Skepthical review: *Geometric Structure of PINN Latent Space for Burger's Equation: Low-Dimensional Manifolds and Initial Condition Encoding*

Summary

The manuscript analyzes the geometry of a Physics-Informed Neural Network (PINN) latent space for a Burgers-equation setup across 25 initial conditions (ICs). From a pretrained network, the authors extract a 10D latent vector at each sampled space–time point, yielding (per IC) a point cloud in $\mathbb{R}^{10}$. They then apply a structured linear-geometry pipeline—global PCA over all latent vectors, per-IC PCA, PCA on per-IC centroids, and subspace comparison via principal-angle-based similarity—to argue that (i) the global latent distribution is effectively ~ 6D ($> 99\%$ variance), (ii) each IC's latent cloud is effectively ~ 3D ($\sim 97\%$ variance in the first three PCs), (iii) IC variation is primarily encoded as translation of these ~ 3D clouds along an almost 1D trajectory (centroid PC1 explains $\sim 99.9\%$), and (iv) the per-IC 3D subspaces are strongly aligned (mean similarity ~ 0.986) (Secs. 2.1–2.3, 3.1–3.5). The analysis is clearly organized and the reported patterns are striking; however, the paper currently lacks essential PDE/PINN specification and validation, contains a core ambiguity about whether the problem is "2D" or actually (x, t) only, and relies heavily on PCA-based variance explanations without robustness checks, nonlinear intrinsic-dimension sanity checks, or quantitative tests of the central 'translation-dominates-IC' interpretation (Secs. 1–4, Figs. 1–12). Addressing these points would substantially improve reproducibility, interpretability, and the credibility/generalizability of the geometric conclusions.

Strengths

- Clear, coherent analysis goal: characterize how a single PINN organizes multiple initial conditions in its latent space, and test this via a consistent multilevel PCA/subspace pipeline (Secs. 1–2).
- Well-structured empirical workflow: global PCA, per-IC PCA, centroid PCA, and principal-angle-based subspace similarity together provide a hierarchical view of latent organization (Secs. 2.2–2.3, 3.1–3.5).
- Reported patterns are internally consistent and visually compelling: ~ 6D global structure, ~ 3D per-IC structure, ~ 1D centroid variation, and high subspace alignment (Secs. 3.1–3.4).
- Interpretation is potentially impactful for PINN interpretability and parameter encoding (IC effects vs spatiotemporal dynamics), and the manuscript generally reads clearly and systematically.
- Dimensional bookkeeping for sampling counts and reshaping is mostly consistent (101×103 points per IC; 25 ICs; 10D latent vectors), which supports confidence in the basic data handling described in Sec. 2.1.

Major issues

1. **Core problem-specification ambiguity/inconsistency: the manuscript repeatedly says "2D Burgers' equation," but the described sampling grid and notation appear to be only (x, t) with a single spatial coordinate x (Secs. 1, 2.1–2.2; pp.2–3). The mention of predicted solution components "u and v" further suggests either a vector-valued field or a different PDE than what is described. This ambiguity undermines interpretation of all geometric findings because the IC family and solution structure depend strongly on the PDE definition.**

   Recommendation: In Sec. 2.1 (and briefly in Sec. 1), state the exact PDE being solved: scalar vs vector Burgers, spatial dimensionality, full equation(s), viscosity and any forcing, domain (x, t) (or (x, y, t)), and boundary/initial conditions. Then make terminology consistent throughout: either correct all "2D" mentions to "1D-in-space Burgers on a 2D spatiotemporal domain (x, t)," or explicitly define the 2D spatial setting and show how y is sampled/stored.

