



Sparse identification of truncation errors

Stephan Thaler¹, Ludger Paehler^{*,1}, Nikolaus A. Adams

Institute of Aerodynamics and Fluid Mechanics, Technical University of Munich, 85748 Garching, Germany

ARTICLE INFO

Article history:

Received 3 April 2019

Received in revised form 16 July 2019

Accepted 20 July 2019

Available online 29 July 2019

Keywords:

Sparse regression

Truncation error

Modified differential equation analysis

Data-driven scientific computing

Preconditioning

ABSTRACT

This work presents a data-driven approach to the identification of spatial and temporal truncation errors for linear and nonlinear discretization schemes of Partial Differential Equations (PDEs). Motivated by the central role of truncation errors, for example in the creation of implicit Large Eddy schemes, we introduce the *Sparse Identification of Truncation Errors* (SITE) framework to automatically identify the terms of the modified differential equation from simulation data. We build on recent advances in the field of data-driven discovery and control of complex systems and combine it with classical work on modified differential equation analysis of Warming, Hyett, Lerat and Peyret. We augment a sparse regression-rooted approach with appropriate preconditioning routines to aid in the identification of the individual modified differential equation terms. The construction of such a custom algorithm pipeline allows attenuating of multicollinearity effects as well as automatic tuning of the sparse regression hyperparameters using the Bayesian information criterion (BIC). As proof of concept, we constrain the analysis to finite difference schemes and leave other numerical schemes open for future inquiry. Test cases include the linear advection equation with a forward-time, backward-space discretization, the Burgers' equation with a MacCormack predictor-corrector scheme and the Korteweg-de Vries equation with a Zabusky and Kruska discretization scheme. Based on variation studies, we derive guidelines for the selection of discretization parameters, preconditioning approaches and sparse regression algorithms. The results showcase highly accurate predictions underlining the promise of SITE for the analysis and optimization of discretization schemes, where analytic derivation of modified differential equations is infeasible.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

When constructed, modified differential equations (MDEs) provide valuable insight into the properties of every discretization scheme including spatial and temporal truncation errors. However, with increasing nonlinearity of the discretization scheme or the underlying Partial Differential Equation (PDE), the analytic approach becomes increasingly intractable. Recent advances in the data-driven learning of differential equations may now allow us to overcome the drawbacks of modified differential equation analysis (MDEA) by reformulating the discovery process as symbolic regression.

MDEA has its roots in von Neumann's stability analysis, which was developed in the 1940s with its first discussion in O' Brien et al. in 1950 [39]. Realizing the potential of von Neumann's approach, Hirt proposed a method to connect the

* Corresponding author.

E-mail addresses: stephan.thaler@web.de (S. Thaler), ludger.paehler@tum.de (L. Paehler), nikolaus.adams@tum.de (N.A. Adams).

¹ Shared first author.

stability of nonlinear difference equations with the form of the truncation error [21]. Building on Hirt's results, Warming and Hyett then established a direct connection between von Neumann's stability analysis and the symbolic form of the MDE as presented by Richtmyer and Morton [44], which is the MDEA as we understand it today. Based on this insight, they showed that the first few terms of the MDE dominate the properties of the numerical discretization. Applications of MDEA developed subsequently include increasing accuracy orders by elimination of leading order truncation error terms [26,27], enhancing stability using a nonlinear numerical viscosity term [33] and adaptive mesh refinement [26]. In light of more widespread application of MDEA, Griffiths and Sanz-Serna examined the limits of MDEA in 1986. They discovered stability criteria for MDEA and the fundamental insight that by constructing an MDE we only use a limited amount of information and the MDE can hence not fully represent the initial PDE discretization [19]. This established clear boundaries on the insight which can be derived from MDEA.

In the past decade the development of implicit Large Eddy Simulation (ILES) turbulence modeling, e.g. Adams et al. [1], led to a renewed interest in MDEA. Based on the inherent turbulence modeling capability of the truncation error as shown by Margolin [35], ILES approaches tune the discretization scheme to model the subgrid-scale stress-tensor using the truncation error. At the center of this approach is the MDEA, but when considering complex flow configurations, the construction of the MDE becomes onerous. An automatic construction of the MDE may overcome this drawback.

The recent advent of data-driven approaches to the discovery of symbolic forms may now provide the toolset to construct MDEs for previously intractable cases in an automatic fashion. Building on breakthrough results of Hod Lipson and collaborators [9,8,49] multiple approaches for the data-driven discovery of PDEs were developed. The core of our ansatz can be summarized as symbolic regression applied to MDEs. We hence constrain the introduction to the three approaches which best satisfy these applicational requirements. These ansätze are the sparse regression approaches of Kutz and Brunton, the physics-informed machine learning of Karniadakis and Dong's PDE-Net.

In 2016, Brunton et al. proposed their initial sparse regression framework called *Sparse Identification of Nonlinear Dynamics* (SINDy) [10], which is a general method for the data-driven identification of dynamical systems. Using finite differences and polynomial interpolation to approximate partial derivatives in time and space, this framework was subsequently extended to PDEs. Said algorithm is named the *PDE functional identification of nonlinear dynamics algorithm* (PDE-FIND) [46] and is explained more closely in section 2.2. Recently, SINDy has seen more generalizations with extensions to time-evolving parametric PDEs [45] and model selection using information criteria [34].

Starting with *Machine Learning of Linear Differential Equations using Gaussian Processes* [41] in 2017, Raissi, Karniadakis and collaborators began to develop a diverse array of techniques for the discovery of coefficients for symbolic terms. This line of inquiry began with a Gaussian process (GP) construction in which the unknown coefficients are recast as GP kernel hyperparameters, which can then be learned through optimization of the marginal likelihood. Notable extensions to this initial framework are the encoding of time integration schemes in the GP kernel as *Numerical GPs* [42] and a reformulation of the framework based on Neural Networks [43]. Raissi et al.'s key insight when working with neural networks was the use of the symbolic form as an additional loss function of the network, hence forcing the neural network to obey the physical structure of the system [43].

A third distinct approach are the *PDE-Nets* of Long, Lu and Dong [31,30]. Here, convolutional neural networks are constructed with partially constrained filters, which approximate the differential operators. This approach rests on breakthroughs of Cai et al. [11] and Dong et al. [16] in which a direct connection between filters and finite difference approximations of differential operators was established. The candidate filters are combined to form the unknown PDE and then predict the function value at the next time step using Forward Euler. Learning the filters by minimizing the loss function, we can then rediscover the exact form of the PDE.

The requirements of the task constrain applicable approaches to the sparse regression framework of Brunton and Kutz [46,61]. *Numerical GPs* have not seen extension to inverse problems, Neural Network approaches lack accuracy and the symbolic identification of the *PDE-Net* cannot be automated. A detailed explanation of the requirements for the ansatz can be found in section 2.2.

Building on the results of Rudy et al. [46], we present a proof of concept showing that MDEs can be identified from simulation data with high accuracy and minimal prior knowledge. Such proof of concept is intended as a stepping stone towards problems for which the analytic derivation of MDEs is intractable. Our approach, the *Sparse Identification of Truncation Errors* (SITE) framework could allow for the discovery of the MDE in these cases. In line with the vision of a recent report by the National Academy of Sciences [15], we understand the current paper as a first development towards fully data-driven MDEA tools for the analysis and optimization of truncation errors. Implicit LES modeling is a direct application where the proposed algorithm would allow an optimal utilization of the truncation errors in the construction of the subgrid-scale model. A second application are small-scale dependent shocks [28] in which the SITE approach may enable more accurate modeling via the construction of better-controlled dissipation schemes. Note that the outlined procedure is independent of the discretization method. We chose finite-difference schemes as a framework in which tractable problems can be defined and analytic MDEs in series form can be derived. Other discretization methods like the finite element method or the finite volume method can be analyzed in a similar manner using the SITE approach. For this proof of concept, we consider one-dimensional test cases only. Extension of SITE to multiple dimensions is straightforward following the extension of the PDE-FIND algorithm to multiple dimensions [46].

The paper is structured as follows. Chapter 2 gives an overview of the theory of analytic MDEA, the preliminaries of the data-driven identification framework and outlines the major challenges. Chapter 3 summarizes the proposed workflow of the

SITE approach before we validate the numerical solvers used for data generation in chapter 4. Numerical test cases in chapter 5 demonstrate the applicability of the procedure to linear and nonlinear PDEs of interest to fluid dynamics: The advection equation is discretized using a forward-time backward-space (FTBS) scheme, a MacCormack predictor-corrector scheme [32] is used for Burgers' equation and the Korteweg-de Vries equation (KdV) is discretized with the Zabusky and Kruskal scheme [57]. Hereby, we assess the applicability of several sparse regression algorithms to the problem of data-driven identification of MDEs, the impact of discretization parameters and the effect of preconditioning. In chapter 6, we summarize our key findings and give guidelines for practical applications before presenting an outlook on future work in chapter 7.

2. Preliminaries

We will present the required theory, starting with a discussion of the analytic derivation of MDEs necessary to assess the quality of predictions in our test cases. This is followed by a coherent exposition of the sparse identification framework and one of its major challenges, multicollinearity. An overview of the investigated sparse regression algorithms at the core of the proposed procedure concludes this preliminaries section.

2.1. Modified differential equation analysis

For demonstration purposes, we begin by considering the linear advection equation

$$u_t + au_x = 0, \quad (1)$$

with a forward-time, backward-space (FTBS) discretization scheme, which is first order accurate in space and time.

$$\frac{u_i^{j+1} - u_i^j}{\Delta t} + a \frac{u_i^j - u_{i-1}^j}{\Delta x} = 0 \quad (2)$$

In order to construct the MDEs, existence of a continuously differentiable function $v(x, t)$ is presumed, which coincides with the numerical solution obtained from eq. (2) at the gridpoints $v(x, t) = v(i\Delta x, j\Delta t) = u_i^j$ [55]. Substitution in eq. (2)

$$\frac{v(x, t + \Delta t) - v(x, t)}{\Delta t} + a \frac{v(x, t) - v(x - \Delta x, t)}{\Delta x} = 0 \quad (3)$$

and Taylor expansion of each term around $v(x, t)$ yields

$$v_t + av_x + v_{tt}\Delta t/2 + v_{ttt}\Delta t^2/6 + \dots - v_{xx}a\Delta x/2 + v_{xxx}a\Delta x^2/6 + \dots = 0, \quad (4)$$

dropping the argument of $v(x, t)$. Villatoro et al. [54] refer to this form of the MDE in eq. (4) as the first modified equation, stressing the non-uniqueness of MDEs, since any linear combination of derivatives of eq. (4) is again a MDE. Following the procedure of Warming et al. [55], the third modified equation [54] can be derived by substitution of higher order time derivatives with spatial derivatives, using a symbolic mathematics package (e.g. *sympy* [37]).

$$\begin{aligned} &v_t + av_x + v_{xx}\Delta x(-a + a^2h)/2 + v_{xxx}\Delta x^2(a - 3a^2h + 2a^3h^2)/6 \\ &+ v_{xxxx}\Delta x^3(-a + 7a^2h - 12a^3h^2 + 3a^4h^3)/24 \\ &+ v_{xxxxx}\Delta x^4(a - 15a^2h + 50a^3h^2 - 60a^4h^3 + 24a^5h^4)/120 \\ &+ v_{xxxxxx}\Delta x^5(-a^5 + 31a^2h - 180a^3h^2 + 390a^4h^3 - 360a^5h^4 + 120a^6h^5)/720 \\ &+ \mathcal{O}(\Delta x^6) = 0, \end{aligned} \quad (5)$$

where $h = \Delta t/\Delta x$. While the procedure of Warming et al. [55] only applies to linear equations, Lerat et al. [29] previously already derived third MDEs for nonlinear equations. When truncating the Taylor series in time and space in the first MDE (4), higher order initial conditions (ICs) and boundary conditions (BCs) are necessary in order to obtain a well-posed problem. There are no higher order time derivatives in a truncated version of the third MDE (5) and the second challenge, higher order BCs, might be handled by enforcing periodic BCs. Yet, there is still criticism directed at the forward numerical solution of MDEs, because a smooth function $v(x, t)$ coinciding with the numerical solution on its gridpoints in general may not satisfy eq. (5) [12]. However, one is usually not interested in solutions of the MDE for specific ICs and BCs, but rather in the form of the MDE itself that contains information about the discretization scheme [12]. Considering this, we demonstrate that reconstruction of MDEs from data is possible without referring to any specific forward solution of the respective MDE.

