

Flow-consistent identification of governing equations from sparsely sampled measurements

Ye Yuan

yue@hust.edu.cn

Huazhong University of Science and Technology <https://orcid.org/0000-0001-7858-0437>

Zhexuan Zeng

Huazhong University of Science and Technology <https://orcid.org/0009-0004-9194-9912>

Dongming Xie

Kaiqi Fang

Huazhong University of Science and Technology

Jiahao Shi

Huazhong University of Science and Technology

Ruoying Yan

Huazhong University of Science and Technology

Xiaoxiao Lu

Huazhong University of Science and Technology

Zhou Deng

Huazhong University of Science and Technology

Yiming Meng

University of Waterloo

Jun Liu

University of Waterloo

Article

Keywords: Continuous-time dynamical system, System identification, Sparse sampling, Koopman operator

Posted Date: March 2nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-8856625/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

Flow-consistent identification of governing equations from sparsely sampled measurements

Zhexuan Zeng^{†1}, Dongming Xie^{†1,2}, Kaiqi Fang¹, Jiahao Shi¹,
Ruoying Yan¹, Xiaoxiao Lu¹, Zhou Deng¹, Yiming Meng³,
Jun Liu³, Ye Yuan^{1*}

¹*Department of Automatic Control, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan, Hubei, China.

²Shanghai Innovation Institute, 699 Huafa Road, Shanghai, China.

³Department of Applied Mathematics, University of Waterloo, 200 University Ave W, Waterloo, ON, Canada.

*Corresponding author(s). E-mail(s): yue@hust.edu.cn;

Abstract

Extracting continuous-time governing equations from sparsely sampled measurements remains a fundamental challenge across science and engineering. Under severe temporal sparsity, conventional data-driven approaches often lose physical consistency, leading to distorted dynamics and unreliable qualitative behavior. Here we propose Hybrid Analytic Neural Dynamics Identification (HANDI), a flow-consistent framework that enables accurate identification of governing equations from sparsely sampled time series. Rather than relying on noisy numerical differentiation or integration, HANDI operates directly on the flow map, bridging discrete measurements and continuous-time dynamics through a Koopman-based formulation. By constructing a hybrid observable space that integrates interpretable analytic structures with neural representations, HANDI enables a sampling-robust linearization of the underlying dynamics. Across canonical nonlinear systems (e.g., limit cycles and bistability) and real-world datasets (e.g., inverted-flag flapping and wheel shimmy dynamics), HANDI consistently attains high mechanistic fidelity under extreme temporal sparsity, reliably uncovering the underlying physics and dynamical geometry from limited measurements.

Keywords: Continuous-time dynamical system, System identification, Sparse sampling, Koopman operator

¹This work is supported by the National Natural Science Foundation of China under Grant No. T2525012.

^{2†}These authors contributed equally to this work.

1 Introduction

Data-driven discovery of governing equations has become a central paradigm for understanding and predicting complex dynamical phenomena across science and engineering. The sampling interval of collected data strongly influences the ability to recover the underlying equations [1]. As the sampling interval increases, transient behaviors become progressively obscured, making accurate identification inherently difficult. This challenge is exacerbated in many real-world scientific settings, where experiments unavoidably produce widely spaced temporal samples due to physical, logistical, or destructive constraints. Such sparsely sampled measurements are common across ecology [2, 3], epidemiology [4–6], neuroscience [7, 8], and biochemical reactions [9–11]. When the sampling is too sparse to resolve transient dynamics, the identified model may exhibit qualitatively incorrect behaviors, which poses a fundamental barrier to domain-specific conclusions and downstream decision-making across diverse scientific systems. For instance, in chemical systems driven by positive feedback, sparse sampling can smooth out rapid self-amplifying responses, falsely suggesting benign or saturating kinetics and thereby underestimating runaway-reaction risks. In sociocultural or ecological systems, sparse measurements may obscure recovery or stabilization phases, leading to the mistaken conclusion and misguided preservation policies. In infectious-disease dynamics, insufficient temporal resolution can blur the fast growth phase of transmission, resulting in biased estimates of contagion and recovery and, consequently, misguided intervention strategies. Ensuring that identified equations faithfully reflect the system’s mechanistic structure is therefore essential, particularly when sparsity is intrinsic to the experimental setting.

Early efforts to identify governing equations from data were pioneered by symbolic regression approaches, notably those of Bongard and Lipson [12] and Schmidt and Lipson [13]. These works demonstrated that nonlinear physical laws can, in principle, be recovered without prior specification of model structure. However, such approaches rely on searching an extremely large combinatorial space of candidate expressions, resulting in rapidly escalating computational complexity and poor scalability with system dimension. To improve tractability, subsequent frameworks focused on identifying governing equations by selecting sparse combinations of candidate terms from a predefined library [14, 15]. Among these, Sparse Identification of Nonlinear Dynamics (SINDy) emerged as a prominent and widely adopted methodology [15], motivating a range of variants aimed at improving robustness to noise and data imperfections. For example, SR3 [16] introduced a relaxed sparse regression formulation to enhance numerical stability, while Weak SINDy (WSINDy) [17] reformulated the problem in a weak integral form to mitigate sensitivity to noisy measurements by avoiding explicit numerical differentiation. Together with numerous related developments [18–21], these methods have enabled the successful identification of governing equations across diverse application domains, including fluid mechanics [22, 23], robotic locomotion and control [24], and biological oscillatory systems [25]. Despite algorithmic differences, most sparse regression-based identification methods and modern symbolic regression frameworks [26, 27] share a common paradigm: they identify continuous-time dynamics from local temporal information extracted from discretely sampled

trajectories. This information is typically obtained either through numerical differentiation or through weak-form integral constraints evaluated over sampling intervals. When sampling is sparse, however, the local temporal structure of trajectories is fundamentally under-resolved. Large sampling intervals obscure curvature and fast transients, rendering derivative-based estimates unstable, while weak-form integration effectively projects the dynamics onto coarse, time-averaged representations. As a result, identified models may interpolate the sampled data yet fail to recover the underlying continuous-time flow.

To address this challenge, we build on a basic physical fact: even when the system is sampled sparsely, its evolution between two measurements is always generated by the consistent underlying flow that maps a state $\mathbf{x}(t_k)$ to its future position $\mathbf{x}(t_k + T_s)$. Then we aim to recover the governing equations by learning this consistent flow and linking this flow to the underlying continuous dynamics over each sampling period, without requiring $T_s \rightarrow 0$. This eliminates both numerical differentiation and integration, thereby preventing the degradation induced by large sampling intervals and enabling reliable recovery of governing equations under sparse sampling. For linear systems, this principle admits an exact realization: the continuous-time system can be uniquely recovered from sparse measurements via an exponential relationship, and the identification error remains essentially invariant to the sampling interval (Supplementary Fig. 1). Sparse sampling is therefore not a numerical obstacle, but a modeling choice, provided that flow consistency is enforced. This sharply contrasts with SINDy and its variants, whose accuracy rapidly degrades under identical sparsity.

Extending this principle to nonlinear systems is substantially more challenging, because the evolution between sparse measurements can no longer be linked to the underlying continuous-time dynamics through a simple, closed-form relationship. Koopman operator theory offers a possible route by lifting the dynamics into a high-dimensional space where the evolution becomes approximately linear [28, 29]. This can be realized by popular data-driven methods such as Dynamic Mode Decomposition (DMD) [30, 31] and Extended DMD (EDMD) [32], which aim to learn the flow directly from data. However, commonly used representations rarely capture an invariant structure of the dynamics in practice, causing identification errors that grow rapidly with the sampling interval (Supplementary Section 3) and leading to a loss of identifiability well before theoretical limits are reached [33, 34].

