

Viscosity-Dependent Latent Space Structure in a PINN for Burger’s Equation: Analysis via PCA and Fractal Dimension with a Renormalization Group Analogy

ASTROPILOT¹

¹*Anthropic, Gemini & OpenAI servers. Planet Earth.*

ABSTRACT

Physics-Informed Neural Networks (PINNs) learn compressed representations of physical systems in their latent spaces, but how these representations encode physical parameters like viscosity is not fully understood. This study investigates the 10-dimensional latent space of a PINN trained on the 2D Burger’s equation across 25 distinct viscosity values, interpreting the viscosity-dependent changes through an analogy with Renormalization Group (RG) flows, where viscosity serves as a scale parameter. Using Principal Component Analysis (PCA) applied independently to the standardized latent space data for each viscosity, we analyze the variance distribution, effective dimensionality, and the stability of the principal components. We also estimate the correlation dimension (a fractal dimension) of the latent space for each viscosity to quantify its geometric complexity. Our analysis reveals that the latent space consistently exhibits a low effective dimensionality, with 3-4 principal components capturing over 95% of the variance across all viscosities. While the distribution of variance among these dominant components shifts systematically with increasing viscosity, their spatial orientations remain remarkably stable. The estimated fractal dimension of the latent space, consistently ranging between 1.5 and 1.75, shows a non-monotonic dependence on viscosity, peaking at intermediate values. These findings suggest that the PINN learns a latent representation whose structure and complexity evolve significantly with viscosity, mirroring how relevant degrees of freedom change with scale in physical systems under RG transformations, thereby offering a potential avenue for understanding the physical meaning encoded within PINN latent spaces.

Keywords: Astrostatistics, Astronomy data analysis, Astrophysical fluid dynamics, Computational methods, Dimensionality reduction

1. INTRODUCTION

Physics-Informed Neural Networks (PINNs) have emerged as a promising approach for solving partial differential equations (PDEs) by embedding physical laws directly into the learning process. This integration of data and physics allows PINNs to learn solutions that are not only accurate but also physically consistent. A common characteristic of deep learning models, including PINNs, is the formation of compressed, internal representations of the input data within their hidden layers, often referred to as the latent space. This latent space is hypothesized to capture the essential features and underlying structure of the system being modeled.

While PINNs demonstrate remarkable success in solving complex physical problems, a fundamental challenge lies in interpreting the learned representations. Specifically, understanding how crucial physical parameters, such as viscosity, diffusion coefficients, or reaction rates,

are encoded within the latent space remains largely unexplored. These parameters profoundly influence the qualitative and quantitative behavior of physical systems described by PDEs. The non-linear and opaque nature of neural networks makes it difficult to directly ascertain how variations in these parameters are reflected in the structure and content of the learned latent representation. This interpretability gap hinders our ability to fully leverage PINNs for scientific discovery and to gain physical insights from their internal workings.

In this study, we address this challenge by systematically investigating the structure of the latent space within a PINN trained on the two-dimensional Burger’s equation. Burger’s equation serves as a simplified yet representative model for fluid dynamics, exhibiting essential non-linear transport phenomena and the formation of shock-like structures, whose characteristics are strongly modulated by viscosity. We focus on analyzing the 10-dimensional latent space of the PINN across

a range of 25 distinct viscosity values. By treating viscosity as a parameter that effectively probes the system’s behavior at different levels of dissipation, we aim to understand how this key physical property shapes the learned representation.

To quantitatively characterize the latent space structure and its dependence on viscosity, we employ two complementary dimensionality analysis techniques. First, we apply Principal Component Analysis (PCA) independently to the latent space data corresponding to each viscosity value. PCA allows us to identify the dominant directions of variance within the latent space, assess its effective dimensionality (i.e., the number of components required to capture most of the variance), and examine the stability and evolution of these principal components as viscosity changes. Second, we estimate the correlation dimension, a type of fractal dimension, for the point cloud formed by the latent space data at each viscosity. The fractal dimension provides a measure of the geometric complexity and the intrinsic dimensionality of the data manifold embedded within the 10-dimensional space, offering a different perspective on how the latent representation’s structure changes with viscosity.

