

Differentiable Sparse Identification of Lagrangian Dynamics

Zitong Zhang, Hao Sun*

Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China
{zhangzitong, haosun}@ruc.edu.cn

Abstract

Data-driven discovery of governing equations from data remains a fundamental challenge in nonlinear dynamics. Although sparse regression techniques have advanced system identification, they struggle with rational functions and noise sensitivity in complex mechanical systems. The Lagrangian formalism offers a promising alternative, as it typically avoids rational expressions and provides a more concise representation of system dynamics. However, existing Lagrangian identification methods are significantly affected by measurement noise and limited data availability. This paper presents a novel differentiable sparse identification framework that addresses these limitations through three key contributions: (1) the first integration of cubic B-Spline approximation into Lagrangian system identification, enabling accurate representation of complex nonlinearities, (2) a robust equation discovery mechanism that effectively utilizes measurements while incorporating known physical constraints, (3) a recursive derivative computation scheme based on B-spline basis functions, effectively constraining higher-order derivatives and reducing noise sensitivity on second-order dynamical systems. The proposed method demonstrates superior performance and enables more accurate and reliable extraction of physical laws from noisy data, particularly in complex mechanical systems compared to baseline methods.

Introduction

Nonlinear dynamics plays a pivotal role across scientific and engineering domains, such as chaotic weather systems, celestial mechanics, neural oscillations, etc. Deriving governing equations directly from data has emerged as a transformative approach for understanding and predicting such complex systems. The autonomous discovery of physical laws from data has been a long-standing goal in scientific research. Deep learning provides powerful tools for modeling dynamical systems and addressing complex challenges (Battaglia et al. 2016; Lenz, Knepper, and Saxena 2015; Sahoo, Lampert, and Martius 2018). These models excel at approximating some fundamental physical principles, including Hamiltonians (Greydanus, Dzamba, and Yosinski 2019) and Lagrangians (Cranmer et al. 2020), as well as sophisticated mathematical frameworks such as

Koopman eigenfunctions (Lusch, Kutz, and Brunton 2018). These eigenfunctions transform nonlinear dynamics into a linear framework within an infinite-dimensional Hilbert space (Koopman and Neumann 1932; Kaiser, Kutz, and Brunton 2021), offering a powerful tool for the analysis and control of systems that are otherwise computationally intractable. Recent work has even combined GNNs with Koopman analysis to model molecular dynamics in an unsupervised manner (Xie et al. 2019). While deep learning models often lack interpretability, and function as black boxes, obscuring the underlying relationships between variables and providing limited insight into the mechanisms governing these relationships, interpretable models provide a compelling alternative. By offering transparent and causally interpretable insights into system behavior, these models enable reliable predictions and a profound understanding of complex dynamics. These approaches bridge a critical gap by providing methods that combine accuracy with interpretability.

Symbolic learning methods (Bongard and Lipson 2007; Quade et al. 2016) seek to extract interpretable mathematical models from data. A significant advancement in this field was achieved by (Schmidt and Lipson 2009), who introduced symbolic regression to distill physical laws from experimental data while balancing accuracy and interpretability. This approach provides a principled framework for discovering concise representations of natural phenomena. However, it scales poorly to high-dimensional systems due to the exponential growth of the search space, limiting its practical applicability. An alternative direction explores symbolic neural networks (Martius and Lampert 2016; Sahoo, Lampert, and Martius 2018; Long, Lu, and Dong 2019; Petersen et al. 2021), which employ mathematical operators as activations to create interpretable models (Mundhenk et al. 2021; Sun et al. 2023). While weight pruning enables parsimony, their reliance on user-defined thresholds compromises robustness. Furthermore, numerical differentiation of system responses exacerbates sensitivity to noise and sparse data, limiting practical applicability.

The Sparse Identification of Nonlinear Dynamics (SINDy) (Brunton, Proctor, and Kutz 2016; Rudy et al. 2017) method has emerged as a powerful framework for discovering governing equations from data. By constructing a library of candidate nonlinear terms and applying

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

sparse regression techniques. Using techniques like sequential threshold ridge regression (STRidge), SINDy iteratively refines sparse solutions through adaptive thresholding, enabling the discovery of interpretable and parsimonious models. While SINDy has demonstrated broad applicability, it struggles with dynamics involving rational functions due to the exponential growth of the candidate function library, complicating sparse regression. SINDy-PI (Kaheman, Kutz, and Brunton 2020) improves noise robustness through convex optimization but is limited to low noise levels and risks instability if incorrect denominator terms are identified. Similarly, RK4-SINDy (Goyal and Benner 2022), which combines Runge-Kutta integration with sparse identification, has shown promise for simple rational systems but may fail in complex scenarios due to potential errors in denominator selection, leading to unstable predictions. These limitations call for more robust methods to handle rational functions. Physics-informed learning helps recover system states and estimate derivatives from sparse, noisy data for equation discovery. Methods like neural networks (Sun et al. 2022) and cubic splines (Sun, Liu, and Sun 2021; Zhang, Liu, and Sun 2024) are commonly used.

Lagrangian provides a compact representation of system dynamics, encapsulating all predictive information within a single scalar function, unlike the more complex structure of differential equations. Lagrangian-SINDy (Chu and Hayashibe 2020) formulates the total energy as a Koopman eigenfunction with zero eigenvalue and applies sparse regression to align the time derivative of this energy with the system’s net power input, enabling accurate recovery of the Lagrangian structure. xL-SINDy (Purnomo and Hayashibe 2023), an extension of Lagrangian-SINDy, integrates the SINDy framework with the proximal gradient method to derive sparse Lagrangian representations. However, current Lagrangian-based identification methods remain limited by their susceptibility to noise or their dependence on substantial data availability due to the chain rule differentiation required for library functions.

