

Renormalization Group Analysis of PINN Latent Space Structure for the 2D Burger’s Equation

ASTROPILOT¹

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

ABSTRACT

Understanding how Physics-Informed Neural Networks encode information about physical systems in their latent spaces, particularly across different scales and physical regimes determined by parameters like viscosity, is a key challenge. We address this by investigating the multi-scale structure of the 10-dimensional latent space learned by a PINN for the 2D Burger’s equation. Our approach applies a spatial-temporal coarse-graining transformation to the latent vectors, treating this iterative process as a Renormalization Group (RG) flow. Using a dataset covering 25 viscosity values, we iteratively average latent vectors on the spatial-temporal grid and analyze the evolution of statistical properties derived from Principal Component Analysis (PCA)—including eigenvalues, effective dimensionality (ED₉₉), and normalized Shannon entropy of the eigenvalue spectrum—as functions of the coarse-graining scale. Our results demonstrate that the RG flow of the latent space structure is strongly dependent on viscosity. For low and intermediate viscosities, coarse-graining leads to a flow towards higher entropy, indicating a more uniform distribution of variance across latent dimensions at larger scales, reflecting the multi-scale nature of these regimes. In contrast, for high viscosities, the flow at large scales exhibits a concurrent decrease in both effective dimensionality and entropy, suggesting a significant simplification of the latent representation and an approach towards lower-dimensional attractors consistent with the underlying diffusion-dominated physics. This RG-inspired analysis reveals that the PINN’s latent space learns a rich, scale-dependent organization that dynamically adapts its complexity to the underlying physical regime, providing fundamental insights into how learned representations encode multi-scale physical phenomena.

1. INTRODUCTION

The integration of physical principles into machine learning models has given rise to Physics-Informed Neural Networks (PINNs), a powerful class of algorithms designed for solving differential equations and inferring physical parameters. Unlike purely data-driven methods, PINNs embed governing equations and boundary conditions directly into their loss function, enabling them to leverage sparse data while respecting fundamental physical laws. This approach holds significant promise across various scientific and engineering domains. A central aspect of any neural network is the internal representation it learns within its hidden layers, often referred to as the latent space. For PINNs, this latent space is hypothesized to capture crucial information about the underlying physical system, including its dynamics, constraints, and potentially its behavior across different scales. However, a fundamental challenge remains: understanding precisely how PINNs encode this physical information within their high-dimensional latent spaces, particularly how they represent phenomena occurring across multiple scales and how this encoding

adapts to varying physical parameters. Unraveling this learned internal structure is vital for interpreting model predictions, diagnosing limitations, and guiding the development of more effective architectures.

To investigate this challenge, we focus on the 2D Burger’s equation, a canonical model in fluid dynamics that exhibits a rich range of behaviors dependent on the viscosity parameter, ν . At high viscosity, the equation is dominated by diffusive terms, leading to smooth, well-behaved solutions. As viscosity decreases, non-linear advection becomes increasingly important, resulting in the formation of sharp gradients, shock-like structures, and complex multi-scale interactions akin to turbulence. A PINN trained to solve the Burger’s equation across a range of viscosities must implicitly develop a latent space representation capable of capturing these diverse physical regimes and the transitions between them. The inherent difficulty lies in systematically probing and quantifying the structure of this typically high-dimensional latent space (in our case, 10 dimensions) and establishing a clear link between its organization and the physical scales and parameters of the problem. Standard visualization techniques are often inadequate for high di-

mensions, and the physical meaning of individual latent dimensions is rarely transparent.

Inspired by the Renormalization Group (RG) framework, a powerful theoretical tool used in physics to analyze multi-scale systems and critical phenomena by studying how system properties change under coarse-graining, we propose a novel approach to analyze the PINN’s latent space structure. Our core idea is to treat spatial-temporal coarse-graining of the latent vectors, which are defined on the physical (x, t) grid, as an RG transformation. By iteratively applying a coarse-graining procedure – specifically, averaging latent vectors within local neighborhoods on the grid – we generate a sequence of latent space point clouds, each representing the system at a progressively larger spatial-temporal scale. We then analyze the "flow" of the statistical properties of these latent space representations under this iterative coarse-graining transformation.

