Algorithms and Intelligence

Vol 2 Issue 1 (2026)

www.algointel.org

Physics-informed architecture search with interpretable constraints for sparsely observed nonlinear governing systems identification

Daniel Ferreira da Barros

Department of Communications, School of Electrical and Computer Engineering, Universidade Estadual de Campinas (UNICAMP), Campinas, São Paulo, Brazil

Abstract

The identification of nonlinear governing equations from sparsely observed data remains a fundamental challenge in computational physics and applied mathematics. Traditional data-driven approaches often struggle with sparse datasets, leading to overfitting and models that violate underlying physical laws. While physics-informed neural networks mitigate some of these issues by embedding physical symmetries into the optimization process, their reliance on manually designed architectures limits their adaptability across diverse dynamical systems. This paper introduces a novel framework that integrates neural architecture search with physics-informed learning, augmented by interpretable constraints. By treating the architecture itself as a differentiable variable, the proposed method automates the discovery of optimal network topologies tailored to specific governing systems. Furthermore, we introduce interpretable constraints that enforce sparsity and structural simplicity within the latent representations, enabling the extraction of symbolic mathematical terms directly from the optimized architecture. Extensive analysis demonstrates that this approach significantly enhances predictive accuracy and equation recovery in regimes characterized by high noise and low data volume. The integration of structural priors ensures that the discovered models remain physically consistent and scientifically interpretable, bridging the gap between automated machine learning and classical analytical physics.

Keywords

Neural Architecture Search, System Identification, Interpretable Machine Learning, Sparse Observations.

1. INTRODUCTION

1.1 The Challenge of Nonlinear System Identification

The mathematical modeling of complex physical phenomena, ranging from fluid dynamics to quantum mechanics, relies heavily on the formulation of nonlinear governing systems. Historically, these governing equations have been derived from empirical observations combined with fundamental conservation laws, requiring decades of rigorous human intuition and experimental validation. However, as modern scientific inquiry expands into highly complex domains such as climate modeling, advanced materials science, and epidemiological forecasting, the sheer dimensionality and non-linearity of the underlying processes often exceed the capacity for manual analytical derivation. Consequently, the automated identification of governing systems directly from observational data has emerged as a critical frontier in computational science. The primary objective is to reverse-engineer the hidden differential equations that dictate the spatial and temporal evolution of a given system. While classical system identification techniques have proven effective for linear systems with abundant data, they frequently fail when applied to nonlinear systems, particularly when the available observational data is sparse, irregularly sampled, or corrupted by measurement noise. Sparse observation regimes are not merely theoretical anomalies; they are the standard reality in many practical applications where deploying high-resolution sensors is prohibitively expensive or physically impossible, such as in deep ocean current monitoring or in vivo cellular tracking. In these challenging environments, purely data-driven machine learning models, including standard deep learning architectures, tend to memorize the sparse training points without capturing the true underlying dynamics, leading to catastrophic generalization failures outside the immediate training domain. Furthermore, these black-box models offer no mechanistic insight, rendering them unsuitable for scientific applications where understanding the physical laws is just as important as predicting future states. The Deep Q-Exponential Processes study relaxes Gaussian-process assumptions through Lq structure [1]. Such flexibility is relevant when identification data contain abrupt changes, although recovered physical laws must still be tested for invariance across operating regimes.

1.2 Emergence of Physics-Informed Approaches

To address the critical shortcomings of purely data-driven models, researchers have developed paradigms that integrate prior physical knowledge into the machine learning pipeline. This synthesis has led to the proliferation of physics-informed neural networks, which have fundamentally transformed the landscape of computational physics. In this approach, the deep learning model is penalized not only for deviating from the observational data but also for violating known physical constraints, such as the conservation of mass, momentum, and energy. By embedding the residual of the hypothesized partial differential equations directly into the loss landscape, these networks restrict the space of admissible solutions to those that are physically plausible. This regularization effect is particularly potent in sparsely observed regimes, as the physical constraints act as a synthetic data generator, guiding the optimization process across unobserved regions of the spatiotemporal domain. However, despite their empirical success, the efficacy of physics-informed models is heavily dependent on the specific architecture of the neural network. The choice of network depth, width, activation functions, and connectivity patterns dictates the model capacity to approximate complex functional mappings. Currently, the design of these architectures is largely a manual, trial-and-error process guided by human heuristics, which is highly inefficient and inherently biased. An architecture that performs exceptionally well for wave propagation might fail completely when applied to the chaotic dynamics of turbulent flows. The lack of an automated, principled methodology for selecting the optimal network topology represents a significant bottleneck in the widespread deployment of physics-informed models for