Contributions: The key contributions of this work are as follows: (1) We bridge this gap by introducing a novel framework that combines Cubic B-Spline approximation with sparse regression for Lagrangian system identification, providing a flexible framework for accurately modeling complex nonlinearities. (2) We address critical limitations of existing Lagrangian-based methods by significantly reducing their sensitivity to noise and reliance on large data volumes. This is achieved through a novel formulation that minimizes the adverse effects of chain rule differentiation on library functions. (3) We develop a robust equation discovery framework that efficiently leverages limited measurements while incorporating physical constraints, enabling precise reconstruction of governing equations via sparse regression. By leveraging the recursive properties of B-spline basis functions, we enable robust derivative computation and noise mitigation, particularly for second-order dynamical systems. This advancement opens new possibilities for extracting physical laws from measurements, with potential applications in robotics, biomechanics, and beyond.

Methodology

We introduce a framework for differentiable sparse identification of lagrangian dynamics. Figure 1 depicts the overall architecture of our method.

Cubic B-Spline Approximation

B-splines are piecewise polynomial basis functions with inherent differentiability. Cubic B-splines provide a differentiable surrogate framework, offering a flexible approach for curve fitting. Since the first derivative of a B-spline curve is another B-spline curve, higher-order derivatives can be computed recursively by applying this property. Consequently, the first and second derivatives of each basis function can be efficiently obtained (see supplementary material for derivations).

To ensure the curve possesses continuous and smooth tangent directions at the starting and ending nodes, meeting the first derivative interpolation requirement, we use clamped B-spline curves for fitting. This approach enforces the curve to pass through the endpoints by repeating the boundary knots $p + 1$ times, where p is the degree of the B-spline, ensuring both endpoint interpolation and smoothness across the curve. In the context of B-spline curves, the term P denotes the control points, which are a set of points used to define the shape of the curve. The first and last control points are set to the start and end points of the data, respectively: $P_0 = D_0, P_s = D_m$, where m and s denote the number of data points and control points, and D_0 and D_m are the first and last data points. For a B-spline of degree p , the first and last knots are repeated. For example, in the case of a cubic B-spline ($p = 3$), the knot vector U is constructed as: $U = [u_0, u_0, u_0, u_0, u_1, u_2, \dots, u_{r-1}, u_r, u_r, u_r, u_r]$. This ensures sufficient support for the boundary basis functions. Mathematically, a B-spline curve $C(u)$ of degree p is expressed as a weighted sum of control points P_i and basis functions $N_{i,p}(u)$:

$$C(u) = \sum_{i=0}^s N_{i,p}(u) \cdot P_i, \quad (1)$$

where P_i represents the i -th control point, $N_{i,p}(u)$ is the i -th B-spline basis function of degree p , u is the parameter within the knot vector U . The control points P_i play a crucial role in determining the geometry of the curve. By adjusting their positions, the shape of the B-spline curve can be precisely controlled. Additionally, the use of clamped boundary conditions ensures that the curve interpolates specific control points, such as the start point P_0 and the end point P_s .

Network Architecture

Formulation of Lagrangian Acquisition for Dynamical Systems: For a multidegree-of-freedom nonrelativistic system, the Lagrangian \mathcal{L} is given by:

$$\mathcal{L} = \sum_i (T_i - V_i), \quad (2)$$

where T_i and V_i denote the kinetic and potential energies of the i -th component, respectively. Given that the total Lagrangian is the sum of the Lagrangians of its individual components, any nonlinear terms appearing in the Lagrangian

