

# Mathematical Interpretation of PINN Latent Space for Burger’s Equation: Learned Dynamics and Geometric Structure

ASTROPILOT<sup>1</sup>

<sup>1</sup>*Anthropic, Gemini & OpenAI servers. Planet Earth.*

## ABSTRACT

Interpreting the internal representations learned by Physics-Informed Neural Networks (PINNs) remains a significant challenge. This study provides a mathematical interpretation of the 10-dimensional latent space,  $L(x, t)$ , learned by a PINN trained to solve the 2D Burger’s equation. We analyze the geometric structure and learned dynamics of this latent space by examining the latent variables themselves and their spatial and temporal derivatives,  $\mathbf{V}_x = \partial L / \partial x$  and  $\mathbf{V}_t = \partial L / \partial t$ , using a dataset of the learned latent space over a 100x100 spatial-temporal grid. Derivatives are computed via finite differences, followed by analysis of descriptive statistics, vector magnitudes, and cosine similarities between  $L, \mathbf{V}_x, \mathbf{V}_t$ . We assess the local dimensionality of the tangent space spanned by  $\mathbf{V}_x$  and  $\mathbf{V}_t$  using singular value decomposition. Finally, sparse regression is employed to discover a system of differential equations governing the latent space evolution,  $\partial L / \partial t = f(L, \mathbf{V}_x, \mathbf{V}_{xx})$ . Our results show that latent variables exhibit significant correlations and heterogeneous statistics. Geometrically, the latent space manifold is structured: spatial gradients  $|\mathbf{V}_x|$  are typically larger than temporal gradients  $|\mathbf{V}_t|$ , and  $\mathbf{V}_x$  and  $\mathbf{V}_t$  vectors are often anti-aligned. The local tangent space is frequently nearly one-dimensional, suggesting a strong constraint on simultaneous spatial and temporal variation. Sparse regression successfully identifies a coupled system of nonlinear partial differential equations for the latent dynamics with high accuracy. Crucially, these learned latent PDEs contain terms structurally analogous to the nonlinear advection ( $L_j \mathbf{V}_{x,j}$ ) and diffusion ( $\mathbf{V}_{xx,j}$ ) operators of the original Burger’s equation, demonstrating that the PINN has encoded key physical principles within its internal representation. This work offers a novel mathematical formalism for interpreting the learned internal models of PINNs, moving beyond black-box function approximation.

*Keywords:* Computational methods, Algorithms, Astronomical techniques, Analytical mathematics, Regression

## 1. INTRODUCTION

The accurate and efficient solution of partial differential equations (PDEs) stands as a fundamental challenge across numerous scientific and engineering domains. Traditional numerical methods, such as finite differences, finite elements, or spectral methods, often necessitate discretizing the physical domain onto complex meshes, which can be computationally expensive and difficult to adapt to intricate geometries or high-dimensional problems. The need for alternative paradigms has spurred the development of data-driven approaches, particularly those leveraging deep learning, for solving and analyzing PDEs.

Physics-Informed Neural Networks (PINNs) represent a significant advancement in this direction. Unlike purely data-driven methods that require large datasets of paired inputs and outputs, PINNs embed the govern-

ing physical laws, expressed as PDE residuals, directly into the neural network’s loss function alongside data constraints. This unique formulation allows PINNs to learn solutions to forward problems, discover unknown parameters in inverse problems, and even approximate operators, often requiring only sparse data measuring the system state.

Despite their empirical success in solving a wide range of PDE problems, including those in fluid dynamics, heat transfer, and quantum mechanics, PINNs largely remain “black boxes.” The internal workings, specifically the mathematical or physical meaning of the learned parameters and the intermediate representations within the network layers, are opaque (Dahlbüdding et al. 2024). Understanding precisely *how* a PINN encodes the complex, often nonlinear dynamics and structure of

a physical system described by a PDE is a critical open problem (Clemente et al. 2025,?).

This lack of interpretability limits our ability to fully trust the model’s predictions, diagnose failures, guide model improvements, or extract novel scientific insights from the learned representation itself (Dahlbüding et al. 2024; Hao et al. 2025). This study addresses this challenge by providing a detailed mathematical interpretation of the internal representation learned by a PINN.

We focus on the 10-dimensional “latent space,” denoted  $L(x, t) \in \mathbb{R}^{10}$ , which represents a learned embedding of the spatial-temporal coordinates  $(x, t)$  into a higher-dimensional vector space by a PINN trained to solve the 2D Burger’s equation:

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} - \nu \frac{\partial^2 u}{\partial x^2} = 0 \quad (1)$$

(McDuffee 2023; Chiodaroli & Feireisl 2024) where  $u(x, t)$  is the dependent variable and  $\nu$  is the viscosity. Burger’s equation serves as a valuable testbed due to its balance of nonlinear advection and diffusion, exhibiting phenomena like wave steepening and shock formation.

Our central hypothesis is that the essential physics captured by the PINN is profoundly imprinted not just in the output layer  $u(x, t)$ , but also in the intrinsic mathematical structure and dynamics of the learned latent space  $L(x, t)$ . To test this, we propose a novel approach: analyzing the geometric properties and searching for the governing differential equations of this latent space by examining the latent variables themselves,  $L(x, t)$ , and, crucially, their spatial and temporal gradient vector fields,  $\mathbf{V}_x(x, t) = \partial L / \partial x$  and  $\mathbf{V}_t(x, t) = \partial L / \partial t$ . These vector fields represent the local directions and rates of change of the latent mapping in  $\mathbb{R}^{10}$  with respect to the physical coordinates, providing insights into the local structure and evolution of the learned representation.

We analyze a dataset of the learned latent space  $L(x, t)$  obtained from a trained PINN over a spatial-temporal grid. We numerically compute the spatial derivatives  $\mathbf{V}_x$  and  $\mathbf{V}_{xx} = \partial^2 L / \partial x^2$ , and the temporal derivatives  $\mathbf{V}_t$  using finite differences. We then quantitatively characterize the latent space and its gradients through statistical analysis of their magnitudes and the geometric relationships (e.g., alignment via cosine similarity) between the vectors  $L$ ,  $\mathbf{V}_x$ , and  $\mathbf{V}_t$  at each point in the  $(x, t)$  domain.

We further investigate the local geometric structure of the learned manifold embedded in  $\mathbb{R}^{10}$  by analyzing the tangent space spanned by  $\mathbf{V}_x$  and  $\mathbf{V}_t$ , quantifying its effective dimensionality using Singular Value Decomposition (SVD). Finally, and most significantly, we employ sparse regression techniques to systematically discover

a system of differential equations that governs the evolution of the latent space, seeking relationships of the form  $\partial L / \partial t = f(L, \mathbf{V}_x, \mathbf{V}_{xx})$ . This step aims to uncover the underlying mathematical rules that dictate how the PINN’s internal representation changes over time based on its current state and spatial configuration.

Through this comprehensive mathematical analysis, we demonstrate that the learned latent space is highly structured and encodes key aspects of the underlying physics. Our results reveal significant correlations and heterogeneous distributions among the latent variables. Geometrically, the spatial gradient vectors  $\mathbf{V}_x$  are typically larger in magnitude than temporal gradient vectors  $\mathbf{V}_t$ , and these vectors are frequently found to be nearly anti-aligned. The local tangent space spanned by  $\mathbf{V}_x$  and  $\mathbf{V}_t$  is often close to one-dimensional, suggesting strong constraints on the simultaneous variation of the latent representation in space and time.

Crucially, our sparse regression analysis successfully identifies a coupled system of nonlinear partial differential equations that accurately describe the latent space dynamics. We show that these learned latent PDEs contain terms structurally analogous to the nonlinear advection ( $L_j \mathbf{V}_{x,j}$ ) and diffusion ( $\mathbf{V}_{xx,j}$ ) operators present in the original Burger’s equation, providing compelling evidence that the PINN has internally learned and encoded key physical principles within its abstract, high-dimensional representation. This work thus offers a novel mathematical formalism for interpreting the learned internal models of PINNs, moving beyond their traditional treatment as black-box function approximators and paving the way for extracting deeper physical and mathematical insights from their learned representations.

