

# Sparse identification of nonlinear dynamics for model predictive control in the low-data limit

Eurika Kaiser<sup>a</sup>, J. Nathan Kutz<sup>b</sup>, Steven L. Brunton<sup>a</sup>

<sup>a</sup>*Department of Mechanical Engineering, University of Washington, Seattle, WA, 98195*

<sup>b</sup>*Department of Applied Mathematics, University of Washington, Seattle, WA, 98195*

---

## Abstract

The data-driven discovery of dynamics via machine learning is currently pushing the frontiers of modeling and control efforts, and it provides a tremendous opportunity to extend the reach of model predictive control. However, many leading methods in machine learning, such as neural networks, require large volumes of training data, may not be interpretable, do not easily include known constraints and symmetries, and often do not generalize beyond the attractor where models are trained. These factors limit the use of these techniques for the online identification of a model in the low-data limit, for example following an abrupt change to the system dynamics. In this work, we extend the recent sparse identification of nonlinear dynamics (SINDY) modeling procedure to include the effects of actuation and demonstrate the ability of these models to enhance the performance of model predictive control (MPC), based on limited, noisy data. SINDY models are parsimonious, identifying the fewest terms in the model needed to explain the data, making them interpretable, generalizable, and reducing the burden of training data. We show that the resulting SINDY-MPC framework has higher performance, requires significantly less data, and is more computationally efficient and robust to noise than neural network models, making it viable for online training and execution in response to rapid changes to the system. SINDY-MPC also shows improved performance over linear data-driven models, although linear models may provide a stopgap until enough data is available for SINDY.

*Key words:* Model predictive control, nonlinear dynamics, sparse identification of nonlinear dynamics (SINDY), system identification, control theory, machine learning.

---

## 1 Introduction

The data-driven modeling and control of complex systems is currently undergoing a revolution, driven by the confluence of big data, advanced algorithms in machine learning, and modern computational hardware. Model-based control strategies, such as model predictive control, are ubiquitous, relying on accurate and efficient models that capture the relevant dynamics for a given objective. Increasingly, first principles models are giving way to data-driven approaches, for example in turbulence, epidemiology, neuroscience, and finance [24]. Although these methods offer tremendous promise, there has been slow progress in distilling physical models of dynamic processes from data. Moreover, many modern techniques in machine learning (e.g., neural networks) rely on access to massive data sets, have limited ability to generalize beyond the attractor where data is collected, and do not readily incorporate known physical constraints. The challenges associated with data-driven discovery limit its use for real-time control of strongly nonlinear, high-dimensional, multi-scale systems, and prevent online recovery to abrupt changes in the dynamics. Fortunately, a new paradigm of sparse and parsimonious modeling is enabling interpretable models in the low-data limit. In this work, we extend the recent sparse identification of

nonlinear dynamics (SINDY) framework to identify models with actuation, and combine it with model predictive control (MPC) for effective and interpretable data-driven, model-based control. We apply the proposed SINDY-MPC method to control several nonlinear systems and demonstrate improved control performance in the low-data limit, compared with other leading data-driven methods, including linear response models and neural networks.

Model-based control techniques, such as MPC [12] and optimal control [48, 15], are cornerstones of advanced process control, and are well-positioned to take advantage of the data-driven revolution. Model predictive control is particularly ubiquitous in industrial applications, as it enables the control of strongly nonlinear systems with constraints, which are difficult to handle using traditional linear control approaches [17, 38, 35, 40, 39, 18, 25, 33, 16]. MPC benefits from simple and intuitive tuning and the ability to control a range of simple and complex phenomena, including systems with time delays, non-minimum phase dynamics, and instability. In addition, it is straightforward to incorporate known constraints, intrinsic compensation for dead time, multiple operating conditions, and it provides the flexibility to formulate and tailor a control objective. The major drawback of model-based control, such as MPC, is in the development of a suitable model via existing system identification or model reduction [8], which may require expensive and time-consuming data collection and computations.

Nearly all industrial applications of MPC rely on

---

*Email addresses:* eurika@uw.edu (Eurika Kaiser), kutz@uw.edu (J. Nathan Kutz), sbrunton@uw.edu (Steven L. Brunton).

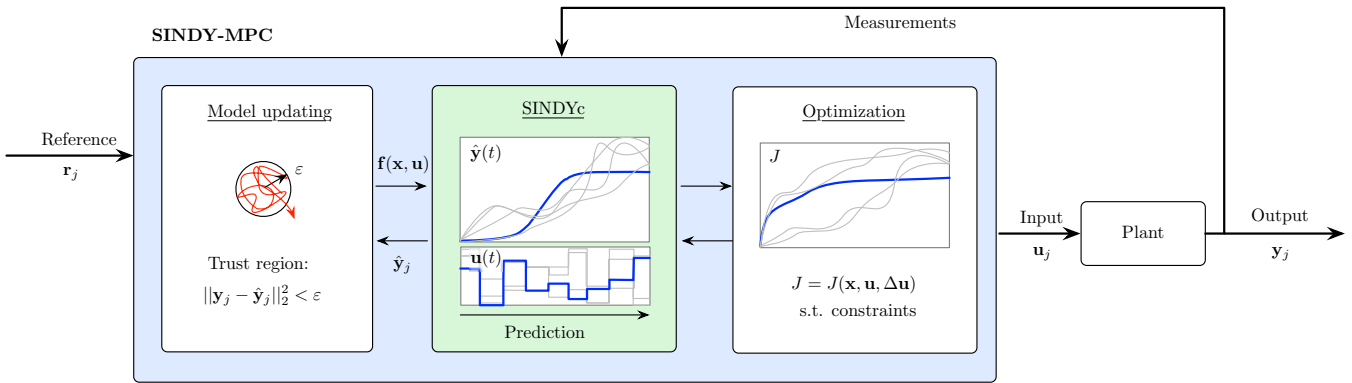


Fig. 1. Schematic overview of the proposed SINDY-MPC framework, using sparse nonlinear models for predictive control.

empirical models, and increasing plant complexity and tighter performance specifications require models with higher accuracy. There are many techniques to obtain data-driven models, including state-space models from the eigensystem realization algorithm (ERA) [21] and other subspace identification methods, Volterra series [6, 5, 30], autoregressive models [2] (e.g., ARX, ARMA, NARX, and NARMAX [3] models), and neural network models [27, 14, 54, 1], to name only a few. These procedures all tend to yield black-box models, with limited interpretability, physical insights, and ability to generalize. More recently, linear representations of nonlinear systems using extended dynamic mode decomposition [55] have been successfully paired with MPC [23], exhibiting improved performance. Nonlinear models based on machine learning, such as neural networks, are increasingly used due to advances in computing power, and recently deep reinforcement learning has been combined with MPC [36, 56], yielding impressive results in the large-data limit. However, large volumes of data are often a luxury, and many systems must be identified and controlled with limited data, for example in response to abrupt changes. Current efforts are focused on *rapid* learning based on minimal data.

There are many important open challenges associated with data-driven discovery of dynamical systems for real-time control. The foremost challenge is the reliance on large quantities of training data to generate models. When abrupt changes occur in the system, an effective controller must rapidly characterize and compensate for the new dynamics, leaving little time for discovery based on limited data. A second challenge is the ability of models to generalize beyond the training data, which is related to the ability to incorporate new information and quickly modify the model. Machine learning algorithms often suffer from overfitting and a lack of interpretability, although the application of these algorithms to physical systems offers a unique opportunity to incorporate known symmetries and constraints. These challenges point to the need for *parsimonious* and interpretable models [4, 47, 9] that may be characterized from limited data and in response to abrupt changes. Whereas traditional methods require unrealistic amounts of training data, the recently proposed SINDY framework [9] relies on sparsity-promoting optimization to identify parsimonious models from lim-

ited data, resulting in interpretable models that avoid overfitting. It has also been shown recently [28] that it is possible to enforce known physics (e.g., constraints, conservation laws, and symmetries) in the SINDY algorithm, improving stability and performance of models.

In this work, we combine SINDY with MPC for enhanced data-driven control of nonlinear systems in the low-data limit. First, we extend the SINDY architecture to identify interpretable models that include nonlinear dynamics and the effect of actuation. Next, we show the enhanced performance of SINDY-MPC compared with linear data-driven models and with neural network models. The linear models are identified using dynamic mode decomposition with control (DMDc) [37, 24], which is closely related to SINDY and traditional state-space modeling techniques such as ERA. SINDY-MPC is shown to have better prediction accuracy and control performance than neural network models, especially for small and moderate amounts of noisy data. In addition, SINDY models are less expensive to train and execute than neural network models, enabling real-time applications. SINDY-MPC also outperforms linear models for moderate amounts of data, although DMDc provides a working model in the extremely low-data limit for simple problems. Thus, in response to abrupt changes, a linear DMDc model may be used until a more accurate SINDY model is trained.