2. **Reproducibility gap: essential details of the PINN architecture, latent-vector definition, IC conditioning, training objective, and validation are missing (Secs. 2.1–2.2, 3). The paper states a pretrained model yields a 10D latent vector and a 13D feature vector, but does not specify (i) where the latent layer is (before/after activation/normalization), (ii) whether IC is an explicit input, an embedding, or handled by training separate networks, (iii) the loss terms and weights (PDE residual vs BC/IC/data), (iv) collocation/data sampling, optimizers/schedule, and (v) whether accuracy/generalization was validated across the 25 ICs. Without these, it is difficult to assess whether latent geometry reflects physics, architecture, or training artifacts.**

   Recommendation: Add a dedicated "PDE + PINN setup" subsection expanding Sec. 2.1/2.2 that includes: network diagram (inputs, layers, activations), explicit definition of the 10D latent vector (exact tensor taken, location), how IC information enters the model (conditional input, parameterization, or otherwise), the full loss with weights, training procedure (sampling, optimizer, epochs/stopping), and quantitative validation (e.g., PDE residual and/or error vs a reference solver) aggregated over ICs. If the model/data are from prior work, still summarize the essentials and cite precisely.

3. **Over-interpretation risk: claims of "3D affine manifolds," "low-dimensional manifold structure," and (implicitly) "disentanglement" are currently supported primarily by PCA variance-explained statistics (Secs. 3.1–3.3, 3.6–4). High explained variance by a few PCs does not rule out curved manifolds, regime mixtures (e.g., time-dominated variance), or nonlinear structure; and "disentanglement" has specific meanings in representation learning that are stronger than "dominant centroid direction + aligned subspaces."**

   Recommendation: In Secs. 3.2–3.3 and 3.6–4, soften language to "approximately low-dimensional linear structure under PCA" unless additional tests are added. Strengthen the claim with at least two simple diagnostics: (i) reconstruction error curves vs number of PCs (per-IC and global), and (ii) a curvature/heterogeneity check such as time-sliced PCA (early vs late time, or several t-bins) to see whether the '3D' claim holds uniformly. Optionally add one nonlinear intrinsic-dimension estimator (e.g., TwoNN, Levina–Bickel MLE, participation ratio) to corroborate that the effective dimension is not purely a linear artifact.

4. **Central "IC is mainly encoded as translation" interpretation is not directly tested (Secs. 3.3–3.6). The current evidence (near-1D centroid PCA + high subspace similarity) is suggestive but does not quantify how much of the between-IC variation is explained by centroid shift vs subspace rotation vs within-manifold shape changes.**

   Recommendation: Add a quantitative effect-size test in Sec. 3.3/3.6: (i) variance decomposition (ANOVA-style) into within-IC covariance vs between-IC centroid covariance in latent space, and report fractions along global PCs; and (ii) a "translated template" reconstruction experiment: choose a reference 3D basis V_{ref} (e.g., from one IC or pooled within-IC covariance), model each IC as $C_k + V_{\text{ref}} z$, report reconstruction error; then allow per-IC bases V_k and report improvement. This directly validates whether translation dominates and how much rotations matter.

5. **Subspace similarity metric and PCA methodology are underspecified, limiting reproducibility and interpretation (Secs. 2.2.1–2.2.2, 2.3.2, 3.4–3.5; Fig. 7 and related). It is unclear whether latent dimensions are only mean-centered or also scaled/standardized, which PCA algorithm is used (SVD vs covariance eigendecomposition), how dimensionality thresholds are applied, and the exact scalar similarity formula underlying the reported mean ~ 0.986 (e.g., $|U^T V|_F^2/d$, mean $\cos^2\theta_i$, product, min, etc.).**

   Recommendation: Expand Sec. 2.2 (and 2.3.2 if needed) with precise definitions: preprocessing (mean-centering; any scaling), PCA implementation (library + SVD/covariance), intrinsic-dimensionality definition $d = \min d : \text{cvar}(d) \geq \tau$ with explicit τ values, and an explicit mathematical formula for the subspace similarity statistic with how principal angles are computed and aggregated. Ensure all reported similarity numbers (Secs. 3.4–3.5) reference this single definition.

