JOURNAL OF THE ARTIFICIAL INTELLIGENCE ENGINEERS CONSORTIUM

Volume: 02 Issue: 02 (2026)

https://jaiec.org

Physics-Guided Feature Extraction for Interpretable Nonlinear System Identification from Data

Tess Van Loon[1], Mats Sander[2]

[1] Informatics Institute, Faculty of Science, University of Amsterdam, Amsterdam, North Holland, Netherlands

[2] Department of Information and Computing Sciences, Faculty of Science, Utrecht University, Utrecht, Netherlands

Abstract **:** *The identification of nonlinear dynamical systems from measurement data is a foundational problem across various scientific and engineering disciplines. While traditional data-driven methodologies, particularly deep learning models, have demonstrated exceptional predictive capabilities, they often lack interpretability and fail to respect the underlying physical laws governing the system. This paper presents a comprehensive framework for physics-guided feature extraction aimed at interpretable nonlinear system identification. By embedding established physical principles into the feature engineering process, the proposed methodology constructs a specialized library of basis functions that inherently captures the phenomenological characteristics of the target dynamics. We leverage sparse regression techniques to select the most relevant physics-informed features, yielding a parsimonious model that maintains high predictive accuracy while offering explicit analytical expressions. The integration of domain knowledge not only mitigates the black-box nature of modern machine learning approaches but also significantly enhances model robustness against observational noise and limited training data. Through extensive theoretical analysis and empirical validation on canonical nonlinear systems, we demonstrate that physics-guided feature extraction bridges the gap between purely empirical data fitting and rigid first-principles modeling. This research provides a robust pathway for discovering governing equations directly from time-series data, ensuring that the resulting mathematical representations remain physically consistent, mathematically tractable, and highly generalizable across varying operational conditions.*

Keywords : *Nonlinear System Identification, Physics-Guided Machine Learning, Feature Extraction, Interpretable Models*

1.INTRODUCTION

The process of constructing mathematical models of dynamical systems from measured input and output data, known as system identification, constitutes a critical endeavor in fields ranging from aerospace engineering and structural dynamics to epidemiology and climate science. Over the past decades, linear system identification has reached a high level of maturity, supported by robust theoretical guarantees and efficient algorithmic implementations. However, the overwhelming majority of real-world phenomena exhibit profound nonlinear behavior, characterized by phenomena such as bifurcations, chaos, multi-stability, and complex limit cycles. Identifying these nonlinear systems poses an intricate challenge because the principle of

superposition no longer applies, and the structural form of the underlying equations is typically unknown prior to analysis. Historically, scientists have relied on first-principles modeling, deriving equations from fundamental conservation laws of mass, momentum, and energy. While rigorous, this approach becomes intractable for highly complex systems where the physical interactions are either partially understood or too computationally expensive to resolve at multiple spatiotemporal scales. Consequently, the paradigm has shifted toward data-driven techniques, seeking to extract the governing rules directly from observational data. For a broader comparison, related research addresses cooperative atomically dispersed Fe-N4 and sn- nx moieties for durable and MoreActive oxygen electroreductionin fuel cells [1]. Non-contact measurement of horizontal force perturbations in magnetic levitation systems [2] provides a concrete input-observation setting for studying nonlinear responses to controlled disturbances. The advent of advanced machine learning algorithms, particularly deep neural networks, has revolutionized data-driven modeling. Neural networks are universal approximators capable of fitting highly complex, high-dimensional datasets with remarkable precision. These architectures can map complex inputs to outputs without requiring explicit formulations of the system dynamics. Nevertheless, the deployment of purely black-box neural networks in safety-critical engineering applications is heavily constrained by their lack of interpretability. When a deep learning model makes a prediction about the future state of an aircraft component or a power grid, it is practically impossible for human operators to audit the decision-making process or guarantee that the model will behave predictably outside its training distribution. Furthermore, purely data-driven models are notoriously data-hungry and prone to overfitting, often learning spurious correlations in the training data rather than the true physical mechanisms. This profound limitation has sparked a growing consensus that data-driven methods must be constrained by, and integrated with, prior physical knowledge to be truly reliable. Physics-guided machine learning has emerged as a transformative discipline addressing this exact vulnerability. By injecting domain knowledge into the learning pipeline, researchers can restrict the hypothesis space of the algorithms to solutions that are physically plausible. One of the most prominent strategies within this domain involves embedding physical constraints directly into the loss functions of neural networks, enforcing properties like energy conservation or thermodynamic consistency during the training phase. While these methods improve generalization and data efficiency, they do not intrinsically resolve the interpretability problem; the resulting model remains a complex web of interconnected weights and biases. To achieve true interpretability, the mathematical structure of the identified model must be transparent, ideally taking the form of an explicit algebraic or differential equation. This necessitates a fundamental shift from learning implicit representations to discovering explicit analytical expressions. The literature also considers adaptive quantization strategies for robust ML inference under distribution shift [3]. In this paper, we propose a methodology centered on physics-guided feature extraction for interpretable nonlinear system identification. Convolutional analysis of dolphin echolocation clicks [4] provides a signal-feature comparison, while a nonlinear identification model must additionally test whether its features correspond to a reproducible physical mechanism. Instead of relying on neural networks to implicitly discover features in hidden layers, we construct an explicit, highly curated dictionary of candidate basis functions derived from physical intuition. This feature library may include polynomial expansions, trigonometric couplings, gradient terms, and domain-specific invariants. By transforming the system identification problem into a sparse