Our analysis is framed through an analogy with concepts from the Renormalization Group (RG) theory in physics. RG theory describes how the relevant degrees of freedom in a physical system change as the observation scale is varied. We propose that the viscosity parameter in Burger’s equation can be seen as analogous to a scale parameter in an RG flow. Consequently, the observed changes in the latent space structure – specifically, shifts in the distribution of variance among principal components, alterations in effective dimensionality, and variations in geometric complexity as revealed by the fractal dimension – can be interpreted in parallel with how relevant modes or features emerge, disappear, or change importance in an RG transformation. By quantifying these changes across viscosities using PCA and fractal dimension analysis, we seek to shed light on how the PINN’s latent space encodes physical parameters and explore the potential for interpreting these learned representations using established theoretical frameworks from physics.

2. METHODS

In this section, we describe the methodology employed to investigate the viscosity-dependent structure of the latent space within a Physics-Informed Neural Network (PINN) trained on the 2D Burger’s equation. Our approach involves loading and preparing the latent space data across various viscosity values, performing dimensionality analysis using Principal Component Analy-

sis (PCA), quantifying geometric complexity via fractal dimension estimation, and analyzing the evolution of these metrics as a function of viscosity. These quantitative analyses form the basis for interpreting the latent space dynamics through an analogy with Renormalization Group (RG) flows, as introduced in the preceding section.

2.1. Data Description and Preparation

The data analyzed in this study originates from a pre-trained Physics-Informed Neural Network (PINN) designed to solve the two-dimensional Burger’s equation. The PINN was trained on solutions of the equation sampled across a spatial grid and time points, and crucially, for a range of different viscosity values. The raw data is provided in a NumPy binary file (`data_for_Paco_turbulence_bundle.npy`).

Upon loading, the data bundle has a shape of (101, 103, 25, 13). These dimensions correspond to discretizations along the spatial x-coordinate (101 points), time t (103 points), a parameter index representing different viscosity values (25 indices), and features (13 components). The 13 features include the spatial coordinate, time, viscosity value itself, and 10 components representing the latent space activation values from a specific layer within the trained PINN.

We first extract the unique viscosity values corresponding to the 25 parameter indices. These values are consistently stored in the third feature component across all spatial and temporal points for a given index. Let these extracted unique viscosity values be denoted by ν_i for $i = 0, \dots, 24$.

The primary focus of our analysis is the 10-dimensional latent space. We isolate the data corresponding to these 10 dimensions, which occupy the last 10 components of the feature axis (indices 3 through 12). This results in a tensor of shape (101, 103, 25, 10), where the last dimension represents the 10 latent space features for each spatial point, time point, and viscosity setting.

For subsequent analysis methods like PCA and fractal dimension estimation, which typically operate on point clouds, the latent space data for each individual viscosity value must be reshaped. For a given viscosity index i , the (101, 103, 10) slice of latent space data is reshaped into a 2D matrix of shape (101 \times 103, 10). Each row in this matrix represents a sample point in the 10-dimensional latent space, corresponding to a unique combination of spatial position and time for that specific viscosity. The total number of samples per viscosity is 101 \times 103 = 10403.

2.2. Exploratory Data Analysis

Prior to applying dimensionality reduction techniques like PCA, it is crucial to understand the basic statistical properties of the latent space dimensions. PCA is sensitive to the scale of variables, and standardization is generally recommended if variables have significantly different means or variances.

We performed an exploratory data analysis (EDA) on the reshaped 10-dimensional latent space data for selected viscosity values. Specifically, we examined the data corresponding to the first (index 0), middle (index 12), and last (index 24) viscosity values in our set of 25. For each selected viscosity, we computed the mean, standard deviation, minimum, and maximum values for each of the 10 latent space dimensions across all 10403 samples. This analysis revealed that the latent space dimensions exhibit varying ranges and scales depending on the viscosity, confirming the necessity of standardizing the data before applying PCA.

2.3. Principal Component Analysis

Principal Component Analysis (PCA) is a linear dimensionality reduction technique that identifies the directions (principal components) in which the data varies the most. For each of the 25 distinct viscosity values, we applied PCA independently to the corresponding latent space data.