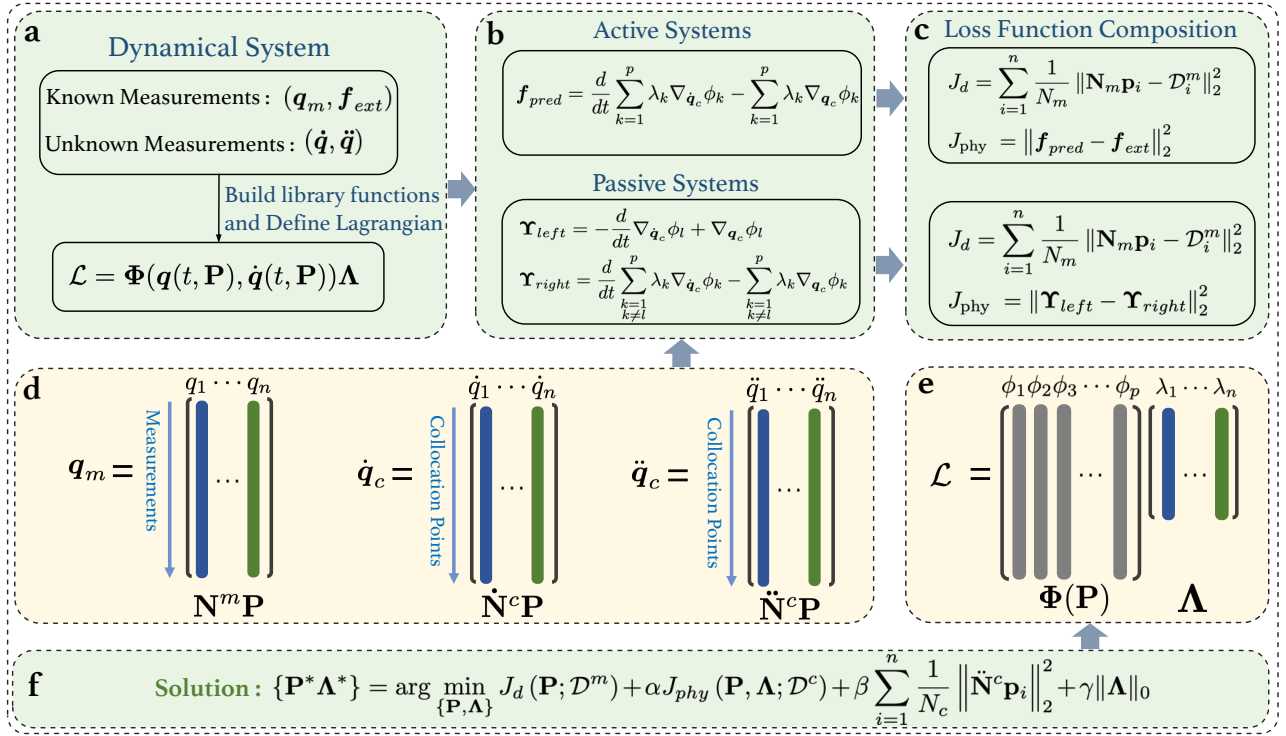


Figure 1: Schematic architecture of differentiable sparse identification of Lagrangian dynamics. (a) Measurement of the dynamical system and definition of the Lagrangian equations; (b) Description of equation formulations across different systems; (c) Loss functions under various system configurations; (d) Representation of values and their derivatives (first and second order) using cubic B-splines for each degree of freedom; (e) Schematic illustration of the sparse representation of the Euler-Lagrange equations; (f) Network is equivalent to solving optimization problem.

of a component will inherently be reflected in the total Lagrangian of the system.

For a system with n degrees of freedom, we formulate the Lagrangian as a linear combination of nonlinear candidate functions. Let $\mathbf{q} = (q_1, q_2, \dots, q_n)$ denote the generalized coordinates of the system. $q_i(t)$ is a function describing the time evolution of the i -th generalized coordinate. The Euler-Lagrange equations govern the dynamics of the Lagrangian system, as described by:

$$\mathcal{L} = \sum_{k=1}^p \lambda_k \phi_k(\mathbf{q}, \dot{\mathbf{q}}) = \Phi(\mathbf{q}, \dot{\mathbf{q}})\Lambda. \quad (3)$$

Here, $\Phi(\mathbf{q}, \dot{\mathbf{q}})$ is the library of nonlinear candidate functions, with each column being a nonlinear candidate function $\phi_k(\mathbf{q}, \dot{\mathbf{q}})$. The vector $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)^\top$ contains coefficients weighting each function in the Lagrangian \mathcal{L} .

We investigate two distinct scenarios: (1) active systems with external input \mathbf{f}_{ext} , and (2) passive systems where some prior knowledge of the Lagrangian is available. Schematic architecture of our approach and problem formulation for each scenario is summarized in Figure 1.

Active dynamical systems Active systems are characterized by non-conservative forces (e.g., control inputs or driving forces) through external input $\mathbf{f}_{ext} = (f_1, f_2, \dots, f_n)^\top$.

The equation for active systems is:

$$\frac{d}{dt} \nabla_{\dot{\mathbf{q}}} \mathcal{L} - \nabla_{\mathbf{q}} \mathcal{L} = \mathbf{f}_{ext}, \quad (4)$$

where $\nabla_{\dot{\mathbf{q}}} \equiv \frac{\partial}{\partial \dot{\mathbf{q}}}$ and $\nabla_{\mathbf{q}} \equiv \frac{\partial}{\partial \mathbf{q}}$. Incorporating Eq. (3) into Eq. (4) for active systems gives:

$$\mathbf{f}_{pred} = \frac{d}{dt} \sum_{k=1}^p \lambda_k \nabla_{\dot{\mathbf{q}}} \phi_k - \sum_{k=1}^p \lambda_k \nabla_{\mathbf{q}} \phi_k, \quad (5)$$

where \mathbf{f}_{pred} denote the predicted value of the external input \mathbf{f}_{ext} based on a coefficient vector $(\lambda_1, \lambda_2, \dots, \lambda_n)$. The time derivative $\frac{d}{dt}$ can be further expanded using the chain rule, yielding terms that involve $\dot{\mathbf{q}}$ and $\ddot{\mathbf{q}}$ as follows:

$$\mathbf{f}_{pred} = \sum_{k=1}^p \lambda_k (\nabla_{\dot{\mathbf{q}}}^\top \nabla_{\dot{\mathbf{q}}} \phi_k \ddot{\mathbf{q}} + \nabla_{\mathbf{q}}^\top \nabla_{\dot{\mathbf{q}}} \phi_k \dot{\mathbf{q}} - \nabla_{\mathbf{q}} \phi_k). \quad (6)$$

Passive dynamical systems Passive systems rely solely on internal energy (e.g., potential and kinetic energy) without external input. They are conservative systems where energy is conserved, and the generalized force f_i is zero:

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{\mathbf{q}}} \right) - \frac{\partial \mathcal{L}}{\partial \mathbf{q}} = 0. \quad (7)$$

For passive systems, certain terms in the Lagrangian equation can still be determined based on prior knowledge of the

system, such as the mass of a component, gravitational constant, elastic coefficient, length, and mass of objects. Furthermore, as indicated by Eq. (2), the Lagrangian of the system is constructed as the sum of the kinetic energy minus the potential energy for all components in the system. Since the Lagrangian is not unique, we normalize the coefficients of the known terms to unity for simplicity. Equation (3) can be expressed as follows:

$$\mathcal{L} = \phi_l(\mathbf{q}, \dot{\mathbf{q}}) + \sum_{\substack{k=1 \\ k \neq l}}^p \lambda_k \phi_k(\mathbf{q}, \dot{\mathbf{q}}). \quad (8)$$

Therefore, we can extract prior knowledge of a constituent of the system and focus on identifying the coefficients of the remaining terms.