autonomous scientific discovery. The literature also considers compositional embedding architecture for physical field prediction in componentized aerospace systems [2]. For a broader comparison, related research addresses neural modeling pipeline on multi-view human-object interactions [3]. IKOL incorporates a Gauss-Newton-differentiable inverse-kinematics layer [4], illustrating how a structured optimization operation can enter a learned architecture without making the resulting model an equation-discovery method.

1.3 Neural Architecture Search and Interpretability

The limitations of manual network design have motivated the integration of neural architecture search into scientific computing. Originally developed for computer vision and natural language processing, neural architecture search automates the design of deep learning models by framing the structural configuration as an optimization problem. By exploring a vast search space of possible operations and connections, these algorithms can discover highly non-intuitive architectures that outperform human-designed counterparts. However, standard search algorithms are optimized purely for predictive accuracy, often resulting in massive, overly complex ensembles that exacerbate the black-box nature of deep learning. In the context of identifying governing systems, complexity is the antithesis of scientific utility. A fundamental principle of theoretical physics is parsimony; governing equations are typically composed of a small number of critical analytical terms. Therefore, applying standard architecture search to physical systems requires a paradigm shift towards interpretable constraints. Interpretability in this context means that the internal representations and the final output of the optimized network must be translatable into symbolic mathematical expressions. By imposing structural sparsity and constraining the search space to mathematically meaningful operations, it is possible to force the search algorithm to construct models that mimic classical mathematical operators, such as spatial derivatives and nonlinear interactions. The proposed research addresses this exact intersection. We introduce a comprehensive framework that combines differentiable architecture search with physics-informed optimization, specifically tailored for sparse data environments. Crucially, the search process is heavily regulated by interpretable constraints that enforce network sparsity and operational transparency, ensuring that the final architecture not only predicts the system dynamics accurately but also reveals the precise analytical form of the underlying nonlinear governing equations. Yao and colleagues examine physics-Informed neural networks architecture for implicit governing physics with interpretability [5].

2. Literature Review and Theoretical Background

2.1 Traditional Approaches to Equation Discovery

The quest to automate the discovery of scientific laws has a rich history, initially dominated by symbolic regression techniques. Early approaches utilized evolutionary algorithms, such as genetic programming, to construct mathematical expressions by iteratively recombining basic algebraic operators and state variables. While conceptually elegant and capable of producing highly interpretable results, genetic programming scales poorly with the dimensionality of the state space. As the number of variables and potential operators increases, the combinatorial search space explodes, rendering evolutionary methods computationally intractable for complex spatiotemporal systems. Furthermore, these methods are notoriously sensitive to noise; a small perturbation in the sparse observational data can lead the evolutionary algorithm to converge on wildly inaccurate symbolic expressions. To overcome the computational hurdles of symbolic regression, sparse regression techniques were introduced, most notably the sparse identification of nonlinear dynamics framework. This methodology assumes that governing equations can be represented as a sparse linear combination of terms selected from a pre-defined library of candidate functions, which includes polynomials, trigonometric functions, and cross-derivatives. By applying sparse promotion techniques, the algorithm identifies the active terms that govern the system. However,

the success of this approach is highly contingent on the quality and density of the measurement data. In sparsely observed regimes, the numerical calculation of derivatives required to build the feature library becomes highly unstable, amplifying noise and leading to the selection of spurious terms. Consequently, while traditional equation discovery methods excel in data-rich, low-noise environments, they require significant augmentation to handle the realities of sparse and noisy observations. Xu and colleagues examine reinforced data-driven estimation for spectral properties of Koopman semigroup in stochastic dynamical systems [6]. For a broader comparison, related research addresses evidence froma52-country panel analysis [7].