6. **Robustness/generalization is not established: the study uses one pretrained network, one PDE parameter setting, and 25 ICs of unspecified diversity (Secs. 2.1, 3.1–3.7, 4). Yet parts of Sec. 3.6 and Sec. 4 read broadly (as if generic to PINNs). Without reruns across seeds/architectures or variation in PDE parameters (e.g., viscosity) and IC families, it is unclear whether the 6D/3D/1D geometry is stable or idiosyncratic.**

   Recommendation: Narrow claims in Secs. 3.6–4 to the studied configuration unless robustness is added. If feasible, add a small robustness section: retrain at least one additional model with a different seed and/or modest architecture change, and (optionally) vary viscosity or IC family; then report whether key summary statistics (global effective dimension, per-IC dimension, centroid PC1 variance, similarity distribution) remain consistent. If out of scope, state this explicitly as a limitation in Sec. 3.7.

7. **Connection to physics and IC parameterization is currently too indirect to support interpretability claims (Secs. 2.1, 3.3, 3.6). The near-linear centroid ordering along CPC1 is hard to interpret without knowing how ICs are generated/ordered; it may be a trivial reflection of a single scalar IC parameter (e.g., amplitude) or even arbitrary indexing.**

   Recommendation: In Sec. 2.1, describe the 25 ICs concretely (functional form, parameters, ranges, and how IC index 0–24 is assigned). In Sec. 3.3/3.6, correlate centroid coordinates (CPC1 or global PC1 projection of centroids) with physically meaningful IC descriptors (amplitude, wavenumber, phase, energy, etc.), and report correlation/regression results. Optionally relate within-IC PCs to solution statistics (e.g., gradient norms/shock indicators) to interpret what the "3D" variation corresponds to physically.

Minor issues

1. Dataset/feature-vector provenance is unclear (Sec. 2.1): the `data_bundle` is described mainly by shape; it is not explicit whether stored "solutions" are PINN predictions or reference-solver outputs, what all 13 features are, and what preprocessing/normalization is applied. The mention of "u and v" + 10 latent dims accounts for 12 features, leaving 1 unexplained.

   Recommendation: In Sec. 2.1, explicitly list all 13 feature channels (with meaning and any scaling), specify whether solutions are PINN outputs, reference solutions, or both, and describe any normalization. Clarify the extra 13th feature to resolve the bookkeeping inconsistency.

2. Threshold choices for "intrinsic dimensionality" (e.g., 95% vs 99%) are not justified and sensitivity is not shown (Secs. 2.2.1, 3.1–3.2; Fig. 3). Dimensional conclusions can change materially with τ.

   Recommendation: Provide a brief justification for the chosen variance thresholds and add a sensitivity table/plot showing inferred dimensions over $\tau \in [0.90, 0.99]$ for global PCA, per-IC PCA, and centroid PCA.

3. Figure captions and plots frequently lack key quantitative annotations and uncertainty/variability depiction (Figs. 1–12; Secs. 3.1–3.5). Several claims (e.g., "near-linearity," "very high alignment") are visually suggested but not accompanied by confidence intervals or distributions (e.g., of similarity scores).

   Recommendation: Add numeric labels where appropriate (exact explained-variance percentages, sample sizes), and include uncertainty estimates (bootstrap over (x, t) points and/or ICs) or at least distribution summaries ($\text{mean} \pm \text{std}$, min/max) for per-IC variance captures and pairwise subspace similarities. Where a "linearity" claim is made, report a simple metric (e.g., R^2 of centroid fit to a line in the leading centroid-PC plane).

4. PC sign/order ambiguity can affect PC-to-PC comparisons (notably dot-product style plots; Figs. 8–9; Sec. 3.5). When eigenvalues are close, PCA directions can swap; signs are arbitrary. This can create unstable "corresponding PC" narratives even if subspaces are aligned.

Recommendation: When comparing components across ICs, emphasize principal angles/subspace metrics (basis-invariant) over raw PC dot products, or explicitly describe how PCs are aligned (e.g., Procrustes alignment, sign fixing by correlation with a reference) and report eigenvalue gaps to justify stable ordering.

5. Centroid computation/weighting is only implicit (Secs. 2.3, 3.3). If the centroid is a uniform mean over all (x, t) points, then late-time/early-time sampling density and temporal dynamics can influence the centroid strongly.

   Recommendation: State explicitly in Sec. 2.3 how centroids are computed (uniform mean over all sampled (x, t) points) and whether alternative weighting (e.g., per-time-slice averaging, equal weight per time) changes the centroid PCA conclusion.