2.2. Identification framework

The properties of the problem at hand largely determine the choice of an appropriate symbolic identification framework. The framework has to be applicable to nonlinear equations, deal with a moderate to large number of candidate functions




Fig. 1. Example library $\Theta(\mathbf{u}) \in \mathbb{R}^{n \times p}$. In our test cases, the library contains from 49 to 210 candidate terms for the advection equation, 28 terms for the Burgers' equation and 68 terms for the KdV equation.

and yield high accuracy predictions in a low noise environment. High accuracy is essential, given that coefficients in MDEs often span several orders of magnitude, e.g. Table 1. GPs of Raissi et al. [41] do not allow nonlinear equations, *Numerical GPs* [42] are not yet extended to inverse problems and *Physics Informed Neural Networks'* (PINNs) [43] continuous and discrete time models suffer from the accuracy requirement. Using PINNs, we only managed to discover terms up to v_{xxx} in eq. (5), all higher order derivatives were not identified correctly. We attribute this behavior to the large number of trainable parameters in the neural network, since finding the precise optimum in high-dimensional spaces is a difficult task. Including high order derivatives aggravates the issue, because the effective depth of the network increases exponentially with increasing derivative order, compromising neural network training.

For our work, we choose PDE-FIND [46]. With the numerical test cases in chapter 5, we will demonstrate that all of the outlined requirements are met. The algorithm leverages linear sparse regression, which facilitates accurate parameter estimation and sets terms not included in the predicted sparse equation exactly to 0. The design of the library enables flexible construction of fairly general candidate function spaces, including higher order temporal derivatives for identification of first MDEs. Furthermore, the vast amount of literature on (sparse) linear regression provides a firm theoretical basis for our work.

We summarize the PDE-FIND method, following [46]. Our objective is to identify first or third MDEs² of the form

$$u_t = F(u_{tt}, u_{ttt}, \dots, u, u_x, u^k u_{xx}, \dots). \quad (6)$$

Given data $u(x, t)$ obtained from the numerical solver under investigation, eq. (6) is set up at each point on the space-time grid using finite differences to approximate any partial derivatives in space and time. This builds a linear system of equations,

$$\mathbf{u}_t = \Theta(\mathbf{u})\xi + \epsilon; \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (7)$$

where $\mathbf{u} \in \mathbb{R}^n$ is a discretized version of $u(x, t)$ and n is the number of gridpoints included in regression. This system is solved for the unknown weights $\xi \in \mathbb{R}^p$ of p candidate terms using sparse regression methods, given that the majority of candidate terms in the library $\Theta(\mathbf{u}) \in \mathbb{R}^{n \times p}$ are not part of the correct MDE. Fig. 1 illustrates an example library $\Theta(\mathbf{u})$. The noise ϵ is assumed to be Gaussian; note that this is an implicit assumption in most sparse regression algorithms. ϵ can be attributed to higher order terms of the MDE not included in the library as well as truncation and round-off error of the finite difference approximations. Correlated candidate feature vectors have been identified as a major challenge for PDE-FIND [46]. Multicollinearity is a central concern in the construction of $\Theta(\mathbf{u})$, since its amount increases with the size of the candidate term pool.

2.3. Multicollinearity

Multicollinearity has been extensively studied in the context of linear regression. For example, it is well known that ordinary least squares (OLS) regression accuracy suffers from it. In-depth coverage of the topic is widely available; hence we will only provide some exemplary intuition about the ill-posed nature of multicollinearity:

² Beware non-uniqueness issues if multiple forms of MDEs (e.g. first and third) can be represented by $\Theta(\mathbf{u})$.

The sparse regression task in eq. (7) is to find the best approximation to u_t by weighted spatial derivatives of $u(x, t)$. Consider a particularly weak IC choice $\bar{u}(x, 0) = \sin(kx)$. Given two collinear terms $\bar{u}_{xx} = -k^2\bar{u}$, any $\gamma \in [0, 1]$ in the linear combination of both terms $\gamma * \bar{u} - (1 - \gamma)\bar{u}_{xx}/k^2$ yields the exact same result. It is thus impossible to distinguish the importance of one term versus the other. This intuition extends in a straightforward manner to multicollinearity, where a given feature vector can be represented with little error by a linear combination of the other feature vectors. The variance inflation factor (VIF) quantifies multicollinearity by measuring the deviation of this representation. A large VIF indicates a small deviation and therefore high multicollinearity.

$$VIF_i = 1/(1 - R_i^2), \quad (8)$$

where R_i^2 is the coefficient of determination of the linear regression

$$\mathbf{f}_i = \zeta_{i0} + \sum_{j=1; j \neq i}^p \zeta_{ij} \mathbf{f}_j, \quad (9)$$

and \mathbf{f}_i is a column of $\Theta(\mathbf{u})$.

The most prominent approach to obtain a well-posed problem is L_2 regularization, i.e. ridge regression [22]. Generalizing L_2 regularization, L_q regularization of the least-squares minimization problem of eq. (7) is defined by adding a penalty term for the weight vector.

$$\arg \min_{\xi} \left(\frac{1}{2} \|\Theta(\mathbf{u})\xi - \mathbf{u}_t\|_2^2 + \lambda \sum_{i=1}^p |\xi_i|^q \right); \quad q \geq 0 \quad (10)$$

L_2 penalizes large coefficients (note the $-k^2$ scaling in the example), resolving the ill-posedness. The solution of eq. (10) can be interpreted from a Bayesian statistics viewpoint to be the maximum a posteriori estimate of a Gaussian likelihood and a log-prior distribution $\log(p(\xi)) = \lambda \|\xi\|^q$ [20, p. 72]. The L_2 penalty corresponds to a Gaussian prior on weights centered at 0, with scale defined by λ . The prior mean of 0 is not supported by actual prior knowledge, thus artificially biasing non-zero coefficients towards zero.

A proper choice of prior for the MDE identification problem is challenging, since information about the scale, and ideally the mean, of each coefficient is necessary. The order of the discretization scheme provides some information about the scale of the dominant order truncation error terms. However, this value is in general too large for sufficient regularization of higher order terms, given that coefficients in MDEs often span several orders of magnitude.

2.4. Sparse regression algorithms

The aim of sparse regression is to find solutions to eq. (7), such that the fewest features possible are included in the model, while representing the data as faultless as achievable given candidate library $\Theta(\mathbf{u})$. To promote sparsity, L_0 regularization (10) represents a natural choice as it penalizes non-zero elements in ξ . L_0 regularization can be formulated either as an unconstrained (11) or constrained (12) optimization problem:

$$\arg \min_{\xi} \left(\frac{1}{2} \|\Theta(\mathbf{u})\xi - \mathbf{u}_t\|_2^2 + \lambda \|\xi\|_0 \right), \quad (11)$$

$$\arg \min_{\xi} \left(\frac{1}{2} \|\Theta(\mathbf{u})\xi - \mathbf{u}_t\|_2^2 \right), \text{ subject to } \|\xi\|_0 \leq s_0. \quad (12)$$

Using a leaps and bound algorithm [17], the solution of eq. (12) is possible up to $s = 30 \dots 40$ [20, p. 57]. However, due to being a combinatorically large problem, most modern sparse regression algorithms focus on numerically efficient approximations to the L_0 regularized problem [59]. We compare the performance of four sparse regression algorithms in the identification of MDEs task. A desirable property of sparse regression algorithms is *sign consistency* [60]. The estimated weights from a *sign consistent* sparse regression algorithm have the same sign as the true weights in the limit $n \rightarrow \infty$, where 0 is considered a distinct sign. As true weights corresponding to irrelevant candidate features are 0, a *sign consistent* algorithm correctly sets those features to 0 in the limit $n \rightarrow \infty$.

The most widespread algorithm among the considered is the *Least Absolute Shrinkage and Selection Operator* (Lasso) [53], which uses L_1 regularization. This is the smallest q such that eq. (10) is a convex problem, improving numerical efficiency of the optimization over non-convex problems significantly. Unlike ridge regression, Lasso sets irrelevant coefficients exactly to 0 [20, p. 73]. Lasso has two major deficiencies though. The L_1 penalty can be interpreted as a Laplace distribution centered at 0 and its scale defined by λ . However, analogous to ridge regression, this prior is not supported by actual prior knowledge. If the prior mean deviates from the data, which it does for all non-zero coefficients, this introduces bias that can be unsatisfactorily large for the problems considered in this paper (see e.g. Fig. 4). Besides, Lasso is sensitive to correlated feature vectors: A necessary and almost sufficient condition for *sign consistency* of Lasso is the *irrepresentability*

condition, which $\Theta(\mathbf{u})$ has to fulfill. The *irrepresentability condition* bounds the extent to which irrelevant features can be represented by true features. This condition is non-trivial to evaluate for moderate p and unknown true features. However, several sufficient conditions for the *irrepresentability condition* are known, a common property of which is a bound on the maximum correlation between feature vectors of $\Theta(\mathbf{u})$ [60].

To cope with highly correlated feature vectors common to $\Theta(\mathbf{u})$, Brunton et al. [10,46] propose *Sequential Threshold Ridge Regression* (STRidge). Hereby, the L_2 penalty is used for regularization of correlated feature vectors, while sequential thresholding promotes sparsity. Given a tolerance tol and λ corresponding to the L_2 penalty as hyperparameters, analytically tractable ridge estimates are obtained sequentially and weights smaller than tol are set to 0, until the non-zero weights converge and yield a sparsity pattern. In a second step, an OLS estimate for the non-zero weights is calculated, thus avoiding any bias from the L_2 penalty.

The recently published *Sparse Relaxed Regularized Regression* (SR3) method [61] relaxes regularization penalties to decouple the accuracy and sparsity requirements by introducing an auxiliary weight vector \mathbf{w} . Application of SR3 to problem (10) yields

$$\arg \min_{\xi, \mathbf{w}} \left(\frac{1}{2} \|\Theta(\mathbf{u})\xi - \mathbf{u}_t\|_2^2 + \lambda \sum_{i=1}^p |w_i|^q + (\gamma/2) \|\xi - \mathbf{w}\|_2^2 \right); \quad q \geq 0 \quad (13)$$

where the additional hyperparameter γ controls the amount of deviation of ξ , which enforces accuracy, from \mathbf{w} , which enforces sparsity. Minimization of eq. (13) with respect to ξ yields

$$\arg \min_{\mathbf{w}} \left(\frac{1}{2} \|\mathbf{F}_\gamma \mathbf{w} - \mathbf{g}_\gamma\|_2^2 + \lambda \sum_{i=1}^p |w_i|^q + (\gamma/2) \right); \quad q \geq 0 \quad (14)$$

with

$$\kappa(\mathbf{F}_\gamma) = \kappa(\Theta(\mathbf{u})) \sqrt{\frac{\gamma + \sigma_{\min}(\Theta(\mathbf{u}))}{\gamma + \sigma_{\max}(\Theta(\mathbf{u}))}}, \quad (15)$$

where $\kappa = \sigma_{\max}/\sigma_{\min}$ is the condition number. γ controls the amount of reduction in κ . For the definitions of $\mathbf{F}_\gamma \in \mathbb{R}^{n \times p}$ and $\mathbf{g}_\gamma \in \mathbb{R}^n$ see Zheng et al. [61]. SR3 has remarkable similarities to the puffer transformation in section 3.1.3, transforming both $\Theta(\mathbf{u})$ and \mathbf{u}_t and significantly reducing κ in the transformed system.

The Forward-Backward Greedy Algorithm (FoBa) [59] approximates problem (12) by sequential greedy selection of the feature, which reduces the OLS residual most. The rationale is to obtain a sparse solution by maximizing the gain from each added new candidate. After a fixed number of forward steps, backward steps aim to eliminate the least important features obtained from forward steps. Features are only deleted if the residual increases by less than half of the residual decrease of the last forward steps, guaranteeing convergence in a finite number of steps. This overcomes one major weakness of forward greedy algorithms, namely the inability to delete features deemed irrelevant during forward stepping. See Zhang [59] for an example of selected features deemed irrelevant. FoBa terminates once the next best forward step reduces the residual by less than the hyperparameter ϵ . For *sign consistency*, FoBa needs to meet the *sparse eigenvalue condition* [59], which is a weaker assumption than the *irrepresentability condition* [7].

3. Sparse identification of truncation errors

The centerpiece of our SITE approach is preconditioning at multiple stages of the workflow illustrated in Fig. 2. For detailed explanations of the individual steps we refer to the respective subsections in this chapter. We begin by using a *Non-Uniform Rational Basis Spline* (NURBS) as IC, which has been constructed to minimize multicollinearity. This is followed by the solver forward run generating the data. Afterwards, we assemble our candidate library following the PDE-FIND framework, scale the candidate terms and apply a puffer transformation. The preconditioned system is solved with a sparse regression algorithm for various hyperparameter values, resulting in a set of models that encode the predicted MDEs. From this set we choose the best model based on the well known Bayesian information criterion (BIC) [50], which we calculate based on a validation dataset.

3.1. Preconditioning

Given that we cannot rely on prior knowledge to deal with this ill-conditioned regression problem (7), we propose a three step preconditioning procedure to reduce multicollinearity.