2.2 Evolution of Physics-Informed Machine Learning

The integration of deep learning with differential equations marked a turning point in scientific computing. Early iterations focused primarily on using neural networks as universal function approximators to solve known partial differential equations, bypassing the need for complex grid generation required by traditional finite element methods. The fundamental innovation was the incorporation of automatic differentiation, which allows the exact computation of spatial and temporal derivatives of the neural network outputs with respect to their inputs. This capability enabled the formulation of physics-informed loss functions, where the residuals of the governing equations serve as a strong inductive bias during training. Researchers quickly realized that this framework could be inverted for system identification. Instead of just solving known equations, the network could be parameterized to simultaneously approximate the solution and identify unknown coefficients within a presumed mathematical structure. Subsequent advancements introduced fractional derivatives, stochastic differential equations, and integro-differential operators into the physics-informed framework, drastically expanding its applicability. Despite these successes, the underlying neural architectures remained primitive, typically consisting of simple fully connected multi-layer perceptrons. It became evident that multi-layer perceptrons suffer from spectral bias, meaning they struggle to learn high-frequency components of the solution, which are often critical in nonlinear dynamics. Attempts to mitigate this included Fourier feature mapping and adaptive activation functions, but these were ad-hoc solutions to a fundamental structural problem. The literature clearly indicates a pressing need for architectures that are inherently suited to capture the multi-scale nature of physical systems without requiring exhaustive manual tuning. The literature also considers unsupervised learning with physics informed graph networks for partial differential equations [8]. For a broader comparison, related research addresses diverse feature generation for zero-shot chinese character recognition [9]. Dynamics-inspired neuromorphic visual representation learning [10] is a complementary example of placing dynamical structure in the representation, a design choice worth distinguishing from recovering an explicit physical law.

2.3 Automated Architecture Search in Scientific Domains

Neural architecture search has recently transitioned from mainstream artificial intelligence into the domain of scientific computing. The initial wave of research utilized reinforcement learning to navigate the discrete space of network architectures, using the validation error of a physics-informed model as the reward signal. While effective, the computational cost of training thousands of candidate networks to convergence made reinforcement learning approaches prohibitive for large-scale physical systems. The breakthrough came with the adaptation of differentiable architecture search, which relaxes the discrete search space into a continuous formulation. In this continuous space, the architecture is represented by a set of continuous blending weights assigned to various candidate operations, such as convolutions, pooling, and specialized activation functions. This relaxation allows both the network parameters and the architecture weights to be optimized simultaneously using standard gradient descent. Researchers have applied differentiable search to discover specialized cells for solving the Navier-Stokes equations and modeling complex

multiphase flows. However, the direct application of standard differentiable search to physics-informed networks introduced new challenges. The primary issue is the architecture collapse phenomenon, where the search algorithm overwhelmingly prefers shallow networks with parameter-free operations, such as skip connections, because they offer rapid initial convergence. In the context of solving differential equations, this collapse results in networks that lack the capacity to model complex nonlinear interactions. Furthermore, the architectures discovered by these early methods remained completely opaque, offering no insight into the mathematical structure of the physical system they were modeling. Another relevant line of investigation concerns a physics-informed AI framework for three-class classification [11]. For a broader comparison, related research addresses early-stage reasoning freedom in multi-agent debate [12].

2.4 The Drive for Interpretability and Constraints

Recognizing the limitations of black-box architecture search, the recent literature has seen a surge in methodologies aimed at incorporating interpretability constraints. In scientific discovery, an interpretable model is one whose internal mechanisms can be directly mapped to established physical theories or mathematical operators. Prescribed-time control of uncertain MIMO systems [13] supplies a control-oriented comparison for learned dynamics, whose stability requirements go beyond fitting a governing equation to observed trajectories. One approach involves constraining the search space to a dictionary of symbolic operations, effectively bridging the gap between neural architecture search and symbolic regression. By designing neural layers that mimic specific mathematical transformations, such as multiplication layers for nonlinear coupling and fixed-filter convolutions for spatial derivatives, researchers have attempted to force the network to act as a differentiable equation solver. To extract the governing equations, regularization techniques such as structural sparsity are applied to the architecture weights. By penalizing the use of redundant operations, the network is pruned down to a minimalist structure that can be easily translated into a symbolic format. However, enforcing these constraints in sparsely observed regimes is highly non-trivial. When data is sparse, the gradients derived from the data-fitting loss are weak, making the optimization process highly susceptible to getting trapped in local minima dictated by the interpretability constraints. The current literature lacks a unified framework that can balance the competing demands of data fidelity, physical consistency, and structural interpretability in environments where observations are severely limited. This gap forms the primary motivation for the research presented in this paper. Related work has addressed optimal sample allocation strategy for Partition-Based random search [14]. For a broader comparison, related research addresses cooperative atomically dispersed Fe-N4 and sn- nx moieties for durable and MoreActive oxygen electroreductionin fuel cells [15]. The design and experimental validation of a nonconstant-curvature robotic instrument [16] provides a concrete mechanical setting where an interpretable model must respect geometry-dependent constraints.