6. Claims about downstream utility (compression/ROM/transfer) are currently not demonstrated quantitatively (Secs. 1, 3.6, 4).

   Recommendation: Either clearly label these as speculative and narrow the language (Abstract/Secs. 1, 3.6, 4), or add a minimal demonstration: project to 6D (global) or 3D (per-IC) and quantify impact on reconstruction/solution error (if the architecture allows decoding), or at least correlate low-dimensional coordinates with PINN error metrics.

7. Related work is sparse given the goals (Sec. 1). The paper would benefit from positioning relative to PINN interpretability, latent-space geometry, and reduced-order modeling/manifold learning for PDEs.

   Recommendation: Add a short related-work paragraph/subsection in Sec. 1 covering prior analyses of internal representations in PINNs, latent manifolds for dynamical systems/PDE surrogates, and connections to ROM/manifold learning; clarify the specific novelty (cross-IC manifold geometry in a PINN latent space).

8. Computational details are missing: the dataset is large ($\sim \mathbf{260}$k vectors), yet runtime/memory and whether PCA is exact, randomized, or batched is not discussed (Secs. 2.2, 3).

   Recommendation: Add a short implementation note: computational cost, PCA method (full SVD vs randomized), and any subsampling used for figures. This improves reproducibility and helps readers scale the approach.

Very minor issues

1. Typos/formatting inconsistencies: "Burger's" vs "Burgers'", sporadic math-mode artifacts for plain percentages/dimensions, inconsistent hyphenation, and occasional spacing/line-break glitches (Secs. 1–4).

   Recommendation: Proofread and standardize terminology and formatting (prefer "Burgers' equation"), keep percentages/dimensions in consistent style, and remove spacing/line-break artifacts.

2. Keywords and front-matter appear mismatched to the topic (Abstract/keywords): astronomy-related keywords and non-standard affiliation text may confuse indexing and readers.

   Recommendation: Replace keywords with PINN/PDE/latent-geometry terms and ensure author/affiliation formatting follows the target venue's conventions.

3. Acronyms/notation such as PC1, CPC1, "Global PC1" are introduced in multiple places without a compact glossary (Secs. 3.1, 3.3, 3.5).

   Recommendation: Define PCA-related acronyms once (ideally in Sec. 2.2 or early Sec. 3) and consider adding a brief notation table distinguishing global PCA vs per-IC PCA vs centroid PCA.

4. Section heading formatting appears inconsistent in the provided markup (Secs. 2–3), which may harm readability/navigation in the final version.

   Recommendation: Standardize section/subsection formatting to the venue style (LaTeX or template), and ensure internal references to section numbers remain correct after editing.

5. Some small numerical cumulative-variance discrepancies appear consistent with rounding (e.g., global PCA cumulative sums; Sec. 3.1).

   Recommendation: Optionally add a brief note on rounding conventions (round components vs cumulative) or adjust displayed totals to avoid confusion.

Key statements and references

- **• Physics-Informed Neural Networks (PINNs) embed governing physical laws directly into the neural network architecture and loss function, enabling mesh-free solution of PDEs in complex geometries and data-sparse regimes, and have been demonstrated as a significant advancement over traditional numerical techniques in prior work [112233441111121121121212121212121212121212121212121212121212121212121212121212121212121212121212121**
- *Reference(s):* 112233441111121121121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212

- **• The 2D Burger's equation is widely used in the literature as a canonical nonlinear PDE that serves as a simplified model for complex fluid-dynamical phenomena such as turbulence and shock formation, and is known to exhibit rich dynamics with strong sensitivity to initial conditions [112233441111121121121212121212121212121212121212121212121212121212121212121212121212121212121212121**

- *Reference(s):* 112234411111211211212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121
- • **Principal Component Analysis (PCA) is a standard statistical technique that performs an orthogonal transformation of possibly correlated variables into linearly uncorrelated principal components, where the components are eigenvectors of the covariance matrix and their eigenvalues quantify the variance along each direction [112234411111211211212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121**
- *Reference(s):* 112234411111211211212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121
- • **Subspace similarity measures based on principal angles, commonly used in prior work, quantify the similarity between two low-dimensional subspaces by computing the cosines of the principal angles between them, with values near 1 indicating nearly parallel subspaces and values near 0 indicating orthogonality [112234411111211211212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121**
- *Reference(s):* 112234411111211211212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121212121

Mathematical consistency audit

This section audits **symbolic/analytic** mathematical consistency (algebra, derivations, dimensional/unit checks, definition consistency).