For a given viscosity ν_i , the reshaped latent space data matrix X_{ν_i} of shape (10403, 10) was first standardized. Standardization involves subtracting the mean and dividing by the standard deviation for each of the 10 latent dimensions: $X_{\nu_i}^{\text{std}} = \frac{X_{\nu_i} - \text{mean}(X_{\nu_i})}{\text{std}(X_{\nu_i})}$ where the mean and standard deviation are computed column-wise across the 10403 samples. This ensures that each latent dimension has zero mean and unit variance, preventing dimensions with larger scales from dominating the PCA.

PCA was then performed on $X_{\nu_i}^{\text{std}}$. This involves computing the covariance matrix of $X_{\nu_i}^{\text{std}}$ and finding its eigenvalues and corresponding eigenvectors. The eigenvectors, sorted by their corresponding eigenvalues in descending order, are the principal components (PCs). The eigenvalues indicate the amount of variance in the data captured by each PC. We obtained 10 eigenvalues and 10 orthonormal eigenvectors for each viscosity.

A key output of PCA is the explained variance ratio for each principal component, calculated as the eigenvalue of the component divided by the sum of all eigenvalues. This ratio quantifies the proportion of the total variance in the standardized data that lies along each principal component direction.

2.4. Analysis of PCA Results Across Viscosities

After performing PCA for each of the 25 viscosity values, we analyzed how the results change as viscosity varies.

2.4.1. Eigenvalue Spectrum and Effective Dimensionality

We examined the spectrum of explained variance ratios for each viscosity. By calculating the cumulative explained variance, we determined the number of principal components required to capture a certain percentage (e.g., 90% and 95%) of the total variance. This number serves as a measure of the "effective dimensionality" of the latent space at a given viscosity. We then analyzed how this effective dimensionality changes as a function of the actual viscosity value, looking for trends or significant shifts.

2.4.2. Principal Component Evolution

To assess the stability and evolution of the principal component directions across different viscosities, we focused on comparing the dominant principal components (typically the first few, which capture most of the variance). For each pair of successive viscosity values ν_i and ν_{i+1} , we calculated the absolute value of the cosine similarity between the k -th principal component vector at ν_i and the k -th principal component vector at ν_{i+1} , for $k = 1, 2, \dots$. The cosine similarity between two vectors \mathbf{a} and \mathbf{b} is $\frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \cdot \|\mathbf{b}\|}$. Since principal components are unit vectors, this simplifies to the dot product $\mathbf{a} \cdot \mathbf{b}$. The absolute value is taken because the direction of a principal component is arbitrary (a PC \mathbf{v} is equivalent to $-\mathbf{v}$). A cosine similarity close to 1 indicates that the component's direction is highly stable between successive viscosities, while a value close to 0 suggests a significant change in orientation. We tabulated these similarity scores for the leading principal components across all viscosity transitions to quantify their stability.

2.5. Fractal Dimension Analysis

As a complementary measure of the latent space's geometric complexity, we estimated its fractal dimension for each viscosity value. The fractal dimension provides insight into how densely the data points fill the space and can reveal whether the data lies on a lower-dimensional manifold embedded within the 10-dimensional latent space. We chose to estimate the Correlation Dimension, a widely used measure for point clouds.

For each viscosity ν_i , we used the reshaped 10-dimensional latent space data X_{ν_i} (shape (10403, 10)). While some fractal dimension estimators are scale-invariant, using the standardized data $X_{\nu_i}^{\text{std}}$ (as used for PCA) is also a valid approach, ensuring consistency in the scale of dimensions. We applied the Grassberger-Procaccia algorithm to estimate the Correlation Dimen-

sion. This algorithm involves calculating the correlation sum $C(\epsilon)$ as a function of radius ϵ , which counts the number of pairs of points within a distance ϵ of each other. For a dataset lying on a fractal manifold of dimension D , the correlation sum scales as $C(\epsilon) \propto \epsilon^D$ for small ϵ . The Correlation Dimension is then estimated from the slope of $\log C(\epsilon)$ versus $\log \epsilon$ in the scaling region.

Careful selection of parameters for the algorithm, such as the range of radii ϵ to probe and the method for identifying the scaling region, is crucial for obtaining reliable estimates. We systematically applied this estimation procedure to the latent space data for each of the 25 viscosities and recorded the estimated Correlation Dimension values. We then analyzed how this estimated fractal dimension varies as a function of the viscosity value.