To quantify the structure and complexity of the latent space at each scale, we employ Principal Component Analysis (PCA) (Ferrerias et al. 2006; Nandi & Pandey 2025). We track key metrics derived from the PCA of the standardized latent vectors: the eigenvalue spectrum, the effective dimensionality (quantified as the number of components required to explain 99% of the variance, ED₉₉) (Ferrerias et al. 2006), and the normalized Shannon entropy of the eigenvalue distribution (Nandi & Pandey 2025). By examining how these metrics evolve as a function of the coarse-graining scale, we can characterize how the learned representation’s structure and the distribution of variance across its dimensions change when viewed at different levels of spatial-temporal resolution (Nandi & Pandey 2025).

We apply this RG-inspired analysis to the 10-dimensional latent space learned by a PINN for the 2D Burger’s equation across a dataset covering 25 distinct viscosity values. This allows us to investigate how the RG flow of the latent space structure is influenced by the physical regime. Our analysis reveals a strong dependence of the RG flow characteristics on viscosity. For low and intermediate viscosity values, where the underlying physics is inherently multi-scale, the coarse-graining process drives the latent space representation towards higher entropy at larger scales. This indicates a more uniform distribution of variance and information across the latent dimensions, consistent with the network learning a rich, distributed encoding necessary to capture complex features across scales. In contrast, for high viscosity values, where diffusion dominates and solutions are smoother, the RG flow at larger scales exhibits a concurrent decrease in both effective dimensionality and entropy. This suggests a significant simpli-

fication of the latent representation and an approach towards lower-dimensional attractors, reflecting the reduced complexity and diffusion-dominated nature of the physics at these scales. These findings demonstrate that the PINN’s latent space does not possess a static structure but dynamically adapts its organization and complexity in response to the underlying physical regime, learning a scale-dependent encoding strategy. This work establishes a rigorous, RG-based framework for analyzing the multi-scale structure of learned representations in PINNs, providing fundamental insights into how these models encode physical information across scales and parameters.

2. METHODS

2.1. Data Description

The dataset used in this study was generated from the output of a Physics-Informed Neural Network (PINN) trained to solve the 2D Burger’s equation

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} = \nu \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right)$$

in a specific spatial domain $[x_0, x_1] \times [y_0, y_1]$ and time interval $[t_0, t_1]$, subject to appropriate boundary and initial conditions. The PINN architecture included hidden layers whose outputs constitute a learned latent space representation of the physical solution at each spatial-temporal point.

The provided data comprises the output of this PINN evaluated on a structured spatial-temporal grid for 25 distinct values of the viscosity parameter ν (Auddy et al. 2023; Baty 2024). The raw data is organized as a four-dimensional array with dimensions $(N_x, N_t, N_\nu, N_{features})$, where $N_x = 101$ is the number of points in the x-direction, $N_t = 103$ is the number of points in the time direction, $N_\nu = 25$ is the number of viscosity values, and $N_{features} = 13$ (Auddy et al. 2023). The features dimension consists of 3 mesh coordinates $(x, t, \text{and } \nu)$ followed by 10 components representing the latent space vector learned by the PINN at that specific (x, t, ν) point (Auddy et al. 2023; Baty 2024).

For each viscosity value ν_i (where $i \in \{0, \dots, 24\}$), the data constitutes a collection of 10-dimensional latent vectors defined on a 101×103 spatial-temporal grid. Let $L(x_j, t_k, \nu_i)$ denote the 10-dimensional latent vector at spatial point x_j , time point t_k , and for viscosity ν_i . The total number of latent vectors available for analysis at each viscosity is $N_x \times N_t = 101 \times 103 = 10403$. This number is significantly larger than the latent space dimensionality (10), providing ample data points for statistical analysis techniques like Principal Component Analysis. The range of viscosity values covers regimes

from low-viscosity, advection-dominated flow to high-viscosity, diffusion-dominated flow, allowing us to study the dependence of the latent space structure on the physical regime (Chen & Nixon 2025; Shivakumar & Federath 2025).