Maths relevance: light

The paper is primarily descriptive/empirical (PCA and geometric interpretation of latent vectors) and contains very few explicit equations or derivations. The audit therefore focuses on internal consistency of dimensional/shape claims, definitions of PCA-derived objects (centroids, subspaces), and whether reported variance decompositions are arithmetically coherent. The main analytic inconsistency is a mismatch between the claimed PDE dimensionality ("2D Burgers") and the (x, t)-only indexing and grid described.

Checked items

1. ✔ **Latent tensor extraction and shape consistency** (Sec. 2.1, p.2)
   - **Claim:** Extracting the last 10 of 13 features yields latent_space_data with shape $(101, 103, 25, 10)$, where latent_space_data$[i, j, k, :]$ is a 10D latent vector $L(x_i, t_j, IC_k)$.
   - **Checks:** dimension/shape consistency, definition consistency
   - **Verdict:** PASS; confidence: high; impact: moderate
   - **Assumptions/inputs:** Original data_bundle has shape $(101, 103, 25, 13)$. The last 10 feature channels correspond to latent variables.
   - **Notes:** Indexing and resulting tensor shape are consistent with slicing 10 channels from the feature dimension.
2. ✔ **Point count per initial condition** (Sec. 2.1, p.2)
   - **Claim:** For each initial condition k, the latent vectors over $i = 0..100$ and $j = 0..102$ form $101 \times 103 = 10403$ points in $\mathbb{R}^{10}$.
   - **Checks:** arithmetic consistency, dimension/shape consistency
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** There are 101 spatial grid points and 103 time steps.
   - **Notes:** $101 \times 103 = 10403$ is correct.
3. ✔ **Global point count after reshaping** (Sec. 2.2.1 (Global PCA), p.3; reiterated Sec. 3.1, p.4)
   - **Claim:** Reshaping all latent vectors across 25 ICs yields $(101 \times 103 \times 25, 10) = (260075, 10)$.
   - **Checks:** arithmetic consistency, dimension/shape consistency
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Per-IC point count is 10403. There are 25 initial conditions.
   - **Notes:** $10403 \times 25 = 260075$ is correct.
4. ✔ **Affine-manifold approximation via PCA (per-IC)** (Sec. 2.3.2, p.4)
   - **Claim:** Each per-IC point cloud can be approximated as an affine subspace defined by centroid C_k and span of first d_{ic} principal vectors $V_k = [v_{k,1}, \ldots, v_{k,d_{\text{ic}}}]$.
   - **Checks:** linear-algebra consistency, definition consistency
   - **Verdict:** PASS; confidence: high; impact: moderate
   - **Assumptions/inputs:** Per-IC PCA is computed on mean-centered data. Leading eigenvectors are orthonormal (as in standard PCA).
   - **Notes:** Using centroid + span(leading PCs) is the standard affine-subspace model implied by PCA.
5. ✔ **Global cumulative variance arithmetic** (Sec. 3.1, p.4)
   - **Claim:** PC1=60.12%, PC2=23.44%, PC3=12.93% (cumulative 96.48%); adding PC4=1.30%, PC5=1.17%, PC6=0.76% yields cumulative 99.72%.

- **Checks:** arithmetic consistency, sanity check (variance fractions sum)
- **Verdict:** PASS; confidence: high; impact: minor
- **Assumptions/inputs:** Percentages are rounded to two decimals.
- **Notes:** $60.12 + 23.44 + 12.93 = 96.49$ (0.01 rounding discrepancy vs 96.48); adding $1.30 + 1.17 + 0.76$ gives 99.72 exactly.