## 2 SINDY-MPC framework

The SINDY-MPC architecture combines the systematic data-driven discovery of dynamics with advanced model-based control to facilitate rapid model learning and control of strongly nonlinear systems. The overarching SINDY-MPC framework is illustrated in Fig. 1. In the following sections, we will describe the sparse identification of nonlinear dynamics with control and model predictive control algorithms.

We consider the nonlinear dynamical system

$$\frac{d}{dt}\mathbf{x} = \mathbf{f}(\mathbf{x}, \mathbf{u}), \quad \mathbf{x}(0) = \mathbf{x}_0 \quad (1)$$

with state  $\mathbf{x} \in \mathbb{R}^n$ , control input  $\mathbf{u} \in \mathbb{R}^q$ , and continuously differentiable dynamics  $\mathbf{f}(\mathbf{x}, \mathbf{u}) : \mathbb{R}^n \times \mathbb{R}^q \rightarrow \mathbb{R}^n$ .




Fig. 2. Schematic of the SINDYc algorithm. Active terms in a library of candidate nonlinearities are selected via sparse regression.

### 2.1 Sparse identification of nonlinear dynamics with control

Here, we generalize the sparse identification of nonlinear dynamics (SINDY) method [9] to include inputs and control, building on a previous conference paper [10]. SINDY identifies nonlinear dynamical systems from measurement data, relying on the fact that many systems have relatively few terms in the governing equations. Thus, sparsity promoting techniques may be used to find models that automatically balance sparsity in the number of model terms with accuracy, resulting in parsimonious models. In particular, a library of candidate nonlinear terms  $\Theta(\mathbf{x})$  is constructed, and sparse regression is used to identify the few active terms [9].

SINDY with control (SINDYc) is based on the same assumption, that Eq. (1) only has a few active terms in the dynamics. SINDY is readily generalized to include actuation, as this merely requires a larger library  $\Theta(\mathbf{x}, \mathbf{u})$  of candidate functions that include  $\mathbf{u}$ ; these functions can include nonlinear cross terms in  $\mathbf{x}$  and  $\mathbf{u}$ . Thus, we measure  $m$  snapshots of the state  $\mathbf{x}$  and the input signal  $\mathbf{u}$  in time and arrange these into two matrices:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_m \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_m \end{bmatrix}. \quad (2)$$

The library of candidate nonlinear functions,  $\Theta$  may now be evaluated using the data in  $\mathbf{X}$  and  $\mathbf{Y}$ :

$$\Theta^T(\mathbf{X}, \mathbf{Y}) = \begin{bmatrix} \mathbf{1} \\ \mathbf{X} \\ \mathbf{Y} \\ \mathbf{X} \otimes \mathbf{X} \\ \mathbf{X} \otimes \mathbf{Y} \\ \vdots \\ \sin(\mathbf{X}) \\ \sin(\mathbf{Y}) \\ \sin(\mathbf{X} \otimes \mathbf{Y}) \\ \vdots \end{bmatrix}, \quad (3)$$

where  $\mathbf{x} \otimes \mathbf{y}$  defines the vector of all product combina-

tions of the components in  $\mathbf{x}$  and  $\mathbf{y}$ . The library of candidate terms is a crucial choice in the SINDY algorithm. One strategy is to start with a basic choice, such as polynomials, and continually increase the complexity of the library by including other terms (trigonometric functions, etc.). Moreover, it is also possible to incorporate partial knowledge of the physics (fluids vs. quantum mechanics, etc.).

The system in Eq. (1) can thus be written as:

$$\dot{\mathbf{X}} = \Xi \Theta^T(\mathbf{X}, \mathbf{Y}). \quad (4)$$

The coefficients  $\Xi$  are *sparse* for most dynamical systems. Therefore, we employ sparse regression to identify a sparse  $\Xi$  corresponding to the fewest nonlinearities in our library that give good model performance:

$$\xi_k = \operatorname{argmin}_{\xi_k} \|\dot{\mathbf{X}}_k - \xi_k \Theta^T(\mathbf{X}, \mathbf{Y})\|_2 + \lambda \|\xi_k\|_1, \quad (5)$$

where  $\dot{\mathbf{X}}_k$  represents the  $k$ -th row of  $\dot{\mathbf{X}}$  and  $\xi_k$  is the  $k$ -th row of  $\Xi$ . To approximate derivatives from noisy state measurements, the SINDY algorithm uses the total variation regularized derivative [42, 13].

The  $\|\cdot\|_1$  term promotes sparsity in the coefficient vector  $\xi_k$ . This optimization may be solved using the LASSO [50] or the sequentially thresholded least squares procedure [9]. The parameter  $\lambda$  is selected to identify the Pareto optimal model that best balances low model complexity with accuracy. A coarse sweep of  $\lambda$  is performed to identify the rough order of magnitude where terms are eliminated and where error begins to increase. Then this parameter sweep may be refined, and the models on the Pareto front are evaluated using information criteria [32].

Since the original SINDY paper [9], it has been extended to include constraints and known physics [28], for example to enforce energy preserving constraints in an incompressible fluid flow. SINDY has also been extended to high dimensional systems, by identifying dynamics on principal components [9], learning partial differential equations [43, 45], and extracting dynamics on delay coordinates [7]. Robust variants of SINDY have been formulated to identify models despite large outliers and noise [51, 44].

### 2.1.1 Discovering discrete-time dynamics

In the original SINDY algorithm, it was shown that it is possible to identify discrete-time models of the form

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k). \quad (6)$$

It is also possible to extend SINDY to identify discrete-time models with inputs and control:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k). \quad (7)$$

Instead of computing derivatives, we collect a matrix  $\mathbf{X}'$  with the columns of  $\mathbf{X}$  advanced one timestep:

$$\mathbf{X}' = [\mathbf{x}_2 \ \mathbf{x}_3 \ \cdots \ \mathbf{x}_{m+1}]. \quad (8)$$

Then, the dynamics may be written as

$$\mathbf{X}' = \mathbf{\Xi} \mathbf{\Theta}^T(\mathbf{X}, \mathbf{\Upsilon}), \quad (9)$$

and the regression problem becomes

$$\xi_k = \operatorname{argmin}_{\xi_k} \|\mathbf{X}'_k - \xi_k \mathbf{\Theta}^T(\mathbf{X}, \mathbf{\Upsilon})\|_2 + \lambda \|\xi_k\|_1. \quad (10)$$

### 2.1.2 Relationship to dynamic mode decomposition

The SINDY regression is related to the dynamic mode decomposition (DMD), which originated in the fluids community to extract spatiotemporal coherent structures from large fluid data sets [41, 46, 52, 24]. DMD modes are spatially coherent and oscillate at a fixed frequency and/or growth or decay rate. Since fluids data is typically high-dimensional, DMD is built on the proper orthogonal decomposition (POD) [19], effectively recombining POD modes in a linear combination to enforce the temporal coherence. The dynamic mode decomposition has been applied to a wide range of problems including fluid mechanics, epidemiology, neuroscience, robotics, finance, and video processing [24]. Many of these applications have the ultimate goal of closed-loop feedback control.

In DMD, a similar regression is performed to identify a linear discrete-time model  $\mathbf{A}$  mapping  $\mathbf{X}$  to  $\mathbf{X}'$ :

$$\mathbf{X}' = \mathbf{A} \mathbf{X}. \quad (11)$$

Thus, SINDY reduces to DMD if formulated in discrete-time, with linear library elements in  $\mathbf{\Theta}$ , and without a sparsity-promoting  $L_1$  penalty term.

DMD was recently extended to include actuation inputs by Proctor et al [37], to disambiguate the effect of internal dynamics and control. In dynamic mode decomposition with control (DMDc), a similar regression is formed, but with the actuation input matrix  $\mathbf{\Upsilon}$ :

$$\mathbf{X}' = \mathbf{A} \mathbf{X} + \mathbf{B} \mathbf{\Upsilon}. \quad (12)$$

Thus, SINDY with control similarly reduces to DMDc un-

der certain conditions. In this work, we will use DMDc and SINDYc to discover dynamics for model predictive control. The DMDc algorithm has also been shown to be closely related to other subspace identification methods, such as the eigensystem realization algorithm [21], but designed for high-dimensional input-output data.

It is interesting to note that the extended DMD [55] regression is performed on the nonlinear library  $\mathbf{\Theta}(\mathbf{X}') = \mathbf{A} \mathbf{\Theta}(\mathbf{X})$ , and an  $L_1$  penalty may also be added. Extended DMD may also be modified to incorporate actuation inputs, and these models have recently been used effectively for model predictive control [23].

