frontiers
in Applied Mathematics
and Statistics

TYPE Original Research
PUBLISHED 14 April 2026
DOI 10.3389/fams.2026.1807939

OPEN ACCESS

EDITED BY
Adnéne Arbi,
Carthage University, Tunisia

REVIEWED BY
Haci Mehmet Baskonus,
Harran University, Türkiye
Jaouad Dabounou,
Hassan Premier University, Morocco

*CORRESPONDENCE
Imane Koulali
imane.koulali@isikun.edu.tr

RECEIVED 10 February 2026
REVISED 19 March 2026
ACCEPTED 23 March 2026
PUBLISHED 14 April 2026

CITATION
Koulali I, Turan E and Eskil MT (2026) LaSIPDE: Latent-Space Identification of Partial Differential Equations from indirect, high-dimensional measurements.
Front. Appl. Math. Stat. 12:1807939.
doi: 10.3389/fams.2026.1807939

LaSIPDE: Latent-Space Identification of Partial Differential Equations from indirect, high-dimensional measurements

Imane Koulali[1]*, Erhan Turan[2] and M. Taner Eskil[1]

[1]Department of Computer Science Engineering, F.M.V. Işık University, Istanbul, Türkiye, [2]Simulargе A.Ş., Istanbul, Türkiye

Discovering governing equations from data is a central challenge in scientific machine learning, particularly when observations are high-dimensional and the underlying state variables are not directly accessible. In this work, we introduce a framework for data-driven discovery of partial differential equations (PDEs) from indirect high-dimensional observations. The proposed approach combines nonlinear representation learning through an autoencoder with sparse identification of governing equations in the latent space, enabling simultaneous model reduction and PDE discovery while preserving spatial structure. Unlike existing methods that either operate on observable variables or discover latent ordinary differential equations, our framework identifies PDEs directly in the learned latent coordinates. We validate the approach on high-dimensional observations generated from Burgers and Korteweg–de Vries (KdV) systems, where the true state variables are intentionally hidden. In both cases, the method successfully recovers the correct dynamical operators, including diffusion, nonlinear advection, and dispersive terms. Although the recovered coefficients differ due to latent coordinate transformations, we show both theoretically and empirically that the discovered equations are dynamically equivalent to the ground-truth systems up to an affine transformation. These results demonstrate that governing PDEs can be recovered from indirect, high-dimensional data without access to the physical state variables, providing a foundation for interpretable model discovery in realistic measurement settings.

KEYWORDS
autoencoder, equation discovery, latent-space identification, partial differential equations, physics-informed algorithm, reduced-order modeling, sparse regression, system identification

1 Introduction

Modeling and predicting complex dynamical systems remain central challenges across scientific and engineering disciplines. Partial differential equations (PDEs) play a fundamental role in describing physical phenomena, yet analytic solutions are rarely available and numerical methods often become computationally intensive for large-scale or multiscale systems [1–3].

Recent advances in data-driven modeling have enabled the discovery of governing equations directly from observed data, offering interpretable models that capture

underlying physical principles. Methods such as sparse regression and physics-informed learning have demonstrated strong performance when the underlying state variables are directly accessible. However, this assumption is rarely satisfied in practice. In many real-world applications, measurements are indirect, high-dimensional, and often transformed representations of the true state. Examples include image-based observations in remote sensing, medical imaging such as dynamic MRI, dense sensor arrays in structural monitoring, or feature-based representations obtained through spectral or modal decompositions.

In such settings, the relationship between the measured data and the underlying physical variables is unknown or highly nonlinear, making direct application of equation discovery methods challenging. Existing approaches either operate directly on observable variables—potentially leading to ill-posed or unstable models—or rely on dimensionality reduction techniques that do not preserve the spatial structure required for PDE identification. Furthermore, many latent-space approaches are limited to discovering ordinary differential equations (ODEs), thereby neglecting the spatial operators that are essential to PDE-governed systems.

To address these limitations, we introduce a framework for simultaneous model reduction and equation discovery that operates directly on high-dimensional, indirect measurement data. By integrating PDE discovery within the latent space of an autoencoder, the proposed method learns low-dimensional coordinates that preserve spatial structure while enabling identification of governing PDEs. This results in an interpretable reduced-order representation of the system dynamics, even when the underlying physical state variables are unknown or inaccessible.

This work introduces several key novelties. First, it enables the discovery of governing *partial differential equations* directly from *indirect, high-dimensional observations*, without requiring access to the underlying physical state variables. Second, unlike existing latent-space approaches that are restricted to ordinary differential equations, the proposed framework explicitly preserves spatial structure and identifies PDE operators in the latent space. Third, we show both theoretically and empirically that the recovered latent dynamics are *dynamically equivalent* to the true governing equations up to an affine transformation, providing a principled explanation for discrepancies in recovered coefficients.

We validate the proposed framework on high-dimensional datasets generated from canonical PDE systems, including the Burgers and Korteweg–de Vries (KdV) equations, where the true state variables are intentionally hidden. The results demonstrate that the method successfully recovers the correct dynamical operators and yields latent-space equations that are dynamically equivalent to the original systems up to an affine transformation.

Throughout this paper, we refer to this framework as LaSIPDE (Latent Space Identification of PDEs from Indirect, High-Dimensional Measurements).

The remainder of this paper is organized as follows. Section 2 reviews related work on equation discovery, reduced-order modeling, and latent dynamical systems. Section 3 presents the proposed framework, including the autoencoder architecture, latent differentiation strategy, and PDE identification procedure. Section 4 reports experimental results on Burgers and KdV systems, along with numerical validation of dynamical equivalence. Section 5 discusses the implications of latent affine indeterminacy and its impact on the interpretation of discovered equations. Finally, Section 6 summarizes the main contributions and outlines directions for future work. Additional data generation parameters, implementation details, derivations, hyperparameter selection heuristics and extended comparisons are provided in the Supplementary material.

2 Related work

2.1 Intrusive and non-intrusive reduced-order modeling

Classical reduced-order modeling (ROM) techniques [4]—such as Proper Orthogonal Decomposition (POD), Galerkin projection—efficiently simulate PDE systems by projecting dynamics onto low-dimensional subspaces, but require explicit knowledge of the governing equations and state variables. In contrast, non-intrusive, data-driven approaches including deep learning-based ROMs (DL-ROMs) [5–13] do not require access to explicit equations, but typically assume that the true state variables are known. Moreover, these methods do not produce symbolic governing equations and often function as black-box models with limited interpretability [14, 15].

2.2 Symbolic regression and sparse system identification

Symbolic regression techniques [16], such as genetic programming [17], Sparse Identification of Nonlinear Dynamics (SINDy), and its extensions [18–20], have been developed to discover governing equations (ODEs and PDEs) directly from time-series or spatial data. These methods use sparse regression to extract interpretable models, offering valuable scientific insight, but fundamentally require direct access to the true state variables. As a result, their applicability is limited to cases where the system coordinates are directly measured and observed, which is not always possible in real-world applications.

2.3 Latent ODE Modeling from Known State Variables (LaSDI Methods)

Recent advances have leveraged autoencoders to compress high-dimensional PDE state-variable data into a low-dimensional latent space, where the system's evolution is modeled by ODEs (e.g., LaSDI, gLaSDI, mLaSDI, and GPLaSDI) [21–25]. These approaches enable substantial dimensionality reduction and efficient simulation, and yield interpretable ODE models in the latent space. However, because spatially extended PDEs are fundamentally transformed into latent ODEs, explicit spatial structure and direct physical correspondence may be lost. This complicates the enforcement of boundary conditions, spatial

TABLE 1 Comparison of representative modeling and equation discovery approaches.

Method class	Requires knowledge of gov. eqs	Indirect high-dim observations	Symbolic equations	PDEs	Spatial structure preserved
Intrusive ROM (POD, Galerkin)	✓	×	× (projected, non-interpretable)	✓	×
DL-ROMs	×	✓	× (black-box)	✓	×
SINDy	×	×	✓ (physical ODE)	×	×
PDE-FIND	×	×	✓ (physical PDE)	✓	✓
LaSDI family	×	×	✓ (latent ODE)	✓	×
SINDy-AE	×	✓	✓ (latent ODE)	✓	×
PINNs/neural operators	✓	×	×	✓	✓
LaSIPDE (this work)	×	✓	✓ (latent PDE)	✓	✓

constraints, and the extraction of physically meaningful laws tied to the original system. As a result, their utility may be limited in domains where preserving physical structure and direct spatial interpretability is essential.