Network Architecture Description: The network architecture overview, as illustrated in Figure 1. Initially, we define n sets of control points for cubic B-splines, represented as $\mathbf{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n\} \in \mathbb{R}^{s \times n}$. These control points are then combined with the spline basis functions $\mathbf{N}(t)$ to interpolate the system’s n -dimensional state of system: $\mathbf{q}(t, \mathbf{P}) = \mathbf{N}(t)\mathbf{P}$. Similarly, the first and second derivatives of the generalized coordinates can be expressed as: $\dot{\mathbf{q}}(t, \mathbf{P}) = \dot{\mathbf{N}}(t)\mathbf{P}$, $\ddot{\mathbf{q}}(t, \mathbf{P}) = \ddot{\mathbf{N}}(t)\mathbf{P}$. The Lagrangian \mathcal{L} in Eq. 3 can be rewritten in the following form:

$$\mathcal{L} = \Phi(\mathbf{q}(t, \mathbf{P}), \dot{\mathbf{q}}(t, \mathbf{P}))\Lambda. \quad (9)$$

The discovery problem can be formally defined as follows: given the dataset $\mathcal{D}^m = \{\mathbf{q}_m^i\}_{i=1, \dots, n} \in \mathbb{R}^{N_m \times n}$, the goal is to determine the optimal parameters \mathbf{P} and Λ such that Eq. (9) is satisfied for $\forall t$. Here, N_m represents the number of measurement points. $\mathcal{D}^c = \{t_0, t_1, \dots, t_{N_c-1}\}$ represents a set of N_c randomly sampled collocation points, where $N_c \gg N_m$. These points are used to improve the satisfaction of physical constraints. The matrix $\mathbf{N}_m \in \mathbb{R}^{N_m \times s}$ represents the spline basis matrix evaluated at the measured time instances, while $\dot{\mathbf{N}}_c \in \mathbb{R}^{N_c \times s}$ denotes the derivative of the spline basis matrix evaluated at the collocation instances.

The loss function of the proposed network consists of two components: the *data loss* (J_d) and the *physics loss* (J_{phy}). As illustrated in Figure 1c, while the data loss is formulated identically for both cases, the physics loss differs between active and passive systems. For active systems, the physics loss minimizes the difference between the predicted external force \mathbf{f}_{pred} and the known external force \mathbf{f}_{ext} . For passive systems, it minimizes the discrepancy between the terms involving prior knowledge Υ_{left} and the terms to be determined Υ_{right} . Mathematically, training the network is equivalent to solving the optimization problem (see Figure 1f), where α , β , and γ are weighting coefficients. To promote optimization, constrain higher-order derivatives, and mitigate noise amplification caused by chain rule differentiation—which exacerbates noise sensitivity in second-order dynamical systems—we introduce the regularization term $J_{reg} = \beta \sum_{i=1}^n \frac{1}{N_c} \|\ddot{\mathbf{N}}_c \mathbf{p}_i\|_2^2$, ensuring better alignment with physical principles and numerically stable.

Inspired by the STLS (Sequential Thresholded Least-Squares) approach in SINDy, we adopt an iterative proce-

Algorithm 1: Differentiable Sparse Identification of Lagrangian Dynamics

Input: $\mathcal{D}^m, \mathcal{D}^c, \mathbf{f}_{ext}$ or known terms

Parameter: Threshold ϵ , weights α, β, γ

Output: Discovered \mathcal{L}^*

- 1: Construct nonlinear candidate function library Φ . Initialize B-spline control points \mathbf{P} and library coefficients Λ
 - 2: Construct basis matrices $\mathbf{N}_m, \dot{\mathbf{N}}_c, \ddot{\mathbf{N}}_c$
 - 3: **repeat**
 - 4: Compute: $\mathbf{q} = \mathbf{N}_m \mathbf{P}, \dot{\mathbf{q}} = \dot{\mathbf{N}}_c \mathbf{P}, \ddot{\mathbf{q}} = \ddot{\mathbf{N}}_c \mathbf{P}$
 - 5: Evaluate losses:
 - 6: $J = J_d + J_{phy} + J_{reg}$
 - 7: Update \mathbf{P}, Λ via gradient descent on J
 - 8: Apply STLS thresholding: $\Lambda_{ij} \leftarrow 0$ if $|\Lambda_{ij}| < \epsilon$
 - 9: **until** convergence
 - 10: **return** $\mathcal{L}^* = \Phi(\mathbf{q}, \dot{\mathbf{q}})\Lambda^*$
-

dure to identify sparse dynamics. The process begins by fitting the coefficients using ordinary least squares. A thresholding step is then applied to eliminate terms with coefficients below a predefined threshold, effectively enforcing sparsity. The algorithm iterates by refitting the remaining terms and reapplying the threshold until convergence is achieved. This approach ensures that only the most significant terms are retained in the final model, resulting in a parsimonious representation of the system’s dynamics. By significantly reducing the number of candidate functions considered during learning, this step accelerates convergence compared to methods without hard thresholding. Finally, we verify whether the cost function has reached the specified tolerance. By substituting the optimal solutions \mathbf{P}^* and Λ^* into Eq. (9), we obtain the Lagrangian of the dynamical system. This overall training scheme, including initialization, optimization loops, and convergence criteria, is systematically outlined in Algorithm 1.