### 2.1.3 Identification of dynamics with feedback control

If the input  $\mathbf{u}$  corresponds to feedback control, so that  $\mathbf{u} = \mathbf{K}(\mathbf{x})$ , then it is impossible to disambiguate the effect of the feedback control  $\mathbf{u}$  with internal feedback terms  $\mathbf{K}(\mathbf{x})$  within the dynamical system; namely, the SINDY regression becomes ill-conditioned. In this case, we may identify the actuation  $\mathbf{u}$  as a function of the state:

$$\mathbf{\Upsilon} = \mathbf{\Xi}_u \mathbf{\Theta}^T(\mathbf{X}). \quad (13)$$

To identify the coefficients  $\mathbf{\Xi}$  in Eq. (4), we perturb the signal  $\mathbf{u}$  to allow it to be distinguished from  $\mathbf{K}(\mathbf{x})$  terms. This may be done by injecting a sufficiently large white noise signal, or occasionally kicking the system with a large impulse or step in  $\mathbf{u}$ . An interesting future direction would be to design input signals that *aid* in the identification of the dynamical system in Eq. (1) by perturbing the system in directions that yield high-value information.

### 2.2 Model predictive control

In this section, we outline the control problem and summarize key results in MPC, which is shown schematically in Fig. 3. Model predictive control solves an optimal control problem over a receding horizon, subject to system constraints, to determine the next control action. This optimization is repeated at each new timestep, and the control law is updated, as shown in Fig. 4.

The receding horizon control problem can generally be formulated as an open-loop optimization at each step, which determines the optimal sequence of control inputs  $\mathbf{u}(\cdot | \mathbf{x}_j) := \{\mathbf{u}_{j+1}, \dots, \mathbf{u}_{j+k}, \dots, \mathbf{u}_{j+m_c}\}$  given the current measurement  $\mathbf{x}_j$  over the control horizon  $T_c = m_c \Delta t$  that minimizes a cost  $J$  over the prediction horizon  $T_p = m_p \Delta t$ ;  $\Delta t$  is the timestep of the model, which may be different from the sampling time of measurements. The control horizon is generally less than or equal to the prediction horizon, so that  $T_c \leq T_p$ ; if  $T_c < T_p$ , then the input  $\mathbf{u}$  is assumed constant thereafter. The first control value  $\mathbf{u}_{j+1} := \mathbf{u}(\mathbf{u}_{j+1} | \mathbf{x}_j)$  is applied, and the optimization is reinitialized and repeated at each subsequent timestep. This results in an implicit feedback control law

$$\mathbf{K}(\mathbf{x}_j) = \mathbf{u}_{j+1}(\mathbf{x}_j), \quad (14)$$

where  $\mathbf{u}_{j+1}$  is the first in the optimized actuation sequence starting at the initial condition  $\mathbf{x}_j$ .

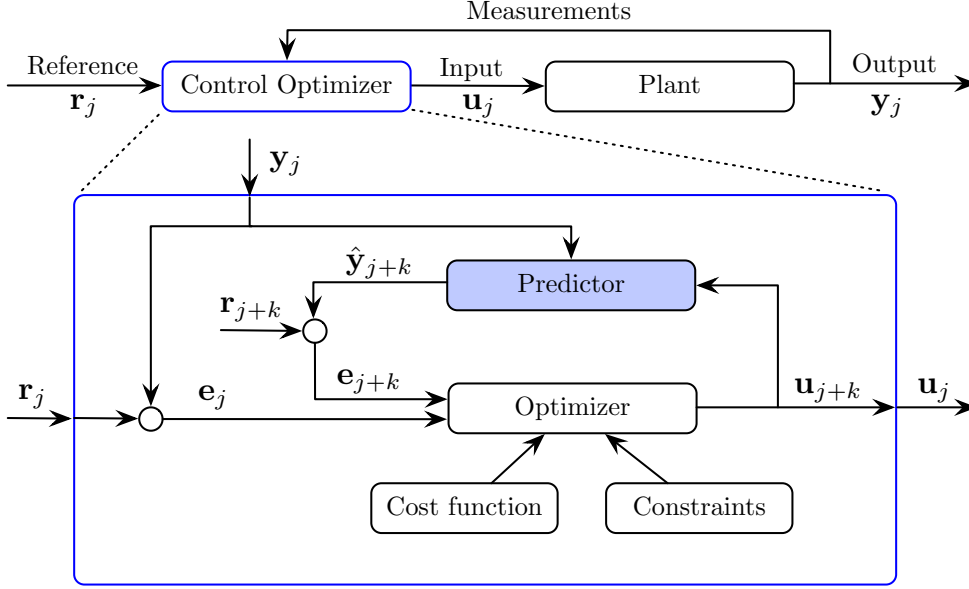


Fig. 3. Schematic for model predictive control. In SINDY-MPC, the output is the full state  $y = x$ , and  $\hat{y} = \hat{x}$  is the model prediction.




Fig. 4. Receding horizon control: The control input sequence (light blue solid) is optimized over the control horizon based on predicted future outputs (red solid) in order to drive the system to the set point (green dashed). The first input (blue star) in the sequence is enacted. Past measurements and control inputs are displayed as purple and blue solid lines, respectively.

The cost optimization at each timestep is given by:

$$\min_{\mathbf{u}(\cdot|\mathbf{x}_j)} J(\mathbf{x}_j) = \min_{\mathbf{u}(\cdot|\mathbf{x}_j)} \left[ \|\hat{\mathbf{x}}_{j+m_p}\|_{\mathbf{Q}_{m_p}}^2 + \sum_{k=0}^{m_p-1} \|\hat{\mathbf{x}}_{j+k}\|_{\mathbf{Q}}^2 + \sum_{k=1}^{m_c-1} \left( \|\hat{\mathbf{u}}_{j+k}\|_{\mathbf{R}_u}^2 + \|\Delta\hat{\mathbf{u}}_{j+k}\|_{\mathbf{R}_{\Delta u}}^2 \right) \right], \quad (15)$$

subject to the discrete-time system model

$$\hat{\mathbf{x}}_{k+1} = \hat{\mathbf{F}}(\hat{\mathbf{x}}_k, \mathbf{u}_k), \quad (16)$$

the input constraints,

$$\Delta \mathbf{u}_{min} \leq \Delta \mathbf{u}_k \leq \Delta \mathbf{u}_{max}, \quad (17a)$$

$$\mathbf{u}_{min} \leq \mathbf{u}_k \leq \mathbf{u}_{max}, \quad (17b)$$

and possibly additional equality or inequality constraints on the state and input. The cost function  $J$  penalizes deviations of the predicted state  $\hat{\mathbf{x}}_k$  along the trajectory and also includes a terminal cost at  $\hat{\mathbf{x}}_{m_p}$ . Expenditures of the input  $\mathbf{u}_k$  and input rate  $\Delta \mathbf{u}_k = \mathbf{u}_k - \mathbf{u}_{k-1}$  are also penalized. Each term is computed as the weighted norm of a vector, i.e.,  $\|\mathbf{x}\|_{\mathbf{Q}}^2 := \mathbf{x}^T \mathbf{Q} \mathbf{x}$ . The weight matrices  $\mathbf{Q} \geq 0$ ,  $\mathbf{Q}_{m_p} \geq 0$ ,  $\mathbf{R}_u > 0$  and  $\mathbf{R}_{\Delta u} > 0$  are positive definite and positive semi-definite, respectively. Note that the model prediction  $\hat{\mathbf{x}}_k$ , which is forecast, may differ from the true measured state  $\mathbf{x}_k$ .

MPC is one of the most powerful model-based control techniques due to the flexibility in the formulation of the objective functional, the ability to add constraints, and extensions to nonlinear systems. The most challenging aspect of MPC involves the identification of a dynamical model that accurately and efficiently represents the system behavior when control is applied. If the model is linear, minimization of a quadratic cost functional subject to linear constraints results in a tractable convex problem. Nonlinear models may yield significant improvements; however, they render MPC a nonlinear program, which can be expensive to solve, making it particularly challenging for real-time control. Fortunately, improvements in computing power and advanced algorithms are increasingly enabling nonlinear MPC for real-time applications.