## 2. METHODS

The primary objective of this study is to perform a detailed mathematical interpretation of the 10-dimensional latent space learned by a Physics-Informed Neural Network (PINN) trained to solve the 2D Burger’s equation. This analysis is based on a dataset representing the learned latent space over a spatial-temporal grid. Our methodology involves several key steps: loading and preparing the data, conducting exploratory analysis of the latent variables, numerically computing spatial and temporal derivatives of the latent space, analyzing the geometric properties and relationships between the latent space vectors and their gradients, assessing the local dimensionality of the learned manifold, and finally, employing sparse regression to discover a system of differential equations governing the latent space dynamics.

### 2.1. Data Acquisition and Preprocessing

The dataset analyzed in this study was generated by a pre-trained PINN solving the two-dimensional Burger’s equation. The dataset is provided as a NumPy array with dimensions  $(100, 100, 12)$ , representing a grid of  $100 \times 100$  points in the  $(x, t)$  domain. The first two channels of the array correspond to the spatial  $(x)$  and temporal  $(t)$  coordinates, respectively, forming coordinate meshes  $X_{\text{mesh}}$  and  $T_{\text{mesh}}$ . The remaining 10 channels represent the 10 dimensions of the learned latent space  $L(x, t) \in \mathbb{R}^{10}$  at each grid point. The dataset was loaded using the NumPy library. The spatial and temporal coordinates were extracted as  $X_{\text{mesh}} = \text{data\_}[:, :, 0]$  and  $T_{\text{mesh}} = \text{data\_}[:, :, 1]$ . The 10-dimensional latent space data was extracted as  $L_{\text{space}} = \text{data\_}[:, :, 2 :]$ , resulting in a  $(100, 100, 10)$  NumPy array. The grid consists of  $N_x = 100$  spatial points and  $N_t = 100$  temporal points. The unique spatial coordinates  $x_i$  ( $i = 0, \dots, N_x - 1$ ) and temporal coordinates  $t_j$  ( $j = 0, \dots, N_t - 1$ ) were extracted and sorted to determine the grid spacing. The grid was found to be uniform, with spatial step size  $\Delta x = x_1 - x_0$  and temporal step size  $\Delta t = t_1 - t_0$ .

### 2.2. Exploratory Data Analysis of the Latent Space

Prior to analyzing the derivatives, an exploratory data analysis (EDA) was performed on the raw latent space data  $L(x, t)$  to understand the basic statistical properties and inter-dependencies among the 10 latent dimensions  $L_k(x, t)$ , where  $k \in \{0, \dots, 9\}$ . For each of the 10 latent dimensions  $L_k(x, t)$ , where  $k \in \{0, \dots, 9\}$ , standard descriptive statistics including mean, variance, minimum value, and maximum value were computed over the entire  $N_x \times N_t$  grid. These statistics provide insights into the distribution and range of values for each individual latent variable (Andrianomena & Hassan 2023; Iwasaki et al. 2023; Horta et al. 2025; Huijse et al. 2025; Lalchand & Eilers 2025). To assess the relationships between different latent dimensions, the Pearson correlation matrix was computed (Trusov et al. 2024; Skowronski et al. 2025). The  $L_{\text{space}}$  array of shape  $(100, 100, 10)$  was reshaped into a 2D array of shape  $(10000, 10)$ , where each row corresponds to a unique  $(x, t)$  grid point and each column represents a latent dimension. The  $10 \times 10$  correlation matrix was then calculated from this reshaped data. High correlation values (positive or negative) indicate that certain latent dimensions covary significantly across the domain, suggesting potential functional dependencies or shared encoding roles (Skowronski et al. 2025).

### 2.3. Numerical Computation of Latent Space Derivatives

To analyze the local structure and dynamics of the learned latent space manifold, we computed the spatial gradient vector field  $\mathbf{V}_x(x, t) = \partial L / \partial x$ , the temporal gradient vector field  $\mathbf{V}_t(x, t) = \partial L / \partial t$ , and the second-order spatial derivative  $\mathbf{V}_{xx}(x, t) = \partial^2 L / \partial x^2$  (Lue et al. 2025,?). These derivatives represent the rate and direction of change of the 10-dimensional latent vector with respect to spatial position and time (Lue et al. 2025,?). The derivatives were computed numerically using finite difference schemes on the uniform spatial and temporal grids (Caunt & Korpi 2001; Minoshima et al. 2019; Balsara et al. 2024). For each latent dimension  $L_k(x, t)$ :

#### 2.3.1. Spatial Derivatives $\mathbf{V}_x$

The first-order spatial derivative  $\partial L_k / \partial x$  was approximated using (Pan et al. 2019; Minoshima et al. 2019).

- Second-order central differences for interior points ( $1 \leq i \leq N_x - 2$ ):

$$\left. \frac{\partial L_k}{\partial x} \right|_{(x_i, t_j)} \approx \frac{L_k(x_{i+1}, t_j) - L_k(x_{i-1}, t_j)}{2\Delta x}$$

- First-order forward difference for the first spatial point ( $i = 0$ ):

$$\left. \frac{\partial L_k}{\partial x} \right|_{(x_0, t_j)} \approx \frac{L_k(x_1, t_j) - L_k(x_0, t_j)}{\Delta x}$$

- First-order backward difference for the last spatial point ( $i = N_x - 1$ ):

$$\left. \frac{\partial L_k}{\partial x} \right|_{(x_{N_x-1}, t_j)} \approx \frac{L_k(x_{N_x-1}, t_j) - L_k(x_{N_x-2}, t_j)}{\Delta x}$$

This computation was performed independently for each of the 10 latent dimensions and each time step, resulting in a  $\mathbf{V}_x$  array of shape  $(100, 100, 10)$  (Fagin et al. 2024).

#### 2.3.2. Temporal Derivatives $\mathbf{V}_t$

The first-order temporal derivative  $\partial L_k / \partial t$  was approximated similarly (Tabatabaei et al. 2023; Xiao et al. 2024).

- Second-order central differences for interior points ( $1 \leq j \leq N_t - 2$ ):

$$\left. \frac{\partial L_k}{\partial t} \right|_{(x_i, t_j)} \approx \frac{L_k(x_i, t_{j+1}) - L_k(x_i, t_{j-1})}{2\Delta t}$$

- First-order forward difference for the first temporal point ( $j = 0$ ):

$$\left. \frac{\partial L_k}{\partial t} \right|_{(x_i, t_0)} \approx \frac{L_k(x_i, t_1) - L_k(x_i, t_0)}{\Delta t}$$

- First-order backward difference for the last temporal point ( $j = N_t - 1$ ):

$$\left. \frac{\partial L_k}{\partial t} \right|_{(x_i, t_{N_t-1})} \approx \frac{L_k(x_i, t_{N_t-1}) - L_k(x_i, t_{N_t-2})}{\Delta t}$$

This yielded a  $\mathbf{V}_t$  array of shape (100, 100, 10) (Price 2024).

### 2.3.3. Second-Order Spatial Derivatives $\mathbf{V}_{xx}$

The second-order spatial derivative  $\partial^2 L_k / \partial x^2$  was approximated using (Haddad & Haddad 2021).

- Second-order central differences for interior points ( $1 \leq i \leq N_x - 2$ ):

$$\left. \frac{\partial^2 L_k}{\partial x^2} \right|_{(x_i, t_j)} \approx \frac{L_k(x_{i+1}, t_j) - 2L_k(x_i, t_j) + L_k(x_{i-1}, t_j)}{(\Delta x)^2}$$

For simplicity in the subsequent regression analysis, values at the spatial boundaries ( $i = 0, N_x - 1$ ) were not used or were handled separately, focusing the analysis on the interior grid points where second-order central differences are applicable (Taruya et al. 2021; Alzahrani et al. 2025). This resulted in a  $\mathbf{V}_{xx}$  array of shape (100, 100, 10), with potential handling for boundary points depending on the specific analysis requiring  $\mathbf{V}_{xx}$ .