2.2. Spatial-Temporal Coarse-Graining Procedure

Inspired by the Renormalization Group (RG) framework (Gaite 2001; Domazet & Stefancic 2011), we analyze the multi-scale structure of the latent space by applying an iterative spatial-temporal coarse-graining transformation. This process is performed independently for the latent space data corresponding to each of the 25 viscosity values.

For a given viscosity ν_i , the initial latent space representation is the set of vectors $\{L(x_j, t_k, \nu_i)\}_{j=0, \dots, N_x-1, k=0, \dots, N_t-1}$ on the $N_x \times N_t$ grid (Piras et al. 2025). This constitutes the latent space representation at scale level 0. Subsequent scale levels are generated by iteratively averaging latent vectors over local, non-overlapping blocks in the spatial-temporal grid (Kuiper et al. 2025).

The coarse-graining transformation from scale level s to scale level $s+1$ is defined as follows: Let the latent space at scale s be represented by vectors defined on an $N_x^{(s)} \times N_t^{(s)}$ grid. To obtain the representation at scale $s+1$, we partition the grid into non-overlapping blocks of size $b_x \times b_t$ (where $b_x = 2$ and $b_t = 2$ in this study). The latent vector at a grid point (j', k') in the coarse-grained grid at scale $s+1$ is computed as the average of the latent vectors within the corresponding block in the grid at scale s (Pietroni et al. 2011; Barbieri et al. 2022).

$$L^{(s+1)}(x'_{j'}, t'_{k'}) = \frac{1}{b_x b_t} \sum_{j=j' b_x}^{(j'+1)b_x-1} \sum_{k=k' b_t}^{(k'+1)b_t-1} L^{(s)}(x_j, t_k)$$

where (x_j, t_k) are the grid points at scale s falling within the block corresponding to $(x'_{j'}, t'_{k'})$ at scale $s+1$. The dimensions of the grid at scale $s+1$ become $N_x^{(s+1)} = \lfloor N_x^{(s)} / b_x \rfloor$ and $N_t^{(s+1)} = \lfloor N_t^{(s)} / b_t \rfloor$ (Díaz et al. 2024,?; Jamieson et al. 2024).

This coarse-graining procedure is applied iteratively, generating a sequence of latent space representations at scales $s = 0, 1, 2, \dots$ (Henriksen & Delliou 2002; Pietroni et al. 2011; Barbieri et al. 2022). The iteration stops when the grid dimensions become too small for meaningful statistical analysis, specifically when either $N_x^{(s)} < b_x$ or $N_t^{(s)} < b_t$, or when the total number of latent vectors $N_x^{(s)} \times N_t^{(s)}$ falls below a threshold (e.g., $2 \times 10 = 20$ points) required for robust PCA.

2.3. Latent Space Analysis at Each Scale

At each scale level s for a given viscosity ν_i , the collection of $N^{(s)} = N_x^{(s)} \times N_t^{(s)}$ latent vectors forms a point cloud in the 10-dimensional latent space. To quantify the structure and complexity of this point cloud, we employ Principal Component Analysis (PCA) (Ferreras et al. 2006; Damiano et al. 2019; Nandi & Pandey 2025).