6. ✔ **Per-IC average cumulative variance arithmetic** (Sec. 3.2, p.5)
   - **Claim:** Average variance explained by first three per-IC PCs is 59.61%, 23.72%, 14.15% with cumulative 97.48%.
   - **Checks:** arithmetic consistency
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Percentages are rounded to two decimals.
   - **Notes:** $59.61 + 23.72 + 14.15 = 97.48$ exactly.

7. ✔ **Intrinsic dimensionality definition and usage** (Sec. 2.2.2, p.3; Sec. 3.2, p.5)
   - **Claim:** d_{ic} is defined as the minimum number of PCs needed to exceed a variance threshold (example 95%), and later all ICs have $d_{\text{ic}} = 3$ under the 95% criterion.
   - **Checks:** definition consistency, logic consistency
   - **Verdict:** PASS; confidence: medium; impact: moderate
   - **Assumptions/inputs:** The threshold used in Results matches the example threshold in Methods.
   - **Notes:** Methods give 95% as the operative example and Results explicitly use 95%. No symbolic inconsistency.

8. ⚠ **Subspace similarity measure definability** (Sec. 2.2.2, p.3; Sec. 3.4, p.7; Fig. 7 caption text on p.7)
   - **Claim:** Subspace similarity between two 3D PCA subspaces is computed using principal angles (heatmap shows average squared cosine), and the reported average similarity is 0.986.
   - **Checks:** definition completeness, invariance/sanity of metric
   - **Verdict:** UNCERTAIN; confidence: medium; impact: moderate
   - **Assumptions/inputs:** Each per-IC principal subspace basis is orthonormal. A scalar similarity statistic is computed from the principal angles.
   - **Notes:** The text references principal angles and also dot products, but does not provide a precise formula for the scalar 'similarity score' (e.g., $\text{mean}_i \cos^2(\theta_i)$ vs other aggregates). Without the explicit definition, internal verification (range, invariance to basis choice, handling of ordering) is not possible.

9. ✔ **Use of absolute dot product between principal vectors** (Sec. 2.2.2, p.3; Fig. 8 description Sec. 3.4, p.7)
   - **Claim:** Alignment between corresponding principal vectors is measured by $|v_{k,1} \cdot v_{j,1}|$ (and similarly for PC2/PC3).
   - **Checks:** range/sanity check, sign ambiguity handling
   - **Verdict:** PASS; confidence: high; impact: minor
   - **Assumptions/inputs:** Principal vectors are unit-norm. Correspondence is by PCA order (largest to smallest eigenvalue).
   - **Notes:** Absolute dot product is in $[0, 1]$ for unit vectors and removes the eigenvector sign ambiguity. (It still assumes meaningful correspondence by rank, which is standard but can be unstable under near-degenerate eigenvalues; not addressed.)

10. ⚠ **Ambiguity of $(L_k - C_k)$ notation for projections** (Sec. 3.5, p.7)
    - **Claim:** The paper projects centered latent vectors $(L_k - C_k)$ for each IC k onto the 6D global principal subspace to measure captured intrinsic variance.
    - **Checks:** notation/shape consistency, definition completeness
    - **Verdict:** UNCERTAIN; confidence: medium; impact: minor
    - **Assumptions/inputs:** For IC k, there are $N = 10403$ latent vectors, forming a matrix $L_k \in \mathbb{R}^{N\times 10}$. C_k is a 10-vector centroid broadcast across rows (or equivalently an $N \times 10$ repeated matrix).
    - **Notes:** As written, L_k is not defined (single vector vs set/matrix). The intended operation is standard, but the paper should specify shapes and whether the projection is applied per sample and aggregated for variance.

11. ✖ **PDE dimensionality vs indexing/domain** (Title/Abstract p.1; Sec. 2.1–2.2 pp.2–3; throughout Results)
    - **Claim:** The study concerns a PINN solving the 2D Burger's equation, while latent vectors are indexed over (x, t) on a grid with only one spatial coordinate.
    - **Checks:** definition consistency, symbol/domain consistency
    - **Verdict:** FAIL; confidence: high; impact: critical
    - **Assumptions/inputs:** A '2D' PDE would require either two spatial coordinates or a clearly defined alternative meaning of '2D'.
    - **Notes:** The paper never defines a second spatial coordinate (y) or a 2D spatial grid; all sampling is over x and t only. If '2D' is intended to mean two velocity components (u, v) in 1D space, this must be stated and the PDE form given; otherwise the claim is inconsistent with the defined domain.