2.4 Latent-Space System Identification from Indirect Observations (ODE Systems)

A different approach, exemplified by SINDy Autoencoder (SINDy-AE) [13, 26], jointly learns intrinsic coordinates and governing equations from indirect, high-dimensional data, eliminating the need for prior knowledge of the state variables. However, SINDy-AE and similar methods are limited to systems governed by ODEs and have not been demonstrated for latent-space PDE discovery or spatiotemporal systems.

2.5 Physics-informed and knowledge-based machine learning for PDEs

Physics-informed machine learning (PIML) [9, 15, 27–29] approaches—including PINNs [1, 30–32], DeepONet [33, 34], and neural operators [35]—incorporate physical constraints into neural network training to approximate solutions and operators for PDEs. However, these methods primarily focus on predictive accuracy rather than symbolic equation discovery and still require direct access to the true state variables.

2.6 Summary and identified research gap

The literature reviewed above highlights substantial progress in reduced-order modeling, symbolic system identification, and physics-informed machine learning. However, these advances address complementary—but largely disjoint—aspects of the broader problem of discovering governing equations from complex data.

Classical and data-driven ROM techniques enable efficient modeling of high-dimensional systems, but either require explicit access to the governing equations or operate as black-box predictors without yielding symbolic models. Sparse regression methods such as SINDy and PDE-FIND provide interpretable equation discovery, yet fundamentally rely on direct access to the true state variables and their derivatives. Autoencoder-based ROM approaches compress high-dimensional PDE data and yield latent dynamical models, but typically reduce spatiotemporal PDEs to latent ODEs, thereby losing explicit spatial structure and limiting physical interpretability. Latent-coordinate discovery methods such as SINDy-AE remove the requirement of known state variables, but have thus far been demonstrated only for ODE systems and not for spatially extended PDEs. Physics-informed learning approaches, including PINNs and neural operators, incorporate physical constraints and achieve high predictive accuracy, but do not aim to recover symbolic governing equations and still assume access to the true state variables.

As a result, there remains no existing framework that simultaneously: (i) operates on high-dimensional, indirect observations without access to the true state variables, (ii) discovers symbolic governing equations, (iii) identifies partial differential equations rather than ordinary differential equations, and (iv) preserves explicit spatial structure in the learned representation. A comparative summary of representative methods along these dimensions is provided in Table 1.

The proposed LaSIPDE (Latent-Space Identification of Partial Differential Equations from Indirect, High-Dimensional Measurements) framework addresses this gap by learning latent coordinates directly from high-dimensional observables and discovering symbolic governing PDEs in the latent space. By integrating autoencoder-based representation learning with sparse PDE identification, the method enables interpretable equation discovery and reduced-order modeling for spatiotemporal systems observed only through indirect, feature-rich measurements.

3 Materials and methods

This section describes the proposed LaSIPDE framework, the synthetic data generation process, and the experimental protocol used to evaluate latent-space PDE discovery from indirect high-dimensional observations.

3.1 Problem formulation

Many physical systems evolve according to spatio-temporal partial differential equations (PDEs). Let $u(x,t)$ denote the (generally unobserved) true state of the system, governed by an unknown dynamical law of the form

$$u_t = \mathcal{N}(u, u_x, u_{xx}, \ldots), \tag{1}$$

where $\mathcal{N}(\cdot)$ represents an unknown nonlinear operator describing the system dynamics.

In classical equation discovery settings, the state variable $u(x,t)$ is assumed to be directly observable. However, in many real-world scenarios, this assumption does not hold. Instead, we have access only to indirect, high-dimensional observations of the system:

$$X(x,t) = \mathcal{H}(u(x,t)) + \Xi, \tag{2}$$

where $\mathcal{H}(\cdot)$ denotes an unknown (possibly nonlinear) observation operator, and Ξ represents measurement noise. Such observations may arise from imaging systems, sensor arrays, or feature-based representations, and typically reside in a high-dimensional space.

Given only these indirect observations $X(x,t)$, the underlying state $u(x,t)$ and its governing equation are both unknown.

The objective of this work is therefore to simultaneously:

1. Learn a low-dimensional latent representation $z(x,t)$ that captures the intrinsic state of the system;
2. Identify an interpretable governing PDE in the latent space of the form

$$z_t = \mathcal{F}(z, z_x, z_{xx}, \ldots), \tag{3}$$

where $\mathcal{F}$ is a sparse operator approximating the true underlying dynamics.

This formulation removes the requirement of direct access to the physical state $u(x,t)$, and instead seeks to recover the governing dynamics from indirect, high-dimensional measurements.

3.2 LaSIPDE framework

Our approach integrates *autoencoder-based dimensionality reduction* with *sparse regression* to jointly learn a *latent representation* and its *governing equation* in an end-to-end manner. By imposing a PDE constraint in the latent space, the framework encourages the learned variables to capture dynamically meaningful structure rather than purely data-driven features. This results in a reduced-order representation that remains interpretable while being consistent with the underlying system dynamics, even when the true state variables are not directly observed.

An overview of the proposed framework is presented in Figure 1. The method consists of (1) an autoencoder that learns a compact latent representation from high-dimensional spatiotemporal observations, and (2) a sparse regression module that identifies governing equations in this latent space.

3.2.1 Learning latent representations via autoencoder

Given high-dimensional spatiotemporal observations $X(x,t) \in \mathbb{R}^m$, we employ a neural autoencoder to learn a low-dimensional latent representation $u \in \mathbb{R}^k$ (with $k \ll m$). The autoencoder consists of:

- *Encoder* $\phi(\cdot;\theta_{\text{enc}})$: maps X to latent state u
- *Decoder* $\psi(\cdot;\theta_{\text{dec}})$: reconstructs X from u

Formally,

$$u = \phi(X;\theta_{\text{enc}}), \quad \hat{X} = \psi(u;\theta_{\text{dec}}) = \psi(\phi(X)), \tag{4}$$

where $\hat{X}$ is the reconstructed approximation of the original input X, and θ_{enc}, θ_{dec} are the trainable parameters of the encoder and decoder respectively.

3.2.2 Latent variable differentiation

To construct a PDE in the latent space, it is necessary to compute temporal and spatial derivatives of the latent variable $u(x,t) = \phi(X(x,t))$.

Time derivatives are computed using the chain rule. Since the encoder ϕ is differentiable, the temporal evolution of the latent variables is obtained as:

$$\frac{\partial u}{\partial t} = \nabla_X \phi(X) \cdot \frac{\partial X}{\partial t}, \tag{5}$$

where $\nabla_X \phi(X)$ denotes the Jacobian of the encoder, and $\frac{\partial X}{\partial t}$ is obtained from the data.

Spatial derivatives are computed directly in the latent space using finite difference schemes applied to $u(x,t)$ over the spatial grid. In practice, first- and higher-order derivatives (e.g., u_x, u_{xx}, u_{xxx}) are approximated using centered finite differences.

This hybrid approach enables efficient and stable computation of latent derivatives without requiring explicit knowledge of the underlying physical state variable. Additional details on the differentiation of the latent variable and its design choice are provided in Supplementary Material S2.

3.2.3 Latent-space governing equation discovery

Using the compact latent representation learned by the autoencoder, the aim is to uncover an explicit governing partial differential equation for the latent state variable(s) $u(x,t)$, assumed to have the general form:

$$\frac{\partial u}{\partial t} = \sum_i \xi_i \Theta_i(u), \tag{6}$$

where $\Theta_i(u)$ are candidate functions of u from a predefined library, and ξ_i are the corresponding coefficients to be identified.