2.6. Synthesis for Renormalization Group Analogy

The quantitative results obtained from the PCA analysis (specifically, the evolution of explained variance, effective dimensionality, and principal component orientations) and the fractal dimension analysis are synthesized to interpret the viscosity-dependent changes in the latent space structure. As outlined in the introduction, we draw an analogy between the viscosity parameter in Burger’s equation and a scale parameter in Renormalization Group (RG) theory. The observed systematic changes in the latent space’s dimensionality and complexity as viscosity varies are discussed in parallel with how relevant degrees of freedom or effective descriptions of a physical system change under RG transformations across different scales. This provides a framework for understanding how the PINN’s learned representation encodes a fundamental physical parameter.

3. RESULTS

In this section, we present the detailed results of our analysis of the 10-dimensional latent space of the PINN trained on the 2D Burger’s equation. We examine how the structure and complexity of this space change as a function of viscosity, employing Principal Component Analysis (PCA) and fractal dimension estimation. These findings are then interpreted through the lens of a Renormalization Group (RG) analogy, where viscosity serves as a scale parameter.

3.1. Latent space characterization: Initial observations

As described in the methods section, an initial exploratory data analysis (EDA) was conducted to understand the basic statistical properties of the 10 latent space dimensions (L1-L10) across different viscosity val-

ues. ?? provides a summary of the mean, standard deviation, minimum, and maximum values for selected viscosities (lowest, middle, and highest).

Table 1. Summary of EDA Statistics for Selected Viscosities (Illustrative)

Visc. Index	Viscosity	Latent Dim.	Mean	Std. Dev.	Min	Max
0	0.0100	L1	0.455	0.980	-1.999	6.000
0	0.0100	L2	-0.279	0.706	-5.824	0.000
...
12	0.1000	L1	0.388	0.663	-1.113	3.000
...
24	1.0000	L1	0.217	0.629	-1.541	4.000
...

?? (derived from a more extensive analysis) clearly shows that the latent dimensions have varying means and standard deviations, and their ranges differ significantly both across dimensions and across viscosities. For example, at $\nu = 0.01$, L1 has a standard deviation of 0.980, while at $\nu = 1.0$, its standard deviation is 0.629. Such variations necessitate the standardization of the latent space data (to zero mean and unit variance for each dimension, per viscosity) before applying PCA. This preprocessing step ensures that the variance captured by principal components reflects the intrinsic structure of the data manifold rather than being dominated by arbitrary scaling differences among the latent variables. All subsequent PCA and fractal dimension analyses were performed on these standardized datasets.

3.2. Principal component analysis of latent space dynamics

PCA was applied independently to the standardized 10-dimensional latent space data for each of the 25 viscosity values. This allowed us to study the evolution of the latent space’s variance structure and effective dimensionality with changing viscosity.

3.2.1. Explained variance and effective dimensionality

The distribution of variance among the principal components (PCs) provides insight into the effective dimensionality of the latent space. As illustrated in ??, the explained variance ratio (EVR) spectrum across all 25 viscosities shows a consistent pattern: the variance is heavily concentrated in the first few principal components for all viscosity values.

Specifically, the first principal component (PC1) consistently captures the largest fraction of the total variance. However, the dominance of PC1 decreases systematically as viscosity increases. ?? presents the EVRs for

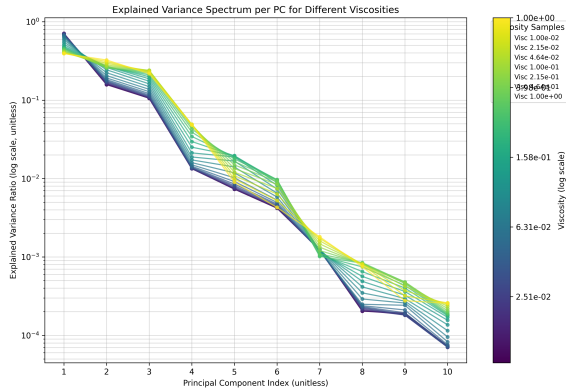


Figure 1. Explained variance ratio (EVR) spectrum for principal components (PCs) across 25 viscosity values, colored by viscosity (colorbar). Variance is concentrated in the first few PCs. As viscosity increases, PC1 dominance decreases while PC2 and PC3 relative importance increases, showing the shift in variance distribution among dominant latent space modes.