Experiments

To evaluate the efficacy of our method, we conducted a series of comprehensive experiments on a wide range of ideal dynamical systems, as illustrated in Figure 2. The datasets capture essential second-order dynamical system behaviors, verifying the proposed modeling methodology. These include both active and passive systems: the active systems consist of single pendulum, double pendulum, and spherical pendulum, while the passive systems include chaotic pendulum, cart-pendulum with a spring, spherical pendulum with a spring and magnetic pendulum. For Euler-Lagrange systems, active systems are characterized by the presence of external forces or control inputs \mathbf{f}_{ext} . These external forces are typically applied to the generalized coordinates (degrees of freedom) of the system, enabling the modification of its dynamical behavior. In contrast, passive systems do not have external forces but leverage prior knowledge, e.g., gravitational potential energy, to model their dynamics.

Experimental Setup

This section details the experimental setup, including the datasets, baselines and metrics.

Datasets. We selected datasets consisting of second-order dynamical systems. All data were generated analytically by predefining ideal dynamical systems and simulating their trajectories using the fourth-order Runge-Kutta method.

Baselines. To demonstrate the effectiveness of proposed method, we compared it with the following baselines: uDSR (Landajuela et al. 2022): A unified symbolic regression framework combines some strategies for complementary benefits. xL-SINDY (Purnomo and Hayashibe 2023): A work integrates SINDy with classical mechanics to identify sparse, interpretable expressions for Lagrangian dynamics.

Evaluation metrics. To evaluate the performance of our method, we adopt several metrics. Our objective is to accurately identify all relevant terms in the governing equations while minimizing the inclusion of irrelevant terms. The relative error, denoted as ℓ_2 , is defined as: $\ell_2 = \|\mathbf{\Lambda}_{\text{id}} - \mathbf{\Lambda}_{\text{true}}\|_2 / \|\mathbf{\Lambda}_{\text{true}}\|_2$, where $\mathbf{\Lambda}_{\text{id}}$ represents the identified coefficients and $\mathbf{\Lambda}_{\text{true}}$ represents the ground truth. To address potential bias caused by significant disparities in coefficient magnitudes, we introduce a non-dimensional measure for a more balanced assessment. The task of discovering governing equations can be viewed as a binary classification problem (Rao et al. 2022), where the goal is to determine the presence or absence of each term in a candidate library. To this end, we employ precision (P) and recall (R) as evaluation metrics, defined as: $P = \|\mathbf{\Lambda}_{\text{id}} \odot \mathbf{\Lambda}_{\text{true}}\|_0 / \|\mathbf{\Lambda}_{\text{id}}\|_0$, $R = \|\mathbf{\Lambda}_{\text{id}} \odot \mathbf{\Lambda}_{\text{true}}\|_0 / \|\mathbf{\Lambda}_{\text{true}}\|_0$, where \odot denotes element-wise product. A term is successfully identified if both $\mathbf{\Lambda}_{\text{id}}$ and $\mathbf{\Lambda}_{\text{true}}$ have non-zero entries.

These systems are analyzed under ideal conditions, where all rods except those in the chaotic pendulum are assumed to be massless. In the chaotic pendulum, the pivots of two rods are positioned differently along their rods, leading to variations in dynamical system’s governing equation. The Euler-Lagrange equations for all dynamical systems discussed in this work are provided in supplementary material. The initial conditions for all experiments are randomly sampled from a uniform distribution within a specified range for each dynamical system, with each condition simulated for a duration of 20 seconds. The candidate library is constructed following the principle of minimal completeness, ensuring that redundant or interfering terms (e.g., simultaneously including $\sin^2 \theta$ and $\cos^2 \theta$) are excluded to maintain a concise and effective representation. We evaluated the proposed method using data corrupted by zero-mean white Gaussian noise $\mathcal{N}(0, \sigma)$ at varying noise levels. The noise was added exclusively to the measured data. For active systems, we assumed that the time-dependent external input is known. For passive systems, we incorporate prior knowledge of a constituent of the system by selecting one of the terms in the total Lagrangian as a known quantity. This allows us to constrain the optimization problem and improve the accuracy of the identified dynamics.

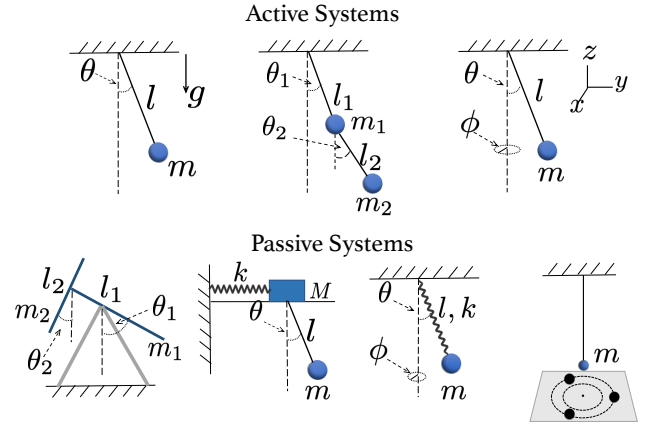


Figure 2: Two types of dynamical systems are considered, as shown from left to right: active systems, including single pendulum, double pendulum, and spherical pendulum; and passive systems, including chaotic pendulum, cart-pendulum with a spring, spherical pendulum with a spring and magnetic pendulum.