12. ⚠ **Feature-count bookkeeping (13 total features)** (Sec. 2.1, p.2)
    - **Claim:** The 13 features consist of predicted solution components (e.g., u and v) and a 10D latent vector.

- **Checks:** dimension/units bookkeeping
- **Verdict:** UNCERTAIN; confidence: high; impact: minor
- **Assumptions/inputs:** Example 'u and v' implies 2 components.
- **Notes:** If there are only two solution components plus 10 latent dimensions, that totals 12, not 13. Either there is an additional predicted quantity or the example is incomplete.

Limitations

- The PDF text provides essentially no explicit equations for the Burgers PDE, the PINN loss, PCA covariance formulas, or the exact subspace-similarity statistic; several checks are therefore limited to definitional/shape consistency rather than step-by-step derivation verification.
- Figures are referenced for many claims; this audit does not validate any numeric outcomes depicted in plots (variance percentages, heatmap values, ordering), only whether the surrounding mathematics/definitions are internally coherent.
- Because no explicit projection/variance-capture formulas are provided, statements about 'percent intrinsic variance captured by a subspace' cannot be symbolically verified beyond noting that such a computation is well-defined if properly specified.

Numerical results audit

This section audits **numerical/empirical** consistency: reported metrics, experimental design, baseline comparisons, statistical evidence, leakage risks, and reproducibility.

13 numeric/logic checks were run: 11 PASS, 0 FAIL, 1 UNCERTAIN due to a checker parsing issue (C8), and 1 PASS that is explicitly a feasibility-only check (C7). Core dimension/product relationships and most reported percentage relationships are internally consistent within stated tolerances; two cumulative-percentage totals differ from summed components by 0.01 percentage points (rounding-consistent).

Checked items

1. ✔ **C1** (Methods §2.1 (page 2))
   - **Claim:** The provided NumPy array 'data_bundle' has dimensions $(101, 103, 25, 13)$.
   - **Checks:** dimension_consistency
   - **Verdict:** PASS
   - **Notes:** Derived counts align with other stated quantities: points_per_IC= $101 \times 103 = 10403$; total_points= $101 \times 103 \times 25 = 260075$; latent_dim stated as 10 with 13 features implies 3 leftover features.
2. ✔ **C2** (Methods §2.1 (page 2))
   - **Claim:** Reshaping per initial condition yields $101 \times 103 = 10403$ latent vectors (points).
   - **Checks:** product_equals_claim
   - **Verdict:** PASS
   - **Notes:** Computed $101 \times 103 = 10403$ matches the claim exactly.
3. ✔ **C3** (Methods §2.2.1 Global PCA (page 3) and Results §3.1 (page 4))
   - **Claim:** All latent vectors across all ICs: $101 \times 103 \times 25 = 260075$ latent vectors; also stated as $10403 \times 25 = 260075$.
   - **Checks:** product_equals_claim_crosscheck
   - **Verdict:** PASS
   - **Notes:** Both products agree: $101 \times 103 \times 25 = 260075$ and $10403 \times 25 = 260075$.
4. ✔ **C4** (Methods §2.1 (page 2))
   - **Claim:** Extracting last 10 components from 13 features yields 'latent_space_data' of shape $(101, 103, 25, 10)$.
   - **Checks:** dimension_difference_consistency
   - **Verdict:** PASS
   - **Notes:** 13 features with a 10D latent vector implies 3 remaining features; consistent with the stated extraction.
5. ✔ **C5** (Results §3.1 (page 4))
   - **Claim:** Global PCA: PC1=60.12%, PC2=23.44%, PC3=12.93%; cumulative first three = 96.48%.
   - **Checks:** percentage_sum_to_cumulative
   - **Verdict:** PASS
   - **Notes:** Sum of reported PCs is 96.49% vs stated 96.48%; within abs tolerance (rounding-consistent).
6. ✔ **C6** (Results §3.1 (page 4))
   - **Claim:** Including PC4=1.30%, PC5=1.17%, PC6=0.76% brings cumulative variance to 99.72%.
   - **Checks:** percentage_sum_to_cumulative
   - **Verdict:** PASS
   - **Notes:** Summed cumulative is 99.71% vs stated 99.72%; within abs tolerance (rounding-consistent).
7. ✔ **C7** (Results §3.1 (page 4))