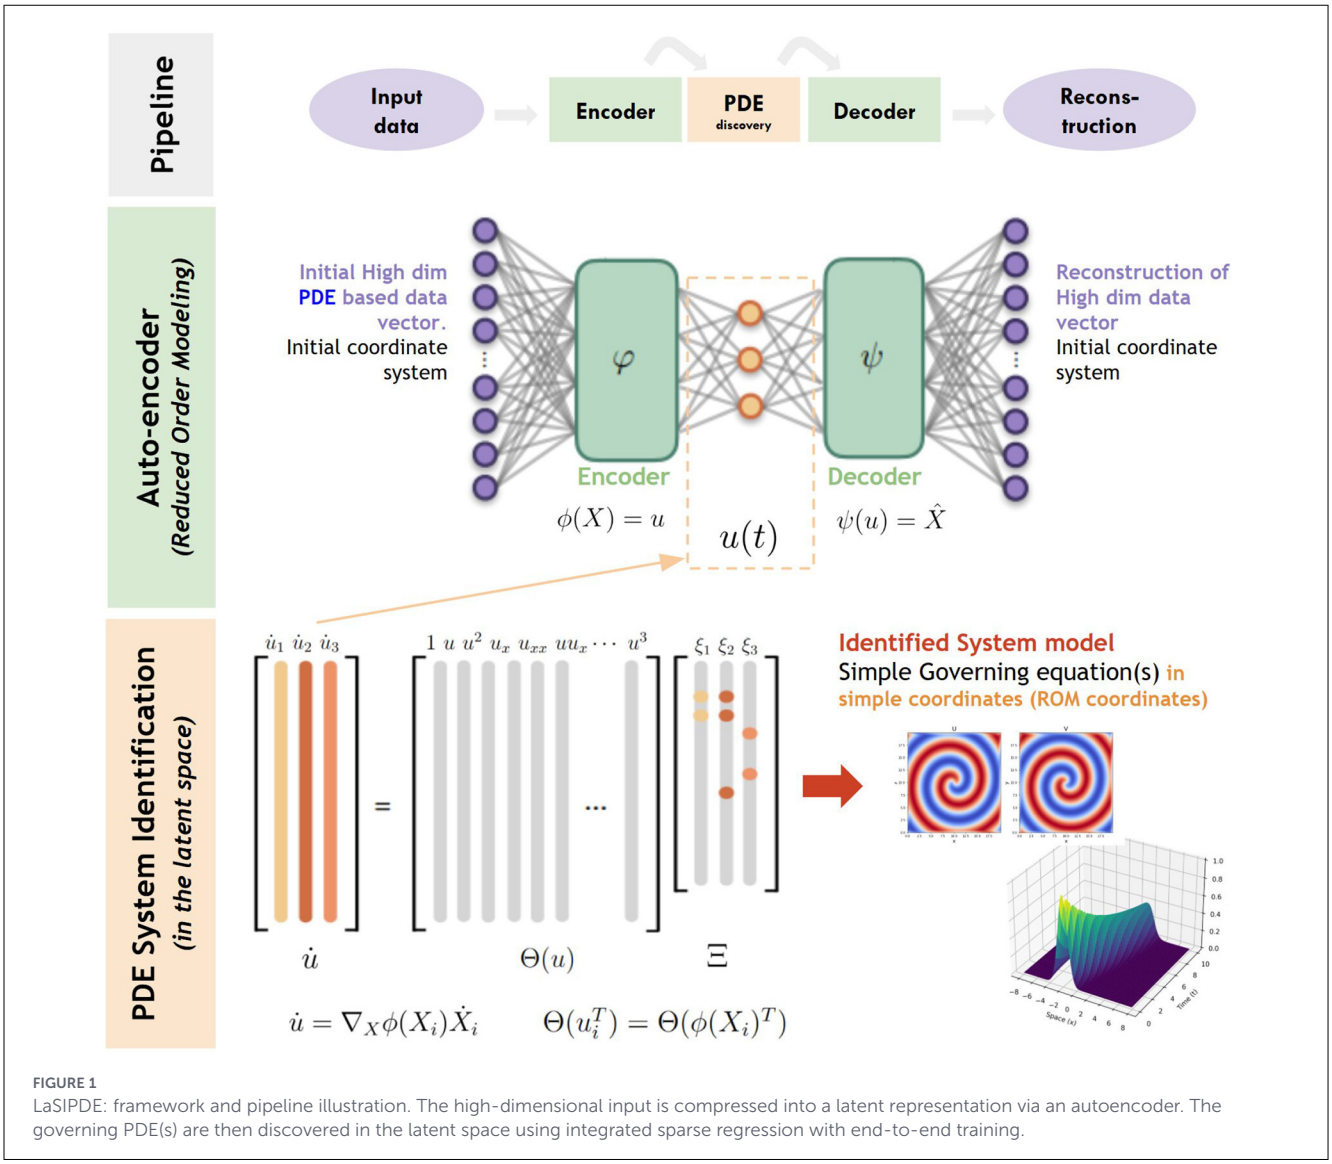

FIGURE 1
LaSIPDE: framework and pipeline illustration. The high-dimensional input is compressed into a latent representation via an autoencoder. The governing PDE(s) are then discovered in the latent space using integrated sparse regression with end-to-end training.

The function library $\Theta(u)$ includes combinations of polynomial terms and spatial derivatives, such as:

$$\Theta(u) = \left[1,\ u,\ u^2,\ \frac{\partial u}{\partial x},\ \frac{\partial^2 u}{\partial x^2},\ u\frac{\partial u}{\partial x},\ \ldots\right] \tag{7}$$

capturing constant, nonlinear, and differential operators relevant for PDE dynamics. Depending on the application, other function types (e.g., trigonometric functions, exponentials or more general nonlinearities) may also be incorporated to enhance expressive power.

The library is evaluated at each latent state, and the coefficients ξ_i are treated as trainable parameters. Rather than performing sparse regression as a separate step, these coefficients are optimized jointly with the neural network weights through a sparsity-promoting loss, enabling end-to-end discovery of the governing equations.

3.2.4 Loss function design

To enable end-to-end discovery of interpretable latent dynamics, we design a composite loss balancing three objectives: (1) accurate reconstruction of input data, (2) consistency between learned and observed dynamics, and (3) sparsity in the discovered PDE. The total training loss is:

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \lambda_1 \mathcal{L}_{\dot{X}} + \lambda_2 \mathcal{L}_{\dot{u}} + \lambda_3 \mathcal{L}_{\text{reg}} \tag{8}$$

where $\lambda_1, \lambda_2, \lambda_3$ control the trade-offs between accuracy, physical consistency, and sparsity, with:

3.2.4.1 Reconstruction loss

Ensures the autoencoder accurately reconstructs the input from the latent space.

$$\mathcal{L}_{\text{recon}} = \frac{1}{m}\sum_{i=1}^{m} \left\| X_i - \hat{X}_i \right\|_2^2 \tag{9}$$

3.2.4.2 Input-space dynamic loss

Encourages the reconstructed trajectory's time evolution to match the predicted latent dynamics mapped back to input space.

$$\mathcal{L}_{\dot{X}} = \frac{1}{m}\sum_{i=1}^{m}\left\|\dot{X}_i - \left(\nabla_u \psi(\phi(X_i))\right)\Theta(\phi(X_i)^T)\Xi\right\|_2^2 \quad (10)$$

3.2.4.3 Latent-space dynamic loss

Ensures the latent variables' time derivatives are consistent with the discovered PDE model.

$$\mathcal{L}_{\dot{u}} = \frac{1}{m}\sum_{i=1}^{m}\left\|\nabla_X \phi(X_i)\dot{X}_i - \Theta(\phi(X_i)^T)\Xi\right\|_2^2 \quad (11)$$

3.2.4.4 Sparsity regularization

Promotes compact, interpretable equations by penalizing nonzero coefficients.

$$\mathcal{L}_{\text{reg}} = \frac{1}{pd}\|\Xi\|_1 \quad (12)$$

To further encourage interpretability, we periodically prune small coefficients in Ξ using sequential thresholding, following the STLSQ strategy.

3.2.5 Training strategy

The entire framework—including the autoencoder, decoder, and sparse regression coefficients—is trained end-to-end via gradient-based optimization. The entire framework—including the autoencoder, decoder, and sparse regression coefficients—is trained end-to-end via gradient-based optimization. Regularization, normalization, and sparsity-promoting strategies are employed to ensure stable training and interpretable equation discovery. Full training details and hyperparameters are provided in Supplementary material S3.