To overcome this barrier, we develop a flow-consistent framework for identifying governing equations from sparsely sampled data. Rather than attempting to reconstruct unresolved transient dynamics, the proposed framework treats the finite-time flow map between measurements as the primary physical object governing identifiability under sparse sampling. Within this perspective, system identification is reframed as enforcing consistency of the flow across sampling intervals. Building on this principle, we construct an expressive and approximately invariant observable space in which the flow map of a nonlinear dynamical system can be accurately represented. This observable space combines analytic basis functions, which promote physical interpretability and equation identification, with neural-network complements that provide the flexibility needed to capture complex nonlinear structure. The resulting hybrid

representation enables the Koopman operator that encodes the flow to act nearly linearly on the learned observables, allowing the consistent flow map to be accurately recovered and, in turn, enabling reliable identification of the underlying continuous-time governing equations even as the sampling interval approaches the theoretical upper bound of identifiability.

2 Results

2.1 Overview of hybrid analytic neural dynamics identification

Ordinary differential equations (ODEs) are ubiquitous in modeling dynamical systems. We consider an autonomous system

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)), \quad \mathbf{x}(t) \in \mathbb{R}^d. \quad (1)$$

Given sparsely sampled trajectories $\{\mathbf{x}_k\}$ with $\mathbf{x}_k := \mathbf{x}(kT_s)$ and a large sampling period T_s , our goal is to recover the governing law $\mathbf{f}(\cdot)$.

Despite sparse measurements, the underlying discrete-time evolution remains rigorously defined by a consistent flow map S^{T_s} for all $\mathbf{x}$, i.e., $\mathbf{x}(t+T_s) = S^{T_s}(\mathbf{x}(t))$, which preserves compositional consistency $S^{T_1+T_2}(\mathbf{x}) = S^{T_1}(S^{T_2}(\mathbf{x}))$ across all intervals. Koopman theory elegantly maps this structure onto a functional space of observables via $K^{T_s}\varphi(\mathbf{x}) = \varphi(S^{T_s}(\mathbf{x}))$, ensuring that the resulting operators inherit the semi-group property. This formulation culminates in the fundamental exponential mapping $K^{T_s} = e^{LT_s}$, where the infinitesimal generator $L = \mathbf{f} \cdot \nabla$ serves as a bridge to the continuous-time dynamical system. By exploiting this intrinsic flow consistency, we reformulate the recovery of the hidden $\mathbf{f}$ from sparse data as a flow-map learning problem. This approach effectively decouples the discovery of governing laws from the temporal resolution of the measurements, bypassing the conventional necessity for the infinitesimal sampling limit ($T_s \rightarrow 0$).

Exploiting this flow consistency hinges on constructing a finite-dimensional observable space that is (approximately) invariant under K^{T_s} ; loss of invariance breaks the identification accuracy (Supplementary Section 3). To address this challenge, we propose a hybrid and sampling-robust framework, termed HANDI, which integrates analytic priors with adaptive neural residuals to identify governing equations from sparse measurements (Fig. 1a). The core of HANDI is an orthogonalized Koopman lifting mechanism that bridges physically interpretable structures with flexible data-driven representations. Specifically, system states are lifted into a hybrid observable space composed of analytic bases $\phi(\mathbf{x})$ and neural observables that account for unmodeled nonlinearities. To ensure identifiability and avoid redundancy between these components, we explicitly project the neural observables onto the orthogonal complement of the analytic subspace (Fig. 1b). This deflation step guarantees that the neural network learns only complementary dynamics not already encoded by the analytic dictionary, preserving interpretability of the identified model.




Fig. 1: Overview of the HANDI framework. **a**, HANDI identifies continuous-time dynamics from sparse measurements (x_k, x_{k+1}) by embedding flow maps into a hybrid observable space. This space integrates analytic bases $\phi(x)$ with orthogonal neural components $G_\perp(x)$ to capture complex dynamics while maintaining the interpretability required to extract governing equations. **b**, To prevent basis redundancy, analytic bases are transformed into an orthonormal set $Q(x)$, while neural bases are projected onto the orthogonal complement to yield the residual components $G_\perp(x)$. **c**, Following lifting to $Z = [Q; G_\perp]$, optimization proceeds in two stages. Stage A optimizes $G_\perp$ and K_g to capture dynamics missing from the analytic bases, providing a robust initialization. Stage B further fine-tunes the model via multi-step trajectory supervision to preserve flow-map consistency. **d**, A principled bridge between the flow map and the governing equation is established via operator inversion without requiring high sampling frequency. The learned operator K_{aug} is transformed into the Koopman generator L_z through a matrix logarithm, and then mapped back to the original basis to extract the governing equation.

Building on this orthogonalized representation, HANDI employs a two-phase training strategy (Fig. 1c): the neural residuals are first aligned to capture discrepancies unexplained by the analytic basis, followed by multi-step rollout optimization of the augmented Koopman operator to enforce long-term flow consistency under sparse sampling. Finally, the learned Koopman operator is mapped back to a generator via operator inversion (Fig. 1d), from which the governing equations $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ are extracted. This step bypasses the classical requirement for the infinitesimal sampling limit ($T_s \rightarrow 0$), enabling accurate identification of continuous-time dynamics from sparse data (see Methods for details).

2.2 Identification of canonical dynamical systems under temporal sparsity

We first evaluate HANDI under increasingly sparse temporal sampling using four canonical nonlinear systems representing fixed point, limit cycle, relaxation, and bistable dynamics. These systems are used to assess whether qualitative invariants of the underlying flows, such as stability, oscillatory structure, and multistability, can be preserved as the measurements become temporally sparse. For each system, datasets are generated with progressively increasing sampling intervals (Supplementary Table 1). All models are presented in non-dimensionalized or rescaled form, with sampling intervals reported in seconds for notational consistency.

Identification performance is assessed using complementary quantitative and qualitative metrics. Quantitative accuracy is evaluated using coefficient normalized root mean squared error (NRMSE) and trajectory mean squared error (MSE) over 100 uniformly sampled initial conditions. Qualitatively, the success rate of attractor consistency (ACR) measures the fraction of trials where the reconstructed model preserves fundamental topological invariants. These invariants are numerically verified through key dynamical signatures including Lyapunov stability for fixed points, orbital frequencies for limit cycles, and manifold geometry for complex attractors. Detailed metric definitions are provided in Supplementary Section 6.1. We compare HANDI with representative identification methods, including SINDy [15], SR3 [16], and PSE [27]; These 3 methods are selected from a broader pool of comparison approaches to ensure clarity of presentation, as they represent widely used identification paradigms and show relatively consistent performance across all considered systems in this study. Comprehensive quantitative evaluations involving all methods are deferred to Section 5.

Fig. 2 summarizes the identification performance across all systems and sampling intervals T_s . Across all cases, HANDI learns Koopman eigenfunctions that faithfully encode the geometric organization of the underlying dynamics, including vanishing structures along limit cycles and basin partitioning in multistable Duffing system (left column of Fig. 2). This structural consistency is reflected by an ACR of 100% over the entire tested sampling range (top middle panels of Fig. 2a–d). By contrast, baseline methods exhibit a rapid loss of qualitative fidelity as T_s increases, with ACR values degrading progressively and frequently collapsing under extreme sparse sampling.




Fig. 2: Identifications of canonical systems under sparse sampling. a–d, Comparison of HANDI against representative baselines on systems with distinct attractors. The learned Koopman eigenfunctions (left) reveal the geometric topology and invariant structures of each attractor. Evaluation across sampling intervals T_s (right) shows that HANDI maintains a 100% ACR and superior phase-portrait fidelity where baselines suffer pronounced topological degradation. Quantitative assessment via trajectory MSE and coefficient NRMSE (bottom right) further demonstrates HANDI's consistent accuracy across sparsity levels. Statistics are summarized over 100 uniformly sampled initial conditions using optimized hyperparameters (Supplementary Section 6.2).