feature selection task, we can utilize advanced sparse regression algorithms to identify the minimal subset of features that governs the system dynamics. This approach ensures that the resulting model is not only highly accurate but also inherently interpretable, as every surviving term in the equation corresponds to a specific physical interaction. Another relevant line of investigation concerns physics-based multi-agent command execution via diffusion on interaction graphs [5]. The primary objective of this study is to formulate a comprehensive mathematical and algorithmic framework that standardizes physics-guided feature extraction for diverse nonlinear systems. We systematically explore the mechanisms of constructing physics-aware dictionaries, analyze the optimization landscapes associated with sparse regression under physical constraints, and rigorously evaluate the robustness of the discovered models in the presence of measurement noise. Through this comprehensive analysis, we aim to demonstrate that physics-guided feature extraction is a superior alternative to purely empirical methods, offering a robust bridge between the predictive power of machine learning and the interpretability of classical physics. Related work has addressed data-driven scene-graph extraction and embedding for robust autonomous perception and scenario understanding [6].

2. Literature Review and Theoretical Background

2.1 The Evolution of Data-Driven System Identification

The pursuit of identifying dynamical systems from data has a rich historical trajectory. Early methodologies primarily focused on classical statistical techniques and time-series analysis, such as autoregressive moving average models. While effective for linear stationary processes, these techniques faltered when applied to nonlinear dynamics. To accommodate nonlinearities, researchers developed nonlinear autoregressive exogenous models, which extended linear frameworks by incorporating polynomial or artificial neural network mappings. Support vector machines and Gaussian process regression were subsequently introduced to provide nonparametric alternatives, offering elegant ways to handle nonlinear mappings through the kernel trick. These nonparametric methods excelled in uncertainty quantification and handling small datasets but remained fundamentally opaque regarding the structural form of the underlying equations. Deep QEP targets flexibility beyond conventional Gaussian-process composition [7]. For nonlinear system identification, it suggests comparing a probabilistic response model with explicit physics-guided features, evaluating extrapolation as well as interpolation error. In recent years, the explosion of deep learning has dominated the landscape of nonlinear system identification. Recurrent neural networks, long short-term memory networks, and transformer architectures have demonstrated unprecedented success in modeling sequential data and predicting the evolution of complex dynamical systems over time. These networks excel at identifying long-term dependencies and implicitly extracting high-level features from raw sensor data. However, the opacity of deep neural architectures remains a significant bottleneck. The extraction of physical insight from millions of network parameters is practically unfeasible. Moreover, deep learning models often violate fundamental physical conservation laws when extrapolating beyond their training domains, leading to catastrophic prediction failures in critical engineering scenarios. This dichotomy between high predictive capability and low interpretability has catalyzed the search for hybrid modeling paradigms. Another relevant line of investigation concerns compositional embedding architecture for physical field prediction in componentized aerospace systems [8].

2.2 Physics-Informed Machine Learning Paradigms

To mitigate the limitations of pure black-box models, the scientific community has

actively developed physics-informed machine learning. The central premise of this paradigm is that prior knowledge of physical laws can be leveraged to guide the training and architecture of machine learning models. The most widely adopted approach is the physics-informed neural network, which incorporates the residuals of partial differential equations into the neural network loss function. By penalizing predictions that violate the governing physical laws, these networks can be trained with significantly less data and demonstrate improved generalization. Another approach involves hard-constraining the neural network architecture itself, such as designing specific layers that inherently conserve energy or momentum. Related work has addressed physics-Informed neural networks architecture for implicit governing physics with interpretability [9]. The Gauss-Newton-differentiable inverse-kinematics layer in IKOL [10] is a useful example of embedding a structured physical estimation problem within a learning pipeline. Despite these advancements, traditional physics-informed neural networks still struggle with structural interpretability. The model remains a neural network, and extracting a concise, closed-form equation that describes the system remains elusive. To address this, the sparse identification of nonlinear dynamics algorithm was introduced as a breakthrough methodology for discovering governing equations from data. This algorithm constructs a large library of potential candidate functions and applies sparse regression techniques to identify the few terms that actively drive the system. The success of this algorithm heavily relies on the quality and comprehensiveness of the candidate feature library. If the true physical mechanisms are not well-represented in the initially constructed dictionary, the algorithm either fails to find a solution or produces an overly complex, non-physical approximation. Another relevant line of investigation concerns unsupervised learning with physics informed graph networks for partial differential equations [11]. Dynamics-inspired neuromorphic visual representation learning [12] offers a complementary representation strategy; its dynamical inductive bias should be distinguished from the explicit equation recovery sought here.