3.3 Experimental setup

3.3.1 Governing canonical PDE systems

To evaluate the proposed framework, we consider two canonical nonlinear partial differential equations: the one-dimensional viscous Burgers equation and the Korteweg–de Vries (KdV) equation. These systems are widely used as benchmark models in scientific computing and represent two distinct classes of nonlinear dynamics.

We denote the system state by $u(x,t)$, where subscripts indicate partial derivatives, e.g., $u_t = \partial u/\partial t$, $u_x = \partial u/\partial x$, $u_{xx} = \partial^2 u/\partial x^2$, and $u_{xxx} = \partial^3 u/\partial x^3$.

3.3.1.1 Example 1: Viscous burgers equation

The viscous Burgers equation is given by:

$$u_t = \nu u_{xx} - u u_x \quad (13)$$

where $\nu > 0$ is the viscosity coefficient. The term uu_x represents nonlinear advection, while νu_{xx} introduces diffusive smoothing. This equation models dissipative transport phenomena and is known for the formation of steep gradients and shock-like structures.

3.3.1.2 Example 2: Korteweg–de Vries (KdV) equation

The Korteweg–de Vries (KdV) equation can be written in evolution form as:

$$u_t = -6uu_x - u_{xxx} \quad (14)$$

where the nonlinear term $6uu_x$ governs advective steepening, and the third-order derivative u_{xxx} introduces dispersive effects. The balance between nonlinearity and dispersion gives rise to stable traveling wave solutions (solitons), making KdV a fundamental model for wave propagation in dispersive media.

These two systems provide complementary test cases: Burgers captures diffusion-dominated dynamics, while KdV captures dispersion-dominated dynamics. In both cases, the true state variable $u(x,t)$ is used only for data generation and is not directly accessible to the model.

3.3.2 Synthetic high-dimensional observations

Let x denote the spatial coordinate and t denote time. The true state of the system is represented by a scalar field $u(x,t)$, which evolves according to the governing PDE. In many practical settings, however, this underlying state variable is not directly observable, and measurements are instead obtained through indirect, high-dimensional transformations of the physical system.

To emulate such conditions, we consider an observation model in which the true state $u(x,t)$ is mapped to a high-dimensional representation $X(u(x,t)) \in \mathbb{R}^m$. This mapping is constructed through nonlinear interactions with fixed basis vectors, producing complex, feature-rich observations that resemble outputs from sensors, modal expansions, or learned feature representations.

Importantly, the proposed framework is trained exclusively on these high-dimensional observations $X(u(x,t))$, without access to the true state $u(x,t)$. The objective is therefore to recover an interpretable representation of the underlying governing dynamics solely from indirect measurements.

Figure 2 illustrates this setting by showing the original solution $u(x,t)$ alongside representative high-dimensional features $X_i(u(x,t))$. While individual features exhibit strong nonlinear distortions, the underlying spatiotemporal structure of the dynamics remains implicitly encoded in the observation space.

All datasets are normalized prior to training, and small Gaussian noise is optionally added to emulate realistic measurement conditions. Full details of the governing equations, numerical simulation procedures, initial conditions, and high-dimensional mapping are provided in Supplementary material S1.

3.3.3 Network architecture and parameters

We employ a symmetric, fully connected autoencoder architecture across all experiments, consisting of linear input

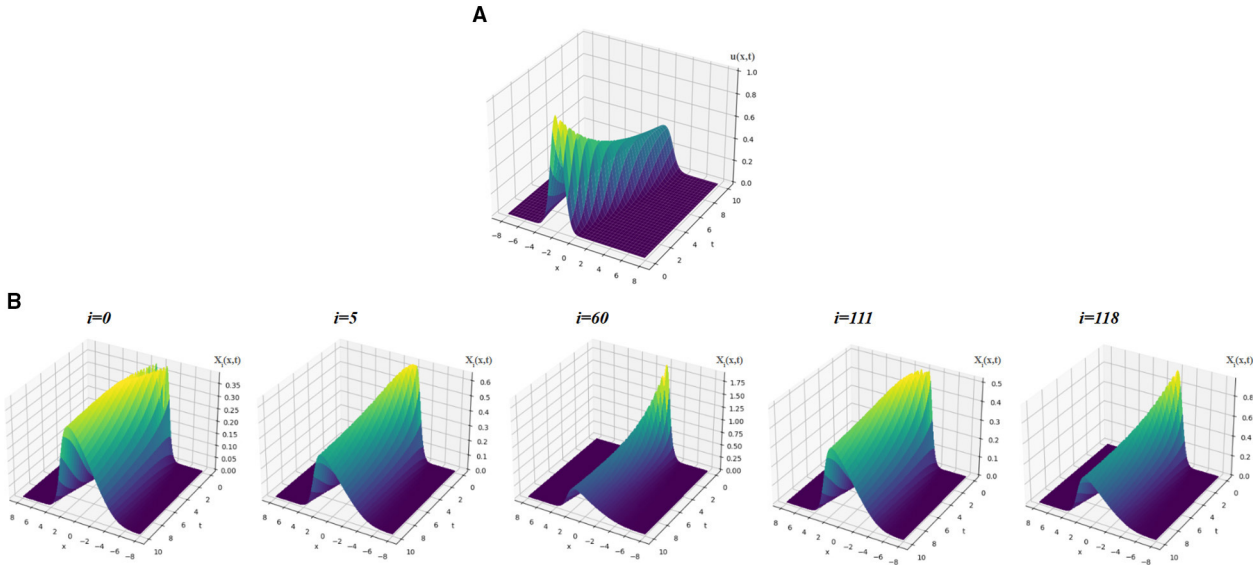

FIGURE 2
Synthetic high-dimensional observation model for the Burgers' system. **(A)** The original physical state $u(x, t)$ representing the hidden governing dynamics. **(B)** Representative channels from the high-dimensional observation space X, illustrating the nonlinear distortions and feature-rich mapping from which the model must autonomously discover the underlying PDE. **(A)** Ground-truth state $u(x, t)$. **(B)** High-dimensional features $X_i(u)$.

and output layers and nonlinear hidden activations. The design is intentionally kept simple to prioritize interpretability and stability in the discovered governing equations, avoiding overparameterization that could obscure the identification of sparse dynamics.

Several architectural and training choices are kept consistent across both systems. In particular, all models use Xavier (Glorot) initialization for weights, while biases and equation discovery coefficients (Ξ) are initialized to zeros and ones, respectively.

Sequential thresholding is applied every 500 epochs, permanently zeroing coefficients with magnitude below 0.1 to promote sparsity in the identified PDEs. The candidate function library includes polynomial terms and spatial derivatives up to third order.

While the overall framework remains unchanged, certain hyperparameters are adapted to account for differences in the underlying dynamical systems and data complexity. In particular, the KdV system requires a deeper encoder–decoder architecture and a longer training duration to accommodate its more complex dispersive dynamics and longer temporal horizon.

A summary of all architectural choices and training hyperparameters for each system is provided in Tables 2, 3. These tables highlight both the shared design principles and the system-specific adjustments made for Burgers and KdV experiments. Further details about the training strategy can be found in Supplementary material S3.

TABLE 2 Summary of experimental parameters for the burgers' equation-based high-dimensional system.

Parameter	Value
Spatial grid points	256
Time steps	100
High-dimensional input size	128
Encoder hidden layers	[64, 32]
Latent dimension	1
Decoder hidden layers	[32, 64]
Activation function	Sigmoid
Weight initialization	Xavier (Glorot)
Optimizer	Adam
Learning rate	1×10^{-4}
Training epochs	6,000
Thresholding interval	Every 500 epochs
Thresholding cutoff	0.1

3.4 Model performance evaluation metrics

To rigorously assess how well the proposed framework captures the structure and evolution of the underlying dynamical system, our evaluation is based on three relative L_2 errors:

$$e_{\mathbf{x}} = \frac{\sum_{i=1}^{m} \left\| \mathbf{x}_i - \psi(\varphi(\mathbf{x}_i)) \right\|_2^2}{\sum_{i=1}^{m} \left\| \mathbf{x}_i \right\|_2^2} \tag{15}$$

which reflects how well the encoder-decoder pair preserves and reconstructs the original input observations; (2) the input time derivative error,

$$e_{\dot{\mathbf{x}}} = \frac{\sum_{i=1}^{m} \left\| \dot{\mathbf{x}}_i - (\nabla_z \psi(\varphi(\mathbf{x}_i))) \left(\Theta(\varphi(\mathbf{x}_i))^T \Xi\right) \right\|_2^2}{\sum_{i=1}^{m} \left\| \dot{\mathbf{x}}_i \right\|_2^2} \tag{16}$$

TABLE 3 Summary of experimental parameters for the KdV' equation-based High-dimensional system.