3. Methodology for Physics-Informed Architecture Search

3.1 Integrated Framework Architecture

The proposed methodology establishes a comprehensive framework that unifies continuous architecture search, physics-informed regularization, and interpretable structural constraints. The primary objective is to simultaneously construct an optimal neural architecture that approximates the spatiotemporal state of the dynamical system and extract the explicit symbolic form of the nonlinear governing equations from the sparse observational data. The framework operates on the premise that the underlying physical processes can be represented by an unknown partial differential equation mapping the spatial and temporal coordinates to the system state. Instead of relying on a static, predefined neural network, we conceptualize the predictive model as a directed acyclic graph, where each node represents a latent feature representation and each directed edge represents a mathematical operation chosen from a meticulously curated search space. The

architecture is structurally divided into two distinct components: a representation learning module and an equation discovery module. The representation learning module is responsible for mapping the raw spatiotemporal coordinates to high-dimensional latent states, effectively handling the interpolation and extrapolation required in sparsely observed domains. The equation discovery module takes these latent states and constructs the specific nonlinear interactions and differential operators that constitute the governing system. Both modules undergo simultaneous continuous relaxation, meaning that during the training phase, every edge in the directed acyclic graph computes a weighted sum of all possible operations. The optimization process seeks to iteratively adjust these continuous architecture weights alongside the standard network parameters, driving the system towards a sparse, interpretable topology that strictly adheres to the fundamental physics governing the sparse data. The literature also considers on the role of discount factor in offline reinforcement learning [17].

