zero (thus removing them from the model with high probability). Although sparsity-inducing priors are shown to be very effective in correctly identifying model terms (by setting superfluous parameters to zero), the method cannot obtain a distributional estimate over the terms in the model independently of their parametrization.⁴ By contrast, the method proposed in this work differs in that prior distributions are placed over both model terms and their parameters independently. Another challenge with sparsity-inducing priors is that the inference is usually intractable and can be expensive to compute. The work in this paper is able to significantly reduce the computational burden by marginalizing the parameters \mathcal{E} from the acceptance ratio (after [23]) such that a new parameter need only be sampled when updating the noise variance.

A related work is the recent contribution of Fung *et al.* [18]. The authors propose a Bayesian formulation of the SINDy method by tracing uncertainty through the nonlinear library and the differential operator in order to account for noise on observed x . The authors then make Gaussian approximations to recover the model evidence in a Bayesian manner. The authors then propose to greedily select terms that maximally increase the model evidence until no more such terms are available. This results in a single model structure estimate with a posterior distribution over its parametrization. The method proposed in this differs in that a joint posterior distribution over both model terms and their parametrization is available explicitly, thanks to the reversible-jump

⁴Because sparsity inducing priors are usually defined continuously over the parameter range it is not possible to marginalize out the parameter values without setting a heuristic threshold on which parameter values should be considered insignificant.

sampler. The proposed approach is also not limited to greedily maximizing the model evidence, which can be important in situations whereby the posterior is multimodal in model space.

The method presented herein is also not the first use of RJMCMC in the Bayesian identification of dynamic models. Troughton & Godsill [23] and Dahlin *et al.* [37] consider an RJMCMC scheme for the identification of linear autoregressive models, and an RJMCMC scheme is employed in [38] to select between two candidate nonlinear models. Recently in [39], an RJMCMC approach is used to select terms in a state-space formulation using Kalman filtering techniques. The paper has the advantage that partial state observation is handled naturally, however, the technique is limited to linear systems only. However, to the author's best knowledge, the present work represents the first application of RJMCMC to the SINDy-type formulation of [1], and the first to explicitly enumerate a full posterior distribution over both model terms and their parametrization jointly in that context.

Another related contribution is the uncertainty quantification enabled by the ensemble-SINDy (E-SINDy) method of Fasel *et al.* [16]. E-SINDy enables heuristic UQ through bootstrapping and (robust) averaging (b(r)agging) on subsets of model terms and data. In this way, the authors are able to give heuristic estimates of quantities such as term inclusion probabilities and bootstrapped distributional estimates of model parameters. While these quantities are undoubtedly useful to practitioners, their heuristic nature makes their theoretical interpretation difficult. Nevertheless, E-SINDy is a natural benchmark comparison for the method proposed in this work. In all forthcoming case studies, both methods are demonstrated.

4. Results

In order to demonstrate the effectiveness of the proposed method, three benchmark case-studies are presented here. Results from both the proposed method and E-SINDy [16] are presented and compared in terms of their confidence in identifying the true dynamics.

The prior structure of the inference is set up initially to be weakly informative. In particular, the values,

$$\mu_m \triangleq \mathbf{0}^{(0)} \quad \text{and} \quad \Sigma_m^{(0)} \triangleq 10^3 \times I \quad (4.1)$$

and

$$a^{(0)} \triangleq 0 \quad \text{and} \quad b^{(0)} \triangleq 0, \quad (4.2)$$

are chosen. The prior distribution over models is selected initially to be flat with $p(m) \propto 1$, $\forall m \in \mathcal{M}$. The authors would note that this choice does not promote model parsimony a priori. The chain is initialized with all terms in Θ present in the model and with the noise variance set to $\sigma^2 = 1$. The Gibbs sampler is run for 6×10^3 steps in total, with the first thousand samples discarded to remove the effects of any transient 'burn-in' behaviour. The authors remark that this length of chain appears more than adequate in terms of convergence to the target distribution (see the appendices for a numerical study).

The E-SINDy results presented hereafter are generated using the publicly available 'pysindy' library [19]. In all cases, 5×10^3 models are sampled (in line with the length of the chains in the RJMCMC) and both data and library ensembling are activated, with the number of candidates dropped in each sample set to 1. In both cases, the sequentially thresholded least-squares (STLSQ) algorithm is used to perform the sparse regression. For more detail on the E-SINDy method and for interpretation of these hyperparameters, the interested reader is directed to the original article [16] and the 'pysindy' documentation [19].

(a) Legendre polynomials

The first case-study is a static sparse polynomial regression problem in which $\Theta(x)$ is static and is set to be the first 10 Legendre polynomials on the interval $x \in [-1, 1]$. In order to demonstrate the recovery of small-valued parameters, a coefficient vector is selected with both zero and near-zero elements. The coefficient vector was generated randomly by the following scheme:

- 10 random parameters were sampled uniformly on $[0,1]$.
- 4 randomly selected parameters (with replacement) were set to 0.
- 2 randomly selected parameters (with replacement) were multiplied by 0.01 to drastically reduce their magnitude.
- All parameters are rounded to 3 decimal places (so that they can be exactly reproduced in this document).

The resulting coefficient vector is given by,⁵

$$\Xi^* \triangleq [0.549, 0, 0.603, 0.545, 0.424, 0.006, 0, 0, 0, 0.004]^T. \quad (4.3)$$

The target of the regression is then generated as

$$y \triangleq \Theta(x)\Xi^* + \epsilon \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 I), \quad (4.4)$$

whereby Gaussian noise is added to the target data with standard deviation equal to 5% of the root-mean-square (RMS) of the noise-free data. The extrapolation performance of the identified models is shown in [figure 2](#).

The performance of both E-SINDy and the proposed method are compared in [figure 1](#). As can be seen in the figure, both the proposed method and E-SINDy methods are able to correctly assign low inclusion probability to the terms in Ξ^* that are equal to zero. The posterior samples generated by BINDy are tightly concentrated around the true values such that they are virtually indistinguishable from their median. However, the STLSQ algorithm is unable to identify terms with parameters lower than the threshold parameter λ , which in this case study is set to the pysindy default value of 0.1. While this is a simplified example, it serves to demonstrate a particular pathology, which may appear when using parameter value as a proxy for term importance. However, the E-SINDy method shows multi-modality, producing some samples in clusters away from the true value. This can be attributed to the library bagging operation. Every subset of the library functions produces a sample (none are rejected—there is no Metropolisization step). This has the effect that some samples are generated with important terms missing, biasing the parameter values of the remaining terms ([figure 2](#)).

(b) Lynx–hare population dynamics

The second case-study considers the well-studied lynx–hare population dynamics dataset. This dataset consists of the number of pelts collected for hares and lynxes annually between 1900 and 1920 by the Hudson Bay Company and is considered to reflect the population level of the two species [16]. As such, the dataset is often used as a canonical example of data generated by the Lotka–Volterra predator–prey population model. Let L and H be the number of lynx and hare in the population, respectively, then their population dynamics are governed by the first-order ordinary differential equation,

$$\left. \begin{aligned} \frac{dL}{dt} &= c_{11}L + c_{12}HL \\ \frac{dH}{dt} &= c_{21}H + c_{22}HL \end{aligned} \right\}, \quad (4.5)$$

and

where the c_{ij} are unknown constants in the model. The dataset represents a significant challenge in that only 21 measurements are available.

Here, the derivatives of the lynx and hare populations are computed numerically in accordance with the approach taken in [16], the columns of Θ are also normalized and a threshold value of 0.19 for λ is used. For $\Theta(x)$, the library of polynomial terms, including interactions up to third order is used. The authors note that this constitutes a more challenging identification task⁶ than the one presented in [16], where only terms up to second order are considered.

⁵For clarity and readability, the results of the case study here considers only a single parametrization of this polynomial. An additional numerical study considering many such parametrizations is presented in the appendices.

⁶The sizes of the respective model spaces increase from 2^6 to 2^{10} when considering terms up to third order.