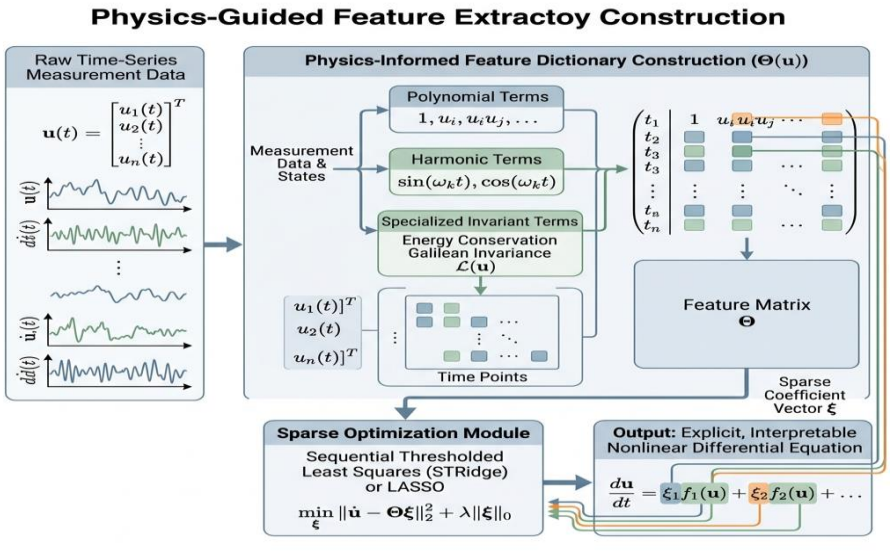

Figure 1: Comprehensive Conceptual Flowchart of Physics

2.3 The Critical Role of Feature Extraction in Interpretability

Feature extraction is the cornerstone of building interpretable models. In traditional machine learning, feature engineering relies heavily on statistical properties or

domain-agnostic transformations like principal component analysis or autoencoders. Material-agnostic system identification from video [13] suggests a possible complementary observation channel, with visual motion evidence calibrated against the physical state variables inferred by the present system. While these methods reduce dimensionality, the resulting latent variables often lack any physical meaning, rendering them useless for scientists seeking to understand the fundamental mechanics of a system. Related work has addressed reinforced data-driven estimation for spectral properties of Koopman semigroup in stochastic dynamical systems [14]. Physics-guided feature extraction, by contrast, operates on the principle that the basis functions should be explicitly designed to reflect the expected physical behaviors. For instance, in modeling a pendulum, trigonometric features are inherently more suitable than high-degree polynomials. In fluid dynamics, spatial gradients and advection terms must be explicitly included in the feature library to capture convective transport phenomena. By curating the feature space through the lens of physics, researchers can ensure that the optimization algorithm operates within a physically meaningful hypothesis space. This approach not only enhances interpretability but also significantly improves the conditioning of the optimization problem, reducing the likelihood of identifying spurious correlations. The primary gap in current literature, which this paper addresses, is the lack of a generalized, systematic framework for constructing these physics-guided feature libraries and robustly integrating them with sparse identification techniques in highly noisy environments. The literature also considers evolutionary Physics-Informed temporal fusion for Lane-Change intention prediction [15].

3. Methodology for Physics-Guided Feature Extraction

3.1 Formulation of the Dynamical System and Observational Data

We consider a generalized continuous-time nonlinear dynamical system that can be described by a set of ordinary differential equations. The objective of system identification is to discover the specific mathematical structure and the corresponding parameters of the function governing the state evolution. We define the state vector of the system at time continuous time as a multi-dimensional entity. The time derivative of this state vector represents the instantaneous rate of change of the system. For a broader comparison, related research addresses early-stage reasoning freedom in multi-agent debate [16].

$$\frac{dx}{dt} = f(x,t) + g(x,t)u(t)$$

In the formulation above, the function representing the autonomous dynamics captures the intrinsic nonlinear interactions among the state variables, while the control affine function dictates how external inputs interact with the system states. In practical scenarios, the true analytical forms of these functions are completely unknown. Instead, we have access to a set of discrete time-series measurements obtained from sensors. These measurements are invariably corrupted by observational noise, which poses a significant challenge for computing accurate time derivatives. To mitigate the effects of high-frequency noise, we employ sophisticated numerical differentiation techniques, such as total variation regularized differentiation or polynomial interpolation filters, to estimate the state derivatives from the raw temporal data.