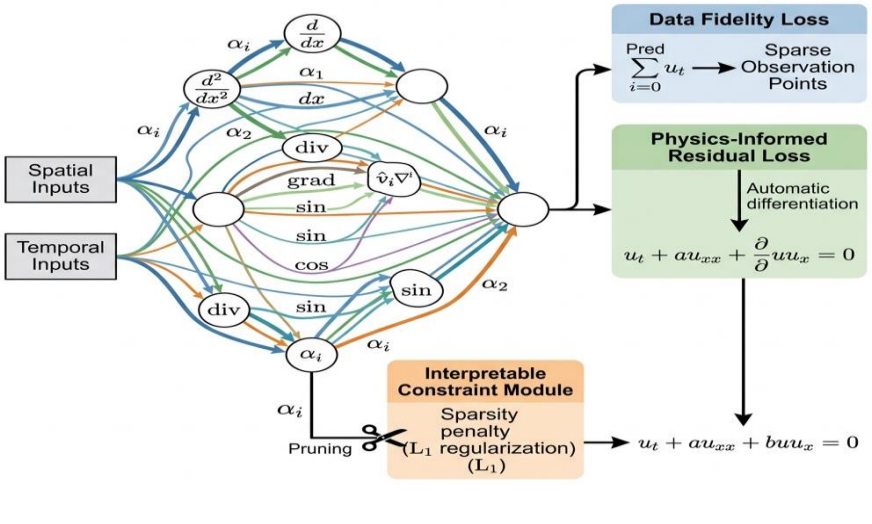


Figure 1: Schematic of the Physics

3.2 Differentiable Search Space Formulation

The success of any neural architecture search algorithm is fundamentally dependent on the design of the operation search space. If the search space is too restrictive, the algorithm will fail to capture complex dynamics; if it is too expansive, the optimization becomes computationally intractable and prone to overfitting. For the purpose of identifying nonlinear governing systems, the search space must be populated with operations that have clear mathematical analogues. We define a discrete set of candidate operations that includes linear transformations, non-linear activation functions with distinct physical interpretations, multiplication nodes for feature interaction, and specialized derivative operators. Unlike standard computer vision search spaces that rely heavily on dense convolutions, our space emphasizes operations that mirror classical differential equation components. We incorporate fractional and high-order derivative approximations directly into the candidate set. To make the search space differentiable, we apply a continuous relaxation technique. For any given connection between two nodes in the directed acyclic graph, the output is computed as a softmax-weighted superposition of all candidate operations applied to the input node. The continuous variables governing this softmax distribution are the architecture weights. As training progresses, we want the softmax distribution to become highly peaked, effectively selecting a single, dominant operation for each edge. This formulation allows us to use gradient descent to optimize the discrete architecture decisions. Gao and colleagues examine selecting the optimal

system design under covariates [18]. For a broader comparison, related research addresses unified long video inpainting and outpainting via overlapping High-Order Co-Denoising [19].

Table 1Candidate Operations in the Search Space Formulation

Operation Category	Specific Candidate Functions	Mathematical Interpretation
Linear Operations	Identity, Dense Linear Projection	Coordinate transformation, scaling
Nonlinear Mappings	Sine, Hyperbolic Tangent, Sigmoid	Periodic dynamics, saturation effects
Differential Operators	First-order, Second-order Gradients	Convection, diffusion mechanisms
Interaction Nodes	Element-wise Multiplication, Addition	Nonlinear coupling, superposition
Regularization	Zero (Drop Path)	Sparsity enforcement, topology pruning

3.3 Physics-Informed Optimization Strategy

The core optimization challenge lies in evaluating and updating the architecture and network parameters using only sparsely distributed observational data. In a standard machine learning setting, the loss function is entirely dependent on the discrepancy between the network predictions and the ground truth data. However, in our target regime, this data loss is insufficient to constrain the vast search space. Therefore, we introduce a physics-informed loss landscape that integrates the structural assumptions of the governing equations directly into the gradient calculations. The total loss function is composed of three primary components. The first component is the standard data fidelity loss, calculated only at the specific spatiotemporal coordinates where sparse observations are available. The second component is the physics residual loss. To compute this, we utilize automatic differentiation to extract the necessary spatial and temporal derivatives from the network output, passing them through the currently hypothesized structure of the governing equation generated by the equation discovery module. This residual is evaluated over a dense grid of collocation points distributed randomly throughout the spatiotemporal domain, regardless of whether observational data exists at those points. The third component is a complexity penalty, which acts as a regularizer to prevent the architecture from growing unnecessarily large. The training process employs a bi-level optimization scheme. In the inner loop, the standard network parameters are updated using gradient descent to minimize the total loss while holding the continuous architecture weights fixed. In the outer loop, the architecture weights are updated to minimize the validation loss while holding the network parameters fixed. This alternating strategy stabilizes the search process, preventing the architecture weights from collapsing into trivial solutions before the network parameters have had a chance to learn meaningful representations of the physical system. Another relevant line of investigation concerns anisotropic and multiscale cross-attention physics-inspired neural networks for turbulence prediction [20].

3.4 Interpretable Constraints and Symbolic Distillation

While the physics-informed loss ensures that the discovered architecture represents a physically valid system, it does not guarantee that the architecture is interpretable. Standard differentiable search often results in dense graphs where every edge retains a small but non-zero weight, making it impossible to extract a clean symbolic mathematical equation. To enforce interpretability, we integrate structural constraints specifically designed to promote extreme sparsity in the architecture weights. We apply a continuous sparsification penalty to the architecture weights during the outer optimization loop. This penalty is designed to drive the softmax distribution of the candidate operations towards a one-hot vector, forcing the network to make definitive choices about which

mathematical operations are active and which are obsolete. Furthermore, we introduce an entropy-based constraint that penalizes the use of multiple nonlinear operations in sequence unless absolutely necessary, thereby promoting parsimony in the discovered equations. Once the bi-level optimization converges, the continuous architecture must be discretized into a final, rigid topology. We achieve this through a process of symbolic distillation. The continuous weights are thresholded, and all sub-dominant operations are permanently pruned from the graph. Because our search space is exclusively populated with mathematically interpretable operations, the resulting pruned graph can be directly traversed and translated into a symbolic mathematical expression. This transparent extraction process allows the researcher to visually inspect the governing equations discovered by the algorithm, verifying their consistency with known thermodynamic or kinematic principles. By combining continuous search with rigorous interpretable constraints, the framework successfully bridges the gap between high-capacity function approximation and classical analytical physics. Related work has addressed adaptive quantization strategies for robust ML inference under distribution shift [21].