Phase portraits further clarify this distinction at the trajectory level (top right panels of Fig. 2a–d). Even at large sampling intervals, HANDI reproduces trajectory evolutions that are consistent with the reference dynamics, preserving key behavioral features such as oscillatory convergence ($T_s = 1.0$ s), sustained periodic motion ($T_s = 1.0$ s), and sharp relaxation transitions ($T_s = 0.5$ s). In contrast, baseline models yield qualitative failures under sparse sampling. For instance, trajectories lose their characteristic spiral structure and collapse into overdamped decay in the fixed point system. For limit cycle dynamics, reconstructed oscillations show substantial frequency attenuation. Such failures indicate violations of flow-level dynamical consistency, undermining physical interpretability and mechanistic insight.

Quantitatively, HANDI consistently achieves lower trajectory prediction and coefficient identification errors across all systems under sparse sampling (bottom of Fig. 2). While baseline methods perform competitively when sampling intervals are sufficiently small, their errors increase sharply as temporal sparsity grows. In contrast, HANDI remains stable across the tested regimes, demonstrating its robustness against severe degradation of local temporal information. Detailed system-specific results are provided in Supplementary Sections 6.3.2–6.3.5.

The maximum sampling interval considered for each system is determined by intrinsic dynamical time scales rather than algorithm-specific tuning. Systems with slower dynamics tolerate substantially larger sampling intervals, whereas systems with fast transitions or sharp basin boundaries require finer temporal resolution. Notably, all methods fail once the sampling interval exceeds a fundamental limit imposed by the continuous-time dynamics, beyond which measurements no longer preserve the geometric structure of the flow [33, 34]. Additional visualizations are provided in Supplementary Section 7.

2.3 Identification of benchmark systems under joint temporal sparsity and limited realizations

To evaluate the limits of system identification under joint temporal sparsity and limited realizations, we benchmarked HANDI on four representative dynamical systems from MD-Bench [35]: chemical autocatalysis, language death, susceptible–infected–recovered (SIR) epidemics, and glycolytic oscillations. These systems exemplify a pervasive challenge in scientific discovery: intrinsic dynamics often evolve rapidly, yet only a small number of experimental realizations are available and measurements are sampled sparsely in time. For instance, millisecond-scale bursts in chemical autocatalysis are typically obscured by the coarse temporal resolution of standard assays [9, 10, 36] (Supplementary Sections 6.3.6–6.3.9). For each system, we restrict the training data to only two trajectories, each sampled sparsely in time (for example, $T_s = 2.25$ s for the autocatalytic system). Under this regime, the available samples are insufficient to resolve fast transients or faithfully trace the geometric structure of the underlying dynamics. This setting therefore constitutes a stringent test of whether an identification method can identify the underlying vector field, rather than merely interpolating between discrete measurements.




Fig. 3: Performance of HANDI on benchmark systems under joint temporal sparsity and limited realizations. HANDI is compared with SINDy [15], SR3 [16], PSE [27], gEDMD [37], and EDMD [32, 38] on four benchmark systems, including chemical autocatalysis (a), language death dynamics (b), the SIR epidemic model (c), and a glycolytic oscillator (d), representing application domains in which temporal sampling and realizations are typically limited. The second column of a–d shows the ground-truth vector fields together with the training measurements (black). Despite sparse sampling and limited trajectory coverage, HANDI accurately recovers both the system trajectories and the underlying vector fields over the state space, achieving consistently low errors E_{L^2} and E_{rel} (third column of a–d). Averaged across all benchmark systems, HANDI yields lower errors in system-specific physical quantities than baseline methods (fourth and fifth columns of a–d and e). All methods are evaluated using optimized hyperparameters and identical datasets, with basis-based approaches sharing the same candidate library. Comparative vector field errors and detailed identification results are provided in Supplementary Sections 6.3.6–6.3.9.

Fig. 3 summarizes the performance of HANDI under sparse temporal sampling and limited trajectory coverage. Across all benchmarks, HANDI accurately identifies the continuous-time vector field within a bounded domain Ω (Fig. 3a–d). This evaluation region Ω is defined by the dynamical range of training trajectories and a modest buffer to assess local extrapolation, while strictly adhering to inherent constraints such as non-negativity. Performance is quantified via the pointwise L^2 error (E_{L^2}) and its normalized relative counterpart (E_{rel}), both of which remain consistently low across considered benchmarks. Fig. 3a–d also demonstrate that the identified models preserve both the correct directionality of motion and coherent geometric structures around equilibria and manifolds, ensuring that the continuous-time dynamics remain physically consistent even far from the training realizations. In contrast, while baseline methods may appear to fit the observed trajectories at sampled time points, their identified vector fields exhibit obvious distortions in regions unconstrained by data (Supplementary Figs. 21, 25, 29, 33). These geometric inconsistencies, manifesting as biased magnitudes or spurious critical structures, directly propagate into erroneous characterizations of stability, basin geometry, and transient behavior. Detailed quantitative comparisons are provided in Supplementary Tables 20, 24, 27, 31.

Furthermore, HANDI achieves consistently low relative errors for key physical quantities across all domains, significantly outperforming or remaining competitive with representative baselines (Fig. 3e). In chemical autocatalysis, HANDI accurately identifies the effective rate k_e and inhibition rate k_{-1} with errors of only 0.19% and 4.37%, respectively. These parameters are essential as they dictate the critical burst-to-saturation profile of the reaction. Similarly, HANDI resolves the enzymatic coupling and substrate influx parameters of the glycolytic oscillator with high precision, thereby accurately capturing the driving regulation of metabolic oscillations. The superior performance of HANDI stems from a fundamental shift in the identification paradigm. This efficacy arises because HANDI reformulates the traditional task of pointwise vector field estimation, which demands dense sampling, into a global identification of linear evolution laws within a latent Koopman subspace. By capturing the underlying flow map that implicitly integrates the system dynamics over finite time intervals, HANDI bypasses the need for numerical differentiation or integration, and remains robust to sparse temporal sampling. Moreover, the operator-inversion process aggregates information across the entire state space, allowing consistent identification even when trajectory coverage is limited. Consequently, the identified evolution laws can be coherently applied across the state space, enabling robust generalization beyond measurements that are sparse in both time and state space.

2.4 Ablation analysis and robustness for system identification

The performance of HANDI framework depends on the architectural integrity of its hybrid modules and the noise level of the given data. Herein, we present an analysis of these factors to validate the necessity of the hybrid design and the model’s robustness against interference.

Model ablation study. We conducted an ablation study to evaluate the role of the neural residual learning and the orthogonal projection strategy. We constructed two types of simplified architectures as ablation baselines: (1) “Residual ablated”,

where the Stage A (neural residual alignment) is removed and retains only Stage B for parameter learning, and (2) “Orthogonal ablated”, where the neural observable is learned without enforcing orthogonality to the analytic subspace.




Fig. 4: Model ablation study and robustness performance of HANDI. a–d, Ablation study on four canonical dynamical systems. The full HANDI model (★) is compared with two ablated variants: ① residual-ablated (residual fitting removed) and ② orthogonality-ablated (orthogonal projection removed). Bar plots report ACR, while box plots show trajectory prediction error (MSE) over 100 initial conditions. Statistical significance between methods is assessed using pairwise Mann–Whitney U tests; asterisks denote significance levels (* $p < 10^{-1}$, ** $p < 10^{-2}$, *** $p < 10^{-3}$, **** $p < 10^{-4}$), and “n.s.” indicates non-significant differences. e, Robustness evaluation under multiplicative measurement noise. Trajectory MSE of HANDI and six baseline methods on four benchmark systems (Chemical autocatalysis, Language death, SIR epidemic, and Glycolytic oscillator) across noise levels ranging from 1% to 5%. HANDI consistently maintains low prediction error across all noise intensities, while several baseline methods exhibit unstable or divergent behavior.