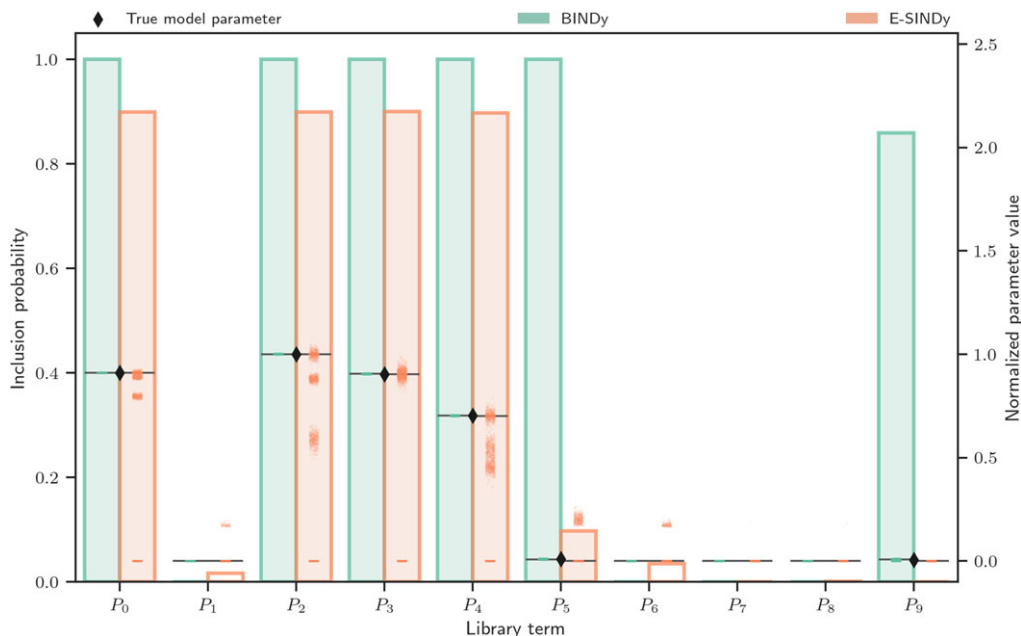


Figure 1. Comparison between RJMCMC and SINDy in quantifying model uncertainty for the toy polynomial data. In the plot, the bar-charts depict model inclusion probability (left axis), point plots depict samples of model parameters (right, log axis). Horizontal bars depict median parameter values, black diamonds depict true values of parameters in the underlying data-generating model.

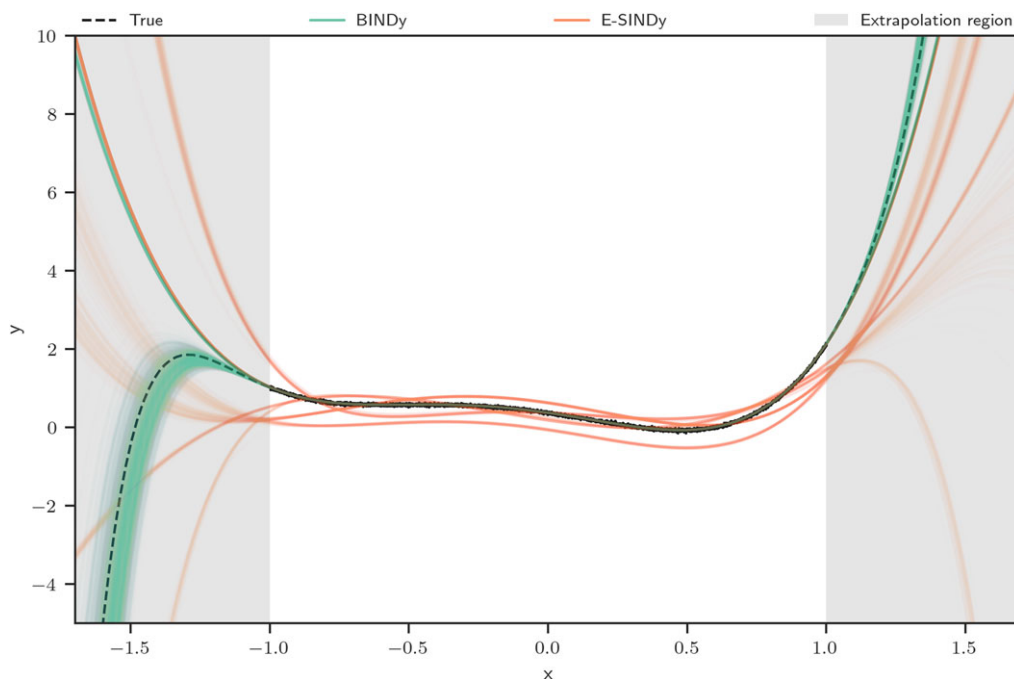


Figure 2. Samples from the identified posterior distribution (BINDy, green) and ensemble (E-SINDy, orange) in both interpolation and extrapolation regimes.

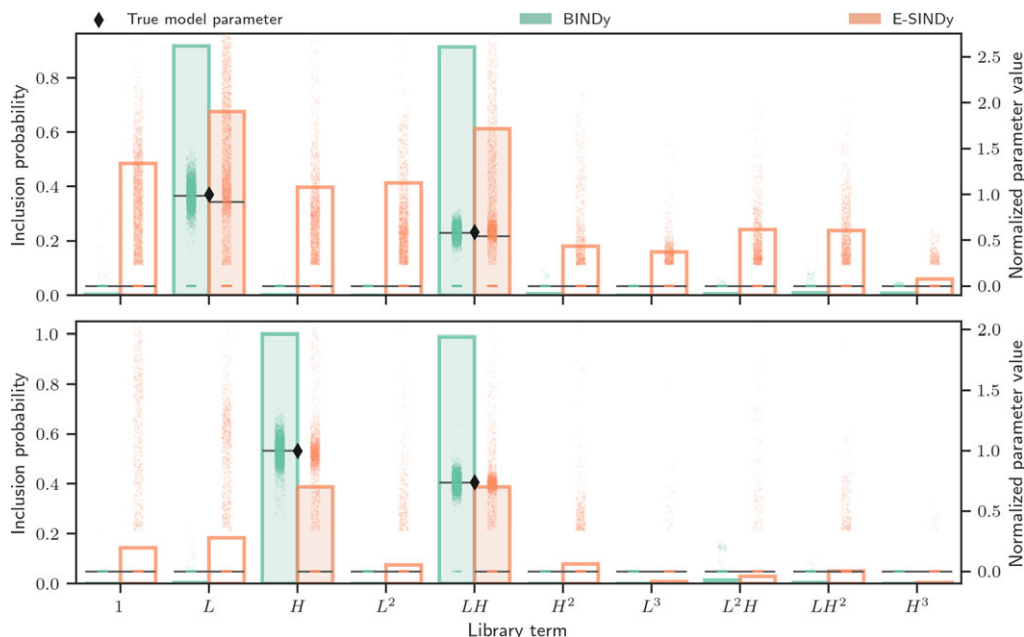


Figure 3. Comparison between RJMCMC and SINDy in quantifying model uncertainty for the lynx–hare population example. In the plot, the bar-charts depict model inclusion probability (left axis), point plots depict samples of model parameters (right, log axis). Horizontal bars depict median parameter values, black diamonds depict maximum-likelihood values of parameters in the Lotka–Volterra dynamics.

The distributions identified by the two methods on the lynx–hare dataset are depicted in figure 3. In the figure, it can be seen that both models are assigning high probability to terms in the Lotka–Volterra model. However, the probability mass is much more concentrated around these terms in the BINDy posterior, compared with those in the E-SINDy ensemble. Once again, it can be seen that the parameter values are tightly distributed around the maximum likelihood values for both methods (assuming Lotka–Volterra dynamics), although the E-SINDy method produces a number of samples close to the cutoff threshold. Particularly concerning is that the median parameter value for all terms in the hare evolution (as identified by ensemble SINDy) is zero. This corresponds to the STLSQ algorithm in E-SINDy removing all parameters and returning a zero model much of the time. These constant dynamics (with no terms in the model) identified by the SINDy ensemble are evident in figure 4 which plots time histories (integrated forward in time numerically, from known initial conditions) of samples from both BINDy and E-SINDy. By comparison, the samples from the BINDy posterior represent the dynamics well, given the scarcity of data available. In extrapolation, neither model is able to closely follow the mean dynamics (computed as the least squares estimate of a, b, c, d , assuming the model structure to be known). However, the uncertainty in the dynamics offered by the BINDy posterior is visually far more compelling, expanding in both the location and magnitude of the population peaks for each species—as might be expected. Access to meaningful uncertainty quantification for real datasets in the low-data regime is of huge importance, for example, when modelling epidemiological dynamics, wherein predicting the height and timing of peaks is critical.