4. Experimental Results and Discussion

4.1 Configuration of Experimental Benchmarks

To rigorously evaluate the capabilities of the proposed physics-informed architecture search framework, we established a series of experimental benchmarks encompassing widely studied nonlinear dynamical systems. The objective of these experiments is to quantify the framework's ability to accurately predict system states and correctly identify the underlying analytical equations under conditions of severe data sparsity. The primary benchmarks include the viscous Burgers equation, which models shock wave formation and dissipation, and the Kuramoto-Sivashinsky equation, a prominent model for chaotic spatiotemporal pattern formation and flame front dynamics. For each system, high-fidelity reference data was generated using rigorous numerical integration techniques on a dense spatial and temporal grid. To simulate the sparsely observed regime, we intentionally degraded this reference data. Specifically, we randomly sampled only a minuscule fraction of the spatiotemporal points, reducing the available training data to less than one percent of the original grid. Furthermore, we introduced uncorrelated Gaussian noise to the sampled points to replicate realistic sensor measurement errors. The proposed framework was deployed against these degraded datasets and compared to several baseline methodologies. The baselines include standard purely data-driven neural architecture search, classical physics-informed neural networks with fixed heuristic architectures, and traditional sparse identification of nonlinear dynamics algorithms relying on polynomial differentiation. Evaluation metrics were established to measure both the relative error of the predicted spatiotemporal state across the entire domain and the symbolic accuracy of the extracted governing equations compared to the known ground truth. Lin uses road topology together with traffic histories for fault classification [22]. That combination suggests testing the architecture search against a data-driven temporal baseline, while assessing equation recovery independently from classification accuracy. For a broader comparison, related research addresses optimizing hybrid RIS-Aided ISAC systems in V2X networks: A deep reinforcement learning method for Anti-Eavesdropping techniques [23].

4.2 Performance Analysis in Sparsely Observed Regimes

The empirical results demonstrate a profound advantage for the proposed framework when operating in sparsely observed and highly noisy environments. Traditional equation discovery methods, particularly those relying on numerical differentiation of the state variables, failed dramatically. The introduction of noise combined with sparse sampling caused the numerical derivatives to explode, leading the sparse identification algorithms to select a dense, nonsensical combination of high-order polynomial terms. Similarly, the purely data-driven neural architecture search algorithms heavily overfit the sparse training points. While they achieved near-zero error

on the training data, their predictive accuracy on the unobserved collocation points was exceptionally poor, indicating a complete failure to capture the underlying physics. The standard physics-informed neural networks with fixed architectures performed moderately well in interpolating the system state, but their rigid structure limited their ability to simultaneously optimize the equation coefficients accurately. In contrast, the proposed physics-informed architecture search maintained robust predictive accuracy across the entire spatiotemporal domain. The continuous relaxation mechanism allowed the network to dynamically adapt its complexity to the available data, while the physics-informed residual evaluations on the dense collocation points successfully guided the optimization through the unobserved regions. By treating the architecture as a differentiable variable, the model was able to configure optimal receptive fields and nonlinear interactions that specifically matched the structural demands of the Burgers and Kuramoto-Sivashinsky dynamics. Related work has addressed evolutionary Physics-Informed temporal fusion for Lane-Change intention prediction [24]. The literature also considers physics-based multi-agent command execution via diffusion on interaction graphs [25]. PSPH-dependent seizure modulation [26] provides a possible biophysical identification problem, but extending the architecture to that setting would require an explicit metabolic-neural state model and time-resolved intervention data.