Ablation studies across four canonical systems (Fig. 4a–d) validate the architectural integrity of HANDI, demonstrating that both residual fitting and basis orthogonalization are indispensable for preserving the framework’s qualitative topological structures and quantitative precision. First, the removal of the residual learning

module (blue) results in a statistically significant increase in identification error, as quantified by the MSE, most notably in the fixed point, limit cycle and Van der Pol systems (Fig. 4a–c). While the MSE for the Duffing oscillator shows only marginal degradation, the corresponding phase portraits reveal a clear loss of qualitative accuracy as quantified by the ACR, failing to capture the precise geometry of the attractor. These results confirm that residual alignment is vital for compensating for the discrete-time evolution of analytical basis and capturing the “missed” physics within the sparse-sampling regime. Second, we evaluated the necessity of the orthogonalization constraint. The variant lacking this projection (green) suffered from severe qualitative collapse; specifically, in the fixed-point and Van der Pol systems, the ACR dropped to zero, indicating a complete failure to recover the system’s governing equation. Quantitatively, removing the orthogonality constraint increased prediction errors by at least an order of magnitude across all four systems. We attribute this degradation to feature redundancy between the analytic and neural bases. Without the orthogonal projection, the optimization landscape becomes ill-conditioned, leading to numerical instability and a failure to isolate the unique contribution of each basis component. In stark contrast, the full HANDI framework consistently achieves superior performance, maintaining mechanistic fidelity and numerical accuracy where ablated variants fail.

Robustness performance. We further evaluate the robustness of HANDI to measurement noise on four representative benchmark systems: Chemical autocatalysis, Language death, SIR epidemic, and the Glycolytic oscillator. Multiplicative noise with intensities ranging from 1% to 5% is applied to the observed trajectories. HANDI is compared against six baseline methods, including SINDy, SR3, PSE, WSINDy, gEDMD and EDMD.

Fig. 4e reports the trajectory MSE of HANDI and baseline models under increasing measurement noise. Across the considered systems and noise levels, HANDI achieves consistently low prediction error and stable long-term simulations, while several baseline methods exhibit rapid performance degradation or simulation failure as noise increases (Supplementary Tables 19, 23, 26, 30). In a small number of 5% noise cases, HANDI performs comparably to the best-performing baseline methods, without exhibiting catastrophic failure. Overall, these results demonstrate that HANDI maintains robust identification performance under moderate measurement noise.

2.5 Identification from experimental data under different temporal sparsity

Real-world experimental data pose a fundamental challenge for system identification: the underlying governing structures are unknown, measurements are noisy, and the data acquisition process varies with the temporal sampling conditions. In this setting, model accuracy can no longer be assessed against a known ground truth. Instead, a meaningful identification should recover dynamical behaviors and physical quantities that are intrinsic to the system and therefore remain consistent as the sampling interval varies. To this end, we apply HANDI to four experimental systems under progressively increasing sampling intervals, including a single pendulum [25], an inverted flag flapping system [39], wheel-shimmy dynamics [39], and a neuronal integration experiment [40].




Fig. 5: Identification of governing equations from experimental data across temporal sparsity. a–d, Performance of HANDI on four experimental systems: pendulum, inverted flag, wheel shimmy dynamics, and neuronal integration, across increasing sampling intervals T_s . Across all regimes, HANDI recovers dynamical behaviors that remain qualitatively and quantitatively consistent even for unseen initial conditions (d), whereas baseline methods (SINDy, SR3, PSE) exhibit characteristic failures such as frequency distortion and simulation divergence. e, Errors of identified quantities, including rope length, natural frequency, steady-state potentials, and oscillation amplitude, demonstrate that HANDI yields accurate, physically plausible estimates (error < 2%) where baselines often produce significant errors or fail to return valid estimates (N/A). See Supplementary Sections 6.3.10–6.3.13 for details.

599 Across all experimental systems, HANDI identifies continuous-time dynamics
600 whose observed dynamical behaviors remain stable as the sampling interval T_s
601 increases (Fig. 5a–d). In contrast, baseline methods exhibit pronounced dependence
602 on temporal sampling sparsity. For the pendulum, inverted flag, and wheel–shimmy
603 systems, which are dominated by rapid energy exchange and oscillatory motion, base-
604 line models suffer from phase drift, frequency attenuation, or outright divergence as
605 T_s increases (Fig. 5a–c). In the neuronal integration experiment, baseline methods fre-
606 quently fail to produce stable long-term simulations under sparse sampling (Fig. 5d;
607 Supplementary Table 36). These failure modes stem from the inherent dependence
608 of local, derivative-based identification methods on sufficiently dense temporal sam-
609 pling, which becomes unreliable when fast transitions are undersampled. By contrast,
610 HANDI identifies a continuous-time Koopman generator that constrains the global
611 flow map, yielding models that reproduce the consistent qualitative dynamical behav-
612 iors, such as oscillatory patterns and steady-state responses, even when identified from
613 sparsely sampled experimental data (fourth column of Fig. 5a–d).

614 This consistency is further supported by the physical quantities inferred from
615 HANDI models across different sampling intervals (Fig. 5e). Despite increases in
616 temporal sparsity, HANDI yields consistent and physically plausible estimates of
617 system-specific parameters. In contrast, baseline methods frequently collapse into non-
618 physical regimes, such as unrealistically large rod length l or oscillation amplitude
619 A , or fail to estimate valid quantities altogether, including the natural frequency ω_n
620 and the mean steady-state membrane potential $\bar{V}$ (reported as N/A). Together, these
621 results indicate that HANDI recovers intrinsic system properties that remain con-
622 sistent across changes in the sampling interval, rather than artifacts induced by a
623 particular level of temporal sparsity.

624 Furthermore, HANDI reliably captures latent dynamical characteristics that are
625 difficult to infer from sparsely sampled trajectories alone. In the wheel–shimmy
626 system, a safety-critical nonlinear vibration, HANDI resolves the slow envelope
627 decay modulating the dominant oscillations, implicitly capturing the effective damp-
628 ing behavior that governs convergence toward the equilibrium manifold (Fig. 5c).
629 Similarly, in the neuronal integration experiment, HANDI consistently recovers char-
630 acteristic relaxation time scales across sampling intervals, despite measurement noise
631 and sparse sampling (Fig. 5d).

632 To assess whether HANDI captures the underlying dynamical structure beyond
633 the training trajectories, we perform a state-space extrapolation test on the inverted
634 flag system, which provides the richest set of experimental realizations among these
635 four systems. The model is identified using only four trajectories acquired at the
636 sparsest sampling interval ($T_s = 0.20$ s). Without any additional fitting, the result-
637 ing continuous-time model is then simulated from two previously unobserved initial
638 conditions. As shown in the lower panel of Fig. 5d, the predicted trajectories closely
639 follow the experimental measurements. Quantitatively, despite being identified from
640 limited and sparsely sampled data, the model preserves the characteristic oscillation
641 amplitudes and dominant frequencies across both unseen initial conditions, with low
642 relative error. These results indicate that HANDI recovers a coherent state-space flow
643

644

that governs inverted flag flapping dynamics, rather than merely interpolating the observed trajectories.