Main Results

Discovery results. Based on our evaluation metrics, a detailed analysis of the experimental results achieved by our method is provided in Table 1. We evaluate our method in comparison with xL-SINDy and uDSR as baseline models. The results listed are averaged over five independent trials. The evaluation metrics computed from the discovered Lagrangians of different dynamical systems under various noise levels are provided. The better performing values under the same metrics are highlighted in bold. Experimental results show that the performance of both models is nearly identical under zero-noise measurement conditions. Although the (ℓ_2) still exhibit small errors, we consider these errors to be within a reasonable range. Under higher noise levels, our model demonstrates significantly lower relative error (ℓ_2) and achieves competitive performance in terms of precision (P) and recall (R). Overall, our method exhibits better robustness and accuracy under noise conditions. Using the magnetical pendulum as an example, we validate our dynamical model by comparing its predictions with the real system (Figure 3ab). Due to chaotic divergence, precise long-term prediction is fundamentally impossible. Although trajectories eventually diverge, our model preserves key statistical properties (amplitude and frequency), maintaining behavior similar to the real system. After equation discovery, we incorporate a damping term to visualize the basin of attraction diagram. It shows a close match in capturing the chaotic dynamics. Figure 3cd compares the true basin of attraction with the basin from our discovered model, showing accurately captured chaotic dynamics. Furthermore, our approach demonstrates exceptional suitability for scenarios involving randomly sampled or incomplete data, a common challenge in real-world applications. To rigorously evaluate the robustness of our method against data imperfections, we conducted a comprehensive set of comparative experi-

Cases	Metrics	ℓ_2 Error $\times 10^{-2}$				P (%)				R (%)			
		Noise level	0%	1%	10%	avg.	0%	1%	10%	avg.	0%	1%	10%
Single pendulum	uDSR	162.44	160.00	159.66	160.7	40	33.33	33.33	35.55	100	100	100	100
	xL-SINDy	0.31	1.53	26.72	9.52	100	100	100	100	100	100	100	100
	Ours	0.41	1.31	10.94	4.22	100	100	100	100	100	100	100	100
Double Pendulum	uDSR	119.40	106.42	105.48	110.43	37.50	27.27	22.22	29.00	60	40	20	40
	xL-SINDy	0.30	45.23	73.42	39.65	100	100	100	100	100	100	100	100
	Ours	1.15	20.40	21.67	14.41	100	100	100	100	100	100	100	100
Spherical pendulum	uDSR	131.57	134.41	165.44	143.81	8.33	7.69	7.14	7.72	33.33	33.33	33.33	33.33
	xL-SINDy	0.51	120.20	83.93	68.21	100	60	50	70	100	100	100	100
	Ours	2.00	31.67	40.51	24.73	100	100	100	100	100	100	100	100
Chaos pendulum	uDSR	223.9	280.83	170.82	225.18	80	40	27.27	49.09	80	80	60	66.67
	xL-SINDy	0.51	5.62	9.12	5.08	100	83.33	87.50	90.28	100	100	85.71	95.24
	Ours	0.16	2.10	3.27	1.84	100	100	100	100	100	100	100	100
Cart Spring Pendulum	uDSR	80.35	81.26	96.92	86.18	75	75	27.27	59.09	75	75	42.86	64.29
	xL-SINDy	1.34	1.55	16.10	6.33	100	87.50	87.50	91.67	100	100	100	100
	Ours	1.42	1.12	2.42	1.65	100	100	87.50	95.83	100	100	100	100
Sphere Spring Pendulum	uDSR	135.28	241.79	333.51	236.86	22.22	1.92	2.50	8.88	57.14	14.29	14.29	28.57
	xL-SINDy	0.29	2.75	3.28	2.11	100	87.50	87.50	91.67	100	71.43	71.43	80.95
	Ours	0.39	2.60	3.11	2.03	100	87.50	87.50	91.67	100	100	85.71	95.23
Magnetical pendulum	uDSR	314.15	322.40	291.81	309.45	44.44	40	30.77	38.40	57.14	57.14	57.14	57.14
	xL-SINDy	1.56	14.33	34.22	16.70	100	100	75	91.67	100	85.71	85.71	90.47
	Ours	1.37	3.60	27.10	10.69	100	100	100	100	100	100	100	100

Table 1: Performance comparison between our proposed framework and other methods in dynamical systems

Cases	Groups	$\ell_2(\times 10^{-2})$	P (%)	R (%)
Single Pendulum	A	0.41	100	100
	B	0.41	100	100
	C	12.36	33.33	100
Double Pendulum	A	1.15	100	100
	B	1.15	100	100
	C	58.65	83.33	60
Spherical Pendulum	A	2.00	100	100
	B	2.04	100	100
	C	91.55	50	100
Chaos Pendulum	A	0.16	100	100
	B	0.16	100	100
	C	18.25	87.50	85.71
Cart Spring Pendulum	A	1.42	100	100
	B	1.40	100	100
	C	47.25	87.50	75
Spherical Spring Pendulum	A	0.39	100	100
	B	1.13	100	100
	C	8.51	85.70	85.70
Magnetical Pendulum	A	1.37	100	100
	B	3.01	100	100
	C	40.55	75	85.71

Table 2: Performance under varying experimental scenarios

ments. The experimental design consisted of two distinct groups: Group A utilized the complete time-series dataset,

while Group B employed a dataset with 5% randomly missing samples, where the missing points were uniformly distributed across the temporal dimension. The experimental results, as detailed in Table 2, reveal that the presence of missing data has a statistically negligible impact on the performance metrics of our method, with performance degradation within an acceptable range across evaluated metrics. acceptable range across evaluated metrics. This reflects the robustness of B-spline curves against noise and their flexibility in handling time-series data without requiring strict dependencies. Specifically, B-splines’ local support property allows them to adapt to irregularities in the data, such as missing points, without significantly affecting the overall fit. Additionally, their smoothness and differentiability ensure that the reconstructed dynamics remain physically plausible, even under incomplete or noisy conditions. These properties make B-splines particularly suitable for applications where data quality may vary or where strict temporal alignment is challenging to achieve.