## 2.4. Analysis of Latent Space Vector Fields

At each grid point  $(x_i, t_j)$ , the latent space value  $L(x_i, t_j)$  (Andrianomena & Tang 2023; Song et al. 2025), spatial gradient  $\mathbf{V}_x(x_i, t_j)$  (Kuiper et al. 2025; Laroche & Speagle 2025), and temporal gradient  $\mathbf{V}_t(x_i, t_j)$  (Kuiper et al. 2025; Laroche & Speagle 2025) are 10-dimensional vectors. We analyzed the magnitudes and relative orientations of these vectors across the  $(x, t)$  domain (Song et al. 2025; Pandya et al. 2025; Kuiper et al. 2025).

### 2.4.1. Magnitude Analysis

The Euclidean norm (magnitude) was computed for each vector at every grid point: (Daglayan et al. 2023, 2024)

$$\begin{aligned} |L(x_i, t_j)| &= \sqrt{\sum_{k=0}^9 (L_k(x_i, t_j))^2} \\ |\mathbf{V}_x(x_i, t_j)| &= \sqrt{\sum_{k=0}^9 (V_{x,k}(x_i, t_j))^2} \\ |\mathbf{V}_t(x_i, t_j)| &= \sqrt{\sum_{k=0}^9 (V_{t,k}(x_i, t_j))^2} \end{aligned}$$

Summary statistics (mean, median, minimum, maximum, and standard deviation) were calculated for the distributions of  $|L|$ ,  $|\mathbf{V}_x|$ , and  $|\mathbf{V}_t|$  over all grid points (Andrianomena & Hassan 2023; Iwasaki et al. 2023; Piras et al. 2025). This analysis helps characterize the typical scale and variability of the latent representation and its rates of change in space and time (Andrianomena & Hassan 2023; Piras et al. 2025).

### 2.4.2. Alignment Analysis

The cosine similarity was computed between pairs of vectors to quantify their relative alignment at each grid point. The cosine similarity between two 10-dimensional vectors  $\mathbf{A}$  and  $\mathbf{B}$  is given by  $\frac{\mathbf{A} \cdot \mathbf{B}}{|\mathbf{A}| |\mathbf{B}|}$ .

- Cosine similarity between  $L$  and  $\mathbf{V}_x$ :  $\cos(\theta_{L, \mathbf{V}_x}) = \frac{L \cdot \mathbf{V}_x}{|L| |\mathbf{V}_x|}$
- Cosine similarity between  $L$  and  $\mathbf{V}_t$ :  $\cos(\theta_{L, \mathbf{V}_t}) = \frac{L \cdot \mathbf{V}_t}{|L| |\mathbf{V}_t|}$
- Cosine similarity between  $\mathbf{V}_x$  and  $\mathbf{V}_t$ :  $\cos(\theta_{\mathbf{V}_x, \mathbf{V}_t}) = \frac{\mathbf{V}_x \cdot \mathbf{V}_t}{|\mathbf{V}_x| |\mathbf{V}_t|}$

Grid points where either vector had zero magnitude were excluded from the similarity calculation for that pair. Summary statistics (mean, median, and standard deviation) were computed for the distributions of these cosine similarity values across the grid. Values close to 1 indicate strong alignment, values close to -1 indicate strong anti-alignment, and values close to 0 indicate near orthogonality.

### 2.5. Local Geometric Structure of the Tangent Field

The spatial and temporal gradient vectors,  $\mathbf{V}_x(x, t)$  and  $\mathbf{V}_t(x, t)$ , locally span a tangent space to the manifold traced by  $L(x, t)$  within the 10-dimensional latent space  $\mathbb{R}^{10}$  (Lalchand & Eilers 2025; Horta et al. 2025). We assessed the effective dimensionality of this local tangent space at each grid point (Cadiou et al. 2025).

At each point  $(x_i, t_j)$ , a  $10 \times 2$  matrix  $M$  was formed with  $\mathbf{V}_x(x_i, t_j)$  and  $\mathbf{V}_t(x_i, t_j)$  as its columns:  $M = [\mathbf{V}_x | \mathbf{V}_t]$ . The Singular Value Decomposition (SVD) of this matrix was computed, yielding two singular values,  $\sigma_1 \geq \sigma_2 \geq 0$  (Rucinski 1998; Creevey et al. 2006; Zuo et al. 2023). The singular values represent the magnitudes of the principal semi-axes of the ellipse formed by mapping the unit circle in the  $(x, t)$  plane to the tangent plane in  $\mathbb{R}^{10}$ .

### 2.6. Discovery of Latent Space Dynamics via Sparse Regression

A key goal was to discover a system of differential equations governing the evolution of the latent space,

specifically seeking relationships of the form  $\partial L_k/\partial t = f_k(L, \mathbf{V}_x, \mathbf{V}_{xx})$  for each latent dimension  $k = 0, \dots, 9$ . We employed sparse regression, inspired by the Sparse Identification of Nonlinear Dynamics (SINDy) framework, to identify the terms and coefficients in these equations from the computed data.

### 2.6.1. Data Preparation for Regression

The computed data arrays  $L_{\text{space}}$ ,  $\mathbf{V}_x$ ,  $\mathbf{V}_t$ , and  $\mathbf{V}_{xx}$ , each of shape  $(100, 100, 10)$ , were reshaped into 2D arrays suitable for regression. This involved flattening the spatial and temporal dimensions, resulting in arrays of shape  $(N_x \times N_t, 10)$ . Let these flattened arrays be denoted  $L_{\text{flat}}$ ,  $\mathbf{V}_{x,\text{flat}}$ ,  $\mathbf{V}_{t,\text{flat}}$ , and  $\mathbf{V}_{xx,\text{flat}}$ . Each row corresponds to a single grid point  $(x_i, t_j)$ , and each column corresponds to a latent dimension  $k$ .

### 2.6.2. Candidate Function Library Construction

For each grid point  $(x_i, t_j)$ , we constructed a vector of candidate functions of  $L$ ,  $\mathbf{V}_x$ , and  $\mathbf{V}_{xx}$  evaluated at that point. These candidate functions form the columns of the library matrix  $\Theta$ . The library was designed to include terms potentially relevant to PDE dynamics, particularly those analogous to terms found in the original Burger’s equation. The library included:

- A constant term (1).
- Linear terms in  $L$ :  $L_0, L_1, \dots, L_9$ .
- Linear terms in  $\mathbf{V}_x$ :  $\partial L_0/\partial x, \dots, \partial L_9/\partial x$ .
- Linear terms in  $\mathbf{V}_{xx}$ :  $\partial^2 L_0/\partial x^2, \dots, \partial^2 L_9/\partial x^2$ .
- Quadratic terms involving  $L$  and  $\mathbf{V}_x$ :  $L_j(\partial L_m/\partial x)$  for all  $j, m \in \{0, \dots, 9\}$ . These terms are structurally analogous to the nonlinear advection term  $u(\partial u/\partial x)$  in Burger’s equation.

Higher-order terms or terms involving  $\mathbf{V}_t$  (except for the target variable) were initially excluded to maintain a manageable library size and prioritize simpler, physically inspired structures. The library matrix  $\Theta$  was constructed by evaluating these candidate functions at all  $N_x \times N_t$  grid points, resulting in a matrix where each row corresponds to a grid point and each column corresponds to a candidate function (Hatipoğlu 2022, 2023).