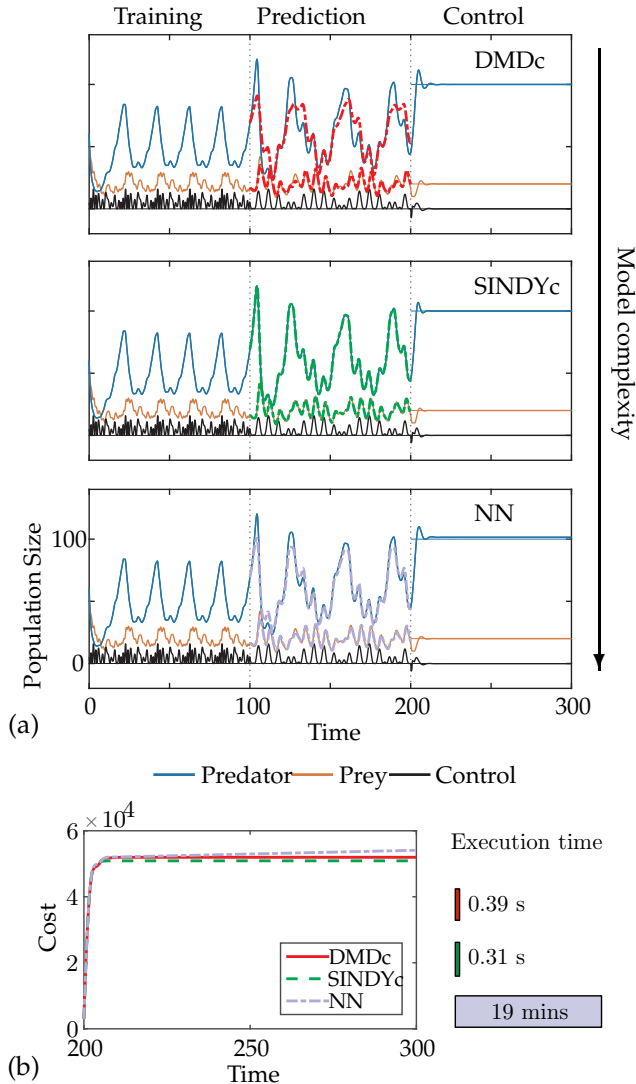


Fig. 5. Prediction and control performance for the Lotka-Volterra system: (a) time series of states and input during training, validation, and control stage, and (b) cumulative cost and execution time of the MPC optimization procedure.

### 3 Example: Lotka-Volterra model

We demonstrate the SINDY-MPC architecture on the Lotka-Volterra system, a two-dimensional, weakly nonlinear dynamical system, describing the interaction between two competing populations. These dynamics may represent two species in biological systems, the dynamic competition in stock markets [26], and they can be modified to study the spread of infectious disease [53].

We consider the system

$$\dot{x}_1 = ax_1 - bx_1x_2 \quad (18a)$$

$$\dot{x}_2 = -cx_2 + dx_1x_2 + u \quad (18b)$$

where  $x_1$  and  $x_2$  are the prey and predator populations, respectively. The constant parameters  $a = 0.5$ ,  $b = 0.025$ ,  $c = 0.5$ , and  $d = 0.005$  represent the growth/death rates, the effect of predation on the prey population, and the growth of predators based on the size of the prey population. The unforced system exhibits a limit cycle behav-

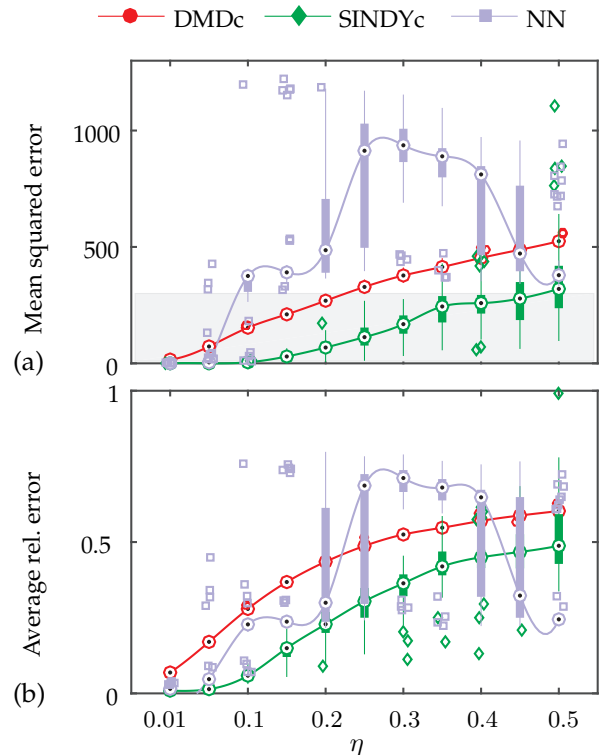


Fig. 6. Prediction performance of the models in dependency on measurement noise for the Lotka-Volterra system: (a) mean squared error and (b) average relative error. Statistics are shown for 50 noise realizations each. Above the shaded region, models in most realizations do not have any predictive power.

ior, where the predator lags the prey, and a critical point  $\mathbf{x}^{crit} = (g/d \ a/b)^T$ , where the population sizes of both species are in balance. The control objective is to stabilize this fixed point. In all examples, the timestep of the system and the models are equal to  $\Delta t = 0.1$ , the weight matrices are  $Q = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$  and  $R_u = R_{\Delta u} = 0.5$ , and the actuation input is limited to  $u \in [-20, 20]$ . The control and prediction horizons are  $m_p = m_c = 10$ . We apply an additional constraint on  $u$ , so that  $x_2$  does not decrease below 10, to enforce a minimum population size required for recovery.

To assess the performance and capabilities of the SINDY-MPC architecture, SINDYc is compared with two representative data-driven models: dynamic mode decomposition with control (DMDc) and a multilayer neural network (NN), which can represent any continuous function under mild conditions [20]. The results are displayed in Fig. 5. The first 100 time units are used to train the models using a phase-shifted sum of sinusoids, a so-called Schroeder sweep [? ], after which the predictive capabilities of these models are validated using sinusoidal forcing with  $u(t) = (2 \sin(t) \sin(t/10))^2$  on the next 100 time units. Different actuation inputs are used during the training and validation stages to assess the models' ability to generalize. Thereafter, MPC is applied for the last 100 time units using a prediction and control horizon of  $m_p = m_c = 5$ . SINDYc shows the best prediction and control performance, followed by the neural network






Fig. 7. Crossvalidated prediction error for increasing length of the training data: (a) time series of the training data, (b) average relative error, and (c) training time in seconds.

and then by DMDc. The neural network has 1 hidden layer with 10 neurons, which is the best trade-off between model complexity and accuracy; increasing the number of neurons or layers has little impact on the prediction performance. It is first trained as a feedforward network and then closed. While the neural network exhibits a similar control performance, execution time of SINDYc is 37 times faster, which is particularly critical in real-time applications.

In practice, measurements are generally affected by noise. We examine the robustness of these models for increasing noise corruption of the state measurements, i.e.  $y = x + n$  where  $n \in \mathcal{N}(0, \sigma^2)$  with standard deviation  $\sigma$ . Crossvalidated prediction performance for different noise magnitudes  $\eta = \sigma / \max(\text{std}(x_i)) \in (0.01, 0.5)$ , where  $\text{std}$  denotes standard deviation, is displayed in Fig. 6. As expected, the performance of all models decreases with increasing noise magnitude. SINDYc generally outperforms DMDc and neural network models, exhibiting a slower decline in performance for low and moderate noise levels. Sparse regression is known to improve robustness to noise and prevent overfitting. The large fluctuation in the neural network performance are due to its strong dependency on the initial network weights.

The amount of data required to train an accurate model is particularly crucial in real-time applications,

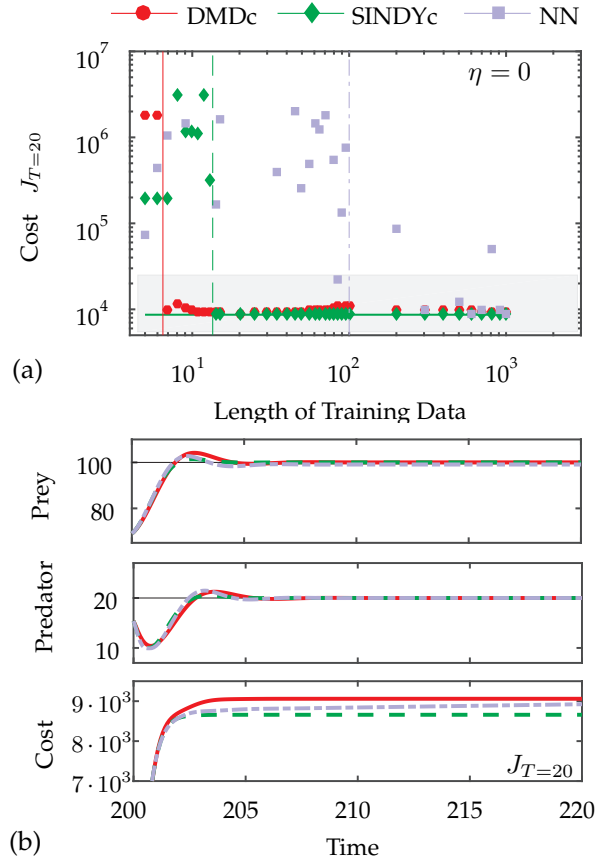


Fig. 8. Control performance for increasing length of the training data: (a) terminal cumulative cost over 20 time units, and (b) time series of states and cost of the best model for each model type ( $m_{\text{train}}^{\text{DMDc}} = 20$ ,  $m_{\text{train}}^{\text{SINDYc}} = 85$ ,  $m_{\text{train}}^{\text{NN}} = 10^3$ ). From  $m_{\text{train}} = 14$  onwards, SINDYc yields comparably well performing models, outperforming all other models. Outside the shaded region, models perform significantly worse or even diverge.

where abrupt changes or actuation may render the model invalid and rapid model updates are necessary. Figure 7 shows the average relative error of the prediction on the 100 time units used for validation, and the time to train a model in seconds for increasing lengths of training data without noise. In the low-data limit, a highly predictive SINDYc model can be learned, discovering the true governing equations within machine precision. Significantly larger amounts of data are required to train an accurate neural network model, although with enough data it outperforms DMDc. DMDc models may be useful in the extremely low-data limit, before enough data is available to characterize a SINDYc model. The training times of SINDYc and DMDc models increase slightly with the amount of data, although they require about two orders of magnitude less time than neural network models.