(c) Lorenz attractor dynamics

The third case study considers the chaotic dynamics of the Lorenz attractor. The dynamics of the Lorenz attractor are well studied in the SINDy literature [1,16], and have become a benchmark system for advancements to the SINDy method. The dynamics are driven by the following system

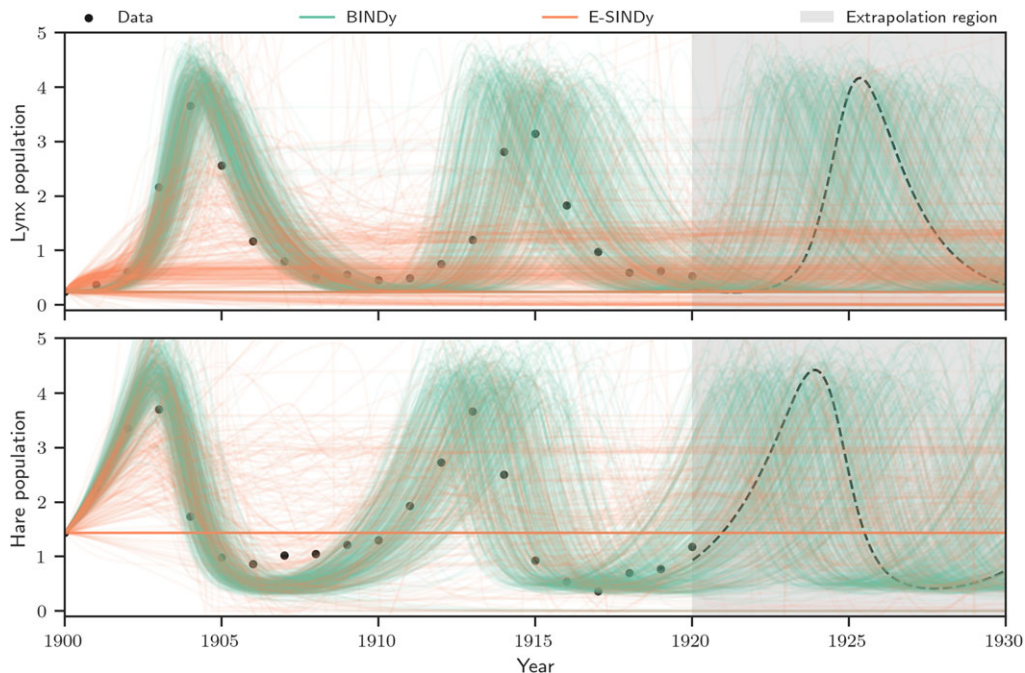


Figure 4. Samples from the identified posterior distribution (BINDy, green) and ensemble (E-SINDy, orange) in both interpolation and extrapolation regimes for the lynx–hare population example.

of ordinary differential equations:

$$\left. \begin{aligned} \frac{dx}{dt} &= 10(x_2 - x_1), \\ \frac{dy}{dt} &= x_1(28 - x_3) \\ \frac{dz}{dt} &= x_1x_2 - \frac{8}{3}x_3. \end{aligned} \right\} \quad (4.6)$$

and

For the case study in this work, the Lorenz equations above are simulated for 10 s (with a further 5 s of unseen data for extrapolation) from an initial condition of $[x_1, x_2, x_3]^T = [-8, 7, 27]^T$ at a sample rate of 10^2 Hz. Overall, 10^3 points are available for training. As before, $\Theta(x)$ is set as a polynomial library, considering terms up to third order. Before training, noise at 2.5% RMS is added to the state variables. The state derivatives are computed numerically using a polynomial smoothed finite-difference scheme with a difference order of 2, a window size of 5 and a polynomial order of 3. In order to reduce the error introduced by the numerical computation of the derivatives close to the edges of the signal, a few time points⁷ are dropped from the signal endpoints. For the SINDy results, the threshold parameter is set to 0.2 (as is done in [16]). All other parameters are unchanged from the previous case-studies.

For this third case-study, the effect of the prior over models $p(m)$ is also considered. In addition to the weakly informative prior used in the previous case-study examples, a regularizing geometric prior over the total number of non-zero elements in the model is also considered. The probability mass function of the regularizing prior is given by,

$$p(m) = (1 - \theta)^d \theta, \quad (4.7)$$

⁷Dropping $2n + 1$ points from each end for an n th-order finite difference scheme is the default used in pysindy [19] and here. After [16], a second-order finite difference scheme is used.

Table 1. Posterior results for the Lorenz attractor case-study. Term inclusion probabilities $p(\Theta_i(x)|\mathcal{D})$ and expected parameter values $E[\mathcal{E}_i|\mathcal{D}]$ (and standard deviations) for each of the three considered methods.

| equation | term | BINDy | | E-SINDy | | BINDy (regularized) | |
|---------------------------|----------|---------------------|-----------------|---------------------|----------------|---------------------|-----------------|
| | | $p(\Theta_i \cdot)$ | mean (s.d.) | $p(\Theta_i \cdot)$ | mean (s.d.) | $p(\Theta_i \cdot)$ | mean (s.d.) |
| $10(x_2 - x_1)$ | x_1 | 1.00 | −10 (0.0759) | 0.95 | −8.74 (3.88) | 1.00 | −10 (0.0742) |
| | x_2 | 1.00 | 10 (0.0655) | 0.95 | 9.33 (2.29) | 1.00 | 10 (0.0641) |
| $x_1(28 - x_3)$ | x_1 | 1.00 | 27.4 (0.314) | 0.95 | 23.9 (9.43) | 1.00 | 27.4 (0.317) |
| | x_2 | 1.00 | −0.872 (0.109) | 0.85 | 1.20 (5.82) | 1.00 | −0.875 (0.109) |
| | x_1x_3 | 1.00 | −0.985 (0.0074) | 0.94 | −0.889 (0.267) | 1.00 | −0.985 (0.0074) |
| $x_1x_2 - \frac{8}{3}x_3$ | x_3 | 1.00 | −2.66 (0.0169) | 0.93 | −2.46 (0.973) | 1.00 | −2.66 (0.0167) |
| | x_1x_2 | 1.00 | 0.995 (0.0038) | 0.94 | 0.946 (0.239) | 1.00 | 0.995 (0.0039) |

where d is the number of terms in model m , and θ is a hyperparameter, set to the strongly regularizing value of 0.99 in this work. The effect of this prior over models is to strongly penalize (by assigning low prior probability to) models with many terms (table 1).

The result of the UQ for the Lorenz equations is compared in figure 5 and table 1 for BINDy (with and without the regularizing prior) and E-SINDy. It can be seen in the figure that all three approaches have been highly successful in assigning high probability to the terms in the data-generating model. As before, sampled model trajectories (integrated forward in time) are plotted for all three methods in figure 6. The chaotic nature of the Lorenz equations makes tracking the true dynamics highly challenging. The well-known ‘lobe-switching’ behaviour of the dynamics means that small errors can quickly cause significant deviations as the dynamics become governed by a different local attractor. This effect is clearly present in the identified distributions, which quickly become multimodal after a few oscillations. Qualitatively, it is observed that both BINDy and regularized BINDy samples remain in step with the true dynamics for the longest, with E-SINDy the first to diverge. This divergence can be seen almost immediately where some orange lines deviate quickly from the observed data. Indeed, this is supported by the numerical results in table 2 which show lower error statistics for the BINDy and regularized BINDy approach. In the extrapolation region, the distributions have become very diffuse, in keeping with what might be expected for uncertain forecasting of a chaotic attractor such as the Lorenz equations.

5. Discussion

In all three case studies, strongly parsimonious models were identified by the proposed methodology. Counterintuitively, these models were assigned high posterior probability despite a flat prior over the model space. It is interesting to consider how this can be the case, for which the authors offer two explanations. On examination of the acceptance ratio in (3.10), one finds that the ratio is proportional to $|\Sigma_m^{(0)}|^{-1/2}/|\Sigma_m^{(0)}|^{-1/2}$. Intuitively, this ratio has a regularizing effect that penalizes larger models that have the same fit to the data due to the differing dimensions of the prior covariance. This effect scales with the parameter covariance meaning that wide priors (like those considered in the case studies here) have a stronger regularizing effect. Another important consideration is the correlations between the columns of $\Theta(x)$ and the state derivatives \dot{x} . These correlations depend on the data, the system under study and chosen library of basis functions. Many, highly correlated terms can complicate the identification procedure with many models able to produce predictions close to the observed data. By contrast, in the case where only a few terms have high correlation, the identification of models with few terms is simplified. This effect could also explain the parsimonious models in the first two case studies.