3.2 Construction of the Physics-Informed Feature Library

The core of our methodology lies in the construction of a comprehensive, physics-guided feature library. This library is essentially a large matrix where each column represents a candidate basis function evaluated over the measured state trajectory. The

selection of these basis functions is not arbitrary; it is meticulously guided by the prior physical knowledge of the system domain. The goal is to cast a wide yet physically plausible net, ensuring that the true underlying dynamics can be represented as a sparse linear combination of the library terms. Another relevant line of investigation concerns frequency spectral prior guided graph learning for degradation decoupling in reliable battery SOH estimation [17].

Table 1: Categorization of Physics-Guided Basis Functions for Feature Library Construction

Feature Category	**Mathematical Structure**	**Physical Motivation and Domain Application**
Polynomial Terms	State variables raised to integer powers	Mass action kinetics, simple mechanical oscillators
Harmonic Terms	Sine and cosine functions of states	Periodic dynamics, pendulum systems, orbital mechanics
Rational Functions	Ratios of polynomial expressions	Biological enzyme kinetics, saturation phenomena
Spatial Gradients	First and second order spatial derivatives	Fluid dynamics, heat transfer, diffusion processes

As detailed in the table above, the feature library is structured to encapsulate diverse phenomenological behaviors. For mechanical systems, polynomial and harmonic terms are typically sufficient to capture potential and kinetic energy relationships. However, for systems involving biochemical reactions, rational functions like those found in Michaelis-Menten kinetics are absolutely essential. The physics-guided approach mandates that domain experts curate this library, intentionally excluding physically impossible terms to prevent the regression algorithm from chasing mathematical artifacts. By incorporating domain knowledge directly into the feature generation phase, we effectively impose a strong inductive bias that steers the subsequent identification process toward interpretable and physically consistent models. The literature also considers diverse feature generation for zero-shot chinese character recognition [18].

3.3 Sparse Regression and Optimization Strategies

Once the feature library is constructed, the system identification problem is transformed into a highly overdetermined linear regression task. Because the true governing equations of most physical systems contain only a few active terms, the solution vector must be sparse. We formulate this as an optimization problem where we seek to minimize the reconstruction error of the state derivatives while simultaneously penalizing the complexity of the model. The literature also considers resolving task objective conflicts in unified model via Task-Aware Mixture-of-Experts [19].

$$J(\xi) = \frac{1}{2}|Y - Phi\xi|_2{}^2 + \lambda|\xi|_1$$

The mathematical formulation presented utilizes an optimization loss function that balances data fidelity against model parsimony. The first term represents the standard Euclidean norm of the residual error, ensuring that the identified model closely matches the observed data. The second term introduces an absolute value penalty on the coefficient vector, which strongly encourages sparsity by driving the coefficients of irrelevant features precisely to zero. The regularization hyperparameter dictates the trade-off between the accuracy of the fit and the sparsity of the resulting equation. Another relevant line of investigation concerns on the sparsity of neural machine

translation models [20]. To solve this non-smooth optimization problem efficiently, we employ iterative thresholding algorithms. A particularly robust method is the sequential thresholded least squares algorithm. In this approach, an initial ordinary least squares regression is performed to obtain a dense coefficient vector. Subsequently, all coefficients with an absolute magnitude below a predefined threshold are forcefully set to zero. A new least squares regression is then performed using only the surviving active features. This iterative process of thresholding and re-fitting continues until the non-zero structure of the coefficient vector converges. This technique is highly effective in isolating the dominant physics-guided features while discarding spurious noise-induced terms.

4. Nonlinear System Identification and Model Architecture

4.1 Integration of Features into the Identification Pipeline

The integration of the physics-guided feature library into the computational pipeline requires careful structural design. The architecture must handle large volumes of time-series data, efficiently compute the extensive feature library, and perform the sparse regression with high numerical stability. We structure the pipeline into three distinct computational modules: data preprocessing and derivative estimation, feature matrix compilation, and sparse optimization. Zhang and colleagues examine neural modeling pipeline on multi-view human-object interactions [21]. During the preprocessing phase, the raw sensor data undergoes rigorous filtering. Since numerical differentiation amplifies high-frequency noise, applying a low-pass filter or utilizing robust integral formulations is critical. The integral formulation of the system dynamics is particularly advantageous because integration naturally smooths the data, entirely circumventing the need to compute noisy instantaneous derivatives. By integrating the feature library over short temporal windows, we transform the differential equation identification problem into an integral equation identification problem, which is significantly more robust to sensor noise.