The effect of the training length on the control performance is shown in Fig. 8. Here, the control performance is evaluated over 20 time units. For small amounts of data, the sparsity-promoting parameter  $\lambda$  in SINDYc is reduced by a factor of 10 until a non-zero entry appears. In the extremely low-data limit, DMDc models perform

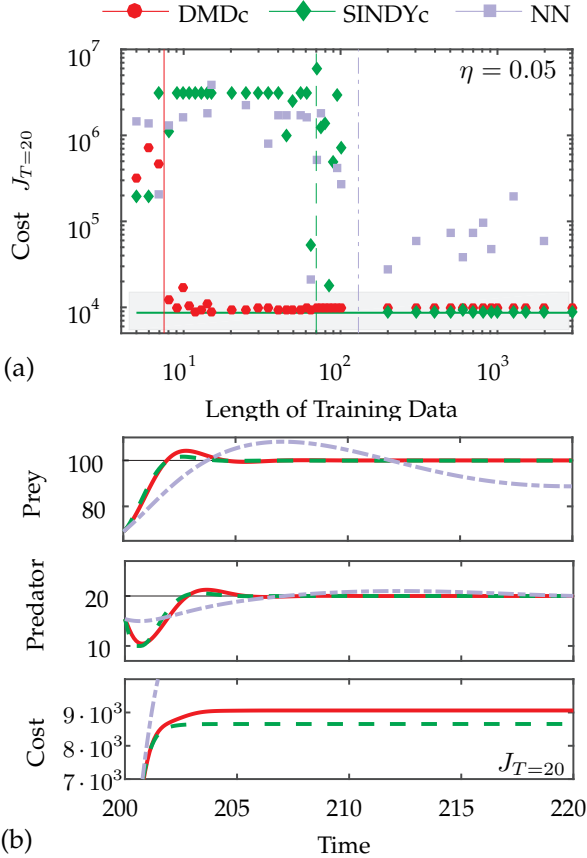


Fig. 9. Control performance for increasing length of noise-corrupted training data with  $\eta = 0.5$ : (a) terminal cumulative cost over 20 time units, and (b) time series of states and cost of the best model for each model type ( $m_{train}^{DMDc} = 12$ ,  $m_{train}^{SINDYc} = 1250$ ,  $m_{train}^{NN} = 65$ ). From  $M_{train} = 200$  onwards, SINDYc yields comparable models, outperforming all other models.

well. SINDYc models require slightly more data than DMDc, but result in the best overall performance. An order of magnitude more data is required to train comparably performing neural network models. SINDYc’s intrinsic robustness to overfitting renders all models from  $m_{train} = 14$  on as having the best control performance compared with the overall best performing DMDc and neural network models. In contrast, DMDc shows a slight decrease in performance due to overfitting and the neural network’s dependency on the initial network weights detrimentally affects its performance. It is interesting to note that the control performance is generally less sensitive than the long-term prediction performance shown in Fig. 7. Even a model with moderately low predictive accuracy may perform well in MPC.

In Fig. 9 we show the same analysis but with noise-corrupted training data. We assume no noise corruption during the control stage. For each training length, the best model out of 50 noise realizations was tested for control. DMDc and SINDYc models both require slightly more data to achieve a similar performance as without noise. Note that neural network models perform significantly

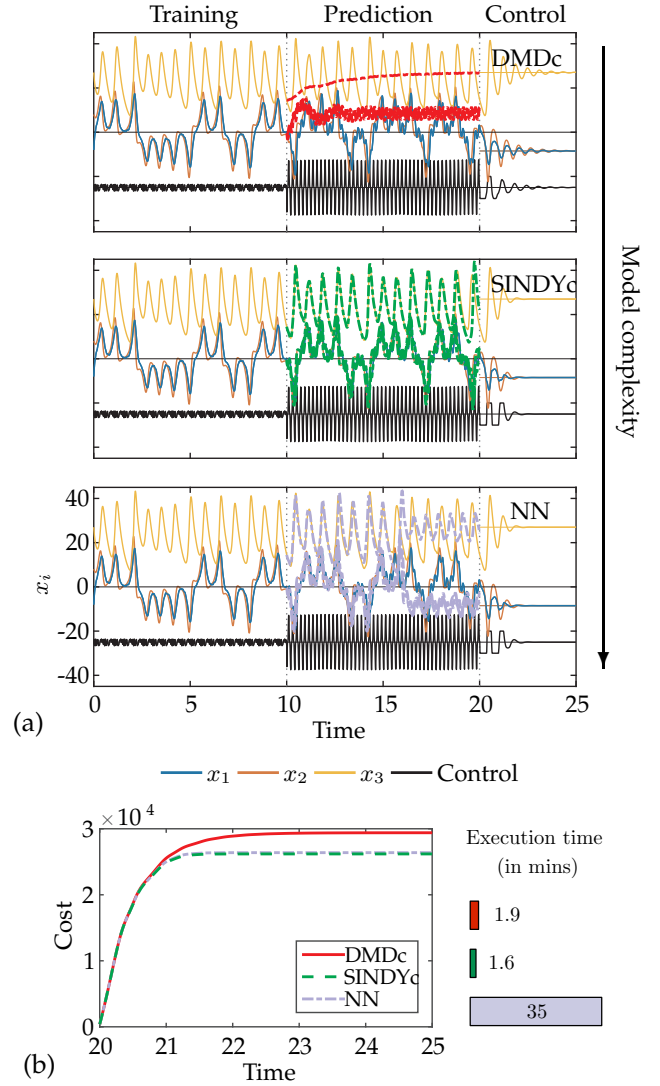


Fig. 10. Prediction and control performance for the chaotic Lorenz system: (a) time series of the states and input (shifted to  $-25$  and scaled by 10 to improve readability) during training, validation, and control stage, and (b) cumulative cost and execution time of the MPC optimization.

worse when trained on noise-corrupted data.

#### 4 Lorenz system

In this section, we demonstrate the SINDY-MPC architecture on the chaotic Lorenz system, a prototypical example of chaos in dynamical systems. The Lorenz system represents the Rayleigh-Bénard convection in fluid dynamics as proposed by Lorenz [29], but has also been associated with lasers, dynamos, and chemical reaction systems [? ]. The Lorenz dynamics are given by

$$\dot{x}_1 = \sigma(x_2 - x_1) + u \quad (19a)$$

$$\dot{x}_2 = x_1(\rho - x_3) - x_2 \quad (19b)$$

$$\dot{x}_3 = x_1x_2 - \beta x_3 \quad (19c)$$

with system parameters  $\sigma = 10$ ,  $\beta = 8/3$ ,  $\rho = 28$ , and control input  $u$  affecting only the first state. A typical trajectory oscillates alternately around the two weakly unsta-



