Response to Authors – Reviewer Comments

Dear authors,

I appreciate the effort in proposing a new metric for dimensionality reduction quality. Your attempt to quantify how well local and global structures are preserved using a coefficient bounded between 0 and 1, and validating it via HDBSCAN clustering quality, is commendable.

However, I would like to highlight several critical points:

1. **Theoretical Justification**:
   Your manuscript does not reference the **Johnson–Lindenstrauss lemma**, which is central to the theoretical foundation for structure preservation in dimensionality reduction. This lemma suggests that to preserve the geometric structure when projecting from high to low dimensions, pairwise distances must be approximately maintained. Your metric appears to align with this goal, but the manuscript lacks a formal connection to this principle. I strongly recommend including a theoretical section or justification referencing this. There are generally two main perspectives regarding the application of this lemma: one that considers it valid, and another that relies on approximations. Although the original UMAP paper claims to approximate this behavior by projecting the data into a lower-dimensional space that preserves the high-dimensional structure, it has already been demonstrated that **multidimensional PCA space** performs better in this task (see 10.1371/journal.pcbi.1011288). However, PCA carries certain risks related to similarity preservation, especially due to its reliance on linear combinations of eigenvectors obtained through eigendecomposition (see 10.3390/biology13070512). In this one from sc-PHENIX (see 10.3390/biology13070512), it has been shown that PCA has the phenomenome of distance concentration issues. All these articles should be cited and discussed.
2. **Curse of Dimensionality and Distance Concentration**:
   The paper does not mention the **curse of dimensionality**, particularly the **concentration of distances**, which affects methods like HDBSCAN in high-dimensional settings (e.g., genomic or transcriptomic data, medical data). In high dimensions, pairwise distances tend to converge, making clustering unreliable. While your current evaluation uses EHR data, this alone is not a sufficient justification for the generalizability of the metric. A discussion of these limitations and their implications is necessary.

3. **Implementation Issues**:
   I attempted to reproduce your metric using your official Python implementation available at PyPI, using the MNIST dataset. Given that MNIST contains ~700 features (pixels) and is commonly used to test both local and global structure preservation, 5,7,8 are together for example a qualitative indication that global structure preservation is in the embedding. And local structure preservation is visualised as well separated classes. sc-PHENIX article mentions that. I expected the metric to work without issue. However, your implementation **fails silently or returns NaN** values, depending on the input. I prepared a Google Colab notebook with code and visualizations demonstrating this issue, which I encourage you to review.

   https://colab.research.google.com/drive/1egkOLNcDhogKIVfmLQGnMmTBQNiJQucT?usp=sharing

   I suggest the following:

   - Provide better **input documentation** and **clear error messages and possible bugs**.

- Host your implementation in a **public GitHub repository**, with both Python and R versions clearly documented.

- Link to any datasets used in the paper and allow reproducibility of figures and results. In case of python use notebooks.

4. **About the Metric's Assumptions**:
   The manuscript assumes that **all structures (local and global)** can and should be preserved. However, this assumption is **not well-defined** and may not hold for all dimensionality reduction tasks, spatially UMAP if it handles. I recommend engaging with the literature that critically discusses this, such as the **sc-PHENIX** paper. That work explains in detail how structure preservation can be subjective and dependent on the manifold geometry, data density, and downstream goals.

5. **PCA vs UMAP**:
   While **PCA** often preserves **global distances** effectively (as discussed in methods like **PHATE** and **sc-PHENIX**), it does not follow the **manifold structure** of the data and typically fails to capture **local neighborhood relationships**. This limitation is particularly critical when working with complex **clinical** or **biological datasets**, such as **single-cell RNA-seq** or **transcriptomic patient data**, where **local structure preservation** is essential for identifying cellular phenotypes or disease subtypes.

   Although **PCA** may yield high values under your proposed **Saturn coefficient**, this could be **misleading** in a real-world clinical context. For example, when I tested your metric using the **synthetic data** provided in your **pip tutorial** (from PyPI), the results indicated that **multidimensional PCA** had a higher **Saturn coefficient** than **UMAP**. This is concerning because **UMAP** is widely regarded as more effective in preserving **local structure**, which is often the most biologically relevant aspect in clinical studies.

   Also, I observed that the **Saturn coefficient consistently diminishes as the number of samples increases**. What is going on? This is concerning — it suggests that the **Saturn coefficient may be sensitive to sample size**, which could undermine its reliability. I am using **random synthetic data**, so ideally, the metric should be **robust to such variation**. I urge the authors to **either justify this behavior explicitly** or consider **revising the implementation to mitigate this effect** in future versions of the method.