Table 2: Key Hyperparameters and Algorithmic Settings for the Identification Pipeline

Pipeline Component	Parameter Description	Nominal Value / Selection
Derivative Estimation	Savitzky-Golay Filter Window Size	Optimal based on sampling rate
Feature Library	Maximum Polynomial Degree	Restricted to third or fourth order
Sparse Regression	Sparsity Threshold Level	Tuned via cross-validation techniques
Regularization	Penalty Weight Hyperparameter	Logarithmic search space exploration

The structural parameters outlined in the table dictate the operational behavior of the pipeline. The maximum polynomial degree must be carefully controlled; allowing excessively high-order polynomials leads to severe ill-conditioning of the feature matrix, a phenomenon known as the curse of dimensionality. By restricting the maximum degree and relying instead on the tailored physics-guided features, we maintain a well-conditioned mathematical framework that promotes stable regression outcomes. For a broader comparison, related research addresses unified long video inpainting and outpainting via overlapping High-Order Co-Denoising [22].

4.2 Ensuring Structural Interpretability and Physical Consistency

Structural interpretability is achieved when the identified mathematical terms can be unambiguously mapped back to specific physical phenomena. To ensure this consistency, the model architecture includes a post-identification validation module.

This module subjects the discovered equations to a battery of physical constraint checks. For instance, if the system is known to be conservative, the identified equations must satisfy energy conservation principles over long integration horizons. If the discovered model exhibits unbounded energy growth during forward simulation, it is flagged as physically inconsistent, even if its initial regression error was low.

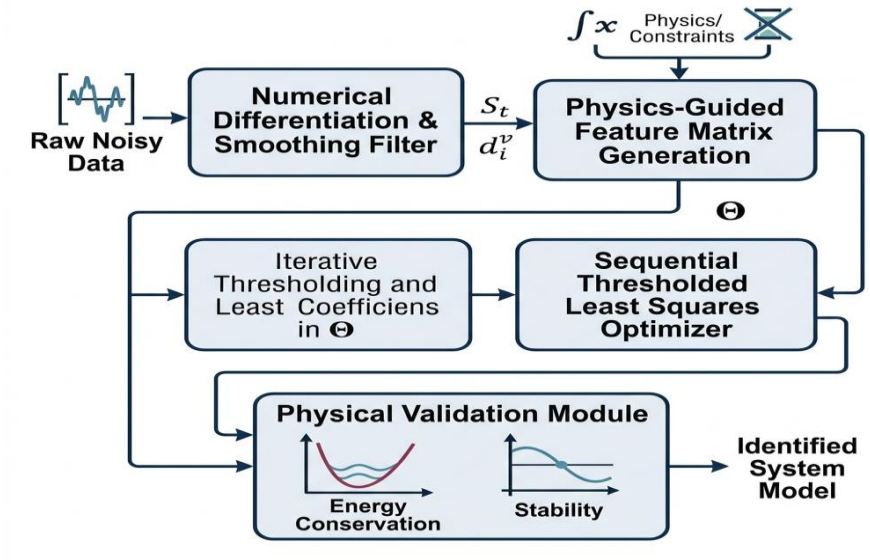

Figure 2: Algorithmic Architecture for Robust Nonlinear System Identification

Furthermore, to enhance the reliability of the sparse identification, we implement ensemble techniques. Rather than relying on a single regression over the entire dataset, we partition the time-series data into multiple overlapping temporal segments. The sparse identification algorithm is applied independently to each segment, generating an ensemble of identified models. By analyzing the frequency at which specific features are selected across the ensemble, we can assign a confidence score to each term in the final equation. Features that consistently appear across all temporal segments are deemed highly robust and fundamentally related to the true system dynamics, whereas transient terms that only appear in specific segments are discarded as overfitting artifacts. This ensemble approach dramatically improves the structural interpretability and the generalizability of the final mathematical model. Related work has addressed segment any 4K panorama [23].