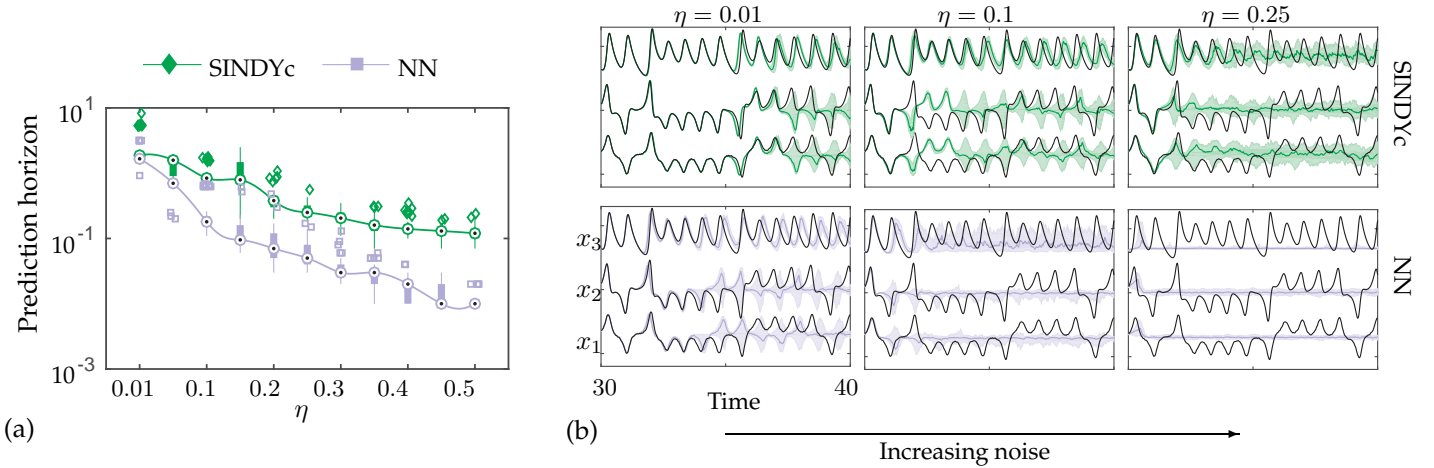


Fig. 11. Crossvalidated prediction performance with increasing measurement noise for the Lorenz system: (a) prediction horizon in time units, and (b) time series with 50 (median as thick colored line) and 25–75 (colored shaded region) percentiles. Statistics are shown for 50 noise realizations each.

ble fixed points  $(\pm\sqrt{72}, \pm\sqrt{72}, 27)^T$ . The chaotic motion of the system implies a strong sensitivity to initial conditions, i.e. small uncertainties in the state will grow exponentially with time. This represents a particularly challenging problem for model identification and subsequent control, as an uncertainty in the measurement of the state will lead inevitably to a completely different behavior of the forecast in the long run, despite having the true model. Discovering the model from (possibly noisy) measurements renders this issue more severe as model uncertainties also affect the forecast accuracy.

The control objective is to stabilize one of these fixed points. In all examples, the timestep of the system is  $\Delta t^{sys} = 0.001$  and the timestep of the model is  $\Delta t^{model} = 0.01$ . The next control input from SINDY-MPC is determined every 10 system timesteps over which the control is kept constant. In general, the timestep of the model is chosen to maximize the control horizon and minimize the length of the sequence of control inputs to be optimized, while assuring the model is as predictive as possible for the given timestep. The weight matrices are  $Q = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ ,  $R_u = R_{\Delta u} = 0.001$ , and the actuation input is limited to  $u \in [-50, 50]$ . The control and prediction horizon is  $m_p = m_c = 10$  and the sparsity-promoting parameter in SINDYc is  $\lambda = 0.1$ , unless otherwise noted. For all models we assume access to full-state information.

We compare the prediction and control performance of the SINDYc model with DMDc and a neural network (NN) model. DMDc is used to model the deviation from the goal state by constructing the regression model based on data from which the goal state has been subtracted. A less naïve approach would partition the trajectory into two bins, e.g. based on negative and positive values of  $x_1$ , and estimate two models for each goal state separately. The neural network consists of 1 hidden layer with 10 neurons and employs hyperbolic tangent sigmoid activation functions. The training of the neural network is performed using the Levenberg-Marquardt algorithm. If the data is corrupted by noise, a Bayesian regularization

is employed, which requires more training time but yields more robust models.

Cross-validated prediction and control performance for the Lorenz system are displayed in Fig. 10. The first 10 time units are used to train with a Schroeder sweep, after which the models are validated on the next 10 time units using a sinusoidally-based high-frequency forcing,  $u(t) = (5 \sin(30t))^3$ . MPC is then applied for the last 5 time units. SINDYc exhibits the best prediction and control performance. The neural network shows comparable performance, although the prediction horizon is much shorter but still sufficient for MPC. Surprisingly, DMDc is able to stabilize the fixed point, despite poor predictions based on a linear model. As DMDc shows negligible predictive capability, we will not present more DMDc results, but instead focus on the comparison between SINDYc and the neural network. As in the previous example, while the neural network exhibits similar control performance, the control execution of SINDYc is 21 times faster.

Figure 11 examines the crossvalidated prediction performance of SINDYc and neural network models based on noisy state measurements for increasing noise magnitude  $\eta = \sigma / \max(\text{std}(x_i)) \in (0.01, 0.5)$ . The performance of both models decreases with increasing noise level, although SINDY generally outperforms the neural network. Unlike the Lotka-Volterra model, the average relative error is misleading in this case. With increasing noise magnitude the neural network converges to a fixed point, having no predictive power, while SINDY still exhibits the correct statistics beyond the prediction horizon; however phase drift leads to a larger average relative error. This can be observed in Fig. 11(b), which shows the median (thick colored line) and the 25–75 percentile region (colored shaded area) of the prediction for three different noise levels. Thus, a better metric for prediction performance is the prediction horizon itself (see Fig. 11(a)). The prediction horizon, measured in time units, is estimated as the time instant when the error ball is larger than a radius of  $\varepsilon = 3$ , i.e. a model is considered predictive if

Property / Model	DMDc	SINDYc	NN
Training with limited data	<u>strong</u> Very few samples are sufficient.	<b>strong</b> Well suited for low and medium amount of data.	<b>weak</b> Requires long time series to learn predictive models.
High-dimensionality	<u>strong</u> Can handle high-dim. data in combination with SVD.	<b>fair</b> Limited by the library size.	<b>strong</b>
Nonlinearities	<b>weak/fair</b> Linear and weakly nonlinear, however with performance loss.	<b>strong</b> Suitable for strongly non-linear systems.	<b>strong</b> Suitable for strongly non-linear systems.
Prediction performance	<b>fair</b>	<u>strong</u>	<b>strong</b>
Control performance	<b>fair</b>	<u>strong</u>	<b>strong</b>
Noise robustness	<b>weak</b> High sensitivity w.r.t. noise.	<u>strong</u> Intrinsic robustness due to sparse regression.	<b>fair</b> Can handle low noise levels.
Parameter robustness	<u>strong</u>	<b>strong</b>	<b>weak</b> High sensitivity w.r.t. initial weights of the network.
Training time	<u>strong</u>	<b>strong</b>	<b>weak</b>
Execution time	<u>strong</u> Fast optimization routines exist for linear systems.	<b>strong</b>	<b>weak</b>

Table 1  
Capabilities and challenges of DMDc, SINDYc and NN models. The model with the strongest performance is underlined.

$\sqrt{\sum_{i=1}^3 (x_i - \hat{x}_i)^2} < \varepsilon$ . This corresponds to roughly 10% error per state variable, considering that the order of magnitude of each state is approximately  $\mathcal{O}(10^1)$ ; this error radius correlates well with the visible divergence of the true and predicted state in Fig. 11(b). For low and moderate noise levels, SINDYc robustly predicts the state with high accuracy. Note that even for  $\eta = 0.25$ , the 1-period prediction would be sufficiently long for a successful stabilization with MPC as we consider a comparably short prediction horizon of  $T_p = 0.1$ .

The effect of the amount of training data on the prediction and control performance is examined in Figs. 12 and 13, respectively. In Fig. 12, we show the average relative error evaluated on the prediction over the next 10 time units, the prediction horizon, and the required training time in seconds for increasing length of noise-free training data. For a relatively small amount of data, SINDYc rapidly outperforms the neural network model with a prediction horizon of 2.5 time units and a significantly smaller error. For a sufficiently large amount of data, SINDYc and the neural network result in comparable predictions. However, SINDYc yields highly predictive models that can be rapidly trained in the low and

moderate data regimes. Models trained on weakly noise-corrupted measurements,  $\eta = 0.05$ , are tested in MPC. For each length of training data, 50 noise realizations are performed and the most predictive model is selected for evaluation in MPC (Fig. 13). Outside the shaded regions, models are generally not predictive or might even diverge. In the noise-corrupted case, it is clear that SINDYc models generally have better control performance than neural network models. For a sufficiently large amount of training data, neural networks can have comparable performance to SINDYc models, although they show a sensitive dependence on the initial choice of the network weights. The control results of the neural network are significantly better here than for the Lotka-Volterra model due to the intrinsic system properties. In chaotic systems, a long enough trajectory will come arbitrarily close to every point on the attractor; thus, measurements of the Lorenz system are in some sense richer than those of the Lotka-Volterra system. A surprising result is that a nearly optimal SINDYc model can be trained on just 8 noisy measurements (compare Fig. 13).