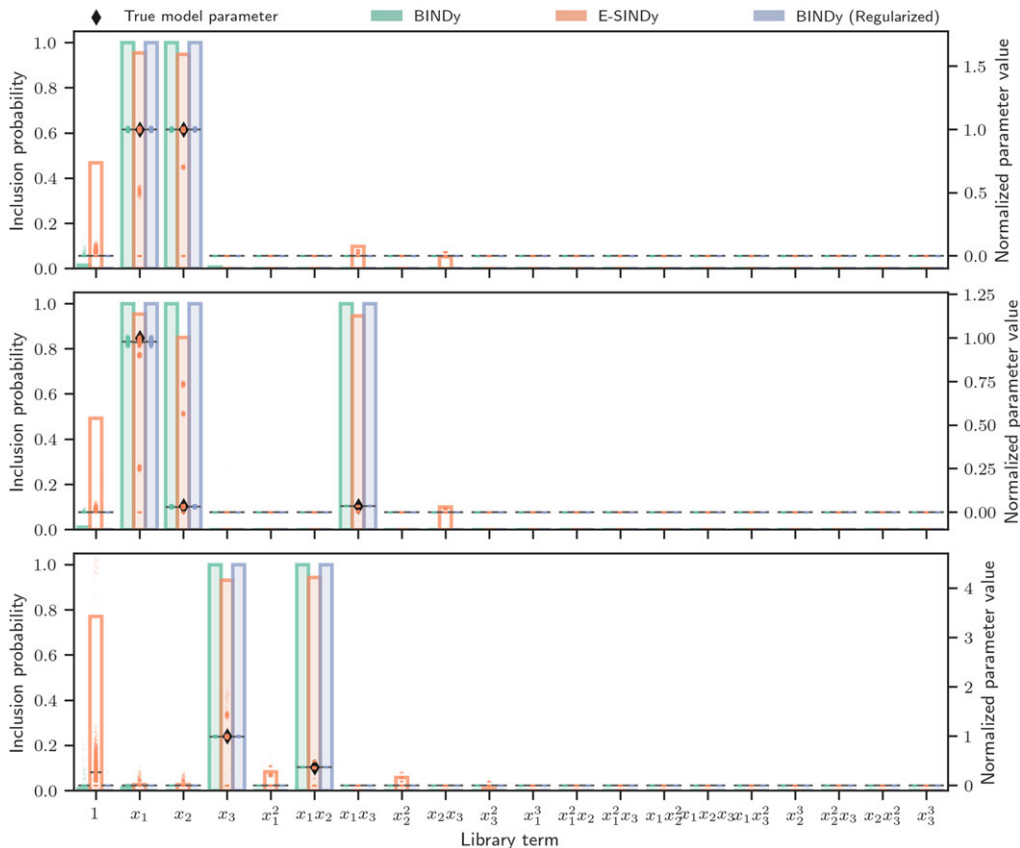


Figure 5. Comparison between RJMCMC and SINDy in quantifying model uncertainty for the Lorenz attractor example. In the plot, the bar-charts depict model inclusion probability (left axis), point plots depict samples of model parameters (right, log axis). Horizontal bars depict median parameter values, black diamonds depict true values of parameters in the underlying data-generating model.

Table 2. Median, mean and standard deviations of sample mean-squared errors computed from the trajectories in figure 5.

| equation | BINDy | | | E-SINDy | | | BINDy (regularized) | | |
|---------------------------|--------|------|------|---------|------|------|---------------------|------|------|
| | median | mean | s.d. | median | mean | s.d. | median | mean | s.d. |
| $10(x_2 - x_1)$ | 73.7 | 73.5 | 18.1 | 82.8 | 83.8 | 19 | 74.8 | 75.2 | 14.4 |
| $x_1(28 - x_3)$ | 100 | 98.1 | 22.9 | 110 | 112 | 25 | 100 | 101 | 18.2 |
| $x_1x_2 - \frac{8}{3}x_3$ | 56.4 | 62.2 | 18.5 | 79.3 | 86.5 | 36.8 | 59 | 65 | 20 |

As with any Bayesian approach, the specification of prior distributions plays an important role. A limitation of the current approach is that the priors over the parameters themselves are required to be conjugate. For most applications, this is not expected to be an issue, although there are several situations (for example bounded parameter spaces or monotonic functions) where this requirement could become restrictive. The authors cannot presently envisage a solution to this issue that does not come at considerable computational cost, should alternative, non-conjugate priors be desired.

The choice of prior over models $p(m)$ may be critical in enforcing parsimony. Here, a strongly regularizing geometric distribution parametrized by the number of terms in the model has been

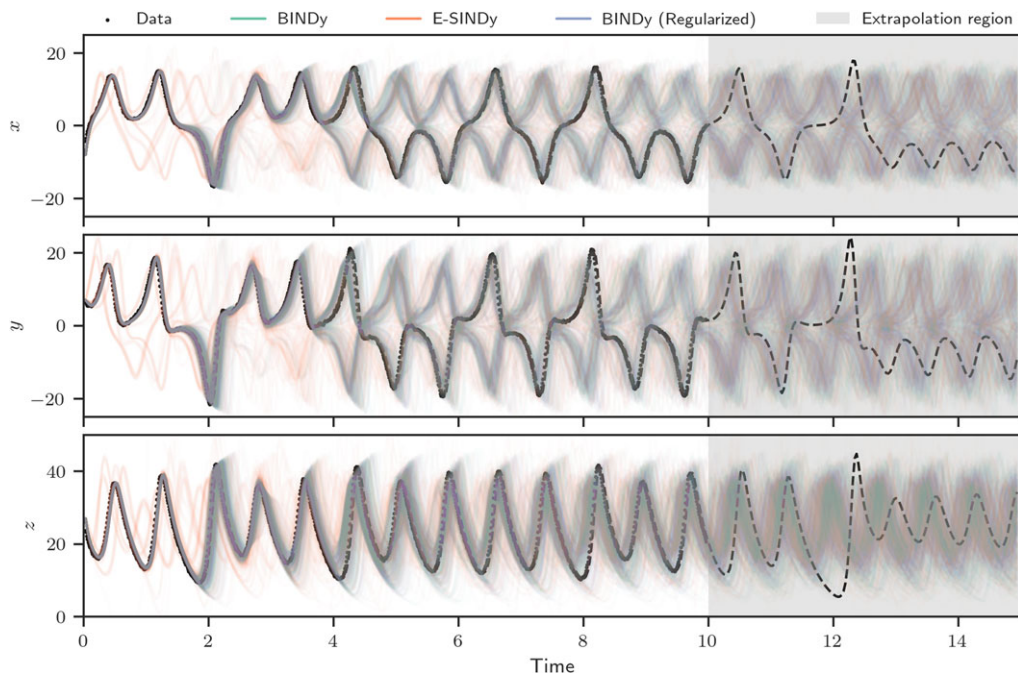


Figure 6. Samples from the identified posterior distribution (BINDy, green) and ensemble (E-SINDy, orange) in both interpolation and extrapolation regimes for the Lorenz attractor example.

applied successfully. However, vastly many more parametrizations are possible. It is expected that prior selection in practice will be driven by expert domain knowledge, wherever possible. For example, one could set the prior over each term directly, assigning high prior probability to terms that are expected to be in the model and low probabilities otherwise. Alternatively, one could promote parsimony in some types of library functions while permitting dense models in other types. The flexible nature of the Bayesian formulation presented in this work allows modellers to encode both domain knowledge and cognitive biases *a priori*. An interesting avenue for further investigation is hierarchical prior structures whereby the model-space prior is permitted to be conditional on the values its parametrization.

One could even imagine a library tempering scheme (similar in some ways to the sequential library and parameter bagging in [16]) whereby model terms with low probability are tempered out of the prior during sampling. Practically, this would enable more efficient sampling (as more model moves are likely to be accepted) as well as permit the operator to start the sampler with a vastly greater (or even infinite) library of terms, with the expectation that all but those in the typical set would be tempered out. Such a scheme would bear some resemblance to well-studied sequential Monte Carlo (SMC) samplers [40], for which a great deal of analysis has already been conducted. However, the application of SMC samplers to domains whereby parameters are permitted to change (such as in the interacting scheme of Jasra *et al.* [41]) has not been as widely studied and remains an open challenge.

A related area for future investigation is the choice of the ‘jump’ kernel $J(m'|m)$. The ‘bit-flipping’ method applied herein has proved to be effective in the case-studies considered so far, but it might be ineffective when significant multi-modality is expected in the posterior, for example, when very differently parametrized models describe the data well, low probability intermediate steps might inhibit exploration of the posterior. Instead, it is interesting to imagine more sophisticated schemes that could better suit problems of this type. Also interesting are the so-called ‘non-reversible’ jump methods [42] that use an adaptive scheme to ensure that the posterior is explored efficiently. Another possibility for the jump kernel is to include an

approach based on the adaptive neighbourhood proposals of [28,29]. For very large model libraries (whereby evaluation of a model with all the terms included may be computationally infeasible), such approaches may give practitioners a method to retain a computationally viable sampler, by considering model moves only to ‘nearby’ models.