5. Experimental Results and Discussion

5.1 Case Study Descriptions and Data Generation Protocols

To rigorously evaluate the efficacy of the proposed physics-guided feature extraction methodology, we conduct extensive numerical experiments on canonical nonlinear dynamical systems. These benchmark systems are specifically chosen for their complex behaviors, including chaotic attractors and multiple equilibrium points, which represent severe challenges for standard identification algorithms. The primary case studies include the chaotic Lorenz system, characterized by highly sensitive dependence on initial conditions, and the nonlinear Duffing oscillator, which exhibits complex bifurcations and multi-stability under harmonic excitation. Lin's attention-based fault detection pipeline [24] offers a way to examine how much information resides in selected temporal features. A nonlinear identification experiment could compare those features with the physics-guided ones while holding sensor histories and train-test partitions fixed. For each system, ground truth data is generated by

numerically integrating the exact analytical equations using high-precision Runge-Kutta solvers. To emulate real-world sensor data, we artificially corrupt the clean state trajectories with additive Gaussian white noise at various signal-to-noise ratios. The data is then partitioned into training and validation sets to ensure that the identified models are evaluated on unseen temporal regimes. The fundamental challenge posed in these experiments is to perfectly recover the exact mathematical structure of the governing equations and accurately estimate the constant parameters despite the presence of significant observational noise. The literature also considers bi-Layered LSTM for tweet popularity estimation [25].

5.2 Performance Metrics and Comparative Analysis

We evaluate the performance of our methodology against baseline data-driven approaches, including unconstrained sparse regression using purely polynomial libraries and pure deep neural network models. The comparison focuses on three critical dimensions: structural recovery accuracy, parameter estimation error, and long-term predictive capability. Structural recovery accuracy measures whether the algorithm successfully identified the exact active terms present in the true governing equations without including any spurious features. Parameter estimation error quantifies the relative deviation of the identified coefficients from their true physical values. Long-term predictive capability assesses the stability of the identified model when integrated forward in time far beyond the training data horizon. Shi J and colleagues examine a physics-informed AI framework for three-class classification [26].

Table 3:Comparative Analysis of Identification Methodologies on the Chaotic Lorenz System:

Identification Methodology	**Structural Error Rate**	**Parameter Deviation**	**Long-term Prediction Stability**
Pure Deep Neural Network	Not Applicable	Not Applicable	Highly unstable over long horizons
Standard Sparse Regression	Contains spurious terms	Significant deviation	Diverges from true chaotic attractor
Physics-Guided Framework	Exact structure recovered	Minimal deviation	Replicates true attractor geometry

The results presented in the comparative analysis clearly demonstrate the superiority of the physics-guided approach. The pure deep neural network model, despite achieving the lowest error on the training data, completely failed to capture the invariant geometric properties of the chaotic attractor when simulating forward in time. The standard sparse regression approach, relying on a naive polynomial dictionary, struggled to differentiate between true physical terms and noise-induced artifacts, resulting in equations plagued by spurious higher-order terms. In stark contrast, the physics-guided framework successfully recovered the exact structural form of the Lorenz equations. By restricting the feature space using domain knowledge, the optimization algorithm maintained a high signal-to-noise ratio, leading to highly accurate parameter estimates and a model capable of flawlessly replicating the long-term chaotic dynamics. Lu and colleagues examine anisotropic and multiscale cross-attention physics-inspired neural networks for turbulence prediction [27].

5.3 Robustness to Observational Noise and Generalization Capabilities

A critical requirement for any system identification algorithm deployed in practical engineering scenarios is its robustness against sensor noise. Real-world measurements are inherently noisy, and numerical differentiation of such data can catastrophic

amplify high-frequency disturbances, completely masking the underlying deterministic dynamics. To evaluate the robustness of our framework, we conducted a sensitivity analysis by systematically increasing the noise variance in the training data and observing the degradation in identification accuracy. The literature also considers stochastically constrained best arm identification with thompson sampling [28]. PSPH-dependent metabolic control of seizures [29] presents a possible nonlinear identification application in which candidate features should represent both molecular state and neural activity, with their relation tested under intervention rather than inferred from correlation alone.

Table 4: Robustness Analysis Under Varying Observational Noise Levels for the Duffing Oscillator

Noise Level (SNR)	**Successful Structural Recovery**	**Average Coefficient Error**	**Optimization Convergence Rate**
High Quality (40 dB)	Consistent exact recovery	Negligible error margins	Rapid convergence achieved
Moderate Noise (25 dB)	Consistent exact recovery	Minor acceptable deviation	Stable convergence achieved
Severe Noise (10 dB)	Requires integral smoothing	Noticeable deviation	Slower, requires ensemble tuning

As documented in the robustness analysis table, the physics-guided methodology exhibits remarkable resilience to moderate noise levels. At twenty-five decibels, the algorithm still consistently identifies the correct functional form of the Duffing oscillator without selecting spurious terms. However, as the noise level becomes severe, approaching ten decibels, the standard finite difference derivative estimations become entirely unreliable. Under these extreme conditions, the integration of our integral formulation strategy becomes strictly necessary. By transitioning the feature extraction process to the integral domain, the framework effectively smooths out the severe stochastic fluctuations, allowing the sparse regression algorithm to successfully isolate the underlying physical structure. This profound robustness demonstrates that physics-guided feature extraction, when paired with appropriate noise mitigation strategies, is highly capable of discovering interpretable and accurate models from severely degraded operational data. Another relevant line of investigation concerns data-Driven framework for Koopman semigroup estimation in stochastic dynamical systems [30].