The aggregated performance across all canonical and experimental systems is summarized in Fig. 6. HANDI consistently achieves high Area Under the Success Curve (AUSC) values, indicating reliable identification performance across a broad range of accuracy thresholds. In several experimental settings, HANDI maintains robust performance under sparse and noisy measurements, whereas baseline methods yield very low AUSC values, reflecting their limited ability to preserve long-term dynamics in such regimes. This clear separation in AUSC highlights the consistent reliability of HANDI across diverse dynamical systems and data conditions.




Fig. 6: Performance of various models across canonical and real-world dynamical systems under sparsity and noise. We compare the performance of HANDI against representative identification approaches, including SINDy, SR3, PSE, WSINDy, gEDMD, and EDMD. Large sampling intervals (e.g., $T_s = 1.0$ s for the fixed-point system) are deliberately selected to simulate scenarios where high-frequency monitoring is technically or economically prohibitive. All methods were evaluated using optimized hyperparameters, with basis-based models utilizing an identical library of candidate functions. Given these consistent constraints, HANDI demonstrates superior global robustness, achieving a significantly higher AUSC across most benchmarks. This demonstrates the framework’s broad applicability and consistent reliability in identifying true governing equations from sparse, corrupted data. Here AUSC quantifies the cumulative success rate across a spectrum of error tolerances $[0, 1]$, where higher values indicate that a larger fraction of predicted trajectories maintain high fidelity across a wide range of accuracy requirements. Statistical results are averaged over 100 uniformly sampled initial conditions for canonical systems and available experimental trajectories for real-world systems.

3 Discussion

This work establishes a flow-consistent perspective for identifying nonlinear dynamical systems from sparsely sampled data. Rather than attempting to reconstruct hidden local dynamics directly, HANDI reframes the identification problem around the intrinsic consistency of the finite-time flow map. By lifting sparse measurements into a Koopman observable space that integrates interpretable analytic structures with an orthogonal neural residual, HANDI captures the action of the underlying flow in a linear form that remains invariant across states. Within this representation, continuous- and discrete-time dynamics are explicitly linked through a matrix exponential, enabling faithful identification even from sparsely sampled measurements.

A key consequence of this formulation is that it preserves not only trajectory-level accuracy but also the global organization of phase space. Through spectral analysis of the learned Koopman generator, HANDI recovers interpretable geometric structures, such as equilibria, bistable potential landscapes, and oscillatory modes, that are closely aligned with the underlying mechanisms. These qualitative invariants are precisely the features that sparse sampling tends to obscure for pointwise or locally constrained identification methods, which may distort the implied vector field away from ground truth. Across canonical benchmarks and real-world systems, the proposed framework consistently identifies governing equations and reconstructs trajectories that remain faithful to the reference dynamics, even as the sampling period approaches theoretical limits of the underlying dynamical systems. Taken together, these results suggest that enforcing flow consistency provides a principled pathway for resolving the ambiguity inherent in sparse sampling identification, shifting the focus from local, pointwise estimation to globally constrained dynamical inference.

Despite these advances, several challenges remain. Chaotic systems, in particular, pose intrinsic difficulties for establishing a correspondence between continuous-time and discrete-time dynamics through a matrix exponential, owing to intermittently forced linear behavior and continuous spectral components that complicate Koopman-based representations in practice [41]. This challenge reflects a fundamental limit of identifiability under sparsity rather than implementation-specific deficiencies, and extending flow-consistent identification into chaotic systems remains an open problem.

High-dimensional systems present a complementary challenge. In settings such as climate dynamics, the dimensionality of the observable space grows rapidly, rendering it difficult to learn compact and interpretable representations of consistent flow maps from limited data. Incorporating stronger physical priors and structural constraints will likely be essential for scaling sparse identification methods to such regimes.

Finally, while this study focused on temporally sparse measurements of ordinary differential equations, many natural processes are inherently spatiotemporal. Extending flow-consistent identification from temporal to spatiotemporal settings represents a promising direction for reconstructing distributed dynamics from limited measurements, and may further clarify the fundamental limits and opportunities of data-driven discovery under extreme sparsity.

4 Methods

Traditional EDMD [32] uses a fixed set of basis functions as Koopman observables. However, when the sampling period T_s changes, a single fixed dictionary may no longer approximate the Koopman invariant subspace well. To address this, we augment the fixed analytic dictionary with a learnable neural dictionary that captures the portion the fixed basis cannot fit.

The method proceeds in two phases. *Stage A coarsely tunes* the neural dictionary to learn components that are (approximately) orthogonal to the fixed observables and to capture the residual left by the fixed dictionary. *Stage B fine-tunes* the same network: we concatenate the fixed and learned observables to form the full feature map, estimate the discrete Koopman operator and its continuous-time generator. Using the block associated with the fixed analytic dictionary, we reconstruct an analytic vector field in the original state coordinates.

4.1 Data preprocessing

We consider N discrete-time trajectories $\{\mathbf{x}_{n,t}\}_{n=1, t=0}^{N, T_n-1}$, with sampling interval $T_s > 0$, $\mathbf{x}_{n,t} \in \mathbb{R}^d$. All snapshots are collected and stacked along the sample dimension:

$$\{(\mathbf{x}_k, \mathbf{y}_k)\}_{k=1}^M = \bigcup_{n=1}^N \{(\mathbf{x}_{n,t}, \mathbf{x}_{n,t+1}) : t = 0, \dots, T_n - 2\}, \quad (2)$$

where $M = \sum_{n=1}^N (T_n - 1)$ denotes the total number of snapshot pairs.

For each state, we build analytic dictionary features $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^m$ and $\mathbf{z}_k = \phi(\mathbf{x}_k)$, where ϕ is formed by selecting or concatenating basis functions from an *analytic* family (we refer to it uniformly as ϕ). These basis functions may include monomial, trigonometric, and exponential functions. We assume that the dictionary is state-inclusive: the first d entries of $\phi(\mathbf{x})$ are exactly the state coordinates, i.e., $\phi(\mathbf{x})_{1:d} = \mathbf{x}$. After lifting, we stack features row-wise as

$$\Phi_x = \begin{bmatrix} \phi^\top(\mathbf{x}_1) \\ \vdots \\ \phi^\top(\mathbf{x}_M) \end{bmatrix} \in \mathbb{R}^{M \times m}, \quad \Phi_y = \begin{bmatrix} \phi^\top(\mathbf{y}_1) \\ \vdots \\ \phi^\top(\mathbf{y}_M) \end{bmatrix} \in \mathbb{R}^{M \times m}. \quad (3)$$

We perform a (no pivot) QR factorization on Φ_x : $\Phi_x = \mathbf{Q}_x \mathbf{R}_\phi$, and define the *orthogonalized basis dictionary*

$$\mathbf{Q}(\cdot) = \phi(\cdot) \mathbf{R}_\phi^{-1}, \quad \mathbf{Q}_x = \Phi_x \mathbf{R}_\phi^{-1}, \quad \mathbf{Q}_y = \Phi_y \mathbf{R}_\phi^{-1}. \quad (4)$$

Note that $\mathbf{Q}(\cdot)$ defines a data-dependent orthogonalized dictionary with respect to the empirical inner product induced by the samples.