The computational complexity of the proposed approach is of interest. The core computational cost in each iteration of the proposed algorithm is the matrix inversion required to perform the conjugate update of the proposed parameters in each iteration of the sampler. Because this cost is (naively) cubic $O(d^3 + Nd^2)$ for an RJMCMC move to a model order of d (from a model order less than d) given N observations, the worst-case analysis is that the method is $O(D^3)$ for a library of size D terms. However, because we expect many of the samples in the posterior to be sparse ($d \ll D$) the actual computational complexity (per sample) is likely to be significantly lower than this. Despite the lack of a concrete computational complexity, the authors remark that empirically the method is extremely fast. The authors are able to run the sampler at over 200 iterations per second on a single core of a Dell XPS laptop with an i7 processor.

Although the case studies in this paper have focused on systems of the first-order ODEs, it is of interest to consider applications to other types of dynamical systems. Since their initial conception, SINDy-type methods have been extended to include PDEs [11], implicit dynamics [12], discrete dynamics and weak PDE solutions [13]. Because the approach here can be thought of as a replacement for the sparse regression at the core SINDy modelling there is no reason that it cannot be applied in any of these cases. However, there are some important considerations. For example in the case on implicit dynamics, care would need to be taken to treat the ill conditioning of the implicit dynamics in a robust probabilistic manner. The authors see these areas as very interesting for future investigation. It is also interesting to imagine the application of the proposed method to systems of very many differential equations. Such high-dimensional problems pose a challenge in that the libraries under consideration would grow combinatorially with the number of degrees of freedom. In this setting, the adaptive neighbourhood proposals in [28,29] may encourage better mixing in the chain.

A key assumption in SINDy methods is that the dynamics are described as $\dot{x} \approx \Theta(x)\mathcal{E}$. Implicit in the least-squares formulation of [1] (as well as almost every other approach for learning sparse \mathcal{E}) is the Gaussian noise model in (2.3). The assumption of Gaussian noise on \dot{x} only, with x (or at least $\Theta(x)$) noise free, is unlikely to hold for many practical applications—especially in the low-data, high-noise regime. The calculation of \dot{x} numerically from noisy x is likely to only compound this issue, particularly for second-order dynamical systems where the numerical derivative must be computed twice. It is certainly possible to envisage more advanced statistical descriptions of the noise model, whereby noise entering the observed x is transformed nonlinearly by Θ and by the numerical differentiation in computing \dot{x} (such an attempt is made in [18]). However, such a formulation is likely to make the inference described in this work impractically expensive. Despite the strong assumption of the noise model in equation (2.3), the authors are encouraged by the strong performance of the proposed method in both numerical and experimental case-studies. One potential explanation for this performance is that the noise variance parameter σ^2 is freely able to inflate during the sampling. This can have the effect of ‘swallowing’ additional variance that may be present in a more accurate likelihood model. The result is a Gaussian approximation to the likelihood function (in a manner similar to [18]).

6. Conclusion

In this work, a novel method for Bayesian identification of nonlinear dynamic systems in a SINDy-like manner has been proposed, considering a joint inference over both model terms and their parametrization. It is the argument of this paper that such a posterior over models is more useful than single (or even an ensemble of) sparse models as it permits the operator to make probability-informed choices, and propagate meaningful uncertainty to further analyses.

Overall the proposed methodology has several advantages in heuristic quantification of uncertainty over ensemble-based SINDy, including prior choice over models and a valid

distributional formulation. The Bayesian formulation has been shown to be effective in three benchmark case studies, correctly assigning high posterior probability to terms in the underlying differential equations.

Data accessibility. All code, data and environment required to generate the results in this study and its appendices are available as electronic supplementary material. Some results in this manuscript were generated using the ‘pysindy’ package (Kaptanoglu *et al.*) available at: <https://joss.theoj.org/papers/10.21105/joss.03994>.

The data are provided in the electronic supplementary material [43].

Declaration of AI use. We have not used AI-assisted technologies in creating this article.

Authors’ contributions. M.D.C.: conceptualization, methodology, software, writing—original draft, writing—review and editing; T.R.: conceptualization, funding acquisition, methodology, software, supervision, writing—review and editing.

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

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

Funding. The authors of this paper gratefully acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) via grant references EP/W005816/1 and EP/W002140/1. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Appendix A. Effect of parametrization in identification of small-valued parameters

In the first case study example, a single parametrization of the Legendre polynomial was considered. Here, the effect of this parametrization is studied. Both BINDy and E-SINDy are applied to 100 parametrizations of the Legendre polynomial in equation (4.3). Parameters are sampled uniformly in the range $[0,1]$ and rounded to three decimal places (as was done in the generation of the coefficients in equation (4.3)). Of the sampled parameters, 4 were set to have zero magnitude and a further two were set to have a small magnitude by multiplication by 0.01. In order to aid in the interpretation of the results, the indices of the zero-valued and small-valued parameters were fixed to be the same as those in the first case study ($[P_1, P_6, P_7, P_8]$ and $[P_5, P_9]$, respectively). Added Gaussian noise was also re-sampled at a level of 5% RMS in every case.

The results of this investigation are plotted in figure 7. As can be seen in the figure, BINDy assigns higher probability to the correct model terms in every case. It is interesting to note that the small-valued parameters (P_5 and P_9) are only identified by BINDy at around 60% of the time on average, less confidence than was seen in the specific parametrization of the results in figure 1. Some of this drop might be explained by the fact that sampling parameters in $[0,1]$, shrinking them by a factor of 100 and then rounding to three decimal places will almost certainly produce some parameters that have zero magnitude or a magnitude so small as to be indistinguishable from noise. However, the authors would note that even this slightly reduced performance represents a significant improvement over the E-SINDy approach that fails to identify the small-valued parameters in almost every case.

Appendix B. Convergence of BINDy

The convergence of RJMCMC methods depends strongly on the choices of the prior distribution and jump kernel as well as the geometry of the posterior. Here we consider the convergence of the proposed scheme to the polynomial case study example in the main body of the paper. In total, 100 samples are drawn from the prior (as defined above) as initial conditions for the Gibbs sampler. The sampler is then run for 1000 iterations, with no samples discarded so that the convergence during the burn in period is visible. The trace of the parameter values in \mathcal{E} is plotted in figures 8 and 9.

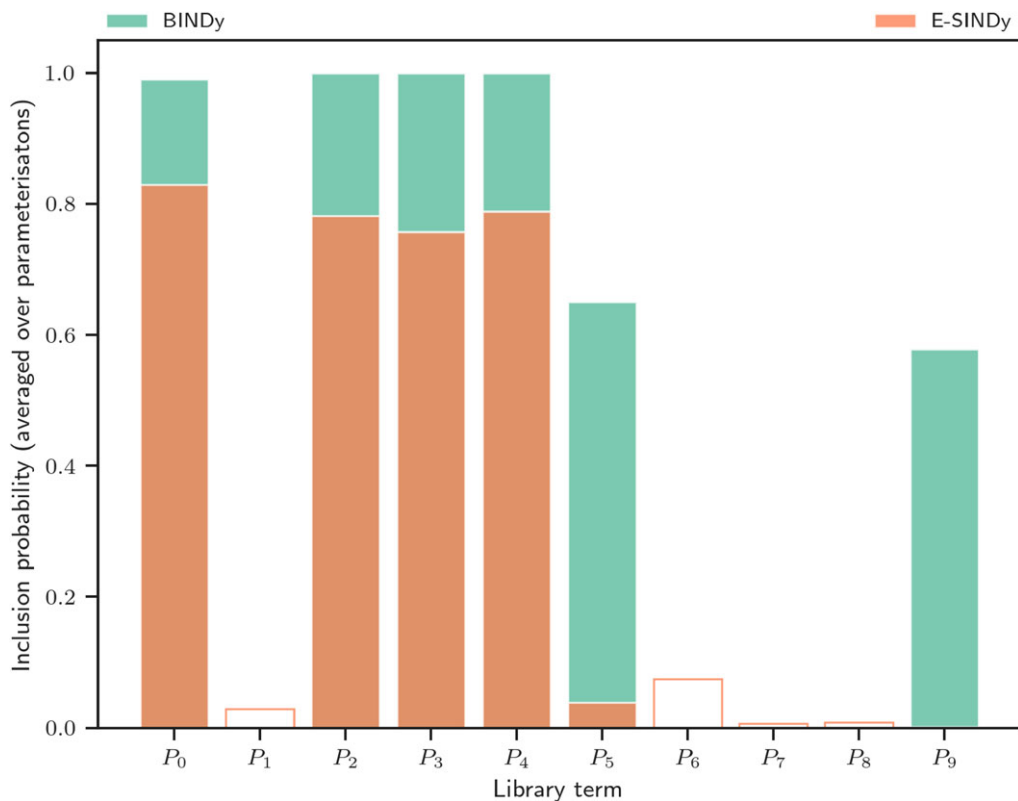


Figure 7. Posterior probabilities of BINDy and E-SINDy of selecting polynomial terms, averaged over 100 parametrizations of the Legendre polynomial in the first case study. Solid blocks correspond to terms that are truly in the data-generating model, hollow bars correspond to spurious terms.