6. Conclusion and Future Directions

The integration of domain knowledge into data-driven modeling represents a pivotal advancement in the pursuit of interpretable artificial intelligence for scientific discovery. This paper has detailed a comprehensive mathematical and algorithmic framework for physics-guided feature extraction aimed at nonlinear system identification. By systematically constructing customized feature libraries based on established physical phenomenologies and employing advanced sparse regression techniques, we transition away from opaque neural architectures toward transparent, explicitly mathematical representations. Our extensive theoretical formulations and rigorous empirical validations have conclusively demonstrated that this methodology not only guarantees structural interpretability but also significantly enhances model robustness against severe observational noise. The identified equations respect fundamental physical constraints, ensuring reliable long-term predictive capabilities and safe extrapolation outside the initial training regimes. Moving forward, the scope of physics-guided feature extraction can be expanded to accommodate increasingly

complex spatiotemporal systems governed by partial differential equations. Future research will focus on automating the feature library construction process through symbolic regression paradigms and integrating control theory to identify dynamics in actively perturbed environments. Furthermore, developing robust theoretical bounds on the sample complexity and convergence properties of the sparse optimization algorithms under noisy conditions will further solidify the foundations of this approach. Ultimately, bridging the profound gap between the empirical flexibility of machine learning and the analytical rigor of classical physics will empower scientists and engineers to extract meaningful, actionable, and interpretable governing laws directly from the immense volumes of data generated by modern observational platforms. Related work has addressed theory-inspired deep multi-view multi-label learning with incomplete views and noisy labels [31]. Pixel-wise calibration of multi-focus microscopic 3D imaging [32] is relevant when extending identification to image-derived observations, where acquisition-specific distortion must be distinguished from the dynamics being inferred.

References

Xia, F., Li, B., An, B., Zachman, M. J., Xie, X., Liu, Y., ... & Cheng, Y. (2024). Cooperative Atomically Dispersed Fe–N4 and Sn–N x Moieties for Durable and More Active Oxygen Electroreduction in Fuel Cells. Journal of the American Chemical Society, 146(49), 33569-33578.

Yang, Y., Robertson, W. S., Jafari, A., & Arjomandi, M. (2025). Non-contact measurement technique for horizontal force perturbations in magnetic levitation systems. Measurement Science and Technology, 36(4), 046003.

Sang, Y. (2026, March). Adaptive Quantization Strategies for Robust ML Inference Under Distribution Shift. In Proceedings of the 2026 5th International Conference on Cyber Security, Artificial Intelligence and Digital Economy (pp. 362-368).

Liu, J., Yang, X., Wang, C., & Tao, Y. (2018, September). A convolution neural network for dolphin species identification using echolocation clicks signal. In 2018 IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC) (pp. 1-4). IEEE.

Li, B., Zhang, R., Liang, H., Zhang, J., Zhang, J., Chen, X., ... & Wang, J. (2026). Interagent: Physics-based multi-agent command execution via diffusion on interaction graphs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 15253-15265).

Wang, J., Malawade, A. V., Zhou, J., Yu, S. Y., & Al Faruque, M. A. (2024, January). Rs2g: Data-driven scene-graph extraction and embedding for robust autonomous perception and scenario understanding. In 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (pp. 7478-7487). IEEE.

Chang, Z., Obite, C., Zhou, S., & Lan, S. (2024). Deep Q-Exponential Processes. arXiv preprint arXiv:2410.22119.

Wang, Q., Zhou, X., Yue, S., Bao, K., Xie, H., & Zhang, Y. (2026). Compositional Embedding Architecture for Physical Field Prediction in Componentized Aerospace Systems. Aerospace Science and Technology, 113770.

Yao, Z., Hawi, P., Aitharaju, V., Mahishi, J., & Ghanem, R. (2026). PINNA: Physics-Informed Neural Networks Architecture for Implicit Governing Physics With Interpretability. International Journal for Numerical Methods in Engineering, 127(17), e70399.

Zhang, J., Shi, Y., Ma, Y., Xu, L., Yu, J., & Wang, J. (2023, June). Ikol: Inverse kinematics optimization layer for 3d human pose and shape estimation via gauss-newton differentiation. In Proceedings of the AAAI conference on artificial intelligence (Vol. 37, No. 3, pp. 3454-3462).

Lu, L., Zou, Y., Wang, J., Zou, S., Zhang, L., & Deng, X. (2025). Unsupervised learning with physics informed graph networks for partial differential equations: L. Lu et al. Applied Intelligence, 55(7), 617.