3.1.1. Spline initialization

A function $u(x, t)$ that cannot be well described by linear combinations of its own derivatives is advantageous for low multicollinearity (section 2.3). Since data from few simulation time steps are sufficient for SITE (see chapter 5): $u(x, 0) \approx u(x, T)$. In cases in which the MDE does not depend on the current function value $u(x, t_{i-1})$, the initial condition $u(x, 0)$ can

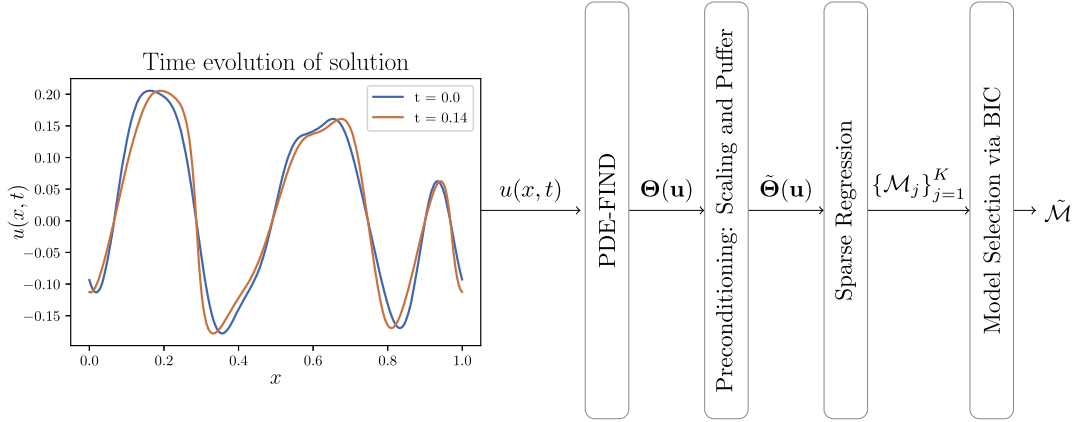


Fig. 2. Illustration of the SITE approach, starting with the solver output data and resulting in the selected model encoding the predicted MDE. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

be chosen freely. This is an opportunity to choose $u(x, 0)$, such that a well-conditioned problem is obtained. We use root mean square-VIF (RMS-VIF) (16) as an objective function to optimize the IC.

$$\text{RMS-VIF} = \sqrt{\frac{1}{p} \sum_{i=1}^p \text{VIF}_i^2} \quad (16)$$

We employ a NURBS s_ϕ as a parametrization for $u(x, 0)$, which provides a flexible definition for $u(x, 0)$ by variation of its weight vector ϕ . This allows an unconstrained optimization, because high order differentiability is achieved through the order of the spline and periodicity can be enforced automatically via additional knots outside Ω , due to local support of the NURBS. The choice of $u(x, 0)$ therefore reduces to finding a weight vector ϕ that minimizes RMS-VIF. We apply a gradient-free particle swarm optimization algorithm [58], as obtaining the gradient of RMS-VIF with respect to ϕ is non-trivial. The procedure is outlined in algorithm 1. Given ϕ , the parametric NURBS s_ϕ is created from which $u(x, 0)$ is obtained. Next, the simulation is run and $\Theta(\mathbf{u})$ is assembled. Lastly, RMS-VIF of $\Theta(\mathbf{u})$ is calculated and fed back into the particle swarm optimization, yielding new proposals for ϕ . Note that the reduction in RMS-VIF and condition number achieved through NURBS initialization comes at no additional cost other than increased preprocessing time.

Algorithm 1 NURBS initialization.

Input: iters, particles

Output: $u(x, 0)$

```

Initialization :  $\{\phi_j^1\}_{j=1}^{\text{particles}} = \text{random}$ 
1: for  $i = 1 \dots \text{iters}$  do
2:   for  $j = 1 \dots \text{particles}$  do
3:      $s_\phi = \text{NURBS}(\phi_j^i)$ 
4:      $u(x, 0) = \text{interpolate}(s_\phi)$ 
5:      $u(x, t) = \text{forwardsolve}(u(x, 0))$ 
6:      $\Theta(\mathbf{u}) = \text{PDE-FIND}(u(x, t))$ 
7:      $\text{RMS-VIF}_j^i = \text{calculateVIF}(\Theta(\mathbf{u}))$ 
8:   end for
9:    $\{\phi_j^{i+1}\}_{j=1}^{\text{particles}} = \text{particleswarm}(\{\text{RMS-VIF}_j^i\}_{j=1}^{\text{particles}})$ 
10: end for
11:  $\phi^{\text{best}} = \arg \min_j \text{RMS-VIF}(\{\phi_j^{\text{iters}}\}_{j=1}^{\text{particles}})$ 
12:  $s_\phi = \text{NURBS}(\phi^{\text{best}})$ 
13:  $u(x, 0) = \text{interpolate}(s_\phi)$ 
14: return  $u(x, 0)$ 

```

3.1.2. Scaling

Higher order derivatives of $u(x, t)$ often span several orders of magnitude. Scaling $\Theta(\mathbf{u})$ such that all features share a common magnitude is a standard preprocessing step in regression analysis [38,20]. We can rewrite eq. (7) using a scaling matrix $\mathbf{S} \in \mathbb{R}^{p \times p}$

$$\mathbf{u}_t = \overline{\Theta}(\mathbf{u})\bar{\xi} + \epsilon; \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) \quad \text{with} \quad \overline{\Theta}(\mathbf{u}) = \Theta(\mathbf{u})\mathbf{S}^{-1}; \quad \bar{\xi} = \mathbf{S}\xi, \quad (17)$$

where we use a default diagonal scaling matrix [38]

$$\mathbf{s}_{kk} = \sqrt{(\boldsymbol{\Theta}(\mathbf{u})^T \boldsymbol{\Theta}(\mathbf{u}))_{kk}}. \quad (18)$$

Even though rescaling does not alter VIF, it reduces the condition number $\kappa(\bar{\boldsymbol{\Theta}}(\mathbf{u})) \leq \kappa(\boldsymbol{\Theta}(\mathbf{u}))$. If $\kappa(\boldsymbol{\Theta}(\mathbf{u}))$ is very large, finding a numerical solution of OLS is problematic [18,40].

3.1.3. Puffer transformation

Puffer transformation [24] is one of our main preconditioning steps. While originally being introduced to extend the applicability of Lasso to problems, where the *irrepresentability condition* is not met, we show that other sparse regression algorithms benefit from the puffer transformation as well (see Fig. 4). Eq. (17) is rewritten by multiplying a precondition matrix $\mathbf{F} \in \mathbb{R}^{n \times n}$ from the left

$$\tilde{\mathbf{u}}_t = \tilde{\boldsymbol{\Theta}}(\mathbf{u})\tilde{\boldsymbol{\xi}} + \tilde{\boldsymbol{\epsilon}}; \quad \text{with} \quad \tilde{\mathbf{u}}_t = \mathbf{F}\mathbf{u}_t; \quad \tilde{\boldsymbol{\Theta}}(\mathbf{u}) = \mathbf{F}\bar{\boldsymbol{\Theta}}(\mathbf{u}); \quad \tilde{\boldsymbol{\epsilon}} = \mathbf{F}\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \tilde{\boldsymbol{\Sigma}}). \quad (19)$$

Jia et al. [24] propose to construct \mathbf{F} from a singular value decomposition (SVD) $\bar{\boldsymbol{\Theta}}(\mathbf{u}) = \mathbf{U}\mathbf{D}\mathbf{V}^T$; $\mathbf{U} \in \mathbb{R}^{n \times p}$; $\mathbf{D} \in \mathbb{R}^{p \times p}$; $\mathbf{V}^T \in \mathbb{R}^{p \times p}$, assuming $\text{rank}(\bar{\boldsymbol{\Theta}}(\mathbf{u})) = p$ with $n > p$ and neglecting zero rows and columns³

$$\mathbf{F} = \mathbf{U}\mathbf{D}^{-1}\mathbf{U}^T \implies \tilde{\boldsymbol{\Sigma}} = \sigma^2 \mathbf{U}\mathbf{D}^{-2}\mathbf{U}^T. \quad (20)$$

Given that $\mathbf{U}^T \mathbf{U} = \mathbf{V}^T \mathbf{V} = \mathbf{I}_d$ from the SVD, it is straightforward to show orthonormality of the preconditioned system matrix $\tilde{\boldsymbol{\Theta}}^T(\mathbf{u})\tilde{\boldsymbol{\Theta}}(\mathbf{u}) = \mathbf{I}_d$, yielding a perfect $\kappa(\tilde{\boldsymbol{\Theta}}(\mathbf{u})) = 1$ (neglecting numerical error). However, the elimination of multicollinearity comes at the cost of inflating the noise variance $\tilde{\boldsymbol{\Sigma}}$ by \mathbf{D}^{-2} . This reduces the signal-to-noise ratio, thus counteracting the benefits from a better conditioned problem [24]. At this point, the two previous preconditioning steps become relevant. Reducing multicollinearity by a good choice of $u(x, 0)$ and scaling $\boldsymbol{\Theta}(\mathbf{u})$ both reduce $\kappa(\bar{\boldsymbol{\Theta}}(\mathbf{u}))$. Therefore, the smallest singular values in \mathbf{D} become larger, hence reducing the noise inflation effect.

The preconditioner \mathbf{F} is generalized by Jia et al. [25], with the goal to bound the inflation of the noise caused by too small singular values: $\hat{\mathbf{D}}^{-1}$ substitutes \mathbf{D}^{-1} in eq. (20),

$$\hat{\mathbf{D}}_{kk}^{-1} = g(\mathbf{D}_{kk}, \tau) / \mathbf{D}_{kk}, \quad (21)$$

requiring a reasonable choice for g and τ . While this approach might be a helpful remedy if noise inflation is a major issue, for the sake of simplicity, we restrict ourselves in this work to the default \mathbf{F} in eq. (20).

3.2. Sparse regression

The purpose of the sparse regression algorithm is to propose a sparse solution to the linear system (7) $\boldsymbol{\xi}$ encoding the predicted MDE. We chose FoBa as the default sparse regression algorithm for SITE due to its high accuracy, numerical efficiency and straightforward hyperparameter tuning (to be shown in chapter 5). Iterating over the hyperparameter ϵ of FoBa gives a set of candidate models $\{\mathcal{M}_j\}_{j=1}^K$, with a different degree of sparsity each. Since the number of terms to be included in the true model⁴ is not known a priori, we employ a data-driven model selection procedure.

3.3. Model selection

We calculate the BIC for all linear models $\{\mathcal{M}_j\}_{j=1}^K$ with respect to a single validation simulation and select the model that maximizes the BIC. A validation dataset $\tilde{\boldsymbol{\Theta}}^{\text{val}}(\mathbf{u})$ can be computed easily by changing $u(x, 0)$. In this paper, we use a NURBS with a smaller number of optimizable knots ϕ resulting in a different initial condition. For a linear regression model, BIC is defined as [5, p. 153]

$$\text{BIC} = -\frac{n_{\text{eff}}}{2} \log \left(\|\tilde{\boldsymbol{\Theta}}^{\text{val}}(\mathbf{u})\hat{\boldsymbol{\xi}} - \tilde{\mathbf{u}}_t^{\text{val}}\|_2^2 \right) - \frac{k}{2} \log(n_{\text{eff}}), \quad (22)$$

where n_{eff} is the effective sample size, k is the number of features included in the model and $\hat{\boldsymbol{\xi}}$ is the sparse regression result from the training set. The BIC aims to find parsimonious models from a balance of the residual sum of squares in the first term and model complexity in the second term of eq. (22). Only for independent, identically distributed data $n_{\text{eff}} = n$ [4]. Clearly, the number of independent samples is smaller than n : Considering the case of ideal advection and periodic boundaries, $u(x, 0)$ is only shifted over time. Therefore, all samples of future time steps are identical to the samples from the first time step (with the simplifying assumption that the advection distance is a multiple of Δx). Due to nonlinear

³ If $n < p$, the design matrix is projected onto the Stiefel manifold, resulting in an empirically verifiable significant reduction in pairwise feature correlation [24].

⁴ We agree that a model can never perfectly represent reality and therefore a “true” model usually does not exist. Since for the problems considered in this paper we can derive the exact solution analytically, we are in fact dealing with the rare case that a true model exists (yet having an infinite number of parameters). We will refer to “true” terms for those terms included in the analytic model.

equations and higher order terms in the MDE, $u(x, t)$ slightly deforms over time, thus increasing the amount of information compared to the ideal advection case. However, correlation between samples within one time step (e.g. neighboring points) decreases the amount of information contained in the data from a single time step. Presuming these two secondary effects roughly balance each other, we take $n_{\text{eff}} = n^x$ as a coarse estimate of n_{eff} , where n^x is the spatial resolution of Ω .

If the puffer transformation is applied, the noise vector $\mathbf{F}\epsilon$ contains statistically dependent terms, which are not accounted for in the assumptions of BIC. Jia et al. [24] therefore propose to compute an OLS estimate without puffer transformation once the set of candidate models with given sparsity pattern has been obtained. In our test cases in chapter 5, we opted to calculate BIC only for models obtained without the puffer transformation in the first place. Hence, we do not provide any formal tests for the procedure proposed by Jia et al. [24].