As can be seen in figure 8, the values of the parameters (green dots) converge quickly to the true values (horizontal black lines) and remain there. Although there is increased variance for the small-valued parameters (corresponding to P_5 and P_9 in the polynomial library), it appears visually to be stationary indicating that the chain has converged. Figure 9 depicts a zoom on the first 100 iterations of the sampler. Here, it can be seen that despite high initial variance, the chain quickly converges and appears well converged after only 100 iterations.

Although the simplistic ‘bit-flipping’ jump kernel used in this study has a number of theoretical limitations, it appears that it is sufficient to produce a well-converged chain in this example.

Appendix C. Robustness of BINDy

An important concept in SINDy-type modelling is the robustness in the high-noise, low-data regime. Indeed, this is a practical consideration for nonlinear system identification approaches more generally. Most SINDy-type approaches return point estimates of the model and its parametrization, given the cognitive bias of sparsity injected into the chosen sparse regression scheme. In this context, one can interpret robustness as choosing the ‘correct’ data-generating model.

The proposed approach targets the posterior probability distribution $p(\mathcal{E}, m | \dot{x}, \Theta(x))$. This distribution explicitly incorporates sources of uncertainty such as low-data and high-noise by way of the data-likelihood model in equation (2.3). Thus, the probability of selecting the correct model is very likely to diminish in the presence of high noise or low data (as the data likelihood

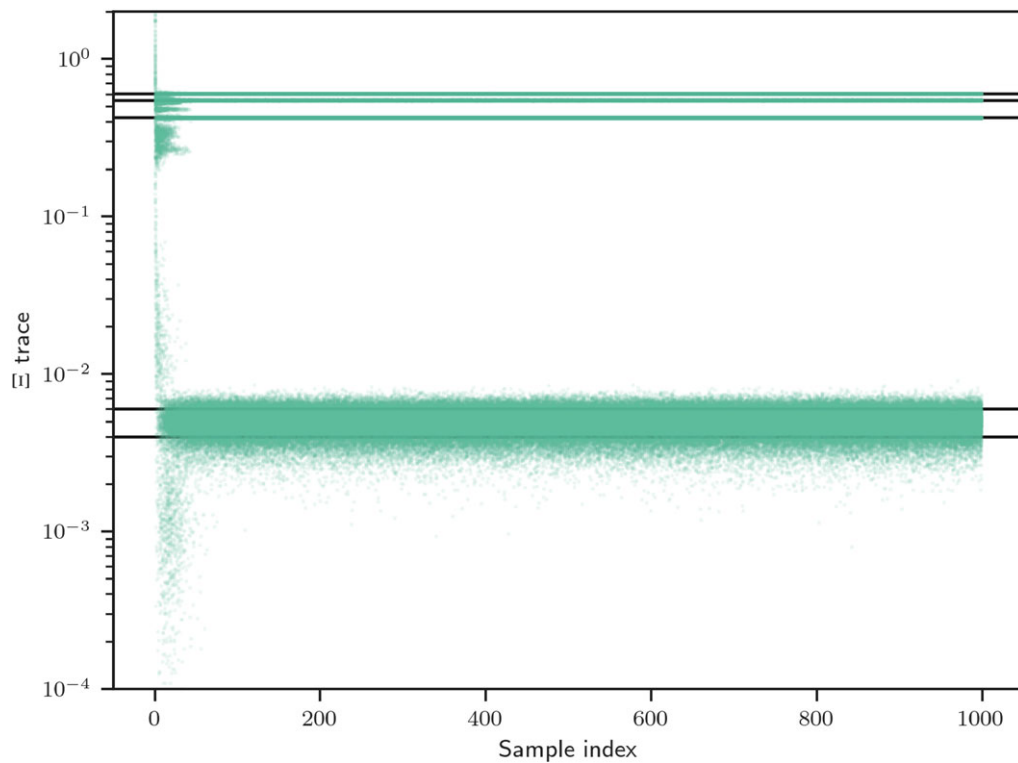


Figure 8. MCMC trace of the elements of \mathcal{E} (fitted to the polynomial case-study). Superimposed trace of 100 chains with initial conditions drawn from the prior.

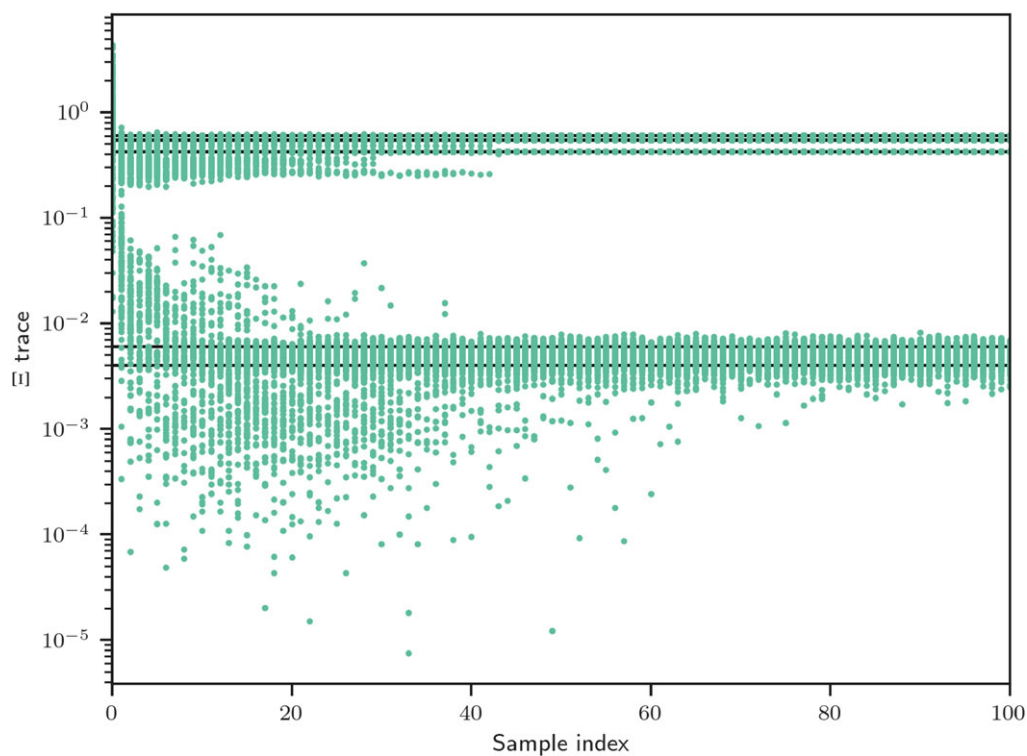


Figure 9. Zoomed view of traces of the elements of \mathcal{E} for the first 100 iterations of the sampler.