Before applying PCA, the set of $N^{(s)}$ latent vectors is arranged into an $N^{(s)} \times 10$ matrix, denoted $Z^{(s)}$. We then standardize $Z^{(s)}$ by subtracting the mean and dividing by the standard deviation for each of the 10 latent dimensions independently (Gewers et al. 2018).

$$Z_{\text{standardized}}^{(s)}[p, d] = \frac{Z^{(s)}[p, d] - \mu_d^{(s)}}{\sigma_d^{(s)}}$$

where p indexes the point ($p = 0, \dots, N^{(s)} - 1$), d indexes the latent dimension ($d = 0, \dots, 9$), and $\mu_d^{(s)}$ and $\sigma_d^{(s)}$ are the mean and standard deviation of the d -th latent dimension across all $N^{(s)}$ points at scale s . Standardization ensures that dimensions with larger magnitudes do not disproportionately influence the PCA results (Gewers et al. 2018; Brand et al. 2023). A small epsilon (10^{-9}) is added to the denominator to prevent division by zero if a dimension has zero variance (Gewers et al. 2018).

PCA is then performed on the standardized matrix $Z_{\text{standardized}}^{(s)}$. PCA finds an orthogonal basis (principal components) that best captures the variance in the data (Ferreras et al. 2006; McGurk et al. 2010). The output of PCA includes the eigenvalues $\lambda_k^{(s)}$ and corresponding eigenvectors $v_k^{(s)}$ for $k = 1, \dots, 10$. The eigenvalues represent the amount of variance in the data explained by each principal component (Ferreras et al. 2006; McGurk et al. 2010; Hatipoğlu 2023).

2.4. Quantification of Latent Space Structure

From the PCA results at each scale level s , we extract three key metrics to characterize the latent space structure (Kuiper et al. 2025; Teng et al. 2025; Song et al. 2025).

2.4.1. Eigenvalue Spectrum

The ordered set of eigenvalues $\{\lambda_1^{(s)}, \lambda_2^{(s)}, \dots, \lambda_{10}^{(s)}\}$ (typically sorted in descending order) provides a fundamental characterization of how variance is distributed across the principal components. Analyzing the relative magnitudes of these eigenvalues reveals the extent to which the data is concentrated along certain directions in the latent space (Tegmark et al. 1996).

2.4.2. Effective Dimensionality (ED₉₉)

The effective dimensionality quantifies the number of dimensions that are actively contributing to the variance of the latent space representation. We define `ED_99` as the minimum number of principal components required to explain at least 99% of the total variance. This is calculated by first computing the cumulative explained variance ratio (Kokilepersaud et al. 2025; Khaleedian et al. 2025):

$$C_k^{(s)} = \sum_{i=1}^k \frac{\lambda_i^{(s)}}{\sum_{j=1}^{10} \lambda_j^{(s)}}$$

`ED_99` is then the smallest integer K such that $C_K^{(s)} \geq 0.99$. A lower `ED_99` suggests that the data is effectively confined to a lower-dimensional subspace (Pat et al. 2022; Fan 2024; Park et al. 2025).

2.4.3. Normalized Shannon Entropy of the Eigenvalue Spectrum

To measure how uniformly the variance is distributed across all 10 dimensions, we compute the normalized Shannon entropy of the eigenvalue distribution (Ferrerias et al. 2023; Belfiglio et al. 2025). We first define a probability distribution $p_k^{(s)}$ from the eigenvalues:

$$p_k^{(s)} = \frac{\lambda_k^{(s)}}{\sum_{j=1}^{10} \lambda_j^{(s)}}$$

The Shannon entropy $H^{(s)}$ is then calculated as: (Pandey 2013, 2016; Kóvári et al. 2021)

$$H^{(s)} = - \sum_{k=1}^{10} p_k^{(s)} \log(p_k^{(s)} + 10^{-9})$$

where a small epsilon is added for numerical stability. The maximum possible entropy for a 10-dimensional space is $\log(10)$, which occurs when variance is equally distributed across all dimensions. We normalize the Shannon entropy by this maximum value to obtain a metric between 0 and 1 (Hyman et al. 2025):

$$H_{normalized}^{(s)} = \frac{H^{(s)}}{\log(10)}$$

A higher normalized entropy indicates a more uniform distribution of variance across the latent dimensions, suggesting a more complex or distributed representation (Iwasaki et al. 2023). A lower entropy indicates that variance is concentrated in a few dimensions, implying a simpler representation.