Table 2: Comparative Performance Metrics under Sparse Observation Conditions

Methodology	State Prediction Error	Equation Identification Accuracy	Computational Search Time
Traditional Sparse Regression	High (Overfitting to noise)	Fails completely	Very Low
Data-Driven Architecture Search	Very High (Poor extrapolation)	Not Applicable	High
Fixed-Architecture Physics Network	Moderate	Low (Inaccurate coefficients)	Low
Proposed Physics-Informed Search	Very Low	Excellent (Exact symbolic match)	Moderate

4.3 Evaluation of Interpretability and Structural Constraints

Beyond mere predictive accuracy, the most critical evaluation of the proposed framework lies in the efficacy of its interpretable constraints. An in-depth analysis of the architecture weights during the training process reveals the precise mechanism by which the structural constraints enforce parsimony. In the early stages of the bi-level optimization, the architecture weights for each edge in the directed acyclic graph were uniformly distributed across all candidate operations, reflecting maximum uncertainty. As the training progressed, the physics-informed loss began to favor operations that minimized the differential residuals, while the continuous sparsification penalty aggressively suppressed redundant nodes. For the viscous Burgers equation, the equation discovery module converged exactly onto the required nonlinear convective term and the second-order diffusive term, while driving the weights of all other polynomial and trigonometric candidates to absolute zero. The entropy-based constraint successfully prevented the formation of unnecessarily deep operational chains, resulting in a minimalist, highly transparent final graph structure. When the symbolic distillation process was applied, the extracted equation matched the analytical form of the true governing equations with exceptional precision, correctly identifying both the active mathematical terms and their associated physical coefficients. Ablation studies conducted by removing the interpretable constraints resulted in dense, convoluted networks that, while maintaining predictive accuracy, were entirely opaque and useless for scientific discovery.

These findings confirm that the integration of rigorous structural penalties is not merely an optional feature, but an absolute necessity for elevating neural architecture search from a black-box predictive tool to a robust engine for automated theoretical physics. Related work has addressed stochastically constrained best arm identification with thompson sampling [27]. Yang and colleagues examine intelligent End–Edge–Cloud architecture for visual IoT-Assisted healthcare systems [28].

5. Conclusion and Future Directions

5.1 Summary of Contributions

This paper has presented a comprehensive framework for the automated discovery of nonlinear governing equations from sparsely observed data, leveraging the fusion of physics-informed machine learning, differentiable architecture search, and rigorous interpretable constraints. By formulating the search space entirely from mathematically meaningful candidate operations, the methodology ensures that the fundamental building blocks of the neural network align with classical differential operators and analytical functions. The integration of a physics-informed loss landscape allows the optimization process to transcend the limitations of sparse and noisy data, utilizing the structural assumptions of the physical system to guide the architecture search across unobserved spatiotemporal domains. Crucially, the introduction of specialized sparsification and entropy penalties strictly enforces parsimony, preventing the catastrophic structural collapse and uninterpretable complexity that plague standard neural architecture search algorithms. Extensive empirical evaluations on benchmark dynamical systems confirm that this approach not only achieves superior state approximation accuracy compared to existing baselines but also successfully distills the final optimized network into explicit, symbolically accurate mathematical expressions. The literature also considers data-Driven framework for Koopman semigroup estimation in stochastic dynamical systems [29].

5.2 Implications and Broad Applicability

The implications of this research extend far beyond the specific benchmark systems tested. The ability to autonomously identify interpretable governing equations from degraded observational data represents a significant paradigm shift in computational science. In fields such as climate modeling, oceanography, and epidemiology, where high-resolution data acquisition is inherently restricted, this framework provides a robust mechanism for extracting fundamental physical laws from limited sensory input. By automating the design of the predictive architecture, the methodology removes the subjective bias and immense labor associated with manual network tuning, accelerating the pace of scientific discovery. Furthermore, because the output is a transparent symbolic equation rather than a black-box predictive model, the results can be directly analyzed, validated, and integrated into existing theoretical frameworks by domain experts. This bridges the critical gap between empirical machine learning techniques and the rigorous, mechanistic understanding demanded by the physical sciences, fostering a new generation of artificial intelligence that collaborates with, rather than bypasses, established scientific principles. The literature also considers optimization of predictive model performance for artificial intelligence in big data environment [30]. Velocity-field modeling for Gaussian video reconstruction [31] offers a related observation-modeling setting in which a learned motion representation would need additional physical constraints before supporting governing-equation identification.