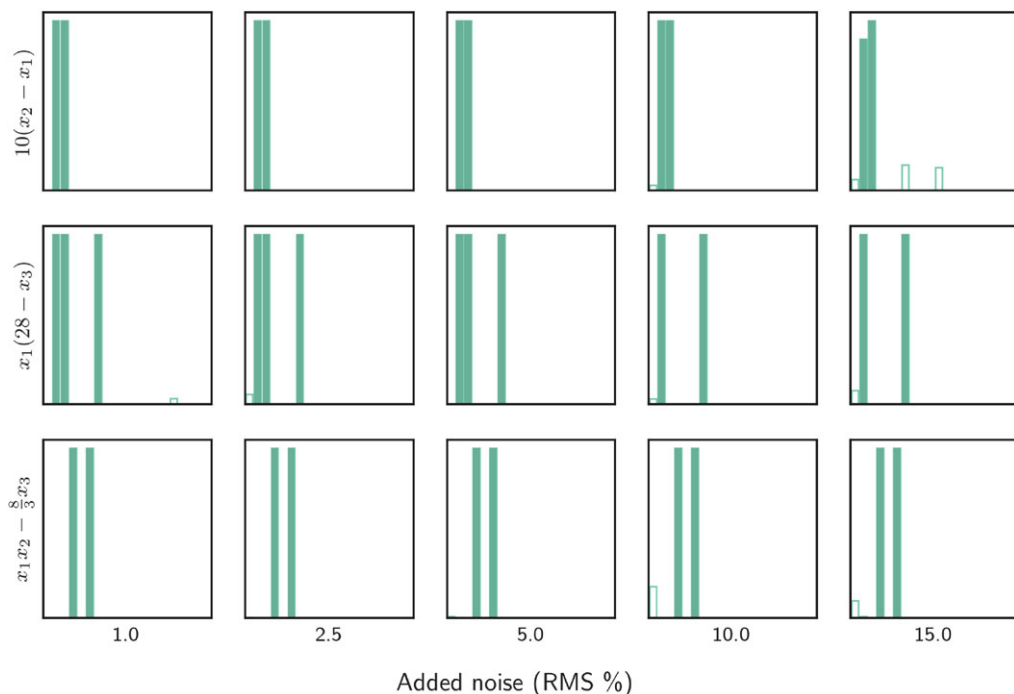


Figure 10. Posterior term inclusion probabilities at a fixed data length for various levels of added Gaussian noise. Estimated by BINDy with a flat model prior. Subplot axes are identical to those of figure 5. Solid bars indicate terms truly in the data generating model, while hollow bars indicate spurious terms.

would become more diffuse). For even very large Gaussian noise, with perfect state observation the conjugate update of the parameters is exact (for a given model order) and so the posterior would exactly reflect the uncertainty due to noise. However, in practice, large noise levels are very likely to introduce non-Gaussian disturbances to the state derivative computation (as is the case in SINDy methods generally) and so this would likely be the limiting factor in practice.

Nevertheless, it is instructive to consider the effect of low-data and high-noise regimes. In figures 10 and 11, the posterior term inclusion probabilities are plotted. BINDy is applied to the Lorenz case-study example with a flat model prior. In figure 10, the length of the observed data is held constant at 10 s (or 1000 observations) and the variance of the added noise is varied from 1% to 15%. As can be seen in the figure, BINDy is very robust to noise and assigns high posterior probability to correct terms in the data-generating model, with some spurious and missed terms assigned low probability at the highest noise levels.

In figure 11, the effect of decreasing amounts of available data is investigated. The level of added noise is held constant at 2.5% RMS and the length of the observed dynamics available for inference is varied between 1.0 s (corresponding to only 100 datapoints) and 10 s (corresponding to 1000 datapoints). The proposed BINDy approach with a flat model prior is applied to the data. Once again BINDy performs excellently, correctly assigning high probability to terms in the Lorenz equations with some spurious and missed terms in the shortest data lengths.

Figure 12 depicts the robustness of BINDy to the effect of data length and noise in a heatmap format. Each element of the heatmap considers the probability of the true data-generating model under the posterior estimated by BINDy. As can be seen in the figure, the probability approaches unity for low noise and long sequences.

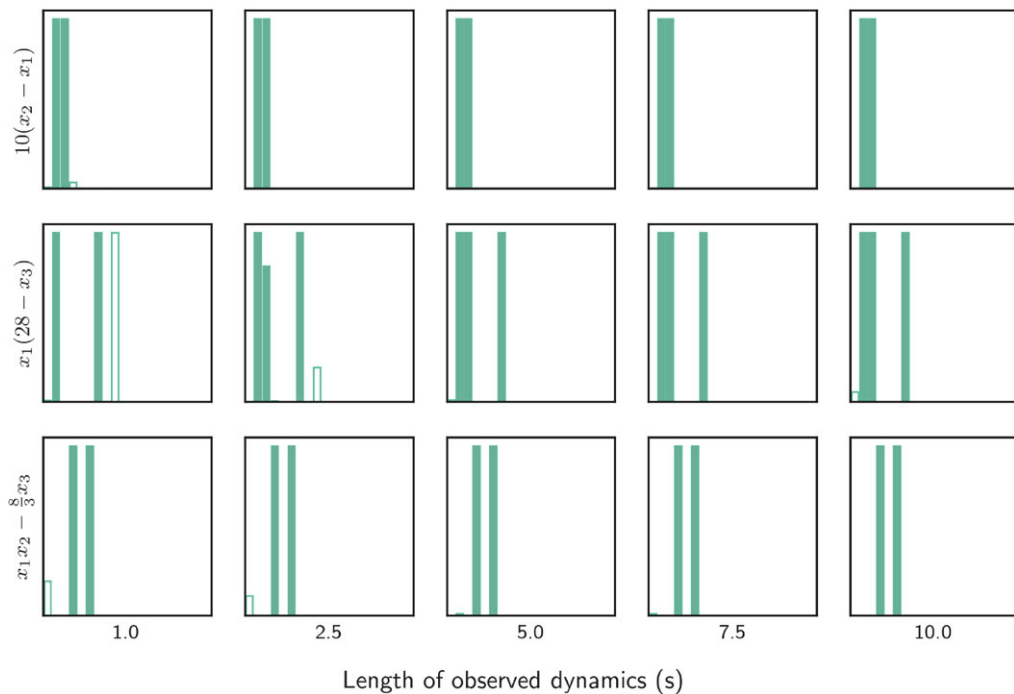


Figure 11. Posterior term inclusion probabilities at a fixed noise level for various amounts of observed dynamics. Estimated by BINDy with a flat model prior. Subplot axes are identical to those of figure 5. Solid bars indicate terms truly in the data generating model, while hollow bars indicate spurious terms.

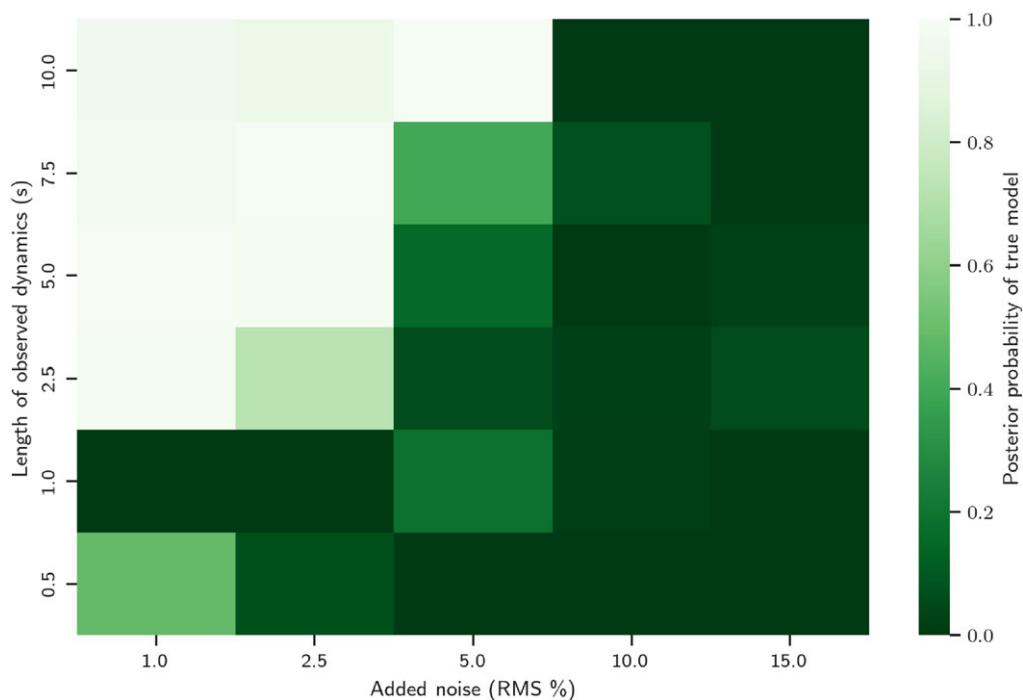


Figure 12. Heatmap depicting the posterior probability of the true model for different noise levels and data lengths.