4. Numerical solvers

To showcase the performance of SITE, we consider three example PDEs with corresponding numerical solvers for which MDEs are to be identified from solver output data, namely the advection, Burgers' and KdV equation. We performed all simulations on a domain $\Omega : x \in [0, 1]$ with periodic BCs $u(0, t) = u(1, t)$.

4.1. Advection equation

We consider the linear advection equation (1) for $a = 1$ with a FTBS discretization scheme (2). The linear accuracy requirement is well known to be

$$\text{CFL} = \frac{a\Delta t}{\Delta x} \leq 1. \quad (23)$$

4.2. Burgers' equation

Burgers' equation is an important model equation in several fields, including fluid dynamics and traffic flow calculations. Due to its similarity to the Navier-Stokes equation and its tendency to develop shock solutions, it is frequently considered as a test case for numerical algorithms in the literature. To demonstrate applicability of SITE to nonlinear equations and more advanced discretization schemes, we discretize the inviscid Burgers' equation

$$u_t + \left(\frac{u^2}{2}\right)_x = 0 \quad (24)$$

by a MacCormack predictor-corrector scheme [32], which is second order accurate in space and time:

$$\begin{aligned} \tilde{u}_i^{j+1} &= u_i^j - h \left(\frac{(u_{i+1}^j)^2}{2} - \frac{(u_i^j)^2}{2} \right), \\ u_i^{j+1} &= u_i^j - \frac{h}{2} \left[\left(\frac{(u_{i+1}^j)^2}{2} - \frac{(u_i^j)^2}{2} \right) + \left(\frac{(\tilde{u}_i^j)^2}{2} - \frac{(\tilde{u}_{i-1}^j)^2}{2} \right) \right]. \end{aligned} \quad (25)$$

The linear stability criterion is given by

$$\text{CFL} = \frac{|u|_{\max} \Delta t}{\Delta x} \leq 1. \quad (26)$$

For derivation of MDEs, the predictor-corrector scheme (25) is rewritten into a single equation [29]

$$\begin{aligned} 0 &= \frac{u_i^{j+1} - u_i^j}{\Delta t} + \frac{(u_{i+1}^j)^2 - (u_{i-1}^j)^2}{4\Delta x} \\ &\quad - \frac{\Delta t}{2} \left(\frac{u_i^j + u_{i-1}^j}{2} \frac{(u_{i+1}^j)^2 - 2(u_i^j)^2 + (u_{i-1}^j)^2}{2\Delta x^2} + \frac{u_i^j - u_{i-1}^j}{\Delta x} \frac{(u_{i+1}^j)^2 - (u_{i-1}^j)^2}{4\Delta x} \right) \\ &\quad + \frac{\Delta t^2}{2} \frac{(u_{i+1}^j)^2 - (u_{i-1}^j)^2}{4\Delta x} \frac{(u_{i+1}^j)^2 - 2(u_i^j)^2 + (u_{i-1}^j)^2}{2\Delta x^2}. \end{aligned} \quad (27)$$

4.3. Korteweg-de Vries equation

The KdV equation describes the asymptotic behavior of one-dimensional waves with small amplitudes. The modeled physical phenomena include shallow water waves and magneto-hydrodynamic waves in a plasma [57]. This test case demonstrates the ability of SITE to identify first MDEs, whose higher order temporal derivatives are not substituted with spatial derivatives. We discretize the KdV equation

$$u_t + 6uu_x + u_{xxx} = 0 \quad (28)$$

using the Zabusky and Kruskal scheme [57], which is second order accurate in space and time:

$$u_i^{j+1} = u_i^{j-1} - 2h(u_{i+1}^j + u_i^j + u_{i-1}^j)(u_{i+1}^j - u_{i-1}^j) - \frac{h}{\Delta x^2}(u_{i+2}^j - 2u_{i+1}^j + 2u_{i-1}^j - u_{i-2}^j). \quad (29)$$

Due to the central in time approximation of (29), an uncentered time discretization has to be used for the first time step [51].

$$u_i^1 = u_i^0 - h(u_{i+1}^0 + u_i^0 + u_{i-1}^0)(u_{i+1}^0 - u_{i-1}^0) - \frac{h}{2\Delta x^2}(u_{i+2}^0 - 2u_{i+1}^0 + 2u_{i-1}^0 - u_{i-2}^0) \quad (30)$$

The linear stability criterion corresponding to the Zabusky and Kruskal scheme is much more restrictive,

$$\left| \frac{\Delta t}{\Delta x} - 2u_{\max} + \frac{1}{\Delta x^2} \right| \leq \frac{2}{3\sqrt{3}}, \quad (31)$$

thus Δt scaling with Δx^3 [51].

4.4. Verification

We utilized the method of manufactured solutions [48] for verification of the custom solvers to eliminate the possibility of incorrect results due to a wrong solver implementation. For all solvers, we chose the manufactured solution $\bar{u}(x, t) = \sin(2\pi(x + t)) + 0.001$. We refined Δx and Δt together with a constant CFL = 0.1 for the advection and Burgers' equation and CFL = 10^{-10} for the KdV equation, due to its more restrictive stability requirement. The solver convergence plots are shown in Fig. A.11 in the Appendix confirming accuracy orders in space and time of 1, 2 and 2 for FTBS, MacCormack and the Zabusky and Kruskal scheme, respectively.

5. Numerical test cases

We investigate the properties of SITE based on the three test cases outlined in chapter 4. The analysis focuses on the default setting of SITE, which uses FoBa with spline initialization and without puffer transformation. All test cases exhibit the same structure: First, we outline the discretization parameters of the data generating simulation as well as the design of the library $\Theta(\mathbf{u})$. Second, the MDE predicted by the default setup of SITE and its accuracy is presented in tabular form. Next, we study the impact of preconditioning steps and its interplay with the choice of sparse regression algorithm. We conclude each test case with an analysis on the impact of the simulation grid on regression accuracy and on BIC model selection for the default setting. BIC model selection is compared to an optimal procedure selecting always the optimal model from the set proposed by FoBa. We defined the optimal model as the model with maximum number of correct terms while not introducing any incorrect terms.

5.1. Algorithmic implementation details

This section summarizes key implementation details of the following test cases. We use *Python* 3.6 [52] with double precision numbers for all computations. In all our test cases we have $n > p$, but the method extends in principle to $n < p$. Note however that we provide no formal tests for this case. Derivatives are approximated by 8th order accurate finite difference stencils from the *findiff* [3] package. Non-centered stencils introduce additional error in the derivative approximations, hence impairing our ability to find the maximal number of candidate terms. Therefore, we only include gridpoints for which centered stencils are available. We pad the data at the boundary of the grid (n^x, n^t) , where n^t is the number of simulation time steps, to ensure that all remaining points have enough neighboring points such that centered stencils are well defined. The minimum padding width depends on the order of the approximated derivative and the accuracy order of the stencil. The spline initialization uses 8th order NURBS with 15 knots within Ω for the training set and 11 knots for the validation set. We optimized NURBS with the particle swarm optimization algorithm implemented in the *pyswarms* [23] package with 50 particles, 100 iterations and default optimization parameters. Fig. 3 illustrates the VIF optimization progress for our test cases.

For comparison, a Gaussian bell curve is considered

$$u(x, 0) = e^{-50(x-0.5)^2} + e^{-50(x+0.5)^2} + e^{-50(x-1.5)^2}; \quad x \in [0, 1], \quad (32)$$

which unlike $\sin(kx)$ from the example in section 2.3 is not inherently collinear to its derivatives.

We compare the maximum performance of the four algorithms from section 2.4. To avoid any bias resulting from model selection, we iterate linearly over the decadic logarithm of their respective hyperparameters and evaluate the accuracy of the predicted set of models. Table B.4 in the appendix summarizes the considered hyperparameters. The same hyperparameters of FoBa are also used to generate $\{\mathcal{M}_j\}_{j=1}^K$ in the analyses of the impact of the discretization parameters. For SR3, we chose




Fig. 3. NURBS particle swarm optimization for the three test cases. Solid lines correspond to the training sets and dashed lines to the validation sets.

L_0 regularization. We stopped the optimization process for Lasso and SR3 after 10000 iterations to keep computational effort within reasonable bounds.

To assess the prediction accuracy of the sparse regression algorithms considered here, we distinguish between correct terms, which are included in the respective analytically derived MDE, and incorrect terms, which are not included. If no incorrect term is included in the predicted model \mathcal{M} , mean absolute error (MAE) and mean relative error (MRE) are calculated from the analytically derived MDE weights ξ ,

$$\text{MAE} = \frac{1}{p_{\mathcal{M}}} \sum_{i=1}^{p_{\mathcal{M}}} |\xi_{i,\mathcal{M}} - \xi_i|; \quad \text{MRE} = \frac{1}{p_{\mathcal{M}}} \sum_{i=1}^{p_{\mathcal{M}}} \left| \frac{\xi_{i,\mathcal{M}} - \xi_i}{\xi_i} \right|. \quad (33)$$

Note that leading order truncation error terms dominate MAE, while the highest order terms dominate MRE (e.g. Table 1). Due to the fact that deviations of even orders of magnitudes in the smallest term would be invisible when adhering to MAE, we deem MRE to be more appropriate to judge the accuracy of sparse regression algorithms. For practitioners who are primarily interested in the leading order truncation error, MAE does however offer valuable information.

We calculate the empirical order of identified terms from $\xi_m \sim \mathcal{O}(\Delta x^k)$ using predictions from two simulations with spatial width Δx_1 and Δx_2 ,

$$k = \frac{\log(\xi_{m1}/\xi_{m2})}{\log(\Delta x_1/\Delta x_2)}. \quad (34)$$

5.2. Advection equation with FTBS

To demonstrate the limits of accuracy of SITE, we study the linear advection equation (1) with a FTBS discretization scheme (2) with the objective to identify its third MDE (5). The discretization parameters are $\text{CFL} = 0.01$ on a grid $(n^x, n^t) = (300, 17)$, yielding 5 time steps after data padding. We build two libraries: a small one contains u and all its spatial derivatives up to order 6. The large one appends all of those combinations of derivatives that add up to a given cumulative order for all cumulative orders up to 6, e.g. augmenting the library by u_x^3 and $u_x u_{xx}$ for a cumulative order of 3. Next, these basis functions are multiplied by u^k for k up to 6 and an intercept is added. This yields $p = 49$ candidate terms in the small library and $p = 210$ in the large one. Note that the maximum number of correct terms included in the libraries is 6. The spline is optimized with respect to the small library. For the large library, $n_{\text{eff}} \approx p$, presuming our coarse estimate $n_{\text{eff}} = n^x$ and considering intra time step correlation.

Table 1 summarizes the predicted MDE of SITE in the default setting for the large library case, demonstrating highly accurate predictions without introducing any incorrect terms. The empirical order is calculated pairwise from eq. (34) with a sequence of $n^x = (200, 300, 400, 500)$. The obtained orders are averaged afterwards.

Fig. 4 compares the considered sparse regression algorithms with respect to MRE and the number of correctly identified terms in the small library case. Lasso consistently predicts less accurate models compared to the other sparse regression algorithms due to the bias from the L_1 regularization. Puffer transformation together with spline initialization (a) improves MRE and the number of identified terms for all algorithms. The only exception is FoBa with 6 terms in the model, exhibiting slightly increased MRE. If STRidge identifies the same sparsity pattern as FoBa, their results are identical, as both rely on OLS to predict ξ . Interestingly, when using puffer, FoBa, STRidge and SR3 yield identical peak accuracy across the range of terms in models. SR3 additionally proposes less accurate models due to the extra degree of freedom from the hyperparameter γ .

Table 1

Summary of SITE default setup prediction for the advection equation.