### 2.6.3. Sparse Regression Algorithm

For each latent dimension  $k \in \{0, \dots, 9\}$ , we sought to find a sparse coefficient vector  $\Xi_k$  such that the temporal derivative  $\mathbf{V}_{t,\text{flat}}[:, k]$  (the  $k$ -th column of  $\mathbf{V}_{t,\text{flat}}$ ) could be approximated by a linear combination of the candidate functions in  $\Theta$  (Horta et al. 2025).

$$\mathbf{V}_{t,\text{flat}}[:, k] \approx \Theta \Xi_k$$

Sparse regression was performed using the Lasso algorithm, which minimizes the objective function:

$$\|\mathbf{y}_k - \Theta \Xi_k\|_2^2 + \alpha \|\Xi_k\|_1$$

where  $\|\cdot\|_2$  is the L2 norm,  $\|\cdot\|_1$  is the L1 norm, and  $\alpha > 0$  is the regularization parameter that promotes sparsity in the coefficient vector  $\Xi_k$ . The parameter  $\alpha$  was selected by balancing the goodness of fit (minimizing the L2 error) with the sparsity of the resulting coefficient vector (Wohlberg & Wozniak 2021; Colgan et al. 2022; Wang et al. 2023).

### 2.6.4. Equation Identification

The output of the sparse regression for each latent dimension  $k$  is a sparse coefficient vector  $\Xi_k$ . The non-zero entries in  $\Xi_k$  identify the specific candidate functions from the library  $\Theta$  that contribute to the dynamics of  $\partial L_k/\partial t$ . The value of a non-zero entry represents the coefficient of the corresponding term in the discovered equation. By examining the non-zero coefficients for each  $k$ , we assembled the system of 10 differential equations governing the latent space dynamics.

### 2.7. Implementation Details

All data loading, preprocessing (Turk & Smith 2011; Mattia et al. 2025; Parsotan et al. 2025), numerical differentiation, statistical analysis (Turk & Smith 2011; Parsotan et al. 2025), geometric analysis (SVD) (Sarkar & Petersen 2025), and sparse regression were implemented using Python, primarily utilizing the NumPy library for numerical operations (Turk & Smith 2011), SciPy for sparse regression solvers (e.g., Lasso), and potentially other libraries for specific tasks if needed.

## 3. RESULTS

This section presents a detailed mathematical interpretation of the 10-dimensional latent space  $L(x, t)$  learned by the PINN trained on the 2D Burger’s equation. We analyze the statistical properties of the latent variables, the characteristics of their spatial and temporal derivatives, the geometric structure of the learned manifold in  $\mathbb{R}^{10}$ , and the governing differential equations discovered for the latent space dynamics.

### 3.1. Statistical characterization of the latent space

An initial exploratory analysis of the 10 latent dimensions,  $L_k(x, t)$  for  $k = 0, \dots, 9$ , was performed across the  $100 \times 100$  spatial-temporal grid. Descriptive statistics for each latent dimension (mean, variance, minimum, maximum) reveal heterogeneity in their distributions. While most means are close to zero,  $L_0$  and  $L_5$  exhibit slightly


larger positive average values. The variances vary considerably across dimensions, ranging from a minimum of approximately 0.48 for  $L_2$  to a maximum of about 1.95 for  $L_3$ . The total range of values (maximum minus minimum) also differs significantly between dimensions, indicating that different latent components activate over different scales and value ranges within the learned representation. This suggests that the network does not treat all latent dimensions identically; they likely encode distinct, though potentially related, features of the solution or the underlying physics.

The Pearson correlation matrix computed between all pairs of latent dimensions highlights significant linear dependencies. As shown in the heatmap in Figure 1, numerous pairs exhibit strong positive or negative correlations, with absolute values  $> 0.8$ . For example,  $L_0$  is highly negatively correlated with  $L_3$  ( $-0.96$ ), while  $L_2$  is strongly positively correlated with  $L_4$  (0.89) and  $L_8$  (0.88). These strong correlations indicate that the 10 latent variables are not independent. Instead, they covary in a highly structured manner across the spatial-temporal domain. This is visually confirmed by pairwise scatter plots, such as the positive correlation between  $L_0$  and  $L_1$  shown in Figure 2, the moderate negative correlation between  $L_0$  and  $L_4$  depicted in Figure 3, and the strong negative correlation between  $L_0$  and  $L_3$  illustrated in Figure 4. The presence of strong correlations suggests that the effective dimensionality of the manifold traced by  $L(x, t)$  within  $\mathbb{R}^{10}$  may be lower than 10, or that the learned representation exploits coordinated changes across multiple latent dimensions to encode features of the Burger’s solution. The presence of strong negative correlations implies antagonistic relationships or balancing effects between certain latent dimensions.

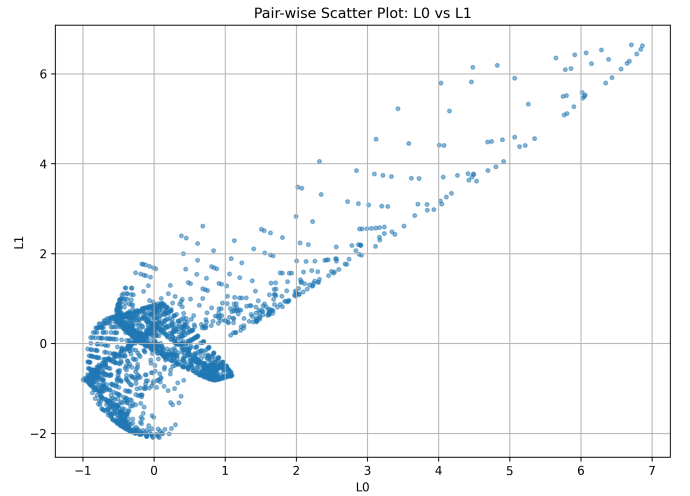
### 3.2. Numerical derivatives of the latent space

To understand how the latent space varies locally in space and time, the first-order spatial gradient  $\mathbf{V}_x = \partial L / \partial x$ , the first-order temporal gradient  $\mathbf{V}_t = \partial L / \partial t$ , and the second-order spatial gradient  $\mathbf{V}_{xx} = \partial^2 L / \partial x^2$  were computed using finite difference approximations on the grid.

Summary statistics for the components of  $\mathbf{V}_x$ ,  $\mathbf{V}_t$ , and  $\mathbf{V}_{xx}$  across the grid reveal distinct characteristics. The components of  $\mathbf{V}_x$  generally have means close to zero (though some are slightly positive) and standard deviations comparable to those of the latent variables themselves. The components of  $\mathbf{V}_t$ , however, have standard deviations and ranges that are notably smaller than those of  $\mathbf{V}_x$  components. This suggests that the rate of change of the latent representation with respect to time is typically less extreme than its rate of change with

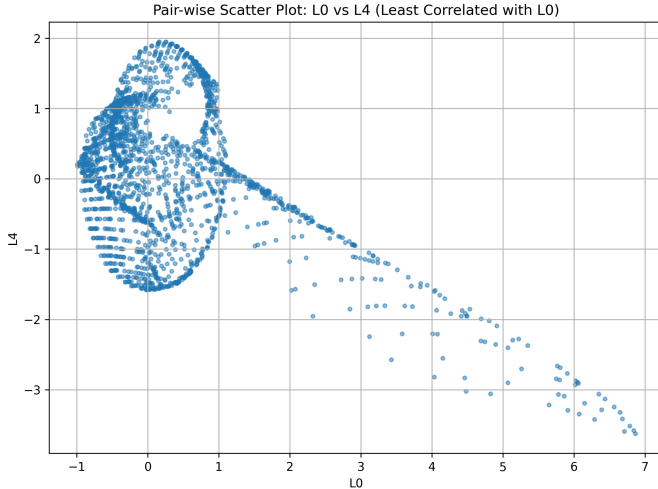


**Figure 1.** Pearson correlation matrix heatmap of the 10 latent dimensions learned by the PINN for the 2D Burger’s equation. The heatmap reveals significant interdependencies between latent dimensions, with many pairs exhibiting strong positive or negative correlations (absolute values  $> 0.8$ ), indicating a highly structured latent space and suggesting a potentially lower effective dimensionality.