Ablation Study. We conducted an ablation study to validate the effectiveness of the module that imposes constraints on higher-order derivatives of the system’s degrees of freedom. This module is designed to mitigate noise amplification caused by chain rule differentiation, which is particularly critical in second-order dynamical systems. To evaluate its impact, we introduced an ablation model that excludes the regularization term on the system’s second-order derivatives. Using the chaotic pendulum as a representative exam-

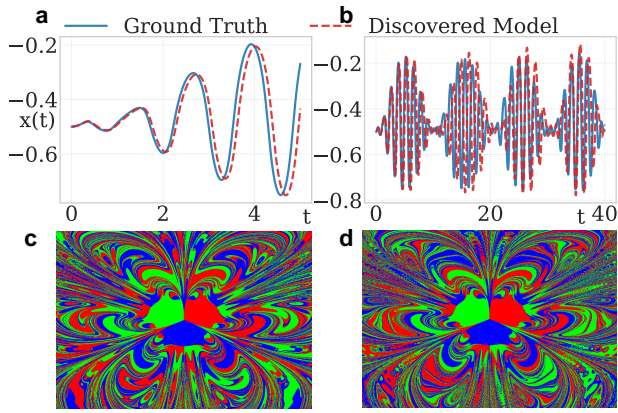


Figure 3: **a** Matches short-term; **b** Diverges long-term but preserves key statistical properties (amplitude range, oscillations). The basins of attraction for three magnets (colored red, blue, and green) are distributed along the edges of an equilateral triangle. **c** shows the true basin of attraction for magnetic pendulum, while **d** displays the basin obtained with the discovered parameters of system.

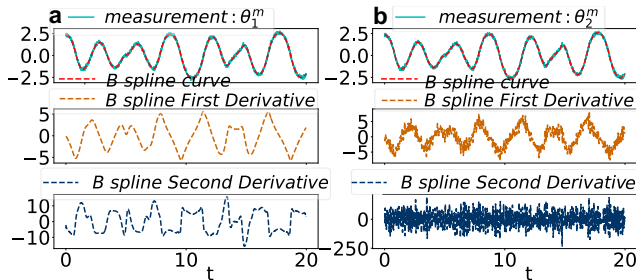


Figure 4: Impact of Second-Order Derivative Regularization on Performance. **a**: Performance of our proposed model in fitting the measurements and their derivatives. **b**: Performance of the model without the regularization term.

ple, we present the training results of the ablation model in Figure 4, where **a** represents our proposed method and **b** represents the model without the regularization term on second-order derivatives. Our observations reveal that, in the absence of the regularization term, the data loss of the system is minimized (the B-spline curve achieves near-perfect fitting of the measurement points). However, the first and second derivatives of the time-dependent functions for the degrees of freedom are poorly approximated, indicating a failure to capture the underlying dynamics accurately. Specifically, the second derivatives exhibit significant deviations from the ground truth, leading to a substantial deterioration in both training speed and model performance. This highlights the importance of constraining higher-order derivatives to ensure physical consistency and robustness. As shown in Table 2, Group A represents the full model, while Group C (without regularization) exhibits degraded performance, validating the necessity of regularization. These results demonstrate that our proposed framework, which incorporates reg-

ularization on higher-order derivatives, is essential for accurate equation discovery, especially under conditions of incomplete or noisy data. By enforcing reasonable constraints on the system’s dynamics, our approach effectively balances data fitting and physical plausibility, enabling reliable identification of governing equations.

Discussion and Limitation. Through our proposed differentiable sparse identification of Lagrangian dynamics, we leverage the differentiability of B-spline basis functions to fit curves and their derivatives by optimizing control points. Using only intuitive measurement data, our method discovers Lagrangian dynamics equations for various systems, effectively mitigating noise amplification caused by chain rule differentiation on candidate function libraries. This enables a more accurate identification of target equations.

Although our method performs well in identifying Lagrangian dynamics equations, it still has some limitations. When dealing with complex systems involving multiple degrees of freedom and high coupling, the size of the candidate function library increases significantly, leading to a sharp rise in computational complexity. This not only prolongs training time but may also make it difficult to converge to the correct solution due to combinatorial explosion. Moreover, although B-splines exhibit strong noise resistance, the computation of higher-order derivatives may still be affected under extreme noise conditions, impacting the model’s stability. Future work could explore more efficient library construction methods, incorporating domain knowledge or adaptive strategies, to further enhance the model’s applicability and robustness in complex systems.

Conclusion

This work introduces a novel differentiable framework for the sparse identification of Lagrangian dynamics from observational data. By integrating cubic B-spline approximation with sparse regression, our method effectively addresses key limitations in existing approaches, particularly in modeling complex nonlinearities and mitigating the impact of noise. The core of our framework is a novel integration of cubic B-splines into the system identification process, enabling flexible and accurate modeling of complex dynamics. This approach incorporates a robust equation discovery mechanism that efficiently utilizes limited measurements and leverages physical constraints. Furthermore, we introduce a recursive derivative computation scheme designed specifically to reduce noise sensitivity in second-order dynamical systems. Experimental results demonstrate that our method significantly outperforms baseline approaches in both accuracy and robustness, particularly in scenarios with high noise levels or sparse data. By leveraging the recursive properties of B-spline basis functions, our framework achieves a precise reconstruction of the governing equations while preserving their physical plausibility. This advancement enhances the reliability of data-driven discovery in complex mechanical systems and paves the way for hybrid models that fuse empirical insights with established physical principles. This synthesis yields actionable understanding of nonlinear dynamics across scientific and engineering fields.