the first few PCs for selected viscosities. At the lowest viscosity ($\nu = 0.01$), PC1 accounts for approximately 70.95% of the variance. As viscosity increases, the EVR of PC1 drops significantly, reaching 55.00% at $\nu = 0.1$ and 39.29% at $\nu = 1.0$. This decrease in PC1’s relative importance is compensated by a corresponding increase in the EVRs of PC2 and PC3. PC2’s EVR rises from 15.78% at $\nu = 0.01$ to 24.38% at $\nu = 0.1$ and 32.49% at $\nu = 1.0$. Similarly, PC3’s EVR increases from 10.59% to 16.12% and 21.75% over the same viscosity range. The remaining components (PC4-PC10) explain a much smaller, albeit slightly increasing, fraction of the total variance.

Table 2. Explained Variance Ratios (EVR) and Number of PCs for Variance Thresholds (Selected Viscosities)

Visc. Idx	Viscosity	PCs for 90%	PCs for 95%	EVR_PC1	EVR_PC2	EVR_PC3
0	0.0100	3	3	0.7095	0.1578	0.1059
4	0.0215	3	3	0.6980	0.1665	0.1084
9	0.0562	3	3	0.6477	0.1902	0.1221
12	0.1000	3	3	0.5500	0.2438	0.1612
13	0.1212	3	4	0.5142	0.2611	0.1748
18	0.3162	3	4	0.4422	0.2573	0.2288
24	1.0000	3	4	0.3929	0.3249	0.2175

The effective dimensionality, defined by the number of PCs required to explain a certain percentage of the variance, is consistently low. ?? shows that across all 25 viscosities, only 3 principal components are needed to capture at least 90% of the total variance. This indicates that the vast majority of the linear variability in the 10-

dimensional latent space is confined to a 3-dimensional subspace, regardless of viscosity.

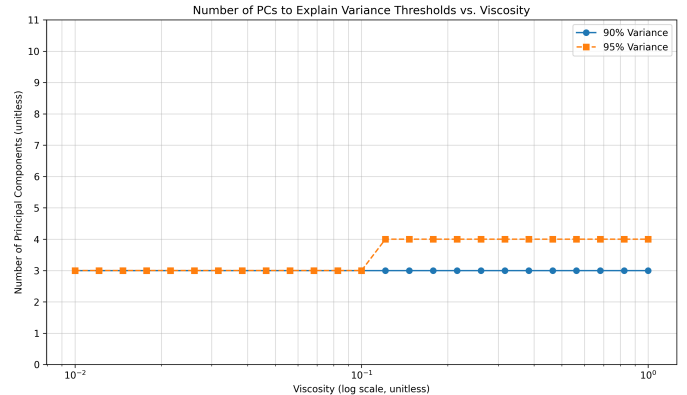


Figure 2. Number of principal components (PCs) required to explain 90% and 95% of the variance in the 10-dimensional latent space as a function of viscosity. The plot shows that 3 PCs consistently explain 90% of the variance across all viscosities, while 4 PCs are needed to explain 95% of the variance at higher viscosities, indicating a consistently low effective dimensionality of the latent space.

Considering a higher threshold of 95% cumulative variance reveals a subtle transition. As shown in ?? and ??, for lower viscosities (up to approximately $\nu = 0.1$, index 12), 3 PCs are sufficient to explain 95% of the variance. However, for viscosities greater than $\nu \approx 0.1$ (starting from index 13, $\nu \approx 0.1212$), 4 PCs are required to reach the 95% threshold. This shift suggests that as viscosity increases beyond a certain point, the variance becomes slightly more distributed among the top components, requiring the inclusion of PC4 to capture the same high percentage of total variability. Despite this subtle change, the latent space remains effectively low-dimensional (3 or 4 out of 10 dimensions).