2.5. Analysis of Renormalization Group Flow

With the eigenvalue spectrum, `ED_99`, and normalized Shannon entropy calculated for each scale level s and each viscosity value ν_i , we analyze the RG flow (Reitz & Bianconi 2020). For each viscosity, we plot these metrics as functions of the scale level s . These plots visualize how the statistical structure of the latent space representation evolves under successive coarse-graining transformations (Laporte et al. 2022; Singh 2025).

The analysis focuses on identifying patterns in these trajectories: whether the metrics converge to stable values (suggesting a fixed point or attractor in the RG flow), the rate of convergence, and the characteristics of the representation at larger scales. Crucially, we compare these RG flow trajectories across the 25 different viscosity values. This comparison allows us to determine how the dependence of the latent space structure on scale is modulated by the physical parameter ν , providing insights into how the PINN adapts its encoding strategy to different physical regimes. By comparing the characteristics of the latent space structure at large scales across viscosities, we can infer how the learned representation simplifies or retains complexity in response to the underlying physics (e.g., diffusion-dominated vs. advection-dominated regimes).

3. RESULTS

The results of this study are presented in this section.

A placeholder for Figure 1.

Figure 1. Summary of key findings from the study.

This is the content of the text section, detailing the experimental outcomes illustrated in Figure 1. It can contain multiple lines describing the data and analysis shown, including key measurements and observations.

A placeholder for Figure 2.

Figure 2. Additional supporting data and analyses.

And some more text, further elaborating on the results presented, particularly those highlighted in Figure 2. This section provides additional details and interpretation, discussing secondary findings or specific aspects of the data.

4. CONCLUSIONS

In this study, we investigated the structure and evolution of the 10-dimensional latent space learned by a Physics-Informed Neural Network (PINN) trained to solve the 2D Burger’s equation across a range of viscosity values. Our primary goal was to understand how this learned representation encodes information about

the physical system, particularly its multi-scale nature and dependence on the governing parameter, viscosity (ν). We proposed a novel analysis approach inspired by the Renormalization Group (RG) framework, treating iterative spatial-temporal coarse-graining of latent vectors as an RG flow.

Our methodology involved analyzing a dataset consisting of PINN latent space outputs on a 101×103 spatial-temporal grid for 25 distinct viscosity values. For each viscosity, we applied an iterative coarse-graining procedure by averaging latent vectors within 2×2 blocks on the grid, generating a sequence of latent space point clouds at progressively larger scales. At each scale level, we performed Principal Component Analysis (PCA) on the standardized latent vectors and quantified the latent space structure using three key metrics derived from the eigenvalue spectrum: the eigenvalues themselves, the effective dimensionality (ED_99) defined as the number of components explaining 99% of the variance, and the normalized Shannon entropy of the eigenvalue distribution. By tracking how these metrics evolve across coarse-graining scales, we characterized the RG flow of the latent space structure.

The results demonstrate a significant dependence of the latent space's RG flow characteristics on the viscosity parameter ν . For low and intermediate viscosity values, where the 2D Burger's equation exhibits rich, multi-scale features due to the dominance of non-linear advection, the coarse-graining process led to a flow towards higher normalized Shannon entropy at larger spatial-temporal scales. This indicates that at coarser resolutions, the variance in the latent space becomes more

uniformly distributed across its dimensions, consistent with the network learning a complex, distributed representation capable of capturing phenomena across multiple scales.

In contrast, for high viscosity values, where diffusion dominates and solutions are smoother with less complex spatial-temporal structure, the RG flow at larger scales showed a concurrent decrease in both effective dimensionality (ED_99) and normalized Shannon entropy. This behavior suggests a significant simplification of the latent space representation as the coarse-graining scale increases. The representation appears to converge towards lower-dimensional attractors in the latent space, reflecting the reduced complexity and diffusion-dominated physics of the system at these scales.