References

1. Brunton SL, Proctor JL, Kutz JN. 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci. USA* **113**, 3932–3937. (doi:10.1073/pnas.1517384113)
2. Schmidt J, Krämer N, Hennig P. 2021 A probabilistic state space model for joint inference from differential equations and data. *Adv. Neural Inf. Process. Syst.* **34**, 12 374–12 385.
3. Cheung SH, Beck JL. 2009 Bayesian model updating using hybrid Monte Carlo simulation with application to structural dynamic models with many uncertain parameters. *J. Eng. Mech.* **135**, 243–255. (doi:10.1061/(ASCE)0733-9399(2009)135:4(243))
4. Wills A, Schön TB, Lindsten F, Ninness B. 2012 Estimation of linear systems using a Gibbs sampler. *IFAC Proc. vol.* **45**, 203–208. (doi:10.3182/20120711-3-BE-2027.00297)
5. Fuentes R, Dervilis N, Worden K, Cross EJ. 2019 Efficient parameter identification and model selection in nonlinear dynamical systems via sparse Bayesian learning. *J. Phys.: Conf. Ser.* **1264**, 012050. (doi:10.1088/1742-6596/1264/1/012050)
6. Longbottom JD, Champneys MD, Rogers TJ. 2024 Probabilistic-numeric SMC sampling for Bayesian nonlinear system identification in continuous time. (<https://arxiv.org/abs/2404.12923>)
7. Wasserman L. 2000 Bayesian model selection and model averaging. *J. Math. Psychol.* **44**, 92–107. (doi:10.1006/jmps.1999.1278)
8. Abdessalem AB, Dervilis N, Wagg D, Worden K. 2018 Model selection and parameter estimation in structural dynamics using approximate Bayesian computation. *Mech. Syst. Signal Process.* **99**, 306–325. (doi:10.1016/j.ymssp.2017.06.017)
9. Abdessalem AB, Dervilis N, Wagg D, Worden K. 2019 Model selection and parameter estimation of dynamical systems using a novel variant of approximate Bayesian computation. *Mech. Syst. Signal Process.* **122**, 364–386. (doi:10.1016/j.ymssp.2018.12.048)
10. Tripura T, Chakraborty S. 2023 A sparse Bayesian framework for discovering interpretable nonlinear stochastic dynamical systems with Gaussian white noise. *Mech. Syst. Signal Process.* **187**, 109939. (doi:10.1016/j.ymssp.2022.109939)
11. Rudy SH, Brunton SL, Proctor JL, Kutz JN. 2017 Data-driven discovery of partial differential equations. *Sci. Adv.* **3**, e1602614. (doi:10.1126/sciadv.1602614)
12. Kaheman K, Kutz JN, Brunton SL. 2020 SINDy-PI: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proc. R. Soc. A* **476**, 20200279. (doi:10.1098/rspa.2020.0279)
13. Messenger DA, Bortz DM. 2021 Weak SINDy for partial differential equations. *J. Comput. Phys.* **443**, 110525. (doi:10.1016/j.jcp.2021.110525)
14. Zheng P, Askham T, Brunton SL, Kutz JN, Aravkin AY. 2018 A unified framework for sparse relaxed regularized regression: SR3. *IEEE Access* **7**, 1404–1423. (doi:10.1109/ACCESS.2018.2886528)
15. Billings SA. 2013 *Nonlinear system identification: NARMAX methods in the time, frequency, and spatio-temporal domains*. New York, NY: John Wiley & Sons.
16. Fasel U, Kutz JN, Brunton BW, Brunton SL. 2022 Ensemble-SINDy: robust sparse model discovery in the low-data, high-noise limit, with active learning and control. *Proc. R. Soc. A* **478**, 20210904. (doi:10.1098/rspa.2021.0904)
17. Hirsh SM, Barajas-Solano DA, Kutz JN. 2022 Sparsifying priors for Bayesian uncertainty quantification in model discovery. *R. Soc. Open Sci.* **9**, 211823. (doi:10.1098/rsos.211823)
18. Fung L, Fasel U, Juniper MP. 2024 Rapid Bayesian identification of sparse nonlinear dynamics from scarce and noisy data. (<https://arxiv.org/abs/2402.15357>)
19. de Silva BM, Champion K, Quade M, Loiseau J-C, Kutz JN, Brunton SL. 2020 Pysindy: a python package for the sparse identification of nonlinear dynamics from data. (<https://arxiv.org/abs/2004.08424>)
20. Mangan NM, Kutz JN, Brunton SL, Proctor JL. 2017 Model selection for dynamical systems via sparse regression and information criteria. *Proc. R. Soc. A* **473**, 20170009. (doi:10.1098/rspa.2017.0009)
21. Green PJ. 1995 Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**, 711–732. (doi:10.1093/biomet/82.4.711)
22. Geyer CJ. 1992 Practical Markov chain Monte Carlo. *Stat. Sci.* **7**, 473–483. (doi:10.1214/ss/1177011137)

23. Troughton PT, Godsill SJ. 1997 A reversible jump sampler for autoregressive time series, employing full conditionals to achieve efficient model space moves. Department of Engineering, University of Cambridge.
24. Brooks SP, Giudici P, Roberts GO. 2003 Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions. *J. R. Stat. Soc. Ser. B: Stat. Methodol.* **65**, 3–39. (doi:10.1111/1467-9868.03711)
25. Murphy KP. 2022 *Probabilistic machine learning: an introduction*. Cambridge, MA: MIT Press.
26. Murphy KP. 2023 *Probabilistic machine learning: advanced topics*. Cambridge, MA: MIT Press.
27. O'Hagan A. 2019 Expert knowledge elicitation: subjective but scientific. *Am. Stat.* **73**, 69–81. (doi:10.1080/00031305.2018.1518265)
28. Liang X, Livingstone S, Griffin J. 2022 Adaptive random neighbourhood informed Markov chain Monte Carlo for high-dimensional Bayesian variable selection. *Stat. Comput.* **32**, 84. (doi:10.1007/s11222-022-10137-8)
29. Caron A, Liang X, Livingstone S, Griffin J. 2023 Structure learning with adaptive random neighborhood informed MCMC. (<https://arxiv.org/abs/2311.00599>)
30. Andrieu C, Thoms J. 2008 A tutorial on adaptive MCMC. *Stat. Comput.* **18**, 343–373. (doi:10.1007/s11222-008-9110-y)
31. Lew J-S, Juang J-N. 2012 Robust generalized predictive control with uncertainty quantification. *J. Guid. Control Dyn.* **35**, 930–937. (doi:10.2514/1.54510)
32. Rogers TJ. 2019 Towards Bayesian system identification: with application to SHM of offshore structures. PhD thesis, University of Sheffield.
33. Fuentes R, Nayek R, Gardner P, Dervilis N, Rogers T, Worden K, Cross E. 2021 Equation discovery for nonlinear dynamical systems: a Bayesian viewpoint. *Mech. Syst. Signal Process.* **154**, 107528. (doi:10.1016/j.ymssp.2020.107528)
34. Nayek R, Fuentes R, Worden K, Cross EJ. 2021 On spike-and-slab priors for Bayesian equation discovery of nonlinear dynamical systems via sparse linear regression. *Mech. Syst. Signal Process.* **161**, 107986. (doi:10.1016/j.ymssp.2021.107986)
35. Chen A, Du Y, Gao L, Lin G. 2021 Bayesian data-driven discovery of partial differential equations with variable coefficients. Available at SSRN 4747393.
36. More KS, Tripura T, Nayek R, Chakraborty S. 2023 A Bayesian framework for learning governing partial differential equation from data. *Physica D* **456**, 133927. (doi:10.1016/j.physd.2023.133927)
37. Dahlin J, Lindsten F, Schön T, Wills A. 2011 Robust ARX models with automatic order determination and student's *t* innovations. Technical report.
38. Tiboaca D, Green P, Barthorpe R, Worden K. 2014 Bayesian system identification of dynamical systems using reversible jump Markov chain Monte Carlo. In *Topics in modal analysis II, volume 8: Proc. of the 32nd IMAC, a Conf. and exposition on structural dynamics, 2014, Orlando, FL, 29 January–2 February 2014*. Berlin, Germany: Springer.
39. Cox B, Elvira V. 2023 Sparse Bayesian estimation of parameters in linear-Gaussian state-space models. *IEEE Trans. Signal Process.* **71**, 1922–1937. (doi:10.1109/TSP.2023.3278867)
40. Chopin N, Papaspiliopoulos O. 2020 *An introduction to sequential Monte Carlo*, vol. 4. Berlin, Germany: Springer.
41. Jasra A, Doucet A, Stephens DA, Holmes CC. 2008 Interacting sequential Monte Carlo samplers for trans-dimensional simulation. *Comput. Stat. Data Anal.* **52**, 1765–1791. (doi:10.1016/j.csda.2007.09.009)
42. Gagnon P, Doucet A. 2020 Nonreversible jump algorithms for Bayesian nested model selection. *J. Comput. Graph. Stat.* **30**, 312–323. (doi:10.1080/10618600.2020.1826955)
43. Champneys MD, Rogers TJ. 2025 BINDy: Bayesian identification of nonlinear dynamics with reversible-jump Markov-chain Monte Carlo. Figshare. (doi:10.6084/m9.figshare.c.7724895)