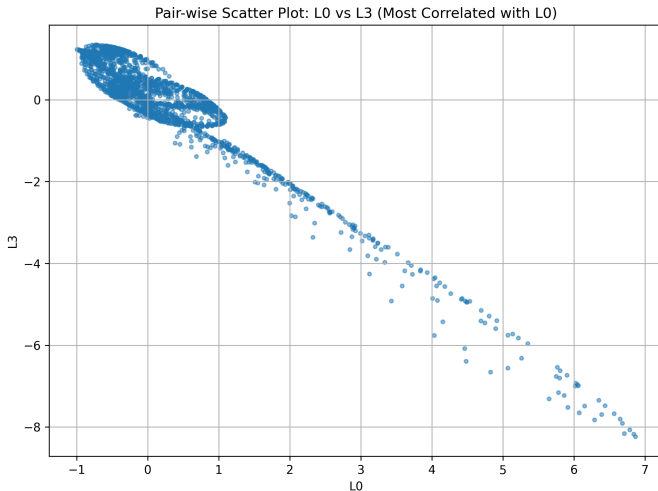


**Figure 2.** Scatter plot of the latent dimension  $L_0$  versus  $L_1$  across the  $(x, t)$  grid. The structured distribution of points visually confirms the positive correlation between these two latent dimensions, indicating a learned relationship within the latent space.

respect to space, when measured along the coordinate axes. In contrast, the components of the second-order spatial derivative  $\mathbf{V}_{xx}$  exhibit much larger standard deviations and significantly wider ranges compared to  $L$ ,  $\mathbf{V}_x$ , and  $\mathbf{V}_t$ . The maximum absolute values of  $\mathbf{V}_{xx}$  components can be very large (e.g., exceeding 100), indicat-



**Figure 3.** Scatter plot of latent dimension  $L_0$  versus  $L_4$  for all  $(x, t)$  grid points, showing the relationship between two dimensions of the PINN’s learned latent space. This plot visually confirms the moderate negative correlation (Pearson correlation coefficient  $\approx -0.45$ ) between  $L_0$  and  $L_4$ , revealing a structured distribution of points rather than random scatter.



**Figure 4.** Pair-wise scatter plot of the latent dimensions  $L_0$  versus  $L_3$ . The data points represent the latent space values  $L(x, t)$  across the entire spatial and temporal grid. This plot visualizes the strong negative correlation (Pearson correlation coefficient  $-0.96$ ) identified between these two latent dimensions, showing a clear linear relationship and indicating a significant dependency in the PINN’s learned representation.

ing regions of very high spatial curvature in the latent space. This is consistent with the PINN learning to represent sharp features or rapid transitions in the solution, such as those associated with shock formation or steep

gradients in Burger’s equation, which would manifest as high curvature in the learned embedding.

### 3.3. Geometric analysis of the latent manifold

The vectors  $L(x, t)$ ,  $\mathbf{V}_x(x, t)$ , and  $\mathbf{V}_t(x, t)$  define the local state and tangent vectors on the 10-dimensional manifold traced by the latent space embedding. We analyzed their magnitudes and relative orientations across the  $(x, t)$  domain.

#### 3.3.1. Vector magnitudes

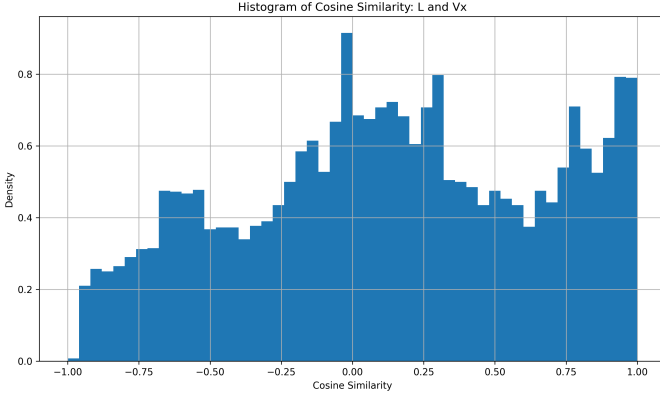
The Euclidean magnitudes  $|L|$ ,  $|\mathbf{V}_x|$ , and  $|\mathbf{V}_t|$  were computed at each grid point. The mean magnitudes are approximately 2.82 for  $|L|$ , 3.52 for  $|\mathbf{V}_x|$ , and 0.93 for  $|\mathbf{V}_t|$ . This confirms the observation from the component statistics: the magnitude of the spatial gradient vector  $|\mathbf{V}_x|$  is, on average, larger than the magnitude of the latent state vector  $|L|$ , and significantly larger than the magnitude of the temporal gradient vector  $|\mathbf{V}_t|$ . This implies that movement along the spatial dimension in the physical domain generally corresponds to a larger displacement in the latent space than movement along the temporal dimension, relative to the current location or unit coordinate change.

#### 3.3.2. Vector alignments

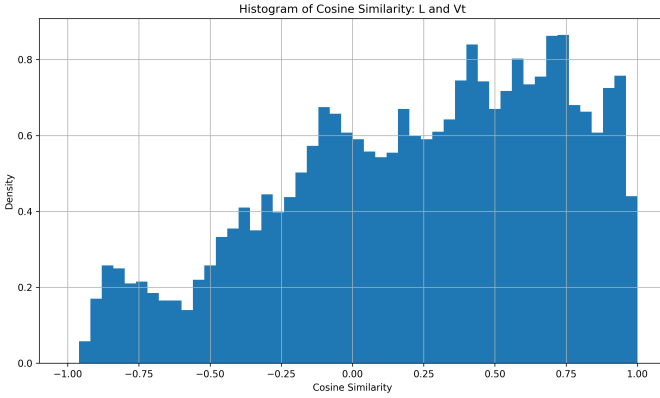
The cosine similarity provides a measure of the angle between pairs of vectors. The distributions of cosine similarities were analyzed. The cosine similarity between  $L$  and  $\mathbf{V}_x$  has a mean near zero (0.125) and, as shown in Figure 5, a broad distribution spanning  $[-1, 1]$ . This broad distribution indicates no strong general alignment between the latent state vector  $L$  and its spatial gradient  $\mathbf{V}_x$  across the entire domain.

Similarly, the cosine similarity between  $L$  and  $\mathbf{V}_t$  has a slightly larger mean (0.237) and a broad distribution, skewed slightly towards positive values, as illustrated in Figure 6. These distributions suggest that there is no strong global preference for the latent state vector  $L$  to be aligned, anti-aligned, or orthogonal to its spatial or temporal gradients  $\mathbf{V}_x$  and  $\mathbf{V}_t$  across the entire domain.

In contrast, the cosine similarity between  $\mathbf{V}_x$  and  $\mathbf{V}_t$  exhibits a mean of  $-0.193$  and a median of  $-0.496$ . The distribution, shown in Figure 7, is broad but strongly skewed towards negative values, with significant density around  $-1$ . This indicates a notable tendency for the spatial and temporal gradient vectors  $\mathbf{V}_x$  and  $\mathbf{V}_t$  to be anti-aligned or have substantial components in opposite directions. Geometrically, this means that the direction of maximal change in the latent space with respect to space is often opposed to the direction of maximal change with respect to time. This geometric feature might reflect the propagation of features in the Burger’s



**Figure 5.** Histogram of cosine similarity between the latent space vector  $L(x, t)$  and its spatial derivative  $\mathbf{V}_x(x, t)$  across the learned spatio-temporal domain. The broad distribution spanning  $[-1, 1]$  indicates no strong general alignment between the latent state and its spatial gradient.