Parameter	Value
Spatial grid points	256
Time steps	1,024
High-dimensional input size	128
Encoder hidden layers	[64, 32, 16]
Latent dimension	1
Decoder hidden layers	[16, 32, 64]
Activation function	Tanh
Weight initialization	Xavier (Glorot)
Optimizer	Adam
Learning rate	1×10^{-3}
Training epochs	15,000
Thresholding interval	Every 500 epochs
Thresholding cutoff	0.1

which quantifies how closely the dynamics predicted by the learned PDE match the observed input dynamics when mapped back to the input space; and (3) the latent time derivative error,

$$e_{\dot{\mathbf{z}}} = \frac{\sum_{i=1}^{m} \left\| \nabla_x \varphi(\mathbf{x}_i)\dot{\mathbf{x}}_i - \Theta(\varphi(\mathbf{x}_i))^T(\Upsilon \odot \Xi) \right\|_2^2}{\sum_{i=1}^{m} \left\| \nabla_x \varphi(\mathbf{x}_i)\dot{\mathbf{x}}_i \right\|_2^2} \tag{17}$$

which measures the consistency between the true latent dynamics (obtained by backpropagating input derivatives) and those predicted by the symbolic model.

Together, these metrics comprehensively characterize reconstruction fidelity and the quality of dynamic modeling in both the input and latent spaces.

These metrics are invariant to affine reparameterizations of the latent variables and therefore provide a meaningful assessment of model performance under the identifiability limits discussed in Section 5.

4 Results

4.1 Example 1: burgers equation

The performance of the proposed framework on the Burgers equation-based high-dimensional system is summarized in Table 4. The quantitative results reported in Table 4 confirm that the model accurately reconstructs both the system state and its temporal evolution. The model achieves very low relative errors for state reconstruction and dynamical consistency, indicating that the learned latent representation preserves both the system state and its temporal evolution with high fidelity.

The discovered latent-space equation correctly identifies the governing structure of the Burgers dynamics, including both the diffusion term and the nonlinear advection term. The diffusion coefficient is recovered with high accuracy (0.099943 compared to the ground-truth value 0.100000), demonstrating the ability of the framework to accurately identify linear operators from indirect observations.

TABLE 4 Summary of results for the burgers equation-based high-dimensional system.

Category	Metric	Value
Model performance evaluation	e_x	6.20×10^{-4}
	$e_{\dot{x}}$	3.961×10^{-2}
	$e_{\dot{z}}$	2.026×10^{-2}
Discovered governing equation	**Ground truth**	$u_t = 0.100000\, u_{xx} - 1.000000\, uu_x$
	Discovered	$z_t = 0.099943\, z_{xx} + 0.468257\, zz_x$

TABLE 5 Summary of results for the KdV equation-based high-dimensional system.

Category	Metric	Value
Model performance evaluation	e_x	1.60×10^{-4}
	$e_{\dot{x}}$	6.70×10^{-3}
	$e_{\dot{z}}$	1.89×10^{-3}
Discovered governing equation	**Ground truth**	$u_t = -6\, uu_x - u_{xxx}$
	Discovered	$z_t = -2.8882\, z_x - 1.3220\, zz_x - 1.0621\, z_{xxx}$

The nonlinear advection term is also correctly selected, although its coefficient differs in both sign and magnitude from the physical value. This discrepancy is consistent across runs and reflects a systematic transformation of the latent variable rather than a failure to recover the underlying dynamics.

Overall, these results indicate that the proposed framework successfully recovers the correct dynamical structure and key physical operators of the Burgers system from high-dimensional, indirect observations.

4.2 Example 2: Korteweg–de Vries (KdV) equation

The performance of the proposed framework on the KdV equation-based high-dimensional system is summarized in Table 5. The quantitative results reported in Table 5 further confirm the ability of the framework to capture both nonlinear and dispersive dynamics from indirect observations.

The discovered latent-space equation correctly identifies the key dynamical components of the KdV system, including the nonlinear advection term and the third-order dispersive term. The coefficient associated with the dispersive operator is recovered with good accuracy (1.0621 compared to the ground-truth magnitude 1.0000), indicating that the framework successfully captures higher-order linear differential operators from indirect observations.

The nonlinear advection term is also correctly selected, although its coefficient differs in magnitude and sign from the physical value. In addition, a linear transport term proportional to z_x appears in the discovered equation. These discrepancies are systematic and consistent across runs, indicating that they arise

from a transformation of the latent variable rather than incorrect identification of the governing dynamics.

Overall, these results demonstrate that the proposed framework successfully recovers the correct dynamical structure of the KdV system from high-dimensional, indirect observations, including both nonlinear and higher-order dispersive effects.

4.3 Numerical verification of dynamical equivalence

To assess whether the discovered latent-space dynamics are dynamically equivalent to the ground-truth systems, we perform a *post hoc* affine alignment between the latent solution $z(x,t)$ and the physical solution $u(x,t)$. We emphasize that this affine alignment is used solely for *post hoc* verification and visualization.

We consider an empirical affine mapping of the form

$$u(x,t) \approx \alpha\, z(x,t) + \beta, \tag{18}$$

where (α, β) are obtained via least-squares regression over the full spatiotemporal domain.

To quantify the discrepancy between the ground-truth solution $u(x,t)$ and the affine-aligned solution $\hat{u}(x,t) = \alpha z(x,t) + \beta$, we consider two standard error metrics.

The root-mean-square error (RMSE) is defined as

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(u_i - \hat{u}_i\right)^2}, \tag{19}$$

which measures the absolute reconstruction error over the spatiotemporal domain.

The relative L_2 error is defined as

$$\mathrm{rel}\, L_2 = \frac{\|u - \hat{u}\|_2}{\|u\|_2}, \tag{20}$$

which quantifies the error relative to the magnitude of the ground-truth solution.

The observed discrepancies are consistent with an affine transformation of the latent variable, whose parameters can be independently estimated from the recovered coefficients and verified numerically (see Supplementary Material S4 for details). The affine parameters obtained from coefficient comparison closely match those obtained from least-squares fitting, providing independent validation of the affine equivalence.

- **Example 1: Burgers equation:** For the Burgers system, the fitted parameters are

$$\alpha = -0.468220, \qquad \beta = 8.69 \times 10^{-6}, \tag{21}$$

with $\mathrm{RMSE} = 2.269 \times 10^{-5}$ and $\mathrm{rel}\, L_2 = 1.067 \times 10^{-4}$.

- **Example 2: KdV equation:** For the KdV system, the fitted parameters are

$$\alpha = -0.220333, \qquad \beta = -4.81 \times 10^{-1}, \tag{22}$$

with $\mathrm{RMSE} = 2.398 \times 10^{-5}$ and $\mathrm{rel}\, L_2 = 8.798 \times 10^{-5}$.

Figures 3, 4 compare the ground-truth solutions, the solutions generated by the discovered latent-space PDEs, the affine-aligned solutions, and the residual differences. In both cases, the residuals remain small relative to the solution amplitude, confirming that the discovered latent-space dynamics reproduce the true system behavior up to an affine transformation of the state variable.

These results provide strong empirical evidence that the identified governing equations are dynamically equivalent to the original physical systems.

4.4 Robustness to measurement noise

To evaluate the robustness of the proposed framework to measurement noise, we repeat the Burgers experiment under increasing noise levels.