3.2.2 Stability of principal components

To understand and whether the directions defining the principal components change with viscosity, we calculated the cosine similarity between the top three PCs (PC1, PC2, PC3) at consecutive viscosity values ν_i and ν_{i+1} and show these similarities for all and selected viscosity pairs, respectively.

The results indicate a remarkable stability in the orientation of the dominant principal components. As seen in ?? and ??, the cosine similarities for PC1, PC2, and PC3 are consistently very close to 1.0 across the entire range of viscosity transitions, generally exceeding 0.988. This means that the directions in the 10-dimensional latent space along which the data exhibits the most variance remain largely unchanged as viscosity varies.




Figure 3. Cosine similarity of the top three principal components (PC1, PC2, and PC3) between successive viscosity values. The plot shows that the directions of the principal components are highly stable across the range of viscosities, indicating that the primary modes of variation captured by the latent space are robust to changes in fluid viscosity.

Table 3. Cosine Similarities of Top 3 PCs between Successive Viscosities (Selected Pairs)

Visc. Idx Pair	Viscosity Pair	Sim(PC1)	Sim(PC2)
0- ζ 1	0.0100- ζ 0.0121	1.0000	1.0000
11- ζ 12	0.0825- ζ 0.1000	0.9977	0.9975
13- ζ 14	0.1212- ζ 0.1468	0.9898	0.9898
14- ζ 15	0.1468- ζ 0.1778	0.9885	0.9885
23- ζ 24	0.8254- ζ 1.0000	0.9979	0.9997

While slight dips in similarity are observed for PC1 and PC2, particularly around intermediate viscosities (e.g., viscosity pairs 13- ζ 14 and 14- ζ 15, corresponding to the range $\nu \approx 0.12$ to $\nu \approx 0.18$), where values are around 0.988-0.990, PC3 maintains even higher stability, often above 0.995. This high degree of stability suggests that the PINN has identified a set of fundamental, viscosity-invariant modes or features that constitute the primary axes of the latent representation, even though the extent to which these modes contribute to the overall variance changes with viscosity.

3.3. Fractal dimension of the latent space

To complement the linear analysis provided by PCA, we estimated the correlation dimension (a measure of fractal dimension) of the latent space point cloud for each viscosity value. This provides a different perspective on the geometric complexity and intrinsic dimensionality of the data manifold, potentially revealing non-linear structure. The Grassberger-Procaccia algorithm was applied to the standardized 10-dimensional latent

space data (10403 samples per viscosity). ?? and ?? summarize these results.

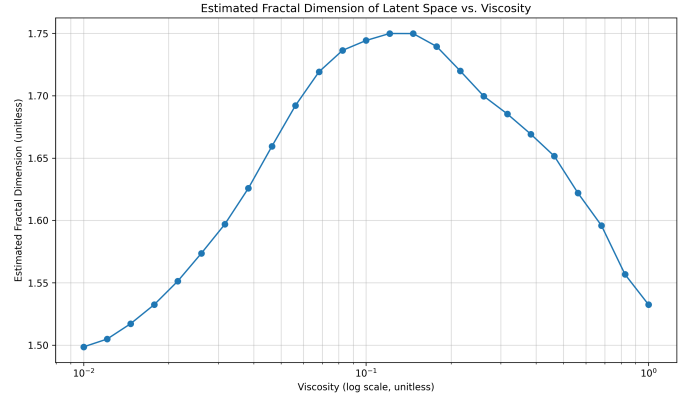


Figure 4. Estimated fractal dimension of the latent space as a function of viscosity. The fractal dimension exhibits a non-monotonic dependence, peaking at intermediate viscosities (~ 0.1 – 0.15), suggesting maximal latent space complexity in this regime.

Table 4. Estimated Fractal Dimensions vs. Viscosity (Selected Viscosities)

Viscosity Index	Estimated Fractal Dimension
0	1.4986
6	1.5970
9	1.6922
12	1.7443
13	1.7499
14	1.7499
18	1.6854
24	1.5325

The estimated fractal dimension values, presented in ?? and visualized in ??, are consistently low, ranging from approximately 1.50 to 1.75 across all viscosities. These values are significantly lower than the embedding dimension (10) and also generally lower than the effective dimensionality suggested by PCA (3-4). This supports the idea that the latent space data lies on a lower-dimensional manifold embedded within the 10-dimensional space, and this manifold is likely non-linear.