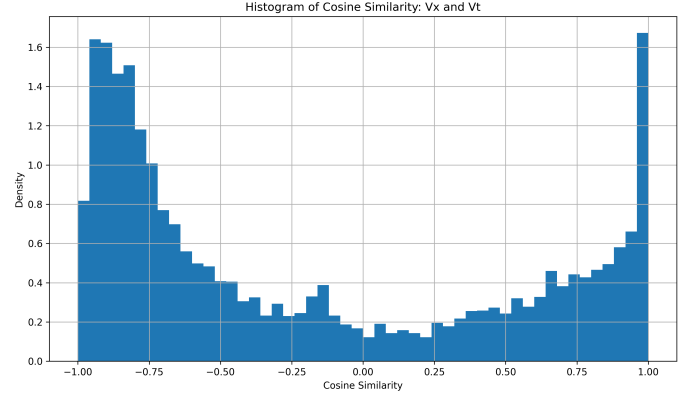
**Figure 6.** Histogram showing the distribution of cosine similarity values between the latent space vector  $L(x, t)$  and its temporal derivative  $\mathbf{V}_t(x, t)$  across the  $(x, t)$  grid. The distribution is broad but skewed towards positive values, indicating a weak tendency for the temporal gradient  $\mathbf{V}_t$  to point in the same general direction as  $L$ .

solution; as a spatial feature moves in time, the change in latent space due to spatial shift ( $\mathbf{V}_x$ ) might be counteracted by the change due to temporal evolution ( $\mathbf{V}_t$ ).

### 3.3.3. Local tangent space dimensionality

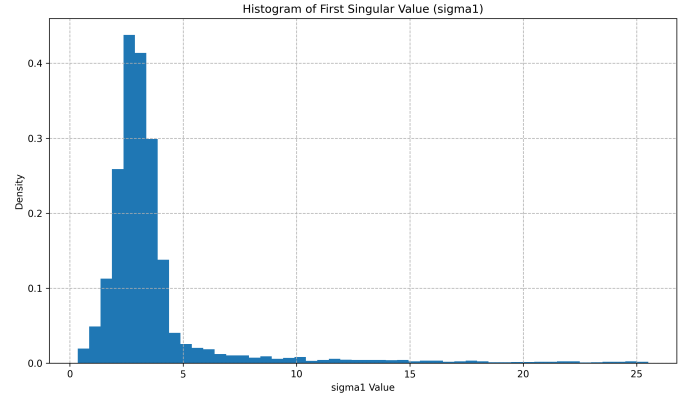
At each point, the local tangent space spanned by  $\mathbf{V}_x$  and  $\mathbf{V}_t$  was analyzed using Singular Value Decomposition (SVD) of the matrix  $[\mathbf{V}_x | \mathbf{V}_t]$ . The singular values  $\sigma_1 \geq \sigma_2$  quantify the extent of variation in the two principal directions within this local tangent plane. The ratio  $\sigma_2/\sigma_1$  indicates how close this plane is to being one-dimensional.

The mean  $\sigma_1$  is 3.71, reflecting the overall scale of variation in the dominant tangent direction, which is



**Figure 7.** Histogram of the cosine similarity between the latent space spatial derivative vector ( $\mathbf{V}_x$ ) and temporal derivative vector ( $\mathbf{V}_t$ ). The distribution is skewed towards negative values, particularly around  $-1$ , indicating that spatial and temporal changes in the latent space are frequently anti-aligned.

comparable to  $|\mathbf{V}_x|$  and  $|L|$ . The distribution of  $\sigma_1$  is shown in Figure 8.




**Figure 8.** Histogram of the first singular value ( $\sigma_1$ ) of the local latent manifold tangent space, derived from the spatial ( $\mathbf{V}_x$ ) and temporal ( $\mathbf{V}_t$ ) derivatives. The distribution shows a wide range of values, indicating varying magnitudes of variation along the dominant tangent direction across the domain.

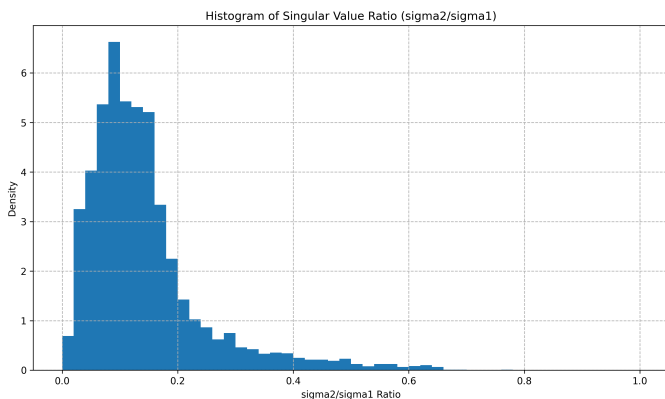
The mean  $\sigma_2$  is 0.45, significantly smaller than  $\sigma_1$ . The distribution of  $\sigma_2$ , presented on a log scale in Figure 9, shows a concentration of values at the lower end.

The ratio  $\sigma_2/\sigma_1$  has a mean of 0.14 and a median of 0.12. As depicted in Figure 10, the distribution of this ratio is heavily concentrated at small values, mostly below 0.3.

This consistently low ratio of singular values indicates that the local tangent space spanned by  $\mathbf{V}_x$  and  $\mathbf{V}_t$  is frequently nearly one-dimensional. While the latent space is 10-dimensional, the surface traced by  $L(x, t)$



**Figure 9.** Histogram of the second singular value ( $\sigma_2$ ) of the matrix  $[\mathbf{V}_x|\mathbf{V}_t]$ , formed by the spatial and temporal derivatives of the latent space  $L(x, t)$ , shown on a log scale. The concentration of values at the lower end of the distribution indicates that  $\sigma_2$  is often small, suggesting the local tangent space spanned by  $\mathbf{V}_x$  and  $\mathbf{V}_t$  is frequently close to one-dimensional.



**Figure 10.** Histogram of the ratio of the second to first singular value ( $\sigma_2/\sigma_1$ ) of the matrix  $[\mathbf{V}_x|\mathbf{V}_t]$ , representing the local dimensionality of the tangent space spanned by the spatial and temporal derivatives of the latent variables. The distribution is heavily skewed towards small values, indicating that the latent manifold is often locally nearly 1-dimensional.

is locally highly constrained, resembling a thin ribbon or curve rather than a flat plane. This suggests that spatial and temporal variations in the latent representation are not independent; they are strongly coupled and predominantly restricted to varying along a single dominant direction locally.

### 3.4. Learned latent dynamics via sparse regression

A central objective was to discover a system of differential equations governing the evolution of the latent space, specifically seeking equations of the form  $\partial L_k/\partial t = f_k(L, \mathbf{V}_x, \mathbf{V}_{xx})$ . Sparse regression was ap-

plied to identify the terms and coefficients in these equations from the computed data.

The temporal derivative  $\mathbf{V}_t$  for each latent dimension  $L_k$  was regressed against a library of candidate functions of  $L$ ,  $\mathbf{V}_x$ , and  $\mathbf{V}_{xx}$  evaluated at each grid point. The library included a constant term, linear terms ( $L_j$ ,  $Vx_j$ ,  $Vxx_j$ ), and quadratic terms ( $L_j^2$ ,  $L_j Vx_j$ ,  $Vx_j^2$ ) for  $j = 0, \dots, 9$ . Lasso regression with  $\alpha = 0.01$  was used to promote sparsity among the candidate terms.

The sparse regression models achieved high accuracy in predicting the temporal derivatives  $\partial L_k/\partial t$ . The  $R^2$  scores for the 10 latent dimensions range from approximately 0.81 to 0.93, indicating that the identified equations capture a large fraction of the variance in the latent space dynamics. The models are moderately sparse, with the number of identified non-zero terms for each equation ranging from 16 to 24 out of 61 total candidate terms, corresponding to sparsity levels between 26% and 39%. The chosen  $\alpha$  value balances fit accuracy and model simplicity; tuning  $\alpha$  could yield sparser or denser models.