Fig. 12. Crossvalidated prediction performance for increasing length of training data (without noise): (a) time series of the training data, (b) average relative error, (c) prediction horizon, and (d) training time in seconds.

## 5 Discussion and Conclusions

In conclusion, we have demonstrated the effective integration of data-driven sparse model discovery for model predictive control in the low-data limit. The sparse identification of nonlinear dynamics (SINDY) algorithm has been extended to discover nonlinear models with actuation and control, resulting in interpretable and parsimonious models. Moreover, because SINDY only identifies the few active terms in the dynamics, it requires less data than many other leading machine learning techniques, such as neural networks, and prevents overfitting. When integrated with model predictive control, SINDY provides computationally tractable and accurate models that can be trained on very little data. The resulting SINDY-MPC framework is capable of controlling strongly nonlinear systems, purely from measurement data, and the model identification is fast enough to discover models in real-time, even in response to abrupt changes to the model. The SINDY-MPC approach is compared with MPC based on data-driven linear models and neural network

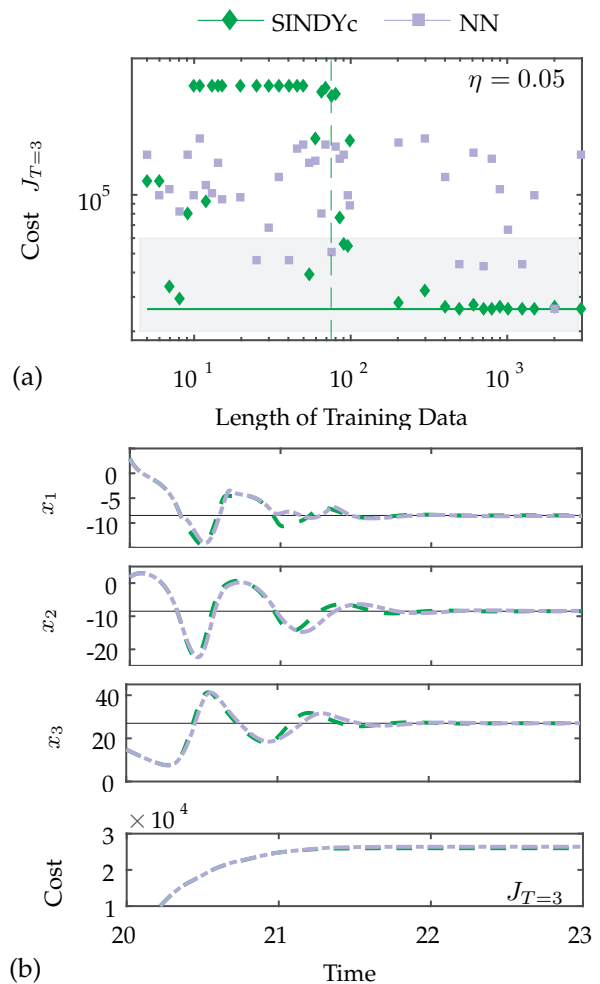


Fig. 13. Control performance for increasing length of noise-corrupted training data: (a) terminal cumulative cost over 3 time units, and (b) time series of states and cost of the best model for each model type ( $m_{train}^{SINDYc} = 38$ ,  $m_{train}^{NN} = 40$ ). Note that from  $m_{train} = 400$  onwards, SINDY identifies the best performing models.

models on two nonlinear dynamical systems.

The relative strengths and weaknesses of each method are summarized in Tab. 1. By nearly every metric, linear DMDc models and nonlinear SINDYc models outperform neural network models (NN). In fact, DMDc may be seen as the limit of SINDYc when the library of candidate terms is restricted to linear terms. SINDY-MPC provides the highest performance control and requires significantly less training data and execution time compared with NN. However, for very low amounts of training data, DMDc provides a useful model until the SINDYc algorithm has enough data to characterize the dynamics. Thus, we advocate the SINDY-MPC framework for effective and efficient nonlinear control, with DMDc as a stopgap after abrupt changes until a new SINDYc model can be identified.

This work motivates a number of future extensions and investigations. Although the preliminary application of SINDYc for MPC is encouraging, this study does not leverage many of the powerful new techniques in sparse

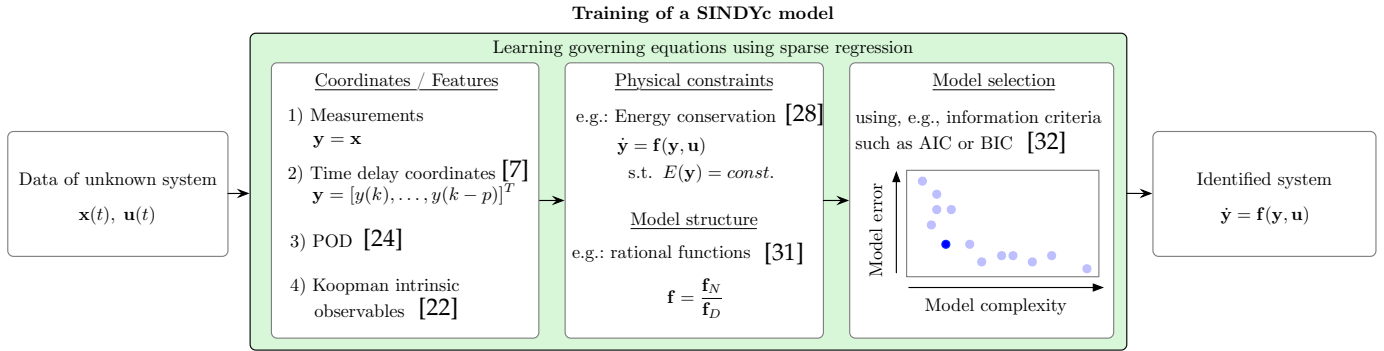


Fig. 14. Illustration of the modular nature of the SINDY with control framework and its ability to handle high-dimensional systems, limited measurements, known physical constraints, and model selection via information criteria.

model identification. Figure 14 provides a schematic of the modularity and demonstrated extensions that are possible within the SINDY framework. In realistic applications, the system may be extremely high-dimensional, and the SINDY library does not scale well with the size of the data. Fortunately, many high-dimensional systems evolve on a low-dimensional attractor, and it is often possible to identify a model on this attractor, for example by identifying a SINDY model on low-dimensional coordinates identified through a singular value decomposition [9] or manifold learning. In other applications, full-state measurements are unavailable, and the system must be characterized by limited measurements. It has recently been shown that delay coordinates provide a useful embedding to identify simple models of chaotic systems [7], building on the celebrated Takens embedding theorem [49]. Delay coordinates also define intrinsic coordinates for the Koopman operator [7], which provides a simple linear embedding of nonlinear systems [34, 11]. Koopman models have recently been used for MPC [23] and have been identified using SINDY regression [22] and subsequently used for optimal control [22]. Similar methods could be used to optimize sensors and exploit partial measurements within the SINDY-MPC framework. All of these innovations suggest a shift from the perspective of *big data* to the control-oriented perspective of *smart data*.

Figure 14 also demonstrates innovations to the SINDY regression to include physical constraints, known model structure, and model selection, which may all benefit the goal of real-time identification and control. Known symmetries, conservation laws, and constraints may be readily included in both the SINDYc and DMDC modeling frameworks [28], as they are both based on least-squares regression, possibly with sequential thresholding. It is thus possible to use a constrained least-squares algorithm, for example to enforce energy conserving constraints in a fluid system, which manifest as anti-symmetric quadratic terms [28]. Enforcing constraints has the potential to further reduce the amount of data required to identify models, as there are less free parameters to estimate, and the resulting systems have been shown to have improved stability in some cases. It is also possible to extend the SINDY

algorithm to identify models in libraries that encode richer dynamics, such as rational function nonlinearities [31]. Finally, incorporating information criteria provides an objective metric for model selection among various candidate SINDY models with a range of complexity.

The SINDY-MPC framework has significant potential for the real-time control of strongly nonlinear systems. Moreover, the rapid training and execution times indicate that SINDY models may be useful for rapid model identification in response to abrupt model changes, and this warrants further investigation. The ability to identify accurate and efficient models with small amounts of training data may be a key enabler of recovery in time-critical scenarios, such as model changes that lead to instability. In addition, for broad applicability and adoption, the SINDY modeling framework must be further investigated to characterize the effect of noise, derive error estimates, and provide conditions and guarantees of convergence. These future theoretical and analytical extensions are necessary to certify the model-based control performance.