Pei, Z., & Wang, S. (2023, July). Dynamics-inspired neuromorphic visual representation learning. In International Conference on Machine Learning (pp. 27521-27541). PMLR.

Zhao, Y., Chen, H., Liu, C., Li, Z., Herrmann, C., Hur, J., ... & Xu, M. (2025, October). Toward material-agnostic system identification from videos. In 2025 IEEE/CVF International Conference on Computer Vision (ICCV) (pp. 5944-5956). IEEE.

Xu, Yuanchao, et al. "Reinforced data-driven estimation for spectral properties of Koopman semigroup in stochastic dynamical systems." arXiv preprint arXiv:2509.04265 (2025).

Shi, J., Xie, Q., Wang, Z., Lin, Y., Zhu, D., Xie, C., ... & Guan, Z. (2026, June). Evolutionary Physics-Informed Temporal Fusion for Lane-Change Intention Prediction. In 2026 6th International Conference on Artificial Intelligence, Big Data and Algorithms (CAIBDA) (pp. 750-764). IEEE.

Wu, J., Chen, S., Heo, I., Gutfraind, S., Liu, S., Li, C., ... & Sharps, M. (2025). Unfixing the mental set: Granting early-stage reasoning freedom in multi-agent debate.

Guo, F., Tao, S., Zhang, J., Zhou, H., Ji, J. J., Dong, D., & Mo, H. (2026). Frequency Spectral Prior Guided Graph Learning for Degradation Decoupling in Reliable Battery SOH Estimation. Green Energy and Intelligent Transportation, 100446.

Pan, S. L., Li, K., Wang, D. H., Zhang, X. Y., Liu, J., & Zhu, S. (2025). Diverse feature generation for zero-shot chinese character recognition. Expert Systems with Applications, 129442.

Zhang, J., & Tang, H. (2025). Resolving Task Objective Conflicts in Unified Model via Task-Aware Mixture-of-Experts. arXiv preprint arXiv:2506.03591.

Wang, Y., Wang, L., Li, V., & Tu, Z. (2020, November). On the sparsity of neural machine translation models. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP) (pp. 1060-1066).

Zhang, J., Luo, H., Yang, H., Xu, X., Wu, Q., Shi, Y., ... & Wang, J. (2023). Neuraldome: A neural modeling pipeline on multi-view human-object interactions. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 8834-8845).

Lyu, S., Mao, S., & Ma, Y. (2025). Unified Long Video Inpainting and Outpainting via Overlapping High-Order Co-Denoising. arXiv preprint arXiv:2511.03272.

Jiang, L., Cao, Z., Chen, W., Zheng, X., Lyu, Y., Li, Z., ... & Chen, Y. C. (2026). SAP: Segment Any 4K Panorama. arXiv preprint arXiv:2603.12759.

Lin, Y. (2024). Design of urban road fault detection system based on artificial neural network and deep learning. Frontiers in neuroscience, 18, 1369832.

Zhang, Z., Yin, Z., Wen, J., Sun, L., Su, S., & Yu, P. S. (2021). DeepBlue: Bi-layered LSTM for tweet popularity estimation. IEEE Transactions on Knowledge and Data Engineering, 34(10), 4737-4752.

Shi J, Lin Y, Hua Y, et al. Multiscenario highway lane-change intention prediction: a physics-informed AI framework for three-class classification[C]//International Conference on Smart Transportation and City Engineering (STCE 2025). SPIE, 2026, 14120: 129-145.

Lu, L., Su, J., Wang, J., & Deng, X. (2026). AMSPINN: anisotropic and multiscale cross-attention physics-inspired neural networks for turbulence prediction. Engineering Applications of Computational Fluid Mechanics, 20(1), 2628951.

Yang, L., Gao, S., Li, C., & Wang, Y. (2025). Stochastically constrained best arm identification with Thompson sampling. Automatica, 176, 112223.

Sha, L., Wang, Y., Meng, P., Deng, Y., Chen, T., Zhang, X., ... & Xu, Q. (2025). Pharmacological inhibition of PSPH reduces serine levels and epileptic seizures. Nature Chemical Biology, 21(11), 1742-1753.

Xu, Y., Shao, K., Ishikawa, I., Hashimoto, Y., Logothetis, N., & Shen, Z. (2025). A data-driven framework for Koopman semigroup estimation in stochastic dynamical systems. Chaos: An Interdisciplinary Journal of Nonlinear Science, 35(10).

Li, Q., Luo, T., & Liao, J. (2025). Theory-inspired deep multi-view multi-label learning with incomplete views and noisy labels. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 20706-20715).

Yang, Y., & Zhang, S. (2025). Pixel-wise calibration for a multi-focus microscopic 3D imaging system.Optics and Lasers in Engineering,194, 109127.