	Analytical weight	Absolute error	Relative error	Empirical order
V_x	-1	$7.89 \cdot 10^{-12}$	$7.89 \cdot 10^{-12}$	0.00
V_{xx}	$1.65 \cdot 10^{-3}$	$8.55 \cdot 10^{-14}$	$5.18 \cdot 10^{-11}$	1.00
V_{xxx}	$-1.80 \cdot 10^{-6}$	$6.48 \cdot 10^{-14}$	$3.61 \cdot 10^{-8}$	2.00
V_{xxxx}	$1.44 \cdot 10^{-9}$	$2.50 \cdot 10^{-16}$	$1.74 \cdot 10^{-7}$	3.00
V_{xxxxx}	$-8.80 \cdot 10^{-13}$	$1.57 \cdot 10^{-16}$	$1.79 \cdot 10^{-4}$	4.00
V_{xxxxxx}	$4.04 \cdot 10^{-16}$	$1.59 \cdot 10^{-19}$	$3.94 \cdot 10^{-4}$	5.00

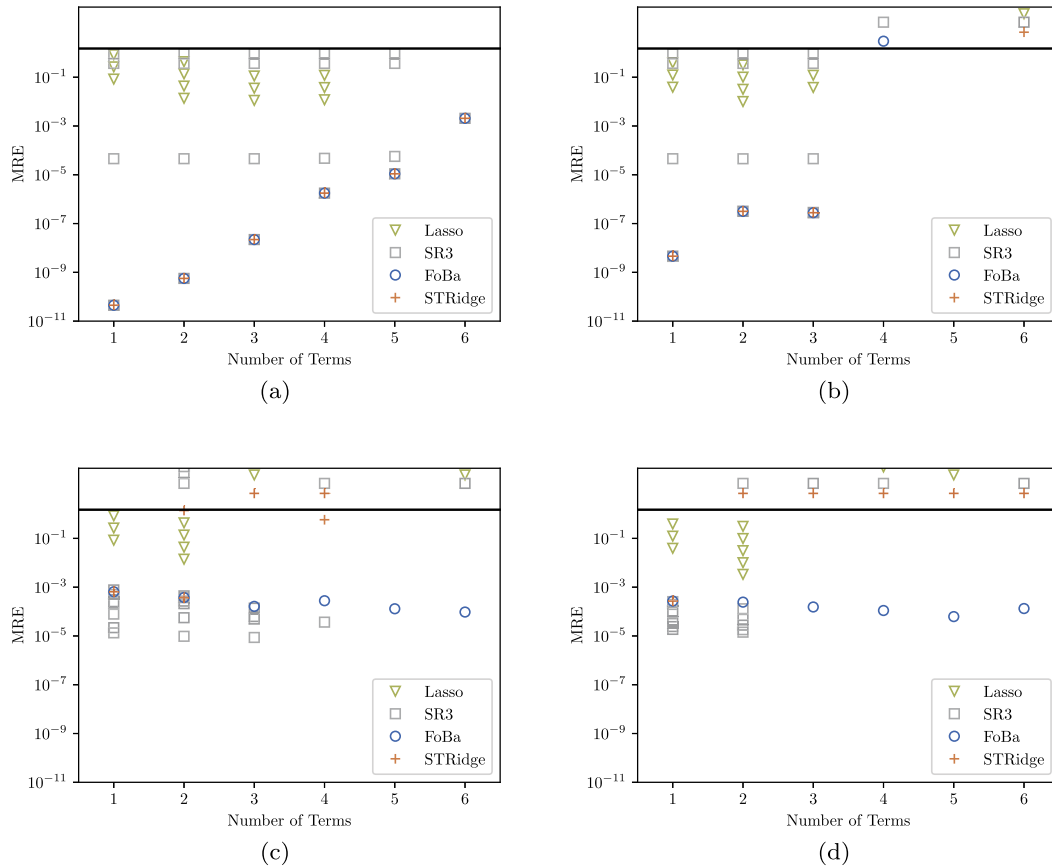


Fig. 4. MRE of sparse regression algorithms as a function of the number of terms included in a model from an iteration over the respective hyperparameters for the advection equation and the small library. Models above the black horizontal line ($y = 10^0$) contain at least one incorrect term and are sorted for visualization purposes. Their respective y-value has no quantitative meaning. The considered setups include spline initialization with puffer transformation (a), Gauss initialization with puffer transformation (b), spline initialization without puffer transformation (c) and Gauss initialization without puffer transformation (d).

Without puffer (c), SR3 is more accurate than FoBa when comparing models with a small number of terms, most likely due to its built-in preconditioning. In contrast, FoBa proposes a model that includes all 6 correct terms, whereas the maximum number of correct terms predicted by SR3 is 4. The impact of the puffer transformation is clearly visible when comparing MRE of FoBa with (a) and without puffer (c). In the latter case, MRE is relatively constant over the number of terms included. In the former case however, the system exhibits a large signal-to-noise ratio and the reduced multicollinearity from puffer improves accuracy, if few, large magnitude terms are included. However, with increasing number of small magnitude terms, the signal-to-noise ratio of the system decreases and the noise inflation from the puffer transformation outweighs the improved multicollinearity, causing larger MRE. Spline initialization yields similar results to the Gauss initial condition (32) if no puffer transformation is applied (d). In conjunction with puffer transformations however (b), spline initialization significantly improves regression results. This behavior stems from smaller multicollinearity that induces less noise by the puffer transformation.

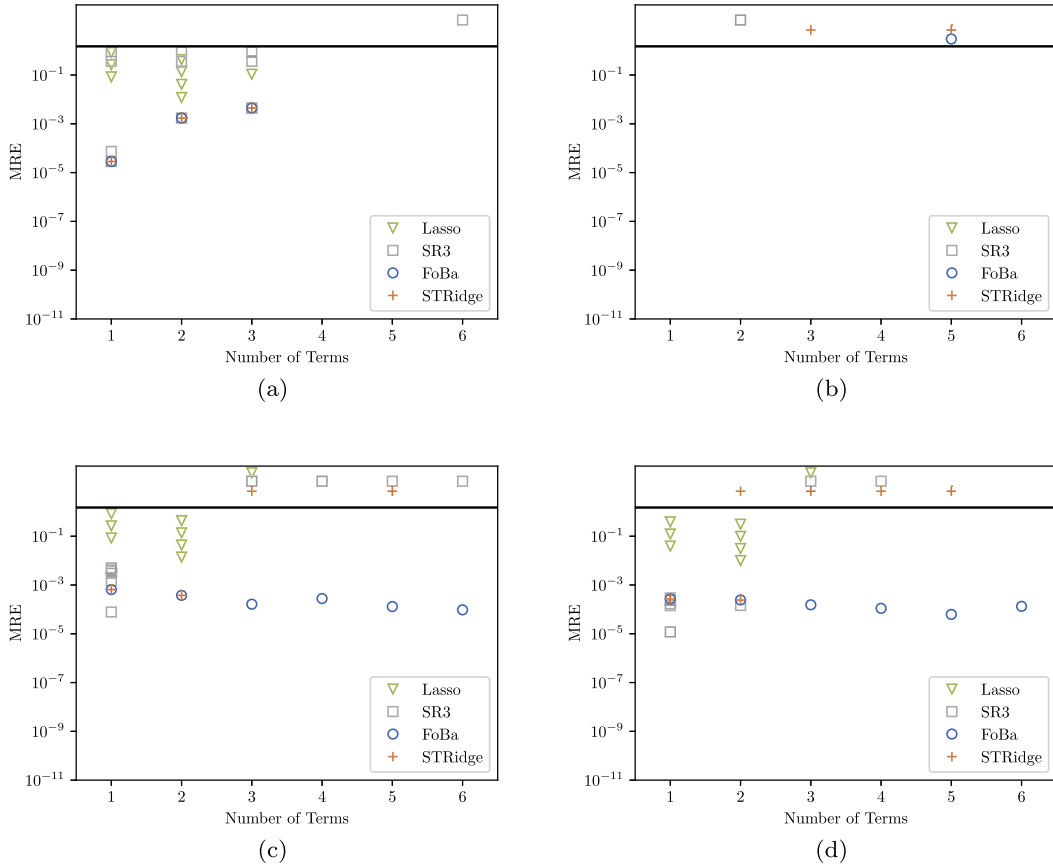


Fig. 5. MRE of sparse regression algorithms as a function of the number of terms included in a model from an iteration over the respective hyperparameters for the advection equation and the small library. Models above the black horizontal line ($y = 10^0$) contain at least one incorrect term and are sorted for visualization purposes. Their respective y-value has no quantitative meaning. The considered setups include spline initialization with puffer transformation (a), Gauss initialization with puffer transformation (b), spline initialization without puffer transformation (c) and Gauss initialization without puffer transformation (d).

When employing the large library (Fig. 5), application of the puffer transformation significantly degrades regression performance. This applies in particular if no spline initialization is used, where none of the algorithms can identify a single correct term. We attribute this behavior to the error inflating property of puffer, which becomes more dominant for increasing VIF due to the additional candidate terms. Larger multicollinearity from the Gauss initialization aggravates this effect. Without a puffer transformation, the results are similar to the small library case, with a slight decline in performance for SR3.

Fig. 6(a) shows the impact of variations in n^x on regression accuracy and on the ability of BIC in selecting the optimal model for 5 time steps in the default setup. The proposed set of models always contains the model with the maximum number of correct terms, however BIC fails to select it for $n^x > 600$. MRE decreases until $n^x = 400$ due to the decreasing importance of higher order terms not included in the model.

These decrease faster in magnitude than the sought terms as the largest correct term not in the library is $\mathcal{O}(\Delta x^6)$. For $n^x > 400$, the benefits of this effect are outweighed by increased noise from the round-off error of the finite difference approximations, resulting in increasing MRE. MAE, which is dominated by large magnitude terms with high signal-to-noise ratio, is more robust to this additional noise. The results for 100 time steps are comparable to the 5 time step case (Fig. 6(b)) with the only noticeable difference being that BIC identifies the correct model up to $n^x = 900$.

5.3. Burgers' equation with MacCormack

After inserting Taylor series into the MacCormack scheme (25) for the Burgers' equation (24) and substituting higher order temporal derivatives with spatial derivatives, the third MDE is obtained [29].

$$v_t + \left(\frac{v^2}{2}\right)_x - \Delta x^2 \left(v_{xxx} \frac{v}{6} (h^2 v^2 - 1) + \frac{v_x v_{xx}}{2} (2h^2 v^2 - h v - 1) + v_x^3 \frac{h}{4} (2h v - 1) \right) + \mathcal{O}(\Delta x^3) = 0 \quad (35)$$

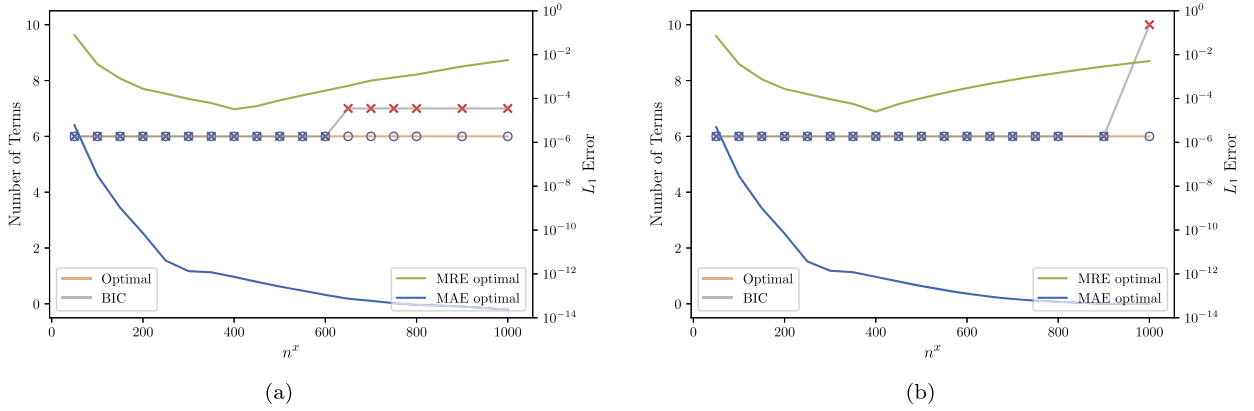


Fig. 6. Number of terms selected by BIC and by optimal choice from the model candidates proposed by FoBa as a function of resolution for the advection equation and 5 time steps (a) or 100 time steps (b). Markers of models only containing correct terms are plotted in blue, models containing at least one incorrect term are plotted in red. BIC is represented by a cross (\times) and the optimal choice by a circle (\circ). MRE and MAE are displayed for the optimal model on the right axis.

Table 2

Summary of SITE default setup prediction for the Burgers' equation.

	Analytical weight	Absolute error	Relative error	Empirical order
vv_x	-1	$5.71 \cdot 10^{-10}$	$5.71 \cdot 10^{-10}$	0.00
$v^3 v_{xxx}$	$9.87 \cdot 10^{-9}$	$-9.36 \cdot 10^{-12}$	$9.49 \cdot 10^{-4}$	2.01
vv_{xxx}	$-1.67 \cdot 10^{-9}$	$2.21 \cdot 10^{-13}$	$1.33 \cdot 10^{-4}$	2.00
$v^2 v_x v_{xx}$	$5.92 \cdot 10^{-8}$	$3.96 \cdot 10^{-11}$	$6.69 \cdot 10^{-4}$	2.00
$vv_x v_{xx}$	$-1.22 \cdot 10^{-8}$	$4.79 \cdot 10^{-12}$	$3.94 \cdot 10^{-4}$	2.00
$v_x v_{vxx}$	$-5.00 \cdot 10^{-9}$	$-7.40 \cdot 10^{-13}$	$1.48 \cdot 10^{-4}$	2.00
vv_x^3	$2.96 \cdot 10^{-8}$	$-7.76 \cdot 10^{-11}$	$2.62 \cdot 10^{-3}$	1.99
v_x^3	$-6.08 \cdot 10^{-9}$	$1.74 \cdot 10^{-12}$	$2.86 \cdot 10^{-4}$	2.00

The discretization parameters are $CFL = 0.5$ on a grid $(n^x, n^t) = (10000, 17)$, yielding 5 time steps after data padding. The larger n^x and CFL number result in a time step size in the same order of magnitude as in the advection case. The library contains u and all its spatial derivatives up to order 3 as well as all of those combinations of derivatives which add up to a given cumulative order for all cumulative orders up to 3. These basis functions are then multiplied by u^k for k up to 3 and an intercept is added, yielding $p = 28$ candidate terms. Since $\Theta(\mathbf{u})$ can represent all second order truncation error terms in the MDE (35), the maximum number of correctly identifiable terms is 8.