To ensure a consistent and interpretable noise model, additive Gaussian noise is introduced into the high-dimensional observations and scaled relative to the standard deviation of the signal. Specifically,

$$\tilde{X} = X + \sigma \cdot \mathrm{std}(X) \cdot \epsilon, \tag{23}$$

where $\epsilon \sim \mathcal{N}(0,1)$ and σ denotes the prescribed noise level. To ensure a controlled and reproducible comparison across the different noise levels, a fixed random seed is used for both the network initialization and the stochastic components of the data generation. In this formulation, σ represents the noise amplitude as a fraction of the signal standard deviation, enabling a consistent interpretation across experiments. For example, $\sigma = 10^{-2}$ corresponds to 1% noise, while $\sigma = 5 \times 10^{-2}$ corresponds to 5% noise relative to the signal magnitude. This normalization ensures that the signal-to-noise ratio (SNR) is controlled consistently across experiments.

The discovered equations and corresponding error metrics are summarized in Table 6. The first row reports the expected latent-space form under affine transformations, which serves as a reference for assessing dynamical equivalence.

At low noise levels ($\leq 10^{-2}$), the method consistently recovers the correct dynamical structure, yielding latent PDEs that remain affinely equivalent to the ground-truth system. Both the diffusion coefficient and nonlinear interaction term are accurately identified, and the resulting dynamics remain stable with low trajectory error.

As the noise level increases beyond 10^{-2}, degradation in the identified model becomes apparent. At 2×10^{-2}, the discovered equation includes an additional transport term proportional to z_x, which is consistent with the expected affine-transformed form. However, the estimated coefficients deviate significantly, and the resulting trajectory errors increase substantially, indicating a loss of dynamical equivalence.

At higher noise levels, essential terms are no longer reliably recovered, leading to structural failure of the identified model. In practice, we observe that noise levels exceeding approximately 3×10^{-2} consistently result in breakdown of the identified dynamics. For noise levels in the range $\sigma \geq 5 \times 10^{-2}$ (i.e., $\geq$ 5% of the signal), the identified models consistently fail to recover the

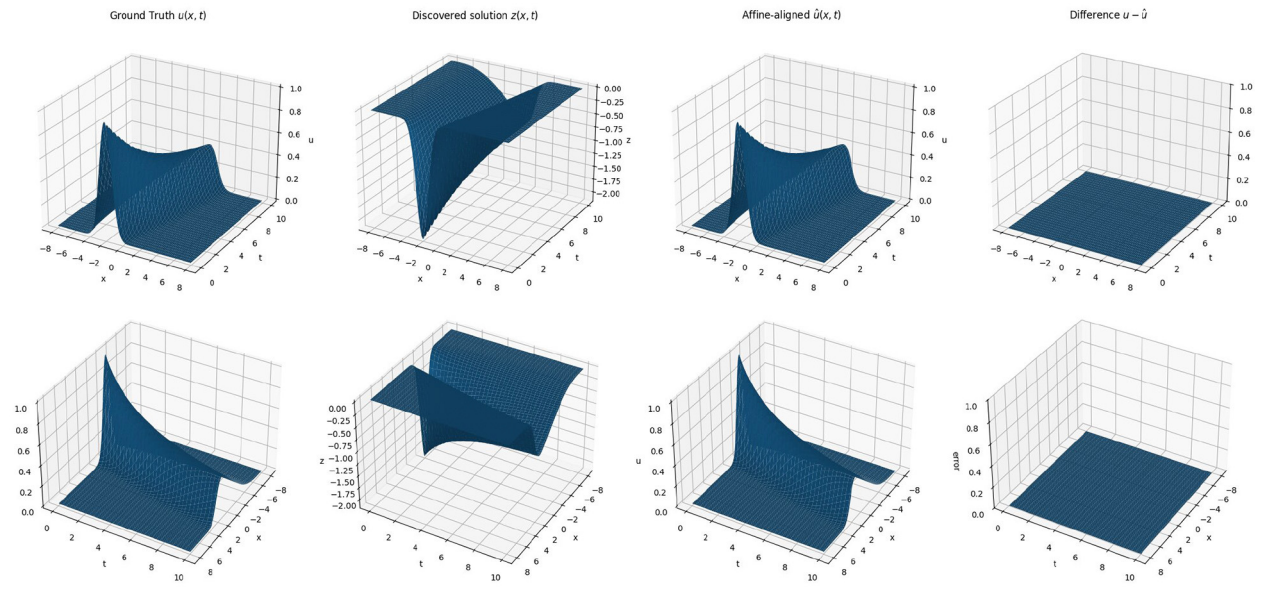

FIGURE 3
Spatiotemporal comparison of the ground-truth Burgers solution $u(x, t)$, the solution $z(x, t)$ generated by the discovered latent-space PDE, the affine-aligned solution $\hat{u}(x, t) = \alpha z(x, t) + \beta$, and the residual difference $u - \hat{u}$. Top and bottom rows show two viewing angles. The difference surface is plotted on the same vertical scale as the ground-truth solution, illustrating that residual discrepancies are negligible relative to the solution amplitude.

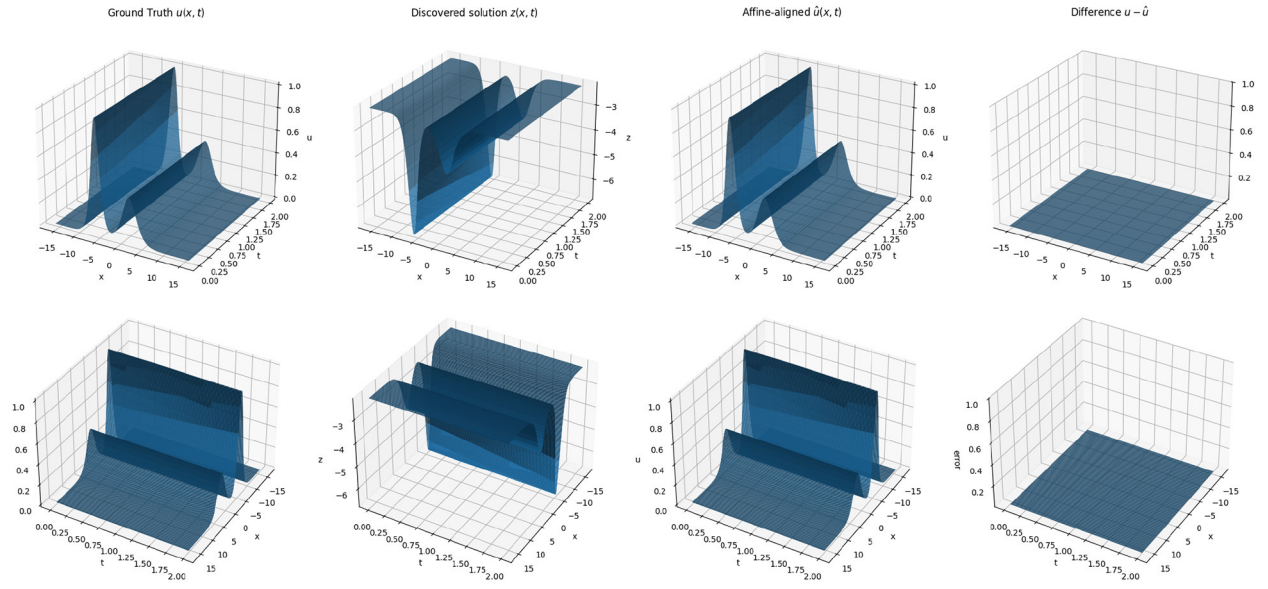

FIGURE 4
Spatiotemporal comparison of the ground-truth KdV solution $u(x, t)$, the solution $z(x, t)$ generated by the discovered latent-space PDE, the affine-aligned solution $\hat{u}(x, t) = \alpha z(x, t) + \beta$, and the residual difference $u - \hat{u}$. Top and bottom rows show two viewing angles. The difference surface is plotted on the same vertical scale as the ground-truth solution, illustrating that residual discrepancies are negligible relative to the solution amplitude.

correct dynamical structure, with missing essential terms and large trajectory errors.

These results indicate that the proposed framework is robust to moderate noise levels (up to approximately 1%–2% of the signal amplitude), but exhibits a clear transition from accurate recovery to structural degradation as noise level increases beyond that.

4.5 Illustrative comparison with existing approaches

To provide additional context for the proposed framework, we evaluated several representative baseline methods under the same indirect high-dimensional observation setting, including *(i)*

TABLE 6 Robustness of latent PDE discovery under increasing noise levels for the Burgers equation. Metrics are reported after affine alignment.