Importantly, the fractal dimension exhibits a non-monotonic dependence on viscosity, as clearly shown in ?. It starts at approximately 1.50 at the lowest viscosity ($\nu = 0.01$), gradually increases as viscosity rises, reaching a peak value of around 1.75 for intermediate viscosities (specifically, between $\nu \approx 0.1$ and $\nu \approx 0.15$, indices 12, 13, 14, see ?). As viscosity continues to increase beyond this peak, the fractal dimen-

sion decreases again, settling around 1.53 at the highest viscosity ($\nu = 1.0$). This non-monotonic behavior suggests that the geometric complexity of the learned latent representation is not a simple function of viscosity but reaches maximum complexity in an intermediate regime.

3.4. Discussion: Viscosity as a scale parameter and the renormalization group analogy

The combined results from PCA and fractal dimension analysis provide a comprehensive picture of how the PINN’s latent space structure evolves with viscosity. These findings can be interpreted using an analogy with Renormalization Group (RG) theory, where viscosity is viewed as a parameter analogous to a physical scale. In this framework, low viscosity corresponds to small scales (UV regime), where dissipative effects are weak and fine structures can develop, while high viscosity corresponds to large scales (IR regime), where dissipation dominates and smooths out solutions.

The observed systematic changes in the latent space structure align conceptually with how relevant degrees of freedom change under an RG flow:

1. **Shift in Variance Distribution (Relative Importance of Modes):** The decrease in PC1’s dominance and the increase in the relative contributions of PC2 and PC3 as viscosity increases from low to high (??, ??) suggest a shift in the primary modes or features that capture the system’s variability. At low viscosity (UV-like), perhaps a single dominant structure (like a sharp shock front) accounts for most of the variance, represented by PC1. As viscosity increases (flowing towards the IR), the solutions become smoother, and the variance becomes more distributed among different features or modes, which are captured by the stable principal components. This is analogous to how different operators or modes become more or less relevant as the scale changes in an RG flow.
2. **Effective Dimensionality:** The consistently low effective dimensionality (3-4 PCs for >90%-95% variance, ??, ??) across all viscosities indicates that the PINN successfully compresses the complex, high-dimensional solution space of the Burger’s equation into a much lower-dimensional latent manifold. The subtle increase from 3 to 4 PCs required for 95% variance around $\nu \approx 0.12$ might suggest that in this intermediate regime, slightly more degrees of freedom are needed to capture the finer details of the latent representation compared to the very low or very high viscosity limits.

3. **Stability of Principal Component Directions (Invariant Features):** The remarkable stability of the top principal component directions (PC1, PC2, PC3) across viscosity changes (??, ??) is a significant finding. It implies that the PINN has learned a set of fundamental, orthogonal directions (features) in its latent space that are robust to variations in viscosity. While the *importance* of these features (measured by EVR) changes, the features themselves (their orientation in the 10D space) remain largely invariant. In the RG analogy, these stable PCs could be seen as corresponding to "relevant operators" whose form persists across different scales, even if their "couplings" (EVRs) run with scale.

4. **Geometric Complexity (Fractal Dimension):** The non-monotonic behavior of the fractal dimension (??, ??) is particularly insightful. The peak complexity at intermediate viscosities ($\nu \approx 0.1 - 0.15$) suggests that the latent manifold is geometrically richest in this regime. At very low viscosity (approaching the inviscid limit where shocks are sharp discontinuities), the solution structure might be simpler in latent space, leading to a lower fractal dimension (≈ 1.5). As viscosity increases, the interplay between non-linearity and diffusion becomes more complex, potentially leading to a richer set of dynamical structures that are encoded as a more complex latent manifold with a higher fractal dimension (≈ 1.75). At very high viscosity (dominated by diffusion), the solutions become very smooth, potentially simplifying the latent manifold structure again, resulting in a decrease in fractal dimension (≈ 1.53). This peak in complexity mirrors phenomena in dynamical systems theory where complexity measures can peak at transitions between different qualitative behaviors.