Acknowledgments

The work is supported by the National Natural Science Foundation of China (No. 92270118 and No. 62276269) and the Beijing Natural Science Foundation (No. 1232009). The Appendix of this manuscript can be found in the preprint version: <https://www.arxiv.org/abs/2511.10706>.

References

- Battaglia, P.; Pascanu, R.; Lai, M.; Jimenez Rezende, D.; et al. 2016. Interaction networks for learning about objects, relations and physics. *Advances in neural information processing systems*, 29.
- Bongard, J.; and Lipson, H. 2007. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24): 9943–9948.
- Brunton, S. L.; Proctor, J. L.; and Kutz, J. N. 2016. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15): 3932–3937.
- Chu, H. K.; and Hayashibe, M. 2020. Discovering interpretable dynamics by sparsity promotion on energy and the Lagrangian. *IEEE Robotics and Automation Letters*, 5(2): 2154–2160.
- Cranmer, M.; Greydanus, S.; Hoyer, S.; Battaglia, P.; Spergel, D.; and Ho, S. 2020. Lagrangian neural networks. *arXiv preprint arXiv:2003.04630*.
- Goyal, P.; and Benner, P. 2022. Discovery of nonlinear dynamical systems using a Runge–Kutta inspired dictionary-based sparse regression approach. *Proceedings of the Royal Society A*, 478(2262): 20210883.
- Greydanus, S.; Dzamba, M.; and Yosinski, J. 2019. Hamiltonian neural networks. *Advances in neural information processing systems*, 32.
- Kaheman, K.; Kutz, J. N.; and Brunton, S. L. 2020. SINDy-PI: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proceedings of the Royal Society A*, 476(2242): 20200279.
- Kaiser, E.; Kutz, J. N.; and Brunton, S. L. 2021. Data-driven discovery of Koopman eigenfunctions for control. *Machine Learning: Science and Technology*, 2(3): 035023.
- Koopman, B. O.; and Neumann, J. v. 1932. Dynamical systems of continuous spectra. *Proceedings of the National Academy of Sciences*, 18(3): 255–263.
- Landajuela, M.; Lee, C. S.; Yang, J.; Glatt, R.; Santiago, C. P.; Aravena, I.; Mundhenk, T.; Mulcahy, G.; and Petersen, B. K. 2022. A Unified Framework for Deep Symbolic Regression. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 33985–33998. Curran Associates, Inc.
- Lenz, I.; Knepper, R. A.; and Saxena, A. 2015. DeepMPC: Learning deep latent features for model predictive control. In *Robotics: Science and Systems*, volume 10, 25. Rome, Italy.
- Long, Z.; Lu, Y.; and Dong, B. 2019. PDE-Net 2.0: Learning PDEs from data with a numeric-symbolic hybrid deep network. *Journal of Computational Physics*, 399: 108925.
- Lusch, B.; Kutz, J. N.; and Brunton, S. L. 2018. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1): 4950.
- Martius, G.; and Lampert, C. H. 2016. Extrapolation and learning equations. *arXiv preprint arXiv:1610.02995*.
- Mundhenk, T. N.; Landajuela, M.; Glatt, R.; Santiago, C. P.; Faissol, D. M.; and Petersen, B. K. 2021. Symbolic Regression via Neural-Guided Genetic Programming Population Seeding. *arXiv preprint arXiv:2111.00053*.
- Petersen, B. K.; Larma, M. L.; Mundhenk, T. N.; Santiago, C. P.; Kim, S. K.; and Kim, J. T. 2021. Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients. In *International Conference on Learning Representations*.
- Purnomo, A.; and Hayashibe, M. 2023. Sparse identification of Lagrangian for nonlinear dynamical systems via proximal gradient method. *Scientific Reports*, 13(1): 7919.
- Quade, M.; Abel, M.; Shafi, K.; Niven, R. K.; and Noack, B. R. 2016. Prediction of dynamical systems by symbolic regression. *Physical Review E*, 94(1): 012214.
- Rao, C.; Ren, P.; Liu, Y.; and Sun, H. 2022. Discovering Nonlinear PDEs from Scarce Data with Physics-encoded Learning. In *The Tenth International Conference on Learning Representations*.
- Rudy, S. H.; Brunton, S. L.; Proctor, J. L.; and Kutz, J. N. 2017. Data-driven discovery of partial differential equations. *Science Advances*, 3(4): e1602614.
- Sahoo, S.; Lampert, C.; and Martius, G. 2018. Learning equations for extrapolation and control. In *International Conference on Machine Learning*, 4442–4450. Pmlr.
- Schmidt, M.; and Lipson, H. 2009. Distilling free-form natural laws from experimental data. *Science*, 324(5923): 81–85.
- Sun, F.; Liu, Y.; and Sun, H. 2021. Physics-informed Spline Learning for Nonlinear Dynamics Discovery. *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, 2054–2061.
- Sun, F.; Liu, Y.; Wang, J.-X.; and Sun, H. 2023. Symbolic Physics Learner: Discovering governing equations via Monte Carlo tree search. In *The Eleventh International Conference on Learning Representations*.
- Sun, L.; Huang, D. Z.; Sun, H.; and Wang, J.-X. 2022. Bayesian Spline Learning for Equation Discovery of Nonlinear Dynamics with Quantified Uncertainty. In *Advances in Neural Information Processing Systems*.
- Xie, T.; France-Lanord, A.; Wang, Y.; Shao-Horn, Y.; and Grossman, J. C. 2019. Graph dynamical networks for unsupervised learning of atomic scale dynamics in materials. *Nature communications*, 10(1): 2667.
- Zhang, Z.; Liu, Y.; and Sun, H. 2024. Vision-based discovery of nonlinear dynamics for 3D moving target. *arXiv preprint arXiv:2404.17865*.