Analysis of the identified equations reveals a consistent structure across the latent dimensions. Each equation for  $\partial L_k/\partial t$  is a sum of a constant term, linear terms involving  $L_j$ ,  $\partial L_j/\partial x$ , and  $\partial^2 L_j/\partial x^2$  for various  $j$ , and quadratic terms such as  $L_j^2$ ,  $L_j(\partial L_j/\partial x)$ , and  $(\partial L_j/\partial x)^2$  for various  $j$ . The equations are coupled, as the dynamics of  $L_k$  depend on other latent variables  $L_j$  and their derivatives.

Crucially, the sparse regression consistently identifies terms of the form  $L_j(\partial L_j/\partial x)$  (i.e.,  $L_j Vx_j$ ) and  $\partial^2 L_j/\partial x^2$  (i.e.,  $Vxx_j$ ). These terms are structurally analogous to the nonlinear advection term  $u(\partial u/\partial x)$  and the diffusion term  $\nu(\partial^2 u/\partial x^2)$  in the original Burger’s equation, respectively. For example, the equation for  $\partial L_0/\partial t$  includes terms like  $-0.18L_3Vx_3$  and  $+0.029Vxx_2$ . The presence and significance of these terms strongly suggest that the PINN has learned to represent the fundamental physical processes of nonlinear advection and diffusion, which govern the Burger’s equation, within its internal latent dynamics. The learned dynamics are not merely a complex black-box mapping but appear to encode the structure of the governing PDE.

### 3.5. Synthesis of findings

The combined statistical, geometric, and dynamic analysis provides a comprehensive mathematical interpretation of the PINN’s learned latent space for the Burger’s equation. The latent variables are not independent but form a correlated, heterogeneous representation, as evidenced by the correlation analysis (e.g.,

Figure 1) and scatter plots (e.g., Figures 2, 3, 4). This representation traces a manifold in  $\mathbb{R}^{10}$  with distinct geometric properties: spatial variations are generally more significant than temporal ones (based on derivative magnitudes), and the directions of maximum spatial and temporal change are often opposed (Figure 7). The manifold is locally highly constrained, as indicated by the low ratio of singular values of the tangent space (Figure 10), suggesting that spatial and temporal variations in the latent representation are tightly coupled and occur primarily along a dominant local direction.

Most importantly, the temporal evolution of this structured latent space is accurately described by a system of learned nonlinear partial differential equations identified via sparse regression. These latent PDEs contain terms that are direct structural analogues of the advection and diffusion operators from the original Burger’s equation. This provides compelling evidence that the PINN has not only learned to approximate the solution  $u(x,t)$  but has also encoded key physical principles of the underlying PDE within the dynamics of its internal, high-dimensional representation. This moves towards demystifying the PINN’s internal workings, demonstrating that they can learn interpretable mathematical structures related to the physics they are trained on.

#### 4. CONCLUSIONS

This study addressed the significant challenge of interpreting the internal representations learned by Physics-Informed Neural Networks (PINNs), moving beyond their conventional treatment as black-box function approximators. Specifically, we focused on providing a mathematical interpretation of the 10-dimensional latent space learned by a PINN trained to solve the 2D Burger’s equation. Our approach involved a comprehensive analysis of the geometric structure and learned dynamics of this latent space by examining the latent variables themselves and their spatial and temporal gradients.

Using a dataset of the learned latent space over a  $100 \times 100$  spatial-temporal grid, we employed numerical finite differences to compute the spatial, temporal, and second-order spatial derivatives of the latent variables. We then performed statistical analysis of the latent variables and their derivative components, analyzed the magnitudes and cosine similarities of the latent state vector  $L$  and its gradient vectors  $\mathbf{V}_x$  and  $\mathbf{V}_t$ , and assessed the local dimensionality of the tangent space

spanned by  $\mathbf{V}_x$  and  $\mathbf{V}_t$  using Singular Value Decomposition. Finally, we applied sparse regression techniques to discover a system of differential equations governing the temporal evolution of the latent space, based on its state and spatial derivatives.

Our results revealed that the learned latent space possesses a rich, interpretable mathematical structure. The 10 latent variables exhibit heterogeneous statistics and significant pairwise correlations, indicating a non-independent, coordinated representation. Geometrically, the latent manifold embedded in  $\mathbb{R}^{10}$  shows distinct characteristics: spatial gradient magnitudes  $|\mathbf{V}_x|$  are typically larger than temporal gradient magnitudes  $|\mathbf{V}_t|$ , and the spatial and temporal gradient vectors  $\mathbf{V}_x$  and  $\mathbf{V}_t$  show a strong tendency towards anti-alignment across the domain. Furthermore, the local tangent space spanned by these gradients is frequently found to be nearly one-dimensional, suggesting a strong constraint on the simultaneous variation of the latent representation in space and time.

Most significantly, the sparse regression analysis successfully identified a coupled system of nonlinear partial differential equations that accurately describe the temporal dynamics of the latent space ( $\partial L_k / \partial t$ ). These learned latent PDEs achieve high  $R^2$  scores in predicting the temporal derivatives and are moderately sparse. Crucially, the identified equations consistently contain terms structurally analogous to the nonlinear advection ( $L_j \mathbf{V}_{x,j}$ ) and diffusion ( $\mathbf{V}_{xx,j}$ ) operators present in the original Burger’s equation.

From these results, we have learned that PINNs do not merely approximate solutions but can encode key physical principles of the governing PDE within the mathematical structure and dynamics of their learned internal representations. The latent space is not an amorphous high-dimensional embedding but a structured manifold whose evolution is governed by interpretable differential relationships. The presence of advection-like and diffusion-like terms in the learned latent dynamics provides compelling evidence that the PINN has internalized core physical mechanisms necessary to solve the PDE. This work demonstrates that mathematical analysis of the latent space, particularly focusing on its gradients and learned dynamics, offers a powerful formalism for demystifying PINNs and extracting deeper insights into how neural networks learn and represent physical systems described by differential equations. Such interpretability is vital for building trust in AI-driven scientific models, diagnosing their failures, and potentially leveraging the learned representations for scientific discovery or model improvement.