Synthesizing these observations, the PINN’s latent space seems to encode viscosity not just as a parameter but as a modulator of the system’s effective description. As viscosity changes, the relative importance of learned features shifts, the effective dimensionality subtly changes how variance is distributed, the core features themselves remain stable, and the geometric complexity of the underlying manifold evolves non-monotonically. This behavior is highly suggestive of an RG-like flow, where viscosity tunes the system’s scale, altering which degrees of freedom are most relevant and how complex the effective description of the system is. The PINN, therefore, appears to have learned a rep-

resentation that naturally reflects the scale-dependent nature of the physical system governed by the Burger’s equation, providing a potential avenue for interpreting such learned representations using established theoretical frameworks from physics.

4. CONCLUSIONS

In this study, we investigated the 10-dimensional latent space of a Physics-Informed Neural Network (PINN) trained to solve the two-dimensional Burger’s equation across a range of 25 distinct viscosity values. The primary goal was to understand how this key physical parameter is encoded within the network’s internal representation and to explore analogies with physical theories like the Renormalization Group (RG). Interpreting the latent spaces of deep learning models, particularly PINNs, remains a significant challenge, hindering our ability to fully leverage them for scientific insight beyond prediction.

The dataset consisted of the activation values from a specific hidden layer of a pre-trained PINN, capturing the latent representation of the Burger’s equation solution at various spatial points, time instances, and for each of the 25 specified viscosity values. For analysis, the 10-dimensional latent space data for each viscosity was treated as a point cloud of over 10,000 samples.

To quantitatively characterize the latent space structure and its dependence on viscosity, we employed two complementary dimensionality analysis techniques. First, Principal Component Analysis (PCA) was applied independently to the standardized latent space data for each viscosity. This allowed us to quantify the distribution of variance, estimate the effective dimensionality, and assess the stability of the dominant principal components across different viscosity settings. Second, we estimated the Correlation Dimension, a type of fractal dimension, for the latent space point cloud at each viscosity using the Grassberger-Procaccia algorithm. This provided a measure of the geometric complexity and intrinsic dimensionality of the underlying data manifold.

Our analysis yielded several key results. PCA consistently showed that the latent space is effectively low-dimensional across all viscosities, with just 3 principal components capturing over 90% of the variance and 3 or 4 components capturing over 95%. While the effective dimensionality remained low, the distribution of variance among the top components shifted systematically with viscosity: the fraction of variance explained by the first principal component decreased as viscosity increased, while the contributions of the second and third components increased. Importantly, the directions of the dominant principal components in the 10-

dimensional space were found to be remarkably stable across viscosity values, exhibiting cosine similarities consistently close to 1.0 between successive viscosity settings. The estimated fractal dimension of the latent space, ranging between 1.5 and 1.75, revealed a non-monotonic dependence on viscosity, peaking at intermediate values before decreasing again at higher viscosities.

From these results, we learned that the PINN’s latent space encodes viscosity not merely as a static input but as a dynamic parameter that shapes the structure and complexity of the learned representation. The consistently low effective dimensionality suggests that the PINN successfully identifies and compresses the essential features of the Burger’s equation solutions onto a low-dimensional manifold. The systematic shift in variance distribution among the stable principal components indicates that while the fundamental ”modes” or ”features” captured by these components remain invariant in their spatial orientation, their relative importance or contribution to the overall variability of the system’s state changes significantly with viscosity. This is analogous to how different degrees of freedom become more or less relevant as a physical system is probed at different scales, as described in Renormalization Group theory. The non-monotonic behavior of the fractal dimension further highlights that the intrinsic geometric complexity of the latent representation is not a simple linear function of viscosity but exhibits richer behavior, potentially reflecting transitions between different physical regimes of the Burger’s equation (e.g., inviscid-like behavior at low viscosity, diffusive-dominated behavior at high viscosity, and complex interplay in between). By drawing an analogy between viscosity and a scale parameter in an RG flow, our findings suggest that the PINN’s latent space learns a representation whose evolution with viscosity mirrors the scale-dependent changes in relevant degrees of freedom seen in physical systems. This work provides a novel perspective on interpreting the internal workings of PINNs, suggesting that concepts from theoretical physics, such as RG, can offer valuable frameworks for understanding how these networks encode physical parameters and capture the multiscale nature of physical phenomena.