## Acknowledgements

EK gratefully acknowledges support by the “Washington Research Foundation Fund for Innovation in Data-Intensive Discovery” and a Data Science Environments project award from the Gordon and Betty Moore Foundation (Award #2013-10-29) and the Alfred P. Sloan Foundation (Award #3835) to the University of Washington eScience Institute. SLB acknowledges support from the Army Research Office through the Young Investigator Program (W911NF-17-1-0422) and a MURI (W911NF-17-1-0306). The authors gratefully acknowledge many valuable discussions with Josh Proctor about sparse model identification and model identification for control.

## References

- [1] Eleni Aggelogiannaki and Haralambos Sarimveis. Nonlinear model predictive control for distributed parameter systems using data driven artificial neural network models. *Computers & Chemical Engineering*, 32(6):1225–1237, 2008.
- [2] Hirotugu Akaike. Fitting autoregressive models for prediction. *Annals of the institute of Statistical Mathematics*, 21(1):243–247, 1969.
- [3] Stephen A Billings. *Nonlinear system identification: NAR-*

- MAX methods in the time, frequency, and spatio-temporal domains. John Wiley & Sons, 2013.
- [4] Josh Bongard and Hod Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948, 2007.
  - [5] Stephen Boyd, Leon O Chua, and Charles A Desoer. Analytical foundations of volterra series. *IMA Journal of Mathematical Control and Information*, 1(3):243–282, 1984.
  - [6] Roger W Brockett. Volterra series and geometric control theory. *Automatica*, 12(2):167–176, 1976.
  - [7] S. L. Brunton, B. W. Brunton, J. L. Proctor, E. Kaiser, and J. N. Kutz. Chaos as an intermittently forced linear system. *Nature Communications*, 8(19):1–9, 2017.
  - [8] S. L. Brunton and B. R. Noack. Closed-loop turbulence control: Progress and challenges. *Applied Mechanics Reviews*, 67:050801–1–050801–48, 2015.
  - [9] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
  - [10] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Sparse identification of nonlinear dynamics with control (SINDYc). *IFAC NOLCOS*, 49(18):710–715, 2016.
  - [11] Marko Budišić, Ryan Mohr, and Igor Mezić. Applied Koopmanism a). *Chaos*, 22(4):047510, 2012.
  - [12] Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer Science & Business Media, 2013.
  - [13] Rick Chartrand. Numerical differentiation of noisy, nonsmooth data. *ISRN Applied Mathematics*, 2011, 2011.
  - [14] Andreas Draeger, Sebastian Engell, and Horst Ranke. Model predictive control using neural networks. *IEEE Control Systems Magazine*, 15(5):61–66, 1995.
  - [15] Geir E. Dullerud and Fernando Paganini. *A course in robust control theory: A convex approach*. Texts in Applied Mathematics. Springer, Berlin, Heidelberg, 2000.
  - [16] Utku Eren, Anna Prach, Başaran Bahadır Koçer, Saša V. Raković, Erdal Kayacan, and Behçet Açıkmeşe. Model predictive control in aerospace systems: Current state and opportunities. *Journal of Guidance, Control, and Dynamics*, 40(7):1541–1566, 2017.
  - [17] Carlos E Garcia, David M Prett, and Manfred Morari. Model predictive control: theory and practice — a survey. *Automatica*, 25(3):335–348, 1989.
  - [18] Jorge L Garriga and Masoud Soroush. Model predictive control tuning methods: A review. *Industrial & Engineering Chemistry Research*, 49(8):3505–3515, 2010.
  - [19] P. J. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge Monographs in Mechanics. Cambridge University Press, Cambridge, England, 2nd edition, 2012.
  - [20] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
  - [21] J. N. Juang and R. S. Pappa. An eigensystem realization algorithm for modal parameter identification and model reduction. *Journal of Guidance, Control, and Dynamics*, 8(5):620–627, 1985.
  - [22] E. Kaiser, J. N. Kutz, and S. L. Brunton. Data-driven discovery of Koopman eigenfunctions for control. *arXiv preprint arXiv:1707.01146*, 2017.
  - [23] Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *arXiv preprint arXiv:1611.03537*, 11 2016.
  - [24] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016.
  - [25] Jay H Lee. Model predictive control: Review of the three decades of development. *International Journal of Control, Automation and Systems*, 9(3):415–424, 2011.
  - [26] S.J. Lee, D.J. Lee, and H.S. Oh. Technological forecasting at the Korean stock market: a dynamic competition analysis using Lotka-Volterra model. *Technol. Forecast. Soc. Chang*, 72:1044–1057, 2005.
  - [27] Richard Lippmann. An introduction to computing with neural nets. *IEEE Assp magazine*, 4(2):4–22, 1987.
  - [28] J.-C. Loiseau and S. L. Brunton. Constrained sparse Galerkin regression. To appear in *Journal of Fluid Mechanics (arXiv preprint arXiv:1611.03271)*, 2017.
  - [29] Edward N Lorenz. Deterministic nonperiodic flow. *J. Atmos. Sciences*, 20(2):130–141, 1963.
  - [30] Bryon R Maner, F. J. Doyle, Babatunde A Ogunnaike, and Ronald K Pearson. A nonlinear model predictive control scheme using second order Volterra models. In *American Control Conference, 1994*, volume 3, pages 3253–3257, 1994.
  - [31] N. M. Mangan, S. L. Brunton, J. L. Proctor, and J. N. Kutz. Inferring biological networks by sparse identification of nonlinear dynamics. *IEEE Transactions on Molecular, Biological, and Multi-Scale Communications*, 2(1):52–63, 2016.
  - [32] N. M. Mangan, J. N. Kutz, S. L. Brunton, and J. L. Proctor. Model selection for dynamical systems via sparse regression and information criteria. *Proceedings of the Royal Society A*, 473(2204):1–16, 2017.
  - [33] D. Q. Mayne. Model predictive control: Recent developments and future promise. *Automatica*, 50(12):2967–2986, 2014.
  - [34] Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325, 2005.
  - [35] Manfred Morari and Jay H Lee. Model predictive control: past, present and future. *Computers & Chemical Engineering*, 23(4):667–682, 1999.
  - [36] Hui Peng, Jun Wu, Garba Inoussa, Qiulian Deng, and Kazushi Nakano. Nonlinear system modeling and predictive control using the RBF nets-based quasi-linear ARX model. *Control Engineering Practice*, 17(1):59–66, 2009.
  - [37] Joshua L Proctor, Steven L Brunton, and J Nathan Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016.
  - [38] S. J. Qin and T. A. Badgwell. An overview of industrial model predictive control technology. In *AICHE Symposium Series*, volume 93, pages 232–256, 1997.
  - [39] S Joe Qin and Thomas A Badgwell. A survey of industrial model predictive control technology. *Control engineering practice*, 11(7):733–764, 2003.
  - [40] James B Rawlings. Tutorial overview of model predictive control. *IEEE Control Systems*, 20(3):38–52, 2000.
  - [41] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D.S. Henningson. Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 645:115–127, 2009.
  - [42] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1):259–268, 1992.
  - [43] S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz. Data-driven discovery of partial differential equations. *Science Advances*, 3(e1602614), 2017.
  - [44] H. Schaeffer and S. G. McCalla. Sparse model selection via integral terms. *Physical Review E*, 96(2):023302, 2017.
  - [45] Hayden Schaeffer. Learning partial differential equations via data discovery and sparse optimization. In *Proc. R. Soc. A*, volume 473, page 20160446. The Royal Society, 2017.
  - [46] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, August 2010.
  - [47] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, 2009.
  - [48] S. Skogestad and I. Postlethwaite. *Multivariable feedback control: analysis and design*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2 edition, 2005.
  - [49] F Takens. Detecting strange attractors in turbulence. *Lecture*



*Notes in Mathematics*, 898:366–381, 1981.

- [50] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. of the Royal Statistical Society B*, pages 267–288, 1996.
- [51] Giang Tran and Rachel Ward. Exact recovery of chaotic systems from highly corrupted data. *arXiv preprint arXiv:1607.01067*, 2016.
- [52] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz. On dynamic mode decomposition: theory and applications. *Journal of Computational Dynamics*, 1(2):391–421, 2014.
- [53] E. Venturino. The influence of diseases on Lotka-Volterra systems. *Rocky Mt. J. Math.*, 24(1):381 – 402, 1994.
- [54] Tong Wang, Huijun Gao, and Jianbin Qiu. A combined adaptive neural network and nonlinear model predictive control for multirate networked industrial process control. *IEEE Transactions on Neural Networks and Learning Systems*, 27(2):416–425, 2016.
- [55] Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley. A data-driven approximation of the Koopman operator: extending dynamic mode decomposition. *J. Nonlin. Sci.*, 25(6):1307–1346, 2015.
- [56] T. Zhang, G. Kahn, S. Levine, and P. Abbeel. Learning deep control policies for autonomous aerial vehicles with MPC-guided policy search. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 528–535. IEEE, 2016.