Table 2 summarizes the predicted MDE of SITE in the default setting. SITE demonstrates consistently accurate predictions of the truncation error terms with relative errors in the order 10^{-4} without predicting any incorrect terms. The empirical order is calculated by a sequence of $n^x = (6000, 8000, 10000, 12000)$.

Fig. 7 compares the accuracy of sparse regression for Burgers' equation. STRidge profits substantially from puffer transformation, then being able to identify all terms correctly using Spline initialization and outperforming FoBa with Gauss initialization. Without puffer, FoBa still detects all terms using spline initialization, while all other algorithms can only detect the term from Burgers' equation, but no truncation error terms. Spline initialization considerably improves results with and without puffer transformation. Interpretations are analogous to the advection test case.

Fig. 8(a) shows the impact of variations in n^x with respect to regression accuracy and the model selection capabilities of BIC for 5 time steps in the default setup. Similarly to the advection case, the model with the maximum number of correct terms is always included in the proposed set. However, BIC can only identify it within a range of $n^x \in [6000, 14000]$, where the optimal error is the smallest. Given that all truncation error terms are of the same order $\mathcal{O}(h^2)$, the increased noise impacts not only MRE, but MAE as well. Since the highest derivative to be approximated is 3, n^x can be chosen much larger than in the advection case before round-off error becomes dominant. Therefore, correct higher order truncation error terms not included in the library can be driven towards 0 effectively.

The results based on the simulation with 100 time steps is shown in Fig. 8(b). In contrast to previous examples, the model with the maximum number of correct terms is often not included in the set of models provided by FoBa outside the range $n^x \in [6000, 14000]$. MRE and MAE increase significantly compared to the 5 time step case, indicating increased noise in the data obtained from later stages of the simulation. Note that both curves lose some meaning in areas where the number of terms in the optimal model changes.

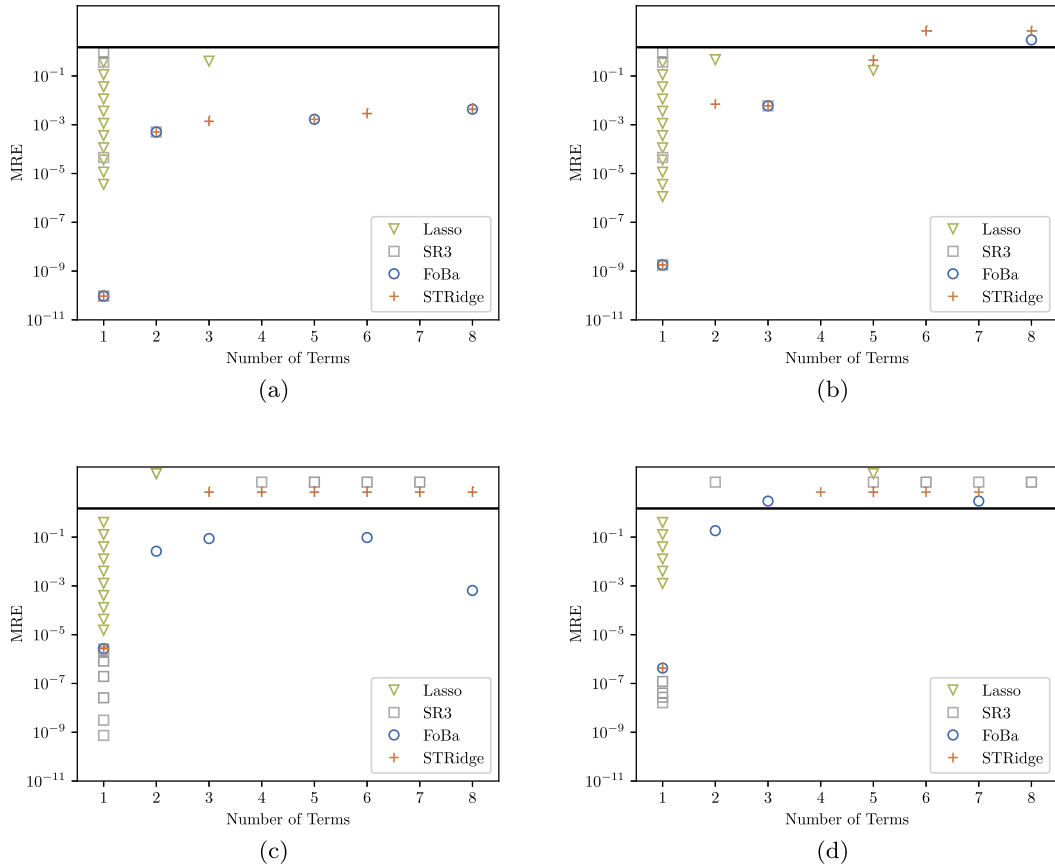


Fig. 7. MRE of sparse regression algorithms as a function of the number of terms included in a model from an iteration over the respective hyperparameters for the KdV equation. Models above the black horizontal line ($y = 10^0$) contain at least one incorrect term and are sorted for visualization purposes. Their respective y-value has no quantitative meaning. The considered setups include spline initialization with puffer transformation (a), Gauss initialization with puffer transformation (b), spline initialization without puffer transformation (c) and Gauss initialization without puffer transformation (d).

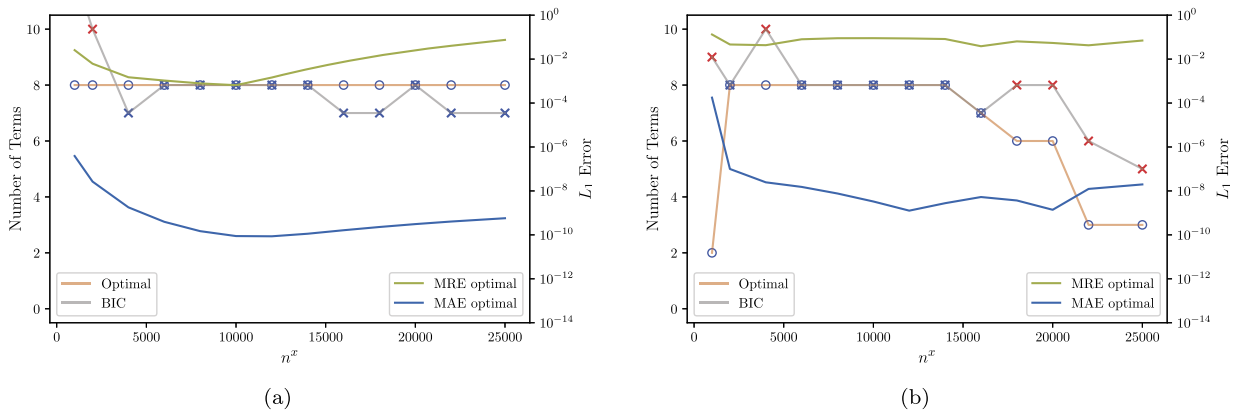


Fig. 8. Number of terms selected by BIC and by optimal choice from the model candidates proposed by FoBa as a function of resolution for the KdV equation and 5 time steps (a) or 100 time steps (b). Markers of models only containing correct terms are plotted in blue, models containing at least one incorrect term are plotted in red. BIC is represented by a cross (\times) and the optimal choice by a circle (\circ). MRE and MAE are displayed for the optimal model on the right axis.

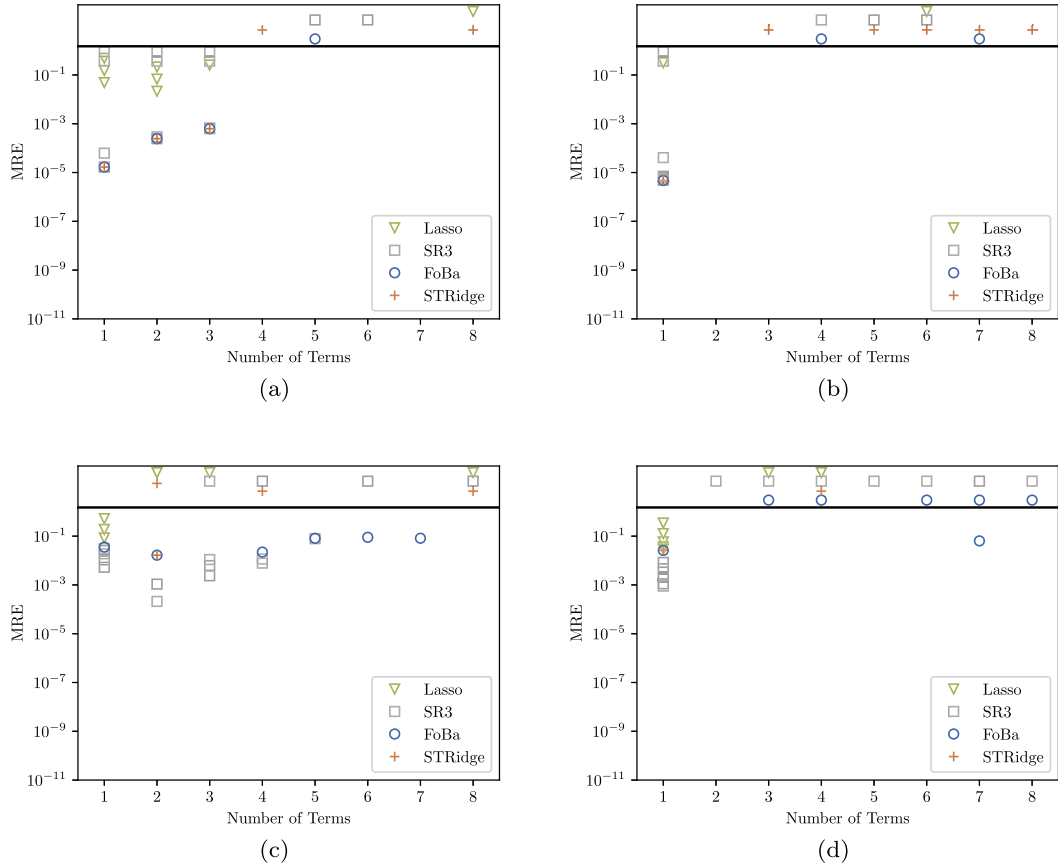


Fig. 9. MRE of sparse regression algorithms as a function of the number of terms included in a model from an iteration over the respective hyperparameters for the KdV equation. Models above the black horizontal line ($y = 10^0$) contain at least one incorrect term and are sorted for visualization purposes. Their respective y-value has no quantitative meaning. The considered setups include spline initialization with puffer transformation (a), Gauss initialization with puffer transformation (b), spline initialization without puffer transformation (c) and Gauss initialization without puffer transformation (d).

5.4. Korteweg-de Vries equation with Zabusky and Kruskal

Inserting Taylor series expansions into the Zabusky and Kruskal [57] discretization scheme (29) yields the first MDE.

$$\begin{aligned} v_t + 6v v_x + v_{xxx} + \Delta x^2 \left(\frac{h^2}{6} v_{ttt} + \frac{1}{4} u_{xxxxx} + v v_{xxx} + 2v_x v_{xx} \right) \\ + \Delta x^4 \left(\frac{h^4}{120} v_{ttttt} + \frac{1}{40} v_{xxxxxx} + \frac{1}{3} v_{xx} v_{xxx} + \frac{1}{6} v_x v_{xxxx} + \frac{1}{20} v v_{xxxxx} \right) + \mathcal{O}(\Delta x^6) = 0 \end{aligned} \quad (36)$$

The discretization parameters are $\text{CFL} = 10^{-6}$ on a grid $(n^x, n^t) = (100, 19)$, yielding 5 time steps after data padding. The library contains u and all its spatial derivatives up to order 7 as well as all of those combinations of derivatives that add up to a given cumulative order for all cumulative orders up to 3. These basis functions are then multiplied by u^k for k up to 3 and an intercept is added. Our goal is to represent a first MDE, thus 2nd and 3rd order time derivatives are appended, yielding $p = 68$ candidate terms. Note the increased padding width due to the higher order time derivatives. $\Theta(\mathbf{u})$ can represent 8 terms from the first MDE up to 4th order (36) and one additional term from the 6th order truncation error, not shown in eq. (36). We constructed $\Theta(\mathbf{u})$ this way to prove that first MDEs can be found with SITE. When including all terms of $\mathcal{O}(h^4)$, we found all spatial derivatives, but could not identify v_{ttt} , which is very small (Table 3) due to the small CFL number enforced by the solver stability criterion.