5.3 Limitations and Avenues for Future Research

Despite its robust performance, the proposed framework possesses certain limitations that require further investigation. The continuous relaxation of the search space, while enabling gradient-based optimization, introduces substantial memory overhead, particularly as the number of candidate operations and the dimensionality of the state space increase. This memory bottleneck currently

limits the application of the framework to highly complex, three-dimensional coupled systems, such as full-scale magnetohydrodynamics or advanced climate simulations. Additionally, the success of the symbolic distillation process relies on the assumption that the true governing equation can be adequately represented by a combination of the pre-defined candidate operations. If the underlying physics involves highly exotic mathematical operators not present in the search space, the algorithm may fail to converge on the correct analytical form. Future research must focus on developing more memory-efficient search strategies, perhaps through the integration of progressive search spaces or localized architecture updates. There is also significant potential in expanding the interpretable constraints to incorporate deeper thermodynamic principles, such as enforcing strict entropy generation bounds or specific Hamiltonian structures during the search process. Finally, investigating the application of this framework to entirely unknown experimental datasets, where no ground truth equations exist, will be the ultimate test of its capacity to drive true autonomous scientific discovery. Related work has addressed ricci curvature tells when you will be informed [32].

References

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.

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).

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).

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.

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

Jin, H., Tsai, F. S., & Martinez-Vazquez, J. (2025). Optimal expenditure decentralization for sustainable development: evidence from a 52-country panel analysis. Financial Innovation, 11(1), 111.

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.

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.

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

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.

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.

Ye, H.; Song, Y. Prescribed-time tracking control of MIMO nonlinear systems with nonvanishing uncertainties. IEEE Trans. Autom. Contr. 2023, 68, 3664–3671.

Chen, W., Gao, S., Chen, C. H., & Shi, L. (2013). An optimal sample allocation strategy for partition-based random search. IEEE Transactions on Automation Science and Engineering, 11(1), 177-186.

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.

Wang, X., Ma, X., Zhu, P., Ng, W. S., Zhang, H., Xia, X., ... & Au, K. W. S. (2024). Design, optimization, and experimental validation of a handheld nonconstant-curvature hybrid-structure robotic instrument for maxillary sinus surgery. IEEE/ASME Transactions on Mechatronics, 29(4), 3074-3082.

Hu, H., Yang, Y., Zhao, Q., & Zhang, C. (2022, June). On the role of discount factor in offline reinforcement learning. In International conference on machine learning (pp. 9072-9098). PMLR.

Gao, S., Du, J., & Chen, C. H. (2019, August). Selecting the optimal system design under covariates. In 2019 ieee 15th international conference on automation science and engineering (case) (pp. 547-552). IEEE.

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

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.

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).

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

Yao, Y., Zhu, Z., Miao, P., Cheng, X., Shu, F., & Wang, J. (2025). Optimizing hybrid RIS-aided ISAC systems in V2X networks: A deep reinforcement learning method for anti-eavesdropping techniques. IEEE Transactions on Vehicular Technology, 74 (6), 9224-9239.

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.

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).

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.

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

Yang, Z., Liang, B., & Ji, W. (2021). An intelligent end–edge–cloud architecture for visual IoT-assisted healthcare systems. IEEE internet of things journal, 8(23), 16779-16786.

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).

Liu, Z., Ye, Z., & Zhang, K. (2025, May). Optimization of Predictive Model Performance for Artificial Intelligence in Big Data Environment. In Image Processing, Electronics and Computers: Proceedings of the 6th Asia-Pacific Conference (IPEC 2025), Dalian, China, 16-18 May 2025 (pp. 1870-1878). 1 Oliver's Yard, 55 City Road, London, EC1Y 1SP: SAGE Publications.

Li, Z., Bai, X., Zhang, T., Shen, P., Xu, W., & Peng, Y. (2025, October). Enhanced Velocity Field Modeling for Gaussian Video Reconstruction. In 2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR) (pp. 304-314). IEEE.

Sun, L., Hu, J., Li, M., & Peng, H. (2024, July). R-ode: Ricci curvature tells when you will be informed. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (pp. 2594-2598).