- **Claim:** The remaining four components individually explain less than 0.3% of the variance each.
- **Checks:** residual_budget_bound
- **Verdict:** PASS
- **Notes:** Feasibility-only check: residual $100 - 99.72 = 0.28\%$ is $< 0.3\%$, so the claim is arithmetically possible, but individual PC7–PC10 values are not verified.

8. ⚠ **C8** (Results §3.2 (page 5))
   - **Claim:** Per-IC average explained variances: PC1=59.61%, PC2=23.72%, PC3=14.15%; average cumulative first three = 97.48%.
   - **Checks:** percentage_sum_to_cumulative
   - **Verdict:** UNCERTAIN
   - **Notes:** Automated checker did not execute this sum due to an unrecognized label/variant for the cumulative field; no arithmetic verification result produced.
9. ✔ **C9** (Results §3.2 (page 5))
   - **Claim:** Average cumulative variance explained by first three per-IC PCs is 97.48% with standard deviation 0.15%.
   - **Checks:** std_dev_nonnegative_sanity
   - **Verdict:** PASS
   - **Notes:** Sanity checks pass: std is nonnegative; mean and mean±std lie within $[0, 100]$.
10. ✔ **C10** (Results §3.3 (page 6))
    - **Claim:** Centroid PCA: CPC1=99.86%, CPC2=0.10%, CPC3=0.02%.
    - **Checks:** partial_percentage_sum_leq_100
    - **Verdict:** PASS
    - **Notes:** Sum is 99.98% ($\leq 100\%$); residual to 100% is $\sim 0.02\%$, consistent with rounding.
11. ✔ **C11** (Results §3.4 (page 7) and Conclusions (page 10))
    - **Claim:** Average subspace similarity is 0.986 with standard deviation 0.014; minimum observed similarity 0.954; conclusion says average exceeds 0.98.
    - **Checks:** inequality_chain_sanity
    - **Verdict:** PASS
    - **Notes:** Inequalities are satisfied ($0 \leq \mathbf{min} \leq \mathbf{mean} \leq 1$, std$\geq 0$) and mean= 0.986 exceeds 0.98; computed mean$-3\times$std= $0.944 \geq 0$.
12. ✔ **C12** (Results §3.4 (page 7))
    - **Claim:** PCA on per-IC principal vectors: for v_{k1} first PC explains 85.45%; for v_{k2} 80.04%; for v_{k3} 97.67%.
    - **Checks:** percentage_range_sanity
    - **Verdict:** PASS
    - **Notes:** All three percentages lie within $[0, 100]$.
13. ✔ **C13** (Results §3.5 (page 7))
    - **Claim:** Global 6D subspace captures 99.72% total variance; on average captures 99.66% of each IC's intrinsic variance; minimum capture 99.24%.
    - **Checks:** inequality_sanity
    - **Verdict:** PASS
    - **Notes:** Sanity checks pass: min (99.24%) $\leq$ average (99.66%), and all percentages are within $[0, 100]$.

Limitations

- Only parsed text from the PDF was used; numeric values embedded solely in plots/figures were not extracted.
- Several results summarize computations (PCA eigenvalues, pairwise similarity matrices) without listing underlying arrays, limiting verification to arithmetic relationships among reported summary numbers.
- No access to the referenced NumPy arrays ('data_bundle', 'latent_space_data') or code outputs, so checks are restricted to internal consistency of stated dimensions, products, sums, and inequality sanity.
- One automated check (C8) returned UNCERTAIN due to an execution/check-parsing issue ("Unrecognized percentage_sum_to_cumulative variant"), so the per-IC average cumulative (97.48%) was not arithmetically verified by the checker.