783 4.2 Stage A: orthogonal residual dictionary learning

784 In the orthogonalized ϕ -space, we assume a discrete linear evolution $\mathbf{z}_{t+1}^\top \approx \mathbf{z}_t^\top \mathbf{K}$ and
 785 estimate the prior one-step operator by the closed form of ridge regression
 786

$$787 \mathbf{K}_Q = \arg \min_{\mathbf{K}} \|\mathbf{Q}_y - \mathbf{Q}_x \mathbf{K}\|_F^2 + \lambda_\phi \|\mathbf{K}\|_F^2 = (\mathbf{Q}_x^\top \mathbf{Q}_x + \lambda_\phi \mathbf{I})^{-1} \mathbf{Q}_x^\top \mathbf{Q}_y, \quad (5)$$

789 where λ_ϕ is a tunable regularization parameter. This induces the residual not captured
 790 by the prior dictionary,
 791

$$792 \mathbf{R}_Q = \mathbf{Q}_y - \mathbf{Q}_x \mathbf{K}_Q \in \mathbb{R}^{M \times m}. \quad (6)$$

793 To explain this residual, we introduce a learnable residual dictionary $\mathbf{g}_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^g$
 794 and stack the corresponding features
 795

$$796 \mathbf{G} = \begin{bmatrix} \mathbf{g}_\theta^\top(\mathbf{x}_1) \\ \vdots \\ \mathbf{g}_\theta^\top(\mathbf{x}_M) \end{bmatrix} \in \mathbb{R}^{M \times g}, \quad (7)$$

797 where θ are network parameters of $\mathbf{g}_\theta$.

798 To avoid the neural dictionary duplicating patterns already represented by the
 799 fixed analytic dictionary, we perform a deflation step that removes from the raw neural
 800 features their least-squares projection onto the samplewise span of the orthogonalized
 801 analytic features, $\text{span}(\mathbf{Q}_x)$. The resulting deflated features $\mathbf{G}_\perp$ lie, up to numerical
 802 tolerances, in the orthogonal complement of $\text{span}(\mathbf{Q}_x)$
 803

$$804 \mathbf{O} = (\mathbf{Q}_x^\top \mathbf{Q}_x + \varepsilon \mathbf{I})^{-1} \mathbf{Q}_x^\top \mathbf{G}, \quad \mathbf{G}_\perp = \mathbf{G} - \alpha \mathbf{Q}_x \mathbf{O}. \quad (8)$$

805 Here, $\alpha \in [0, 1]$ controls the deflation strength, while the small Tikhonov term $\varepsilon > 0$
 806 stabilizes the projection, where α and ε are both tunable. This procedure suppresses
 807 redundancy between the fixed and neural dictionaries, improves the identifiability of
 808 subsequent operator estimates, and yields better numerical conditioning downstream.
 809 Using $\mathbf{G}_\perp$, we define the closed-form linear mapping
 810

$$811 \mathbf{K}_g^{\text{closed}} = (\mathbf{G}_\perp^\top \mathbf{G}_\perp + \lambda_g \mathbf{I})^{-1} \mathbf{G}_\perp^\top \mathbf{R}_Q, \quad \hat{\mathbf{R}}_Q = \mathbf{G}_\perp \mathbf{K}_g, \quad (9)$$

812 where λ_g is tunable, and then train the parameters by minimizing the Stage A objec-
 813 tive, which reduces the discrepancy between the predicted residual $\hat{\mathbf{R}}_Q$ and the target
 814 residual $\mathbf{R}_Q$ while regularizing the learned operator $\mathbf{K}_g$ toward its closed-form ridge
 815 $\mathbf{K}_g^{\text{closed}}$
 816

$$817 \mathcal{L}_1 = \|\hat{\mathbf{R}}_Q - \mathbf{R}_Q\|_F^2 + \beta \|\mathbf{K}_g - \mathbf{K}_g^{\text{closed}}\|_F^2, \quad (10)$$

818 where β is tunable. Optimize θ and $\mathbf{K}_g$ to obtain the trained residual dictionary $\mathbf{g}_\theta$
 819 together with its orthogonalized features $\mathbf{G}_\perp$.
 820
 821
 822
 823
 824
 825
 826
 827
 828

4.3 Stage B: Koopman-based vector field identification

We construct the augmented observable by concatenating the prior analytic features and the orthogonalized neural residual features, $\mathbf{Z}(\mathbf{x}) = [\mathbf{Q}^\top(\mathbf{x}), \mathbf{G}_\perp^\top(\mathbf{x})]^\top \in \mathbb{R}^{m+g}$, and stack all samples row-wise to obtain $\mathbf{Z}_x = [\mathbf{Z}(\mathbf{x}_1), \dots, \mathbf{Z}(\mathbf{x}_M)]^\top$ and $\mathbf{Z}_y = [\mathbf{Z}(\mathbf{y}_1), \dots, \mathbf{Z}(\mathbf{y}_M)]^\top$. The augmented discrete Koopman operator is then estimated in closed form via ridge regression,

$$\mathbf{K}_{\text{aug}} = (\mathbf{Z}_x^\top \mathbf{Z}_x + \lambda_{\text{full}} \mathbf{I})^{-1} \mathbf{Z}_x^\top \mathbf{Z}_y \in \mathbb{R}^{(m+g) \times (m+g)}, \quad (11)$$

where λ_{full} is tunable. To promote temporal coherence beyond a single step, we first define, for any trajectory n and starting time t with $t+S \leq T_n-1$, a contiguous window $\{\mathbf{x}_{n,t}, \dots, \mathbf{x}_{n,t+S}\}$ and set $\mathbf{Z}_0 := \mathbf{Z}(\mathbf{x}_{n,t})$ and $\mathbf{Z}_k := \mathbf{Z}(\mathbf{x}_{n,t+k})$ for $k = 1, \dots, S$. We then enforce one-step and multi-step consistency in the feature space by minimizing

$$\mathcal{L}_2 = \|\mathbf{Z}_1 - \mathbf{Z}_0 \mathbf{K}_{\text{aug}}\|_F^2 + \gamma \cdot \frac{1}{S} \sum_{k=1}^S \|\mathbf{Z}_k - \mathbf{Z}_0 (\mathbf{K}_{\text{aug}})^k\|_F^2, \quad (12)$$

where $\gamma > 0$ controls the multi-step weight and the loss is evaluated in all valid windows, S is the maximum rollout length and γ, S are tunable; each training epoch first refreshes $\mathbf{K}_{\text{aug}}$ on the entire dataset and then updates $\mathbf{g}_\theta$ using $\mathcal{L}_2$. For interpretability and continuous-time analysis, we map the discrete operator back to the native analytic basis by a similarity transform applied only to the ϕ -block, with $\mathbf{T} = \text{diag}(\mathbf{R}_\phi, \mathbf{I}_g)$ and $\mathbf{K}_{\text{aug}}^{\text{nat}} = \mathbf{T}^{-1} \mathbf{K}_{\text{aug}} \mathbf{T}$. The continuous-time generator is obtained via the principal matrix logarithm,

$$\mathbf{L} = \frac{1}{T_s} \log(\mathbf{K}_{\text{aug}}^{\text{nat}}), \quad (13)$$

from which we extract the block associated with the original state coordinates to write the recovered vector field as

$$\dot{\mathbf{x}}(t) = \phi^\top(\mathbf{x}(t)) \mathbf{C}_{\text{hat}}, \quad \mathbf{C}_{\text{hat}} := \mathbf{L}_{1:m, 1:d}. \quad (14)$$

4.4 Uniqueness and numerical stability of operator inversion

The recovery of a continuous-time generator $\mathbf{L}$ from a discrete-time Koopman operator $\mathbf{K}_{\text{aug}}^{\text{nat}}$ relies on the matrix logarithm (13). This inversion is generally not unique due to the multivalued nature of the complex logarithm. A sufficient condition for uniqueness is imposed on the spectrum of the true generator $\mathbf{L}$: if all eigenvalues λ satisfy

$$|\text{Im}(\lambda)T_s| < \pi, \quad (15)$$

then the principal matrix logarithm recovers the unique generator whose spectrum lies within the principal branch. Under this condition, eigenvalues of $\exp(\mathbf{L}T_s)$ remain strictly away from the negative real axis, where the principal logarithm is discontinuous. A detailed theoretical discussion of this identifiability condition is provided in Supplementary Section 7.