Noise	Discovered equation	Dynamical behavior	ν Error	rel L_2	RMSE
Expected	$z_t = \nu z_{xx} - \frac{1}{a} z z_x + \frac{b}{a} z_x$	Affine-equivalent	0.0		
0	$z_t = 0.100000 z_{xx} + 0.375700 z z_x$	Affine-equivalent	0.0	9.59×10^{-8}	2.04×10^{-8}
10^{-6}	$z_t = 0.099943 z_{xx} + 0.468257 z z_x$	Affine-equivalent	5.7×10^{-4}	1.07×10^{-4}	2.27×10^{-5}
10^{-2}	$z_t = 0.099900 z_{xx} + 0.770700 z z_x$	Affine-equivalent	1.0×10^{-3}	1.87×10^{-4}	3.98×10^{-5}
2×10^{-2}	$z_t = 0.1919 z_{xx} + 0.1649 z z_x + 0.7280 z_x$	Degraded	9.19×10^{-1}	8.42×10^{-1}	1.79×10^{-1}
3×10^{-2}	Essential coeffs z_x and $z z_x$ pruned	Structure failure	—	—	—

POD combined with PDE-FIND, *(ii) SINDy-AE* applied to high-dimensional observations, and *(iii) a vanilla autoencoder* trained using reconstruction loss only.

These methods are applied outside their intended regimes, and the comparison is therefore intended to highlight practical limitations rather than to establish a formal performance ranking. In particular, existing approaches either require direct access to the true state variables, operate on projected representations that may not preserve governing operators, or are restricted to ordinary differential equation (ODE) dynamics without explicit spatial structure.

All methods were evaluated on the same KdV-based high-dimensional dataset. In this setting, POD + PDE-FIND achieved low reconstruction error but failed to recover the correct PDE structure, SINDy-AE produced degenerate low-variance polynomial ODE dynamics without spatial terms, and the standard autoencoder provided accurate reconstruction without yielding interpretable dynamics. In contrast, the proposed framework successfully recovered nonlinear and dispersive operators consistent with the underlying PDE while maintaining low dynamical error across all reported metrics.

A detailed description of the experimental setup, discovered models, and quantitative metrics (including the relative L_2 errors e_X, $e_{\dot{X}}$, and $e_{\dot{Z}}$) is provided in Supplementary material S5.

5 Discussion

5.1 Latent affine indeterminacy

Latent representations learned by autoencoder-based models are inherently non-identifiable, in the sense that the learned variable $z(x,t)$ is defined only up to an affine transformation of the underlying physical state $u(x,t)$. In general, the encoder may recover a representation of the form

$$z(x,t) = a\,u(x,t) + b, \qquad a \neq 0, \tag{24}$$

without affecting reconstruction accuracy.

This ambiguity arises because both the encoder and decoder can compensate for global scaling and shifting transformations without altering the reconstructed observations. As a result, the learned latent variable does not necessarily coincide with the physical state, but instead represents an equivalent coordinate system in which the dynamics are expressed.

In this context, the encoder is free to scale ($|a| \neq 1$), sign-flip ($a < 0$), or shift ($b \neq 0$) the latent variable z relative to the ground-truth u without loss of generality. Such transformations do not represent a failure of the model; rather, they yield a dynamically equivalent representation where the decoder symmetrically compensates for the latent reparameterization to ensure accurate reconstruction of the observables. As long as the discovered latent PDE is affinely equivalent to the physical governing law—a property we verify via least-squares mapping between the learned and physical states—the identification is considered a success. For instance, in a system where the underlying dynamics are governed by the Burgers' equation, the essential physics lies in how the wave propagates, steepens, and dissipates over time. If the discovered model accurately reproduces these characteristic behaviors, the identification is physically true within its own latent scale. This flexibility allows the framework to remain strictly *unsupervised*, discovering the essential dynamical structure without requiring prior knowledge of the physical units or orientation of the underlying state.

5.2 Effect on discovered governing equations

This affine indeterminacy has a direct impact on the coefficients of the discovered governing equations. Substituting $z = au + b$ into a PDE reveals that linear differential operators (e.g., diffusion or dispersion terms) remain invariant under affine transformations, while nonlinear terms scale with the latent amplitude parameter a.

Additionally, a nonzero offset b introduces new linear terms that are not present in the original physical equation. In particular, an affine shift can generate an additional transport term proportional to z_x.

For example, applying this transformation to the Burgers equation yields a latent-space equation of the form

$$z_t = \nu z_{xx} - \frac{1}{a} z z_x + \frac{b}{a} z_x, \tag{25}$$

where the diffusion coefficient remains unchanged, while the nonlinear term is rescaled and an additional transport term may appear if $b \neq 0$.

Similarly, for the KdV equation, the transformed dynamics take the form

$$z_t = -\frac{6}{a} z z_x + \frac{6b}{a} z_x - z_{xxx}, \tag{26}$$

where the dispersive term remains invariant, while nonlinear and transport terms are affected by the affine parameters.

As a result, discrepancies in the magnitude and sign of nonlinear coefficients, as well as the appearance of additional linear terms, are expected consequences of latent reparameterization rather than incorrect model identification. Detailed derivations of these transformations and coefficient consistency analysis are provided in Supplementary material S4.

5.3 Interpretation of results

The results obtained for both the Burgers and KdV systems are consistent with the theoretical effects of affine transformations on the governing equations.

In the Burgers experiment, the diffusion coefficient is recovered with high accuracy, consistent with its invariance under affine transformations. In contrast, the nonlinear advection coefficient differs in both sign and magnitude, reflecting the scaling of the latent variable. The absence of an additional transport term indicates that the offset parameter b is negligible in this case, implying that the learned latent representation primarily differs from the physical state by a global scaling and sign change.

In the KdV experiment, the coefficient associated with the dispersive term is similarly recovered with good accuracy, confirming that higher-order linear differential operators are preserved under the learned representation. The nonlinear advection coefficient again differs in magnitude and sign, consistent with latent scaling. In addition, the presence of a linear transport term proportional to z_x indicates that the offset parameter b is non-negligible, as predicted by the affine transformation analysis.

Combined with the numerical verification results in Section 4.3, these observations confirm that the discovered latent-space equations for both systems are affinely equivalent to their respective ground-truth PDEs, and that the underlying dynamics—both diffusive and dispersive—have been correctly identified from indirect high-dimensional observations.

6 Conclusion

This work introduces a data-driven framework for the discovery of governing partial differential equations (PDEs) from indirect high-dimensional observations. The proposed approach integrates nonlinear reduced-order representation learning through an autoencoder with sparse identification of governing equations in the latent space, enabling simultaneous model reduction and equation discovery in an end-to-end manner.

Through experiments on Burgers and Korteweg–de Vries (KdV) systems, we demonstrate that the framework consistently recovers the correct dynamical operators underlying the observed data, including diffusion, nonlinear advection, and higher-order dispersive effects. While the recovered coefficients may differ in magnitude and sign from their physical counterparts, we show both theoretically and empirically that these discrepancies arise from affine transformations of the latent variables. Numerical verification confirms that the identified latent-space PDEs are dynamically equivalent to the original governing equations, as evidenced by low reconstruction errors and close agreement between aligned solutions.

Importantly, the results highlight that accurate recovery of governing dynamics does not require direct access to the physical state variables. Instead, the framework is capable of extracting physically meaningful and interpretable representations directly from indirect, high-dimensional measurements. This establishes a proof of concept that PDE discovery can be performed in realistic settings where observations are noisy, transformed, or only partially informative.

The proposed method also provides a natural bridge between model order reduction and equation discovery, as the learned latent dynamics serve as a compact and interpretable representation of the underlying system. In contrast to approaches that either prioritize reconstruction or latent ODE modeling, the present framework explicitly preserves spatial structure and enables identification of PDE operators in the latent space.