Table 3 summarizes the predicted MDE of SITE in the default setting. The empirical order is calculated by a sequence of $n^x = (88, 100, 112, 125)$. Note that again, SITE did not predict any incorrect terms, but regression accuracy is significantly lower compared to both previous examples.

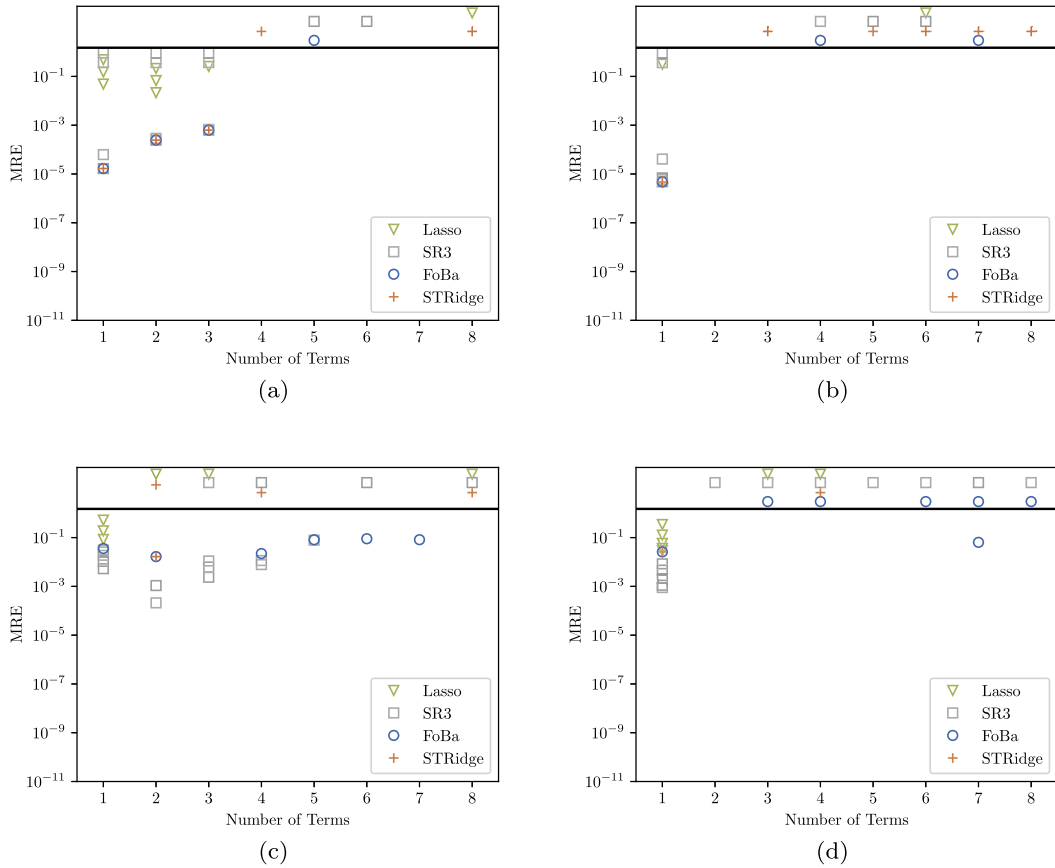


Fig. 10. Number of terms selected by BIC and by optimal choice from the model candidates proposed by FoBa as a function of resolution for the KdV equation and 5 time steps (a) or 100 time steps (b). Markers of models only containing correct terms are plotted in blue, models containing at least one incorrect term are plotted in red. BIC is represented by a cross (\times) and the optimal choice by a circle (\circ). MRE and MAE are displayed for the optimal model on the right axis.

Table 3

Summary of SITE default setup prediction for the KdV equation.

	Analytical weight	Absolute error	Relative error	Empirical order
$v v_x$	-6	$2.25 \cdot 10^{-2}$	$3.75 \cdot 10^{-3}$	-0.01
v_{xxx}	-1	$-5.08 \cdot 10^{-6}$	$5.08 \cdot 10^{-6}$	0.00
v_{ttt}	$-2.09 \cdot 10^{-15}$	$7.24 \cdot 10^{-16}$	$3.47 \cdot 10^{-1}$	2.03
v_{xxxxx}	$-2.5 \cdot 10^{-5}$	$-7.47 \cdot 10^{-9}$	$2.99 \cdot 10^{-4}$	2.00
$v v_{xxx}$	$-1 \cdot 10^{-4}$	$1.85 \cdot 10^{-5}$	$1.85 \cdot 10^{-1}$	1.96
$v_x v_{xx}$	$-2 \cdot 10^{-4}$	$5.96 \cdot 10^{-6}$	$2.98 \cdot 10^{-2}$	1.92
$v_{xxxxxxx}$	$-2.5 \cdot 10^{-10}$	$-2.06 \cdot 10^{-12}$	$8.26 \cdot 10^{-3}$	3.98

The results of the sparse regression comparison in Fig. 9 are mostly analogous to the other test cases. With applied puffer transformation, no algorithm can detect more than 3 terms. This is significantly smaller than the number of terms identifiable without puffer. Gauss initialization yields poor results in this test case: the majority of predicted models only contains one of two terms of the KdV equation.

Fig. 10(a) shows the dependency of SITE predictions on spatial resolution in default setting for 5 time steps. The optimal resolution range $n^x \in [75, 125]$ is very narrow and outside this range, the maximum number of terms in models proposed by FoBa decreases instantly. Even for $n^x \in [75, 125]$, one term within $\mathcal{O}(h^4)$ – which could be represented by $\Theta(\mathbf{u})$ – is not identified. Given that a 7th order spatial derivative is to be approximated, n^x is considerably limited due to round-off error effects. Unlike in the Burgers case, driving higher order terms towards 0 is therefore not possible, compromising the ability to identify all terms within one order of magnitude. If we were not to include the 7th spatial derivative in the library, n^x would still be limited due to the restrictive stability criterion of the Zabusky and Kruskal scheme (31). As Δt scales with Δx^3 , increasing n^x would quickly decrease Δt and therefore result in increasing round-off errors for the time derivative

approximations, for which at least 3rd-order derivatives are to be approximated. Utilizing simulation data from 100 time steps (Fig. 10(b)) instead of 5 slightly improves quality measures in this test case, including an extended optimal resolution range up to $n^x = 137$.

6. Discussion

We discuss a few guidelines obtained from the advection, Burgers and KdV test cases. The predictions of our SITE approach are determined by three factors: Simulation data, choices in the algorithmic procedure and in the model selection step. The simulation data has considerable influence, defining the signal-to-noise ratio of the regression problem. Moderate corruption of the simulation signal by noise can be mitigated by either using total variation regularized derivatives [13], or using signal-noise separating neural networks as in [47]. Δx and Δt should be chosen to balance noise from higher order truncation error terms not included in the library and noise due to round-off error of the finite difference approximations. While we mainly focused on the impact of n^x in chapter 5, we found that prediction results strongly depend on a sensible choice of Δt as well. The range of Δt and Δx for maximum accuracy predictions differs significantly between problems. This range is particularly narrow when high order spatial and temporal derivatives are being approximated as in the KdV case. Approximation of high order derivatives is a limiting factor for SITE, as round-off error substantially limits the range of simulation parameters, curtailing the opportunity to drive high order terms towards 0. Using higher precision numbers than double precision, as it is possible with e.g. FORTRAN [36] or Julia [6], has the potential for significant improvements in both regression accuracy and number of identifiable terms.

The number of simulation time steps impacts regression accuracy, the maximum number of correct terms proposed by FoBa and BIC model selection. However, the extent of impact on these three quality measures differed considerably between test cases and increasing the number of time steps did not reveal a clear positive or negative trend. The mechanisms involved are not yet fully understood and warrant further investigation.

The algorithmic setup is defined by appropriate choices of library, preconditioning and sparse regression algorithm. In library construction, using existing prior knowledge about the form of candidate functions is advisable. Reducing the number of candidates usually leads to smaller multicollinearity, which is beneficial to sparse regression. Performing an iterative approach by constructing the library over multiple successive iterations might be helpful. At this point, a word of caution is in order: In case the algorithm has uncharacteristically poor performance this may be due to an erroneous construction of the candidate term library. If a correct, strong impact term is not included in the library, we have seen that this term is then often approximated by a linear combination of incorrect terms from the library.

The test cases demonstrate that the benefits of preconditioning depend on the problem considered. Puffer transformation can improve both regression accuracy and the number of identified terms if the signal-to-noise ratio of terms included in the model is high. However, the noise inflation property quickly outweighs these benefits if multicollinearity is severe. Using spline initialization together with puffer is essential to decrease the amount of additional noise from multicollinearity. Generalizations of **F**, following Jia et al. [25], might aid in the reduction of the noise inflation effect. Spline initialization without puffer behaves the opposite way. While there is only a minor benefit in low noise problems like the advection case, the benefits are considerable in higher noise problems (Burgers, KdV). If an unconstrained choice of $u(x, 0)$ is infeasible due to problem constraints, spline initialization can be skipped.

Preconditioning choices are also linked to the selected sparse regression algorithm. For practical applications, a robust default setup seems to be FoBa with spline initialization and without the puffer transformation. Apart from robust term detection, FoBa yields highly accurate predictions even without puffer transformation. This is only surpassed by SR3 for models with a small number of terms. FoBa facilitates model selection in comparison to Lasso or SR3. While the former only proposes a few models, the latter yield different models for every set of hyperparameters. If a different algorithm than FoBa is used, puffer transformation often improves the results both in accuracy and number of identified terms. For our test cases, we did not find any evidence that STRidge was more reliable than FoBa.⁵

Model selection using information criteria has previously been introduced to the PDE-FIND framework by Mangan et al. [34]. Their approach is based on a validation set of $N \sim 100$ simulations $\{u^i(x, t)\}_{i=1}^N$ with varying initial conditions $u^i(x, 0)$. For each model from $\{M_j\}_{j=1}^K$, the dynamics of model j are integrated in time for each $u^i(x, 0)$, yielding $u_k^j(x, t)$. For each simulation in the validation set, a residual between $u^i(x, t)$ and $u_k^j(x, t)$ is calculated and the RMS error of these residuals is used as likelihood in the Akaike information criterion (AIC) [2]. Since each simulation is considered as one sample in AIC, a sufficiently large N has to be chosen, resulting in large computational effort for time integration of the candidate models. In order to avoid the computational effort and the issues related to time integration of MDEs discussed in section 2.1, we calculate BIC directly from the linear model (7) with respect to a single validation simulation. A major advantage of the

⁵ We assume that the reduced reliability of FoBa with respect to STRidge reported by Rudy et al. [46] was due to an incorrect implementation of FoBa. For our implementation, we therefore made adjustments according to the work of Zhang [59].

procedure of Mangan et al. [34] is that rather uncorrelated samples can be used in AIC, hence eliminating the need to estimate n_{eff} , which becomes a free parameter of the model selection step of SITE.

Estimating n_{eff} with n^x is very coarse: n_{eff} clearly does not increase linearly with n^x due to increased intra timestep correlation. However, we found our results to be robust with respect to n_{eff} as long as the order of magnitude was roughly correct. For problems where convection is not the dominant effect, a different choice of n_{eff} should be made on a problem specific basis. For alternatives to the proposed rule of thumb, one might consider estimation methods from the literature e.g. [4], or [56]. There is legitimate criticism aimed at the BIC due to its implicit dependence on a prior distribution which can substantially deviate from the prior distribution a considerate investigator would choose [56]. However, we found BIC to be capable of identifying the optimal model for a range of n_x in all test cases. Instead of BIC, AIC could potentially be a natural second choice, yielding similar model selection from our experience.⁶ There exists a vast amount of literature on improvements to the BIC, such as EBIC [14] and MBIC [56], allowing for custom adjustments of the BIC to the individual problem at hand.

7. Conclusion

We presented SITE, a novel data-driven approach to modified differential equation analysis. Its effectiveness in discovering first and third MDEs was demonstrated in various test cases. The current implementation of SITE is by no means optimal, as neither the preconditioning steps nor the model selection procedure have been optimized. We still showed that high quality results can be obtained, underlining the promise of the approach for extending the MDEA toolbox to discretization schemes, where analytic derivation of MDEs is infeasible. Applications that might benefit from SITE include optimization of numerical discretization schemes, the construction of dissipation schemes, and ILES turbulence modeling. A stepping stone for the approach will be a deeper understanding of its limitations, glimpses of which we were able to witness with Burgers' equation and the KdV equation. Other research directions deemed worthy of inquiry are the handling of more nonlinear problems, higher order derivatives, truncation errors in multiple dimensions, application to the finite element method and application to the finite volume method.