875 From a numerical perspective, the sensitivity of operator inversion originates from
876 the computation of the matrix logarithm itself. When eigenvalues of the Koopman
877 operator approach the negative real axis, corresponding to $\text{Im}(\lambda)T_s \approx \pm\pi$, small
878 perturbations in $\mathbf{K}_{\text{aug}}^{\text{nat}}$ may be significantly amplified in $\log(\mathbf{K}_{\text{aug}}^{\text{nat}})$. This behavior
879 reflects an intrinsic property of the logarithm mapping rather than a limitation of
880 the proposed method, and it becomes increasingly pronounced as temporal sampling
881 becomes sparse.

882 In practice, we compute the matrix logarithm by Schur decomposition-based algo-
883 rithm to evaluate the principal matrix logarithm. Specifically, the method employs
884 the inverse scaling and squaring algorithm, providing a numerically stable evaluation
885 in the presence of non-normality and clustering of eigenvalues [42, 43]. While this
886 approach ensures robust numerical computation of the principal logarithm, it does
887 not alter the inherent ill-conditioning of the logarithm near its branch cut.

888 All experiments reported in this study operate in regimes where the generator
889 inversion remains numerically well conditioned. Empirically, when the above spectral
890 condition is violated, continuous-time identifiability deteriorates across all methods
891 considered, indicating system aliasing imposed by sparse temporal sampling rather
892 than a failure of the proposed approach (Supplementary Section 7).
893

894 4.5 Coefficient-level metric

895 We report a coefficient-level NRMSE between the estimated coefficient matrix $\mathbf{C}_{\text{hat}}$
896 and the ground truth $\mathbf{C}_{\text{true}}$:
897

$$\begin{aligned} \text{NRMSE} &= \frac{\sqrt{\frac{1}{md} \sum_{i=1}^m \sum_{j=1}^d (\mathbf{C}_{\text{hat}}(i, j) - \mathbf{C}_{\text{true}}(i, j))^2}}{\frac{1}{|\Omega|} \sum_{(i,j) \in \Omega} |\mathbf{C}_{\text{true}}(i, j)|}, \quad (16) \\ \Omega &= \{(i, j) : |\mathbf{C}_{\text{true}}(i, j)| > 0\}, \end{aligned}$$

898 where the denominator normalizes by the average magnitude of the nonzero ground-
899 truth entries. Lower values of NRMSE indicate better agreement.

900 The trajectory-level MSE for k th initial condition is computed by the squared
901 prediction error over multiple time steps:
902

$$\text{MSE}^{(k)} = \frac{1}{d N_{\text{sim}}} \sum_{j=0}^{N_{\text{sim}}-1} \|\mathbf{x}_k(t_j) - \hat{\mathbf{x}}_k(t_j)\|^2, \quad (17)$$

903 where N_{sim} is the number of time steps. For trajectories that diverge or become
904 numerically unstable, the $\text{MSE}^{(k)}$ is reported as NaN and excluded from averaging.

905 To assess robustness across all trajectories, we also report the AUC based on
906 trajectory-level prediction errors. Given a set of MSE values $\{\text{MSE}^{(k)}\}_{k=1}^N$, we define
907 the success rate at threshold $\delta > 0$ as
908

$$\text{SR}(\delta) = \frac{1}{N} \sum_{k=1}^N \mathbb{I}(\text{MSE}^{(k)} \leq \delta), \quad (18)$$

where $\mathbb{I}(\cdot)$ denotes the indicator function. The AUC is then defined as the normalized area under the success-rate curve over a fixed interval $[0, \delta_{\max}]$:

$$\text{AUC} = \frac{1}{\delta_{\max}} \int_0^{\delta_{\max}} \text{SR}(\delta) d\delta, \quad (19)$$

where $\delta_{\max} = 1.0$ is fixed for all methods and systems. By construction, $\text{AUC} \in [0, 1]$, with larger values indicating that a higher fraction of trajectories achieve low prediction error across a wide range of tolerances.

Data availability

Source data are provided with this paper. All datasets are available via GitHub at <https://github.com/kaixinzeng/HANDI> under the MIT License. Details of data-generation and experimental setup are provided in Supplementary Section 4.

Code availability

All code used for model training and analysis in this study is available via GitHub at <https://github.com/kaixinzeng/HANDI> under the MIT License.

967 References

- 968
- 969 [1] Song, W., Jiang, S., Camps-Valls, G., Williams, M., Zhang, L., Reichstein, M.,
970 Vereecken, H., He, L., Hu, X., Shi, L.: Towards data-driven discovery of governing
971 equations in geosciences. *Communications Earth & Environment* **5**(1), 589 (2024)
- 972
- 973 [2] Colwell, R.K.: Spatial scale and the synchrony of ecological disruption. *Nature*
974 **599**(7886), 8–10 (2021)
- 975
- 976 [3] Zhang, Z., Zhang, Q., Chen, B., Yu, Y., Wang, T., Xu, N., Fan, X., Penuelas, J.,
977 Fu, Z., Deng, Y., *et al.*: Global biogeography of microbes driving ocean ecological
978 status under climate change. *Nature Communications* **15**(1), 4657 (2024)
- 979
- 980 [4] Kaniadakis, G., Baldi, M.M., Deisboeck, T.S., Grisolia, G., Hristopulos, D.T.,
981 Scarfone, A.M., Sparavigna, A., Wada, T., Lucia, U.: The κ -statistics approach
982 to epidemiology. *Scientific Reports* **10**(1), 19949 (2020)
- 983
- 984 [5] Kharazmi, E., Cai, M., Zheng, X., Zhang, Z., Lin, G., Kaniadakis, G.E.: Identi-
985 fiability and predictability of integer-and fractional-order epidemiological models
986 using physics-informed neural networks. *Nature Computational Science* **1**(11),
987 744–753 (2021)
- 988
- 989 [6] Ye, Y., Pandey, A., Bawden, C., Sumsuzzman, D.M., Rajput, R., Shoukat, A.,
990 Singer, B.H., Moghadas, S.M., Galvani, A.P.: Integrating artificial intelligence
991 with mechanistic epidemiological modeling: a scoping review of opportunities and
992 challenges. *Nature Communications* **16**(1), 581 (2025)
- 993
- 994 [7] Zhou, Z., Liu, J., Wu, W.E., Fang, R., Liu, S., Wei, Q., Yan, R., Guo, Y., Tao,
995 Q., Wang, Y., *et al.*: Revealing neurocognitive and behavioral patterns through
996 unsupervised manifold learning of dynamic brain data. *Nature Computational*
997 *Science*, 1–15 (2025)
- 998
- 999 [8] Banerjee, A., Chandra, S., Ott, E.: Network inference from short, noisy, low time-
1000 resolution, partial measurements: Application to *C. elegans* neuronal calcium
1001 dynamics. *Proceedings of the National Academy of Sciences* **120**(12), 2216030120
1002 (2023)
- 1003
- 1004 [9] Howlett, M.G., Fletcher, S.P.: From autocatalysis to survival of the fittest in
1005 self-reproducing lipid systems. *Nature Reviews Chemistry* **7**(10), 673–691 (2023)
- 1006
- 1007 [10] Pu, C., Luan, Y., Wang, Y., Xiao, Z.: Thermal kinetics and flash DSC experimen-
1008 tal validation of conventional nitrocellulose and ladder-structured nitrocellulose
under non-isothermal condition. *Combustion and Flame* **260**, 113216 (2024)
- 1009
- 1010 [11] Alagna, N., Dúzs, B., Dietrich, V., Younesi, A.T., Lehmann, L., Ulbricht, R.,
1011 Köppl, H., Walther, A., Gerber, S.: Deep learning reaction framework (dlrn) for
1012 kinetic modeling of time-resolved data. *Communications Chemistry* **8**(1), 153