## REFERENCES

- Alzahrani, A. Y., Haroon, A., Elsanhoury, W., & Çınar, D. C. 2025, An In-depth Analysis of Photometric and Kinematic Characteristics of SAI 16, SAI 81 and SAI 86 Open Clusters Utilizing Gaia DR3. <https://arxiv.org/abs/2504.15344>
- Andrianomena, S., & Hassan, S. 2023, Latent space representations of cosmological fields. <https://arxiv.org/abs/2311.00799>
- Andrianomena, S., & Tang, H. 2023, Radio Galaxy Zoo: Leveraging latent space representations from variational autoencoder. <https://arxiv.org/abs/2311.08331>
- Balsara, D., Bhoriya, D., Zanotti, O., & Dumbser, M. 2024, Well-balanced high order finite difference WENO schemes for a first-order Z4 formulation of the Einstein field equations, doi: <https://doi.org/10.3847/1538-4365/ad7d0d>
- Cadiou, C., Laigle, C., & Agertz, O. 2025, How complex are galaxies? A non-parametric estimation of the intrinsic dimensionality of wide-band photometric data. <https://arxiv.org/abs/2404.02962>
- Caunt, S. E., & Korpi, M. J. 2001, A 3D MHD model of astrophysical flows: algorithms, tests and parallelisation, doi: <https://doi.org/10.1051/0004-6361:20010157>
- Chiodaroli, E., & Feireisl, E. 2024, On the long time behaviour of solutions to the Navier-Stokes-Fourier system on unbounded domains. <https://arxiv.org/abs/2406.09587>
- Clemente, F. D., Scialpi, M., & Bejger, M. 2025, Explainable autoencoder for neutron star dense matter parameter estimation, doi: <https://doi.org/10.1088/2632-2153/add3bd>
- Colgan, R. E., Yan, J., Márka, Z., et al. 2022, Architectural Optimization and Feature Learning for High-Dimensional Time Series Datasets, doi: <https://doi.org/10.1103/PhysRevD.107.022009>
- Creevey, O. L., Brown, T. M., Jiménez-Reyes, S., & Belmonte, J. A. 2006, Understanding the Relationship Between Observations and Stellar Parameters in an Eclipsing Binary System. <https://arxiv.org/abs/astro-ph/0605177>
- Daglayan, H., Vary, S., Leplat, V., Gillis, N., & Absil, P. A. 2023, Direct Exoplanet Detection Using L1 Norm Low-Rank Approximation. <https://arxiv.org/abs/2304.03619>
- Daglayan, H., Vary, S., Absil, O., et al. 2024, An Alternating Minimization Algorithm with Trajectory for Direct Exoplanet Detection – The AMAT Algorithm, doi: <https://doi.org/10.1051/0004-6361/202451242>
- Dahlbüdding, D., Molaverdikhani, K., Ercolano, B., & Grassi, T. 2024, Approximating Rayleigh Scattering in Exoplanetary Atmospheres using Physics-informed Neural Networks (PINNs). <https://arxiv.org/abs/2408.00084>
- Fagin, J., Park, J. W., Best, H., et al. 2024, Latent Stochastic Differential Equations for Modeling Quasar Variability and Inferring Black Hole Properties, doi: <https://doi.org/10.3847/1538-4357/ad2988>
- Haddad, F., & Haddad, N. 2021, Solar System, Astrophysics, and Cosmology from the Derivative Expansion, doi: <https://doi.org/10.1007/s10714-021-02849-5>
- Hao, B., Braga-Neto, U., Liu, C., Wang, L., & Zhong, M. 2025, Training PINNs with Hard Constraints and Adaptive Weights: An Ablation Study. <https://arxiv.org/abs/2404.16189>
- Hatipoğlu, G. 2022, Utility of PCA and Other Data Transformation Techniques in Exoplanet Research. <https://arxiv.org/abs/2211.14683>
- . 2023, PCA-based Data Reduction and Signal Separation Techniques for James-Webb Space Telescope Data Processing. <https://arxiv.org/abs/2301.00415>
- Horta, D., Price-Whelan, A. M., Hogg, D. W., Ness, M. K., & Casey, A. R. 2025, *Lux*: A generative, multi-output, latent-variable model for astronomical data with noisy labels. <https://arxiv.org/abs/2502.01745>
- Huijse, P., Ridder, J. D., Eyer, L., et al. 2025, Learning novel representations of variable sources from multi-modal *Gaia* data via autoencoders. <https://arxiv.org/abs/2505.16320>
- Iwasaki, D., Cooray, S., & Takeuchi, T. T. 2023, Extracting an Informative Latent Representation of High-Dimensional Galaxy Spectra. <https://arxiv.org/abs/2311.17414>
- Kuiper, D., Contardo, G., Huppenkothen, D., & Hessels, J. W. T. 2025, Representation learning for fast radio burst dynamic spectra. <https://arxiv.org/abs/2412.12394>
- Lalchand, V., & Eilers, A.-C. 2025, Shared Stochastic Gaussian Process Latent Variable Models: A Multi-modal Generative Model for Quasar Spectra. <https://arxiv.org/abs/2502.19824>
- Laroche, A., & Speagle, J. S. 2025, Closing the stellar labels gap: Stellar label independent evidence for  $[\alpha/M]$  information in Gaia BP/RP spectra, doi: <https://doi.org/10.3847/1538-4357/ad9607>

- Lue, A., Genel, S., Huertas-Company, M., Villaescusa-Navarro, F., & Ho, M. 2025, Cosmology with One Galaxy: Auto-Encoding the Galaxy Properties Manifold. <https://arxiv.org/abs/2502.17568>
- Mattia, G., Crocco, D., Fuksman, D. M., et al. 2025, PyPLUTO: a data analysis Python package for the PLUTO code. <https://arxiv.org/abs/2501.09748>
- McDuffee, R. 2023, Deriving the properties of space time using the non-compressible solutions of the Navier-Stokes Equations. <https://arxiv.org/abs/1607.05400>
- Minoshima, T., Miyoshi, T., & Matsumoto, Y. 2019, A high-order weighted finite difference scheme with a multi-state approximate Riemann solver for divergence-free magnetohydrodynamic simulations, doi: <https://doi.org/10.3847/1538-4365/ab1a36>
- Pan, L., Padoan, P., & Åke Nordlund. 2019, Inaccuracy of Spatial Derivatives in Simulations of Supersonic Turbulence, doi: <https://doi.org/10.3847/1538-4357/ab16de>
- Pandya, S., Patel, P., Nord, B. D., Walmsley, M., & Ćiprijanović, A. 2025, SIDDA: Sinkhorn Dynamic Domain Adaptation for Image Classification with Equivariant Neural Networks. <https://arxiv.org/abs/2501.14048>
- Parsotan, T., Palmer, D. M., Ronchini, S., et al. 2025, BatAnalysis – A Comprehensive Python Pipeline for Swift BAT Time-Tagged Event Data Analysis. <https://arxiv.org/abs/2502.00278>
- Piras, D., Herold, L., Lucie-Smith, L., & Komatsu, E. 2025,  $\Lambda$ CDM and early dark energy in latent space: a data-driven parametrization of the CMB temperature power spectrum, doi: <https://doi.org/10.1103/PhysRevD.111.083537>
- Price, D. C. 2024, Reduced-resolution beamforming: lowering the computational cost for pulsar and technosignature surveys. <https://arxiv.org/abs/2402.12723>
- Rucinski, S. 1998, Determination of broadening functions using the Singular Value Decomposition (SVD) technique. <https://arxiv.org/abs/astro-ph/9807327>
- Sarkar, S., & Petersen, M. S. 2025, commensurability: a Python package for classifying astronomical orbits based on their toroid volume, doi: <https://doi.org/10.21105/joss.07009>
- Skowronski, J., Piatti, D., Rapagnani, D., et al. 2025, A Comparative Analysis of R-Matrix Fitting:  $^{12}\text{C}(p, \gamma)^{13}\text{N}$  as Test Case. <https://arxiv.org/abs/2502.14040>
- Song, Y., Villar, V. A., Martinez-Galarza, J. R., & Dillmann, S. 2025, A Poisson Process AutoDecoder for X-ray Sources. <https://arxiv.org/abs/2502.01627>
- Tabatabaei, J., Banihashemi, A., Baghran, S., & Mashhoon, B. 2023, Anisotropic Cosmology in the Local Limit of Nonlocal Gravity, doi: <https://doi.org/10.3390/universe9090377>
- Taruya, A., Nishimichi, T., & Jeong, D. 2021, Grid-based calculations of redshift-space matter fluctuations from perturbation theory: UV sensitivity and convergence at the field level, doi: <https://doi.org/10.1103/PhysRevD.105.103507>
- Trusov, S., Zarrouk, P., Cole, S., et al. 2024, The 2-point correlation function covariance with fewer mocks, doi: <https://doi.org/10.1093/mnras/stad3710>
- Turk, M. J., & Smith, B. D. 2011, High-Performance Astrophysical Simulations and Analysis with Python. <https://arxiv.org/abs/1112.4482>
- Wang, H., Sreejith, S., Lin, Y., et al. 2023, Neural Network Based Point Spread Function Deconvolution For Astronomical Applications, doi: <https://doi.org/10.21105/astro.2210.01666>
- Wohlberg, B., & Wozniak, P. 2021, PSF Estimation in Crowded Astronomical Imagery as a Convolutional Dictionary Learning Problem, doi: <https://doi.org/10.1109/LSP.2021.3050706>
- Xiao, H., Dai, L., & McQuinn, M. 2024, Detecting Dark Matter Substructures on Small Scales with Fast Radio Bursts, doi: <https://doi.org/10.1103/PhysRevD.110.023516>
- Zuo, S., Chen, X., & Mao, Y. 2023, A Semi-blind PCA-based Foreground Subtraction Method for 21 cm Intensity Mapping, doi: <https://doi.org/10.3847/1538-4357/acb822>