Acknowledgements

This project has received funding from the German Research Council (DFG) under grant agreement No. 326472365. All data and codes used in this manuscript are publicly available on Github at <https://www.github.com/tumaer/SITE>.

Appendix A




Fig. A.11. L_0 (solid line) and L_2 error (dashed line) as a function of grid resolution for MMS with empirical orders noted for each refinement step. CFL = 0.1 for FTBS and MacCormack and CFL = 10^{-10} for the Zabusky and Kruskal scheme due to the restrictive stability criterion. The errors are computed at $t_{\text{test}} = 0.1$ for FTBS and MacCormack and $t_{\text{test}} = 10^{-8}$ for the Zabusky and Kruskal scheme.

⁶ We repeated the analysis in Fig. 6 (a) using AIC instead of BIC, which yielded the exact same results.

Appendix B

Table B.4

Considered hyperparameters of sparse regression algorithms: extrema, logarithmic step size and resulting total number of models.

	Lasso	STRidge	SR3	FoBa
min 1	$\log_{10}(\lambda_{\min}) = -15$	$\log_{10}(\lambda_{\min}) = -15$	$\log_{10}(\lambda_{\min}) = -15$	$\log_{10}(\epsilon_{\min}) = -15$
max 1	$\log_{10}(\lambda_{\max}) = 14.5$	$\log_{10}(\lambda_{\max}) = 14$	$\log_{10}(\lambda_{\max}) = 14$	$\log_{10}(\epsilon_{\max}) = 14.5$
step size 1	0.5	1	1	0.5
min 2		$\log_{10}(\text{tol}_{\min}) = -16$	$\log_{10}(\gamma_{\min}) = -5$	
max 2		$\log_{10}(\text{tol}_{\max}) = 9$	$\log_{10}(\gamma_{\max}) = 5$	
step size 2		1	1	
number models	60	780	330	60

References

- [1] N.A. Adams, S. Hickel, S. Franz, Implicit subgrid-scale modeling by adaptive deconvolution, *J. Comput. Phys.* 200 (2004) 412–431.
- [2] H. Akaike, A new look at the statistical model identification, in: *Selected Papers of Hirotugu Akaike*, Springer, 1974, pp. 215–222.
- [3] M. Baer, Findiff, 2018.
- [4] J. Berger, M.J. Bayarri, L.R. Pericchi, The effective sample size, *Econom. Rev.* 33 (2014) 197–217.
- [5] J.O. Berger, L.R. Pericchi, *Objective Bayesian Methods for Model Selection: Introduction and Comparison*, Lecture Notes-Monograph Series, 2001, pp. 135–207.
- [6] J. Bezanson, S. Karpinski, V.B. Shah, A. Edelman, Julia: a fast dynamic language for technical computing, preprint, arXiv:1209.5145, 2012.
- [7] P.J. Bickel, Y. Ritov, A.B. Tsybakov, Simultaneous analysis of Lasso and Dantzig selector, *Ann. Stat.* 37 (2009) 1705–1732.
- [8] J. Bongard, H. Lipson, Automated reverse engineering of nonlinear dynamical systems, *Proc. Natl. Acad. Sci. USA* 104 (2007) 9943–9948.
- [9] J.C. Bongard, H. Lipson, Nonlinear system identification using coevolution of models and tests, *IEEE Trans. Evol. Comput.* 9 (2005) 361–384.
- [10] S.L. Brunton, J.L. Proctor, J.N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proc. Natl. Acad. Sci. USA* 113 (2016) 3932–3937.
- [11] J.-F. Cai, B. Dong, S. Osher, Z. Shen, Image restoration: total variation, wavelet frames, and beyond, *J. Am. Math. Soc.* 25 (2012) 1033–1089.
- [12] S.-C. Chang, A critical analysis of the modified equation technique of Warming and Hyett, *J. Comput. Phys.* 86 (1990) 107–126.
- [13] R. Chartrand, Numerical differentiation of noisy, nonsmooth data, *ISRN Appl. Math.* 2011 (2011).
- [14] J. Chen, Z. Chen, Extended Bayesian information criteria for model selection with large model spaces, *Biometrika* 95 (2008) 759–771.
- [15] N.R. Council, et al., *Assessing the Reliability of Complex Models: Mathematical and Statistical Foundations of Verification, Validation, and Uncertainty Quantification*, National Academies Press, 2012.
- [16] B. Dong, Q. Jiang, Z. Shen, Image restoration: wavelet frame shrinkage, nonlinear evolution PDEs, and beyond, *Multiscale Model. Simul.* 15 (2017) 606–660.
- [17] G.M. Furnival, R.W. Wilson, Regressions by leaps and bounds, *Technometrics* 16 (1974) 499–511.
- [18] G.H. Golub, P.C. Hansen, D.P. O’Leary, Tikhonov regularization and total least squares, *SIAM J. Matrix Anal. Appl.* 21 (1999) 185–194.
- [19] D.F. Griffiths, J.M. Sanz-Serna, On the scope of the method of modified equations, *SIAM J. Sci. Stat. Comput.* 7 (1986) 994–1008.
- [20] T. Hastie, R. Tibshirani, J.H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, second edition, Springer Series in Statistics, Springer, New York, 2009.
- [21] C.W. Hirt, Heuristic stability theory for finite-difference equations, *J. Comput. Phys.* 2 (1968) 339–355.
- [22] A.E. Hoerl, R.W. Kennard, Ridge regression: biased estimation for nonorthogonal problems, *Technometrics* 12 (1970) 55–67.
- [23] L. James, V. Miranda, PySwarms: a research toolkit for Particle Swarm Optimization in Python, *J. Open Source Softw.* 3 (2018) 433.
- [24] J. Jia, K. Rohe, Preconditioning to comply with the irrepresentable condition.
- [25] J. Jia, K. Rohe, Preconditioning the Lasso for sign consistency, *Electron. J. Stat.* 9 (2015) 1150–1172.
- [26] G. Klopfer, D. McRae, The nonlinear modified equation approach to analyzing finite difference schemes, in: *5th Computational Fluid Dynamics Conference*, Reston, Virginia, American Institute of Aeronautics and Astronautics, 1981, p. 429.
- [27] G.H. Klopfer, D.S. McRae, Nonlinear truncation error analysis of finite difference schemes for the Euler equations, *AIAA J.* 21 (1983) 487–494.
- [28] P.G. LeFloch, S. Mishra, Numerical methods with controlled dissipation for small-scale dependent shocks, *Acta Numer.* 23 (2014) 743–816.
- [29] A. Lerat, R. Peyret, Noncentered schemes and shock propagation problems, *Comput. Fluids* 2 (1974) 35–52.
- [30] Z. Long, Y. Lu, B. Dong, PDE-Net 2.0: learning PDEs from data with a numeric-symbolic hybrid deep network, 2018.
- [31] Z. Long, Y. Lu, X. Ma, B. Dong, PDE-Net: learning PDEs from data, in: *International Conference on Machine Learning*, 2018, pp. 3208–3216.
- [32] R. MacCormack, The effect of viscosity in hypervelocity impact cratering, in: *4th Aerodynamic Testing Conference*, Reston, Virginia, American Institute of Aeronautics and Astronautics, 1969, p. 4067.
- [33] A. Majda, S. Osher, A systematic approach for correcting nonlinear instabilities, *Numer. Math.* 30 (1978) 429–452.
- [34] N.M. Mangan, J.N. Kutz, S.L. Brunton, J.L. Proctor, Model selection for dynamical systems via sparse regression and information criteria, *Proc. R. Soc. A* 473 (2017) 20170009.
- [35] L.G. Margolin, W.J. Rider, A rationale for implicit turbulence modelling, *Int. J. Numer. Methods Fluids* 39 (2002) 821–841.
- [36] M. Metcalf, J.K. Reid, M. Cohen, *Fortran 95/2003 Explained*, vol. 416, Oxford University Press, Oxford, 2004.
- [37] A. Meurer, C.P. Smith, M. Paprocki, O. Čertík, S.B. Kirpichev, M. Rocklin, A. Kumar, S. Ivanov, J.K. Moore, S. Singh, T. Rathnayake, S. Vig, B.E. Granger, R.P. Muller, F. Bonazzi, H. Gupta, S. Vats, F. Johansson, F. Pedregosa, M.J. Curry, A.R. Terrel, Š. Roučka, A. Saboo, I. Fernando, S. Kulal, R. Cimrman, A. Scopatz, *SymPy: symbolic computing in Python*, *PeerJ Comput. Sci.* 3 (2017) e103.
- [38] A. Neumaier, Solving ill-conditioned and singular linear systems: a tutorial on regularization, *SIAM Rev.* 40 (1998) 636–666.
- [39] G.G. O’Brien, M.A. Hyman, S. Kaplan, A study of the numerical solution of partial differential equations, *J. Math. Phys.* 29 (1950) 223–251.
- [40] C.C. Paige, M.A. Saunders, LSQR: an algorithm for sparse linear equations and sparse least squares, *ACM Trans. Math. Softw.* 8 (1982) 43–71.
- [41] M. Raissi, P. Perdikaris, G.E. Karniadakis, Machine learning of linear differential equations using Gaussian processes, *J. Comput. Phys.* 348 (2017) 683–693.
- [42] M. Raissi, P. Perdikaris, G.E. Karniadakis, Numerical Gaussian processes for time-dependent and nonlinear partial differential equations, *SIAM J. Sci. Comput.* 40 (2018) A172–A198.

- [43] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* 378 (2019) 686–707.
- [44] R.D. Richtmyer, K. Norton, *Difference Methods for Initial Value Problems*, Interscience, 1967.
- [45] S. Rudy, A. Alla, S.L. Brunton, J.N. Kutz, Data-driven identification of parametric partial differential equations.
- [46] S.H. Rudy, S.L. Brunton, J.L. Proctor, J.N. Kutz, Data-driven discovery of partial differential equations, *Sci. Adv.* 3 (2017) e1602614.
- [47] S.H. Rudy, J.N. Kutz, S.L. Brunton, Deep learning of dynamics and signal-noise decomposition with time-stepping constraints, *J. Comput. Phys.* (2019).
- [48] K. Salari, P. Knupp, *Code Verification by the Method of Manufactured Solutions*.
- [49] M. Schmidt, H. Lipson, Distilling free-form natural laws from experimental data, *Science* 324 (2009) 81–85.
- [50] G. Schwarz, Estimating the dimension of a model, *Ann. Stat.* 6 (1978) 461–464.
- [51] T.R. Taha, M.I. Ablowitz, Analytical and numerical aspects of certain nonlinear evolution equations, III: numerical, Korteweg-de Vries equation, *J. Comput. Phys.* 55 (1984) 231–253.
- [52] P.C. Team, *Python: A Dynamic, Open Source Programming Language*, Python Software Foundation, 2015, 78.
- [53] R. Tibshirani, Regression shrinkage and selection via the Lasso, *J. R. Stat. Soc. B (Methodol.)* (1996) 267–288.
- [54] F.R. Villatoro, J.I. Ramos, On the method of modified equations, I: asymptotic analysis of the Euler forward difference method, *Appl. Math. Comput.* 103 (1999) 111–139.
- [55] R. Warming, B. Hyett, The modified equation approach to the stability and accuracy analysis of finite-difference methods, *J. Comput. Phys.* 14 (1974) 159–179.
- [56] D.L. Weakliem, A critique of the Bayesian information criterion for model selection, *Sociol. Methods Res.* 27 (1999) 359–397.
- [57] N.J. Zabusky, M.D. Kruskal, Interaction of solitons in a collisionless plasma and the recurrence of initial states, *Phys. Rev. Lett.* 15 (1965) 240.
- [58] T. Zeugmann, P. Poupart, J. Kennedy, X. Jin, J. Han, L. Saitta, M. Sebag, J. Peters, J.A. Bagnell, W. Daelemans, G.I. Webb, K.M. Ting, J.S. Shirabad, J. Fürnkranz, E. Hüllermeier, S. Matwin, Y. Sakakibara, P. Flener, U. Schmid, C.M. Procopiuc, N. Lachiche, Particle swarm optimization, in: C. Sammut, G.I. Webb (Eds.), *Encyclopedia of Machine Learning*, Springer reference, Springer, New York, NY, 2011, pp. 760–766.
- [59] T. Zhang, Adaptive forward-backward greedy algorithm for sparse learning with linear models, in: *Advances in Neural Information Processing Systems*, 2009, pp. 1921–1928.
- [60] P. Zhao, B. Yu, On model selection consistency of lasso, *J. Mach. Learn. Res.* 7 (2006) 2541–2563.
- [61] P. Zheng, T. Askham, S.L. Brunton, J.N. Kutz, A.Y. Aravkin, A unified framework for sparse relaxed regularized regression: SR3, *IEEE Access* 7 (2019) 1404–1423.