(2025)	1013
	1014
[12] Bongard, J., Lipson, H.: Automated reverse engineering of nonlinear dynamical systems. <i>Proceedings of the National Academy of Sciences</i> 104 (24), 9943–9948 (2007)	1015
	1016
	1017
	1018
[13] Schmidt, M., Lipson, H.: Distilling free-form natural laws from experimental data. <i>Science</i> 324 (5923), 81–85 (2009)	1019
	1020
	1021
[14] Pan, W., Yuan, Y., Gonçalves, J., Stan, G.-B.: A sparse bayesian approach to the identification of nonlinear state-space systems. <i>IEEE Transactions on Automatic Control</i> 61 (1), 182–187 (2015)	1022
	1023
	1024
	1025
[15] Brunton, S.L., Proctor, J.L., Kutz, J.N.: Discovering governing equations from data by sparse identification of nonlinear dynamical systems. <i>Proceedings of the National Academy of Sciences</i> 113 (15), 3932–3937 (2016)	1026
	1027
	1028
[16] Champion, K., Zheng, P., Aravkin, A.Y., Brunton, S.L., Kutz, J.N.: A unified sparse optimization framework to learn parsimonious physics-informed models from data. <i>arXiv preprint arXiv:1906.10612</i> (2019)	1029
	1030
	1031
	1032
[17] Messenger, D.A., Bortz, D.M.: Weak SINDy: Galerkin-based data-driven model selection. <i>Multiscale Modeling & Simulation</i> 19 (3), 1474–1497 (2021)	1033
	1034
	1035
[18] Champion, K., Lusch, B., Kutz, J.N., Brunton, S.L.: Data-driven discovery of coordinates and governing equations. <i>Proceedings of the National Academy of Sciences</i> 116 (45), 22445–22451 (2019)	1036
	1037
	1038
	1039
[19] Rudy, S.H., Brunton, S.L., Proctor, J.L., Kutz, J.N.: Data-driven discovery of partial differential equations. <i>Science Advances</i> 3 (4), 1602614 (2017)	1040
	1041
	1042
[20] Chen, Z., Liu, Y., Sun, H.: Physics-informed learning of governing equations from scarce data. <i>Nature Communications</i> 12 (1), 6136 (2021)	1043
	1044
	1045
[21] Lu, P.Y., Ariño Bernad, J., Soljačić, M.: Discovering sparse interpretable dynamics from partial observations. <i>Communications Physics</i> 5 (1), 206 (2022)	1046
	1047
[22] Brunton, S.L., Kutz, J.N.: <i>Data-driven Science and Engineering: Machine Learning, Dynamical Systems, and Control</i> . Cambridge University Press, Cambridge, UK (2022)	1048
	1049
	1050
	1051
[23] Loiseau, J.-C., Noack, B.R., Brunton, S.L.: Sparse reduced-order modelling: Sensor-based dynamics to full-state estimation. <i>Journal of Fluid Mechanics</i> 844 , 459–490 (2018)	1052
	1053
	1054
	1055
[24] Sayahkarajy, M., Witte, H.: Temporal evolution of the hydrodynamics of a swimming eel robot using sparse identification: SINDy-DMD. <i>J</i> 8 (1) (2025)	1056
	1057
	1058

- 1059 [25] Prokop, B., Gelens, L.: From biological data to oscillator models using SINDy.
1060 iScience **27**(4), 109316 (2024)
1061
- 1062 [26] Cranmer, M.: Interpretable machine learning for science with PySR and symbol-
1063 icregression. *jl. arXiv preprint arXiv:2305.01582* (2023)
1064
- 1065 [27] Ruan, K., Xu, Y., Gao, Z.-F., Liu, Y., Guo, Y., Wen, J.-R., Sun, H.: Discovering
1066 physical laws with parallel symbolic enumeration. *Nature Computational Science*,
1067 1–14 (2025)
- 1068 [28] Koopman, B.O.: Hamiltonian systems and transformation in Hilbert space. *Pro-
1069 ceedings of the National Academy of Sciences of the United States of America*
1070 **17**(5), 315 (1931)
1071
- 1072 [29] Mezić, I.: Spectral properties of dynamical systems, model reduction and
1073 decompositions. *Nonlinear Dynamics* **41**(1), 309–325 (2005)
1074
- 1075 [30] Schmid, P.J.: Dynamic mode decomposition of numerical and experimental data.
1076 *Journal of Fluid Mechanics* **656**, 5–28 (2010)
1077
- 1078 [31] Tu, J.H.: Dynamic mode decomposition: Theory and applications. PhD thesis,
1079 Princeton University (2013)
1080
- 1081 [32] Williams, M.O., Kevrekidis, I.G., Rowley, C.W.: A data-driven approximation
1082 of the Koopman operator: Extending dynamic mode decomposition. *Journal of*
1083 *Nonlinear Science* **25**, 1307–1346 (2015)
1084
- 1085 [33] Zeng, Z., Liu, J., Yuan, Y.: A generalized Nyquist-Shannon sampling theorem
1086 using the Koopman operator. *IEEE Transactions on Signal Processing* **72**, 3595–
1087 3610 (2024)
- 1088 [34] Zeng, Z., Yue, Z., Mauroy, A., Gonçalves, J., Yuan, Y.: A sampling theorem
1089 for exact identification of continuous-time nonlinear dynamical systems. *IEEE*
1090 *Transactions on Automatic Control* **69**(12), 8402–8417 (2024)
1091
- 1092 [35] Bideh, A.Z., Georgievska, A., Gryak, J.: MDBench: Benchmarking data-driven
1093 methods for model discovery. *arXiv preprint arXiv:2509.20529* (2025)
1094
- 1095 [36] Galanics, C., Sintár, V., Szalai, I.: Autocatalytic flow chemistry. *Scientific Reports*
1096 **13**(1), 9211 (2023)
1097
- 1098 [37] Klus, S., Nüske, F., Peitz, S., Niemann, J.-H., Clementi, C., Schütte, C.: Data-
1099 driven approximation of the Koopman generator: Model reduction, system
1100 identification, and control. *Physica D: Nonlinear Phenomena* **406**, 132416 (2020)
1101
- 1102 [38] Mauroy, A., Goncalves, J.: Koopman-based lifting techniques for nonlinear sys-
1103 tems identification. *IEEE Transactions on Automatic Control* **65**(6), 2550–2565
1104 (2019)

[39] Yang, A., Axås, J., Kádár, F., Stépán, G., Haller, G.: Modeling nonlinear dynamics from videos. <i>Nonlinear Dynamics</i> 113 (10), 10881–10909 (2025)	1105 1106 1107
[40] Brain Science, A.I.: AllenSDK Documentation. https://allensdk.readthedocs.io/en/latest/index.html . Accessed: 26 Oct 2025	1108 1109 1110
[41] Brunton, S.L., Brunton, B.W., Proctor, J.L., Kaiser, E., Kutz, J.N.: Chaos as an intermittently forced linear system. <i>Nature Communications</i> 8 (1), 19 (2017)	1111 1112 1113
[42] Higham, N.J.: <i>Functions of Matrices: Theory and Computation</i> . SIAM, Philadelphia, PA, USA (2008)	1114 1115 1116
[43] Al-Mohy, A.H., Higham, N.J.: Improved inverse scaling and squaring algorithms for the matrix logarithm. <i>SIAM Journal on Scientific Computing</i> 34 (4), 153–169 (2012)	1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133 1134 1135 1136 1137 1138 1139 1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1150

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Supplementary.pdf](#)
- [sparsemeasurements.mp4](#)