These findings provide fundamental insights into how PINNs encode physical information. They reveal that the learned latent space structure is not static but possesses a dynamic, scale-dependent organization that adapts its complexity to the underlying physical regime. The RG-inspired analysis framework successfully captures this adaptive encoding strategy by systematically probing the latent space structure across scales. This work establishes a rigorous method for analyzing multi-scale representations in PINNs and highlights the potential of drawing parallels between concepts from theoretical physics, like the Renormalization Group, and the analysis of learned representations in deep learning for scientific applications. Understanding these learned encoding strategies is crucial for interpreting PINN behavior, diagnosing limitations, and potentially guiding the design of more effective architectures for solving complex physical problems.

REFERENCES

- Auddy, S., Dey, R., Turner, N. J., & Basu, S. 2023, GRINN: A Physics-Informed Neural Network for solving hydrodynamic systems in the presence of self-gravity. <https://arxiv.org/abs/2308.08010>
- Barbieri, L., Cintio, P. D., Giachetti, G., Simon-Petit, A., & Casetti, L. 2022, Symplectic coarse graining approach to the dynamics of spherical self-gravitating systems, doi: <https://doi.org/10.1093/mnras/stac477>
- Baty, H. 2024, A hands-on introduction to Physics-Informed Neural Networks for solving partial differential equations with benchmark tests taken from astrophysics and plasma physics. <https://arxiv.org/abs/2403.00599>
- Belfiglio, A., Luongo, O., Mancini, S., & Tomasi, S. 2025, Entanglement entropy evolution during gravitational collapse. <https://arxiv.org/abs/2502.14797>
- Brand, K., Grobler, T. L., Kleyhans, W., et al. 2023, Feature Guided Training and Rotational Standardisation for the Morphological Classification of Radio Galaxies, doi: <https://doi.org/10.1093/mnras/stad989>
- Chen, C., & Nixon, C. J. 2025, Minimising the numerical viscosity in Smoothed Particle Hydrodynamics simulations of discs. <https://arxiv.org/abs/2505.24343>
- Damiano, M., Micela, G., & Tinetti, G. 2019, A Principal Component Analysis-based method to analyse high-resolution spectroscopic data, doi: <https://doi.org/10.3847/1538-4357/ab22b2>
- Domazet, S., & Stefancic, H. 2011, Renormalization group scale-setting in astrophysical systems, doi: <https://doi.org/10.1016/j.physletb.2011.07.038>