Future work will focus on expanding the scope of the framework to more complex physical domains. While the current study demonstrates the framework on 1D spatiotemporal systems, the architecture is natively extensible to higher-dimensional problems. The use of convolutional autoencoders allows for the straightforward substitution of 1D kernels with 2D or 3D variants to handle multi-dimensional spatial domains. Furthermore, the symbolic discovery layer is agnostic to the spatial dimensionality, as it operates on the latent representation z. Preliminary results on 2D systems suggest that the framework maintains its ability to identify governing PDEs from indirect observations. However, a comprehensive extension involving *multi-field coupling, cross-component interactions*, and the identification of *non-trivial spatial manifold topologies* in higher dimensions is currently the subject of a forthcoming manuscript. These directions are essential for assessing the scalability and robustness of the approach in complex industrial and experimental settings.

Overall, this work demonstrates that governing equations can be reliably recovered from indirect high-dimensional observations, providing a foundation for interpretable, physics-informed modeling in settings where traditional equation discovery methods are not directly applicable.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

IK: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. ET: Conceptualization, Methodology, Resources, Supervision, Writing – review & editing. ME: Conceptualization, Funding acquisition,

Methodology, Project administration, Supervision, Writing – review & editing.

Funding

The author(s) declared that financial support was received for this work and/or its publication. This research was part of project number 118C122 supported by the Industrial Ph.D. Program (2244) of the Scientific and Technological Research Council of Turkey (TÜBITAK).

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fams.2026.1807939/full#supplementary-material

References

1. Chen W. Physics-informed machine learning for reduced-order modeling of nonlinear problems. *J Comput Phys.* (2021) 446:110666–110666. doi: 10.1016/j.jcp.2021.110666

2. Zienkiewicz OC, Taylor RL. *The Finite Element Method*. New York: McGraw-hill (1977).

3. Fornberg B. *A Practical Guide to Pseudospectral Methods*. Cambridge: Cambridge University Press. (1998).

4. Loiseau JC, Brunton SL, Noack BR. *9 From the POD-Galerkin Method to Sparse Manifold Models*. Berlin, Boston: De Gruyter (2021). p. 279–320. doi: 10.1515/9783110499001-009

5. Mallick S, Mittal M. AI-based model order reduction techniques: a survey. *Arch Comput Methods Eng.* (2025) 32:2321–46. doi: 10.1007/s11831-024-10207-2

6. Xiao D. *Non-intrusive reduced order models and their applications*. Doctoral dissertation Diss Imperial College London. (2016).

7. Fresca S, Manzoni A. POD-DL-ROM enhancing deep learning-based reduced order models for nonlinear parametrized PDEs by proper orthogonal decomposition. *Comput Methods Appl Mech Eng.* (2022) 388:114181–114181. doi: 10.1016/j.cma.2021.114181

8. Pant P, Doshi R, Bahl P, Barati Farimani A. Deep learning for reduced order modelling and efficient temporal evolution of fluid simulations. *Phys Fluids.* (2021) 33:107101. doi: 10.1063/5.0062546

9. von Rueden L, Mayer S, Beckh K, Georgiev B, Giesselbach S, Heese R, et al. Informed machine learning—a taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Trans Knowl Data Eng.* (2023) 35:614–33. doi: 10.1109/TKDE.2021.3079836

10. Willard J, Jia X, Xu S, Steinbach M, Kumar V. Integrating physics-based modeling with machine learning: a survey. *arXiv* [preprint]. *arXiv:2003.04919* (2020).

11. Swischuk R. Projection-based model reduction: formulations for physics-based machine learning. *Comput Fluids.* (2019) 179:704–17. doi: 10.1016/j.compfluid.2018.07.021

12. Gin C, Lusch B, Brunton SL, Kutz JN. Deep learning models for global coordinate transformations that linearise PDES. *Eur J Appl Math.* (2021) 32:515–39. doi: 10.1017/S0956792520000327

13. Champion K, Lusch B, Kutz JN, Brunton SL. Data-driven discovery of coordinates and governing equations. *Proc Nat Acad Sci.* (2019) 116:22445–51. doi: 10.1073/pnas.1906995116

14. Brunton SL, Noack BR, Koumoutsakos P. Machine learning for fluid mechanics. *Annu Rev Fluid Mech.* (2020) 52:477–508. doi: 10.1146/annurev-fluid-010719-060214

15. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nat Rev Phys.* (2021) 3:422–40. doi: 10.1038/s42254-021-00314-5

16. Dong J, Zhong J. Recent advances in symbolic regression. *ACM Comput Surv.* (2025) 57:1–37. doi: 10.1145/3735634

17. Sun S, Tian S, Wang Y, Li B. The data-driven discovery of partial differential equations by symbolic genetic algorithm. *Nonlinear Dyn.* (2024) 112:19871–85. doi: 10.1007/s11071-024-10093-0

18. Brunton SL, Proctor JL, Kutz JN. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc Nat Acad Sci.* (2016) 113:3932–7. doi: 10.1073/pnas.1517384113

19. Rudy SH, Brunton SL, Proctor JL, Kutz JN. Data-driven discovery of partial differential equations. *Sci Adv.* (2017) 3:1602614–1602614. doi: 10.1126/sciadv.1602614

20. Kaptanoglu AA, de Silva BM, Fasel U, Kaheman K, Goldschmidt AJ, Callaham JL, et al. PySINDy: a comprehensive Python package for robust sparse system identification. *arXiv* [Preprint]. *arXiv:211108481*. (2021).

21. Fries WD, He X, Choi Y. Lasdi: parametric latent space dynamics identification. *Comput Methods Appl Mech Eng.* (2022) 399:115436. doi: 10.1016/j.cma.2022.115436

22. He X, Choi Y, Fries WD, Belof JL, Chen JS. gLaSDI: parametric physics-informed greedy latent space dynamics identification. *J Comput Phys.* (2023) 489:112267. doi: 10.1016/j.jcp.2023.112267

23. Bonneville C, Choi Y, Ghosh D, Belof JL. Gplasdi: Gaussian process-based interpretable latent space dynamics identification through deep autoencoder. *Comput Methods Appl Mech Eng.* (2024) 418:116535. doi: 10.1016/j.cma.2023.116535

24. Anderson W, Chung K, Choi Y. mLaSDI: multi-stage latent space dynamics identification. *arXiv* [preprint]. *arXiv:250609207*. (2025).

25. Stephany R, Anderson WM, Choi Y. Higher-order LaSDI: reduced order modeling with multiple time derivatives. *arXiv* [preprint]. *arXiv:251215997*. (2025).

26. Bakarji J. Discovering governing equations from partial measurements with deep delay autoencoders. *Proc R Soc A*. (2022) 479:20230422. doi: 10.1098/rspa.2023.0422

27. Krishnapriyan A. Characterizing possible failure modes in physics-informed neural networks. In: *Advances in Neural Information Processing Systems*. (2021). p. 34.

28. Wang R, Yu R. Physics-guided deep learning for dynamical systems: a survey. *ACM Comput Surv*. (2025) 58:1–31. doi: 10.1145/3766887

29. Meng C, Griesemer S, Cao D, Seo S, Liu Y. When physics meets machine learning: a survey of physics-informed machine learning. *Mach Learn Comput Sci Eng*. (2025) 1:20. doi: 10.1007/s44379-025-00016-0

30. Raissi M, Perdikaris P, Karniadakis GE. Physics- informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys*. (2019) 378:686–707. doi: 10.1016/j.jcp.2018.10.045

31. Darbon J, Meng T. On some neural network architectures that can represent viscosity solutions of certain high dimensional Hamilton-Jacobi partial differential equations. *J Comput Phys*. (2021) 425:109907–109907. doi: 10.1016/j.jcp.2020.109907

32. Luo K, Zhao J, Wang Y, Li J, Wen J, Liang J, et al. Physics-informed neural networks for PDE problems: a comprehensive review. *Artif Intell Rev*. (2025) 58:323. doi: 10.1007/s10462-025-11322-7

33. Lu L, Jin P, Pang G, Zhang Z, Karniadakis GE. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat Mach Intell*. (2021) 3:218–29. doi: 10.1038/s42256-021-00302-5

34. Wang S, Wang H, Perdikaris P. Learning the solution operator of parametric partial differential equations with physics-informed DeepONets. *Sci Adv*. (2021) 7:eabi8605. doi: 10.1126/sciadv.abi8605

35. Li Z. Fourier neural operator for parametric partial differential equations. *arXiv* [preprint]. *arXiv:2010.08895*. (2021).