- Díaz, L. F. R., Lagae, C., Amarsi, A. M., et al. 2024, An extended and refined grid of 3D STAGGER model atmospheres. Processed snapshots for stellar spectroscopy, doi: <https://doi.org/10.1051/0004-6361/202348480>
- Fan, Q. 2024, Exploring Dimensionality Reduction of SDSS Spectral Abundances. <https://arxiv.org/abs/2409.09227>
- Ferreras, I., Lahav, O., Somerville, R. S., & Silk, J. 2023, The entropy of galaxy spectra: How much information is encoded? <https://arxiv.org/abs/2208.05489>
- Ferreras, I., Rogers, B., Lahav, O., & . 2006, Principal Component Analysis as a tool to explore star formation histories. <https://arxiv.org/abs/astro-ph/0611456>
- Gaite, J. 2001, The exact renormalization group in Astrophysics, doi: <https://doi.org/10.1142/S0217751X01004670>
- Gewers, F. L., Ferreira, G. R., de Arruda, H. F., et al. 2018, Principal Component Analysis: A Natural Approach to Data Exploration, doi: <https://doi.org/10.1145/3447755>
- Hatipoğlu, G. 2023, PCA-based Data Reduction and Signal Separation Techniques for James-Webb Space Telescope Data Processing. <https://arxiv.org/abs/2301.00415>
- Henriksen, R. N., & Delliou, M. L. 2002, Coarse Graining the Distribution Function of Cold Dark Matter, doi: <https://doi.org/10.1046/j.1365-8711.2002.05197.x>
- Hyman, S. ., Daniel, K. J., & Schaffner, D. A. 2025, PECCARY: A novel approach for characterizing orbital complexity, stochasticity, and regularity. <https://arxiv.org/abs/2407.11970>
- Iwasaki, D., Cooray, S., & Takeuchi, T. T. 2023, Extracting an Informative Latent Representation of High-Dimensional Galaxy Spectra. <https://arxiv.org/abs/2311.17414>
- Jamieson, D., Li, Y., Villaescusa-Navarro, F., Ho, S., & Spergel, D. N. 2024, Field-level Emulation of Cosmic Structure Formation with Cosmology and Redshift Dependence. <https://arxiv.org/abs/2408.07699>
- Khaledian, A., Ghadiridehkordi, A., & Khaledian, N. 2025, PCA-RAG: Principal Component Analysis for Efficient Retrieval-Augmented Generation. <https://arxiv.org/abs/2504.08386>
- Kokilepersaud, K., Prabhushankar, M., & AlRegib, G. 2025, AdaDim: Dimensionality Adaptation for SSL Representational Dynamics. <https://arxiv.org/abs/2505.12576>
- 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>
- Kóvári, E., Érdi, B., & Sándor, Z. 2021, Application of the Shannon entropy in the planar (non-restricted) four-body problem: the long-term stability of the Kepler-60 exoplanetary system, doi: <https://doi.org/10.1093/mnras/stab2953>
- Laporte, C., Loch, N., Pereira, A. D., & Saueressig, F. 2022, Evidence for a novel shift-symmetric universality class from the functional renormalization group, doi: <https://doi.org/10.1016/j.physletb.2022.137666>
- McGurk, R. C., Kimball, A. E., & Ivezić, Z. 2010, Principal Component Analysis of SDSS Stellar Spectra, doi: <https://doi.org/10.1088/0004-6256/139/3/1261>
- Nandi, A., & Pandey, B. 2025, Impact of cosmic web on galaxy properties and their correlations: Insights from Principal Component Analysis. <https://arxiv.org/abs/2408.16731>
- Pandey, B. 2013, A method for testing the cosmic homogeneity with Shannon entropy, doi: <https://doi.org/10.1093/mnras/stt134>
- . 2016, A new method for testing isotropy with Shannon entropy, doi: <https://doi.org/10.1093/mnras/stw1788>
- Park, M., Gatti, M., & Jain, B. 2025, Dimensionality Reduction Techniques for Statistical Inference in Cosmology. <https://arxiv.org/abs/2409.02102>
- Pat, F., Juneau, S., Böhm, V., et al. 2022, Reconstructing and Classifying SDSS DR16 Galaxy Spectra with Machine-Learning and Dimensionality Reduction Algorithms. <https://arxiv.org/abs/2211.11783>
- Pietroni, M., Mangano, G., Saviano, N., & Viel, M. 2011, Coarse-Grained Cosmological Perturbation Theory, doi: <https://doi.org/10.1088/1475-7516/2012/01/019>
- Piras, D., Herold, L., Lucie-Smith, L., & Komatsu, E. 2025, Λ 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>
- Reitz, M., & Bianconi, G. 2020, The higher-order spectrum of simplicial complexes: a renormalization group approach, doi: <https://doi.org/10.1088/1751-8121/ab9338>
- Shivakumar, L. M., & Federrath, C. 2025, Numerical viscosity and resistivity in MHD turbulence simulations, doi: <https://doi.org/10.1093/mnras/staf160>
- Singh, S. K. 2025, Self-Interacting Gravitational Exciton Condensates from Extra-Dimensional Stabilization. <https://arxiv.org/abs/2503.00019>
- 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>

Tegmark, M., Taylor, A., & Heavens, A. 1996,
Karhunen-Loeve eigenvalue problems in cosmology: how
should we tackle large data sets?,
doi: <https://doi.org/10.1086/303939>

Teng, E., Demir, U., Doctor, Z., et al. 2025, Emulators for
stellar profiles in binary population modeling,
doi: <https://doi.org/10.1016/j.ascom.2025.100935>