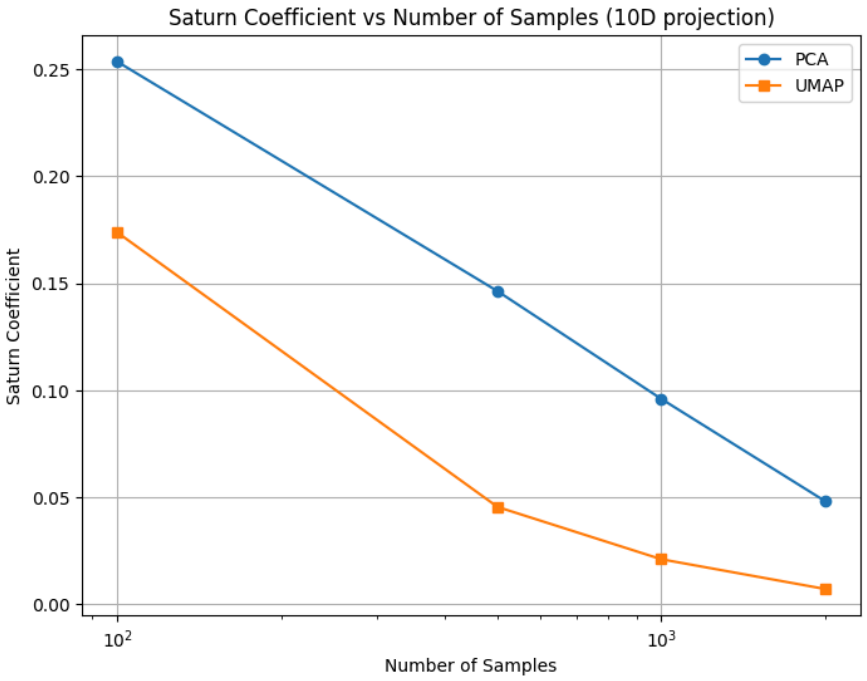

These findings suggest that the **Saturn coefficient** may be biased toward **global distance preservation**, and thus may not fully reflect the capacity of a dimensionality reduction method to retain clinically meaningful patterns, such as **subtype separability** or **patient stratification**.

Furthermore, validation metrics like **HDBSCAN** combined with the **Adjusted Rand Index (ARI)** might not be sufficient or appropriate to benchmark **local structure preservation** in such contexts. I suggest using MELLON and observe where the density of the data is.

To enhance the clinical applicability of your method, I strongly recommend testing it on real-world **clinical datasets**, particularly those where **manifold continuity** and **biological interpretability** are paramount.

6. **Conclusion**:
   In my tests, your metric appears to effectively capture **global relationships**, but I remain unconvinced of its ability to reliably assess **local structure preservation**. This distinction is critical. I strongly encourage the authors to clarify what specific aspect of structure their metric is intended to quantify: is it **neighborhood consistency**, **manifold integrity**, or **global distance preservation**? This should be explicitly stated in the manuscript. In my tests using both **multidimensional PCA** and **UMAP**, the results for your coefficient suggest strong global alignment, but do not provide evidence of reliable local preservation—especially when evaluated using **HDBSCAN** with a ari **clustering metric**. Please demonstrate this distinction, or make the necessary modifications to your method to show how it can effectively capture **global**, **local**, or **combined structure preservation**. This could involve proposing a **new implementation** or an **extension** of your current metric. It would also be valuable to include biologically realistic datasets such as **MNIST** or **scRNA-seq**, where preserving different structural aspects is particularly important. There is a interesting continuum microarray dataset that you could used it is microarray data of different time points (https://www.mdpi.com/2079-7737/13/7/512), probability for continuum structure preservation?

   Additionally, since the method is unsupervised, it remains unclear how well it performs in assessing local structure without external validation. I would also like to know whether the metric is strictly intended for low-dimensional visualization (e.g., 2D or 3D), or if it can also be applied in general dimensionality reduction scenarios beyond visualization tasks. Please elaborate on these points.

I suggest submitting your manuscript to a **preprint server** (if not already done), or at least including **supplemental material** with detailed **reproducibility instructions** and **code examples** (e.g., Python and R). This would greatly facilitate community adoption and validation.

Additionally, I encourage the authors to **clarify or reconsider** their **definition of structure preservation**—whether it refers to **local**, **global**, or potentially **continuum structures**. In this field, the notion of "preserving structure" remains highly **ambiguous** and is often dependent on the assumptions and priorities of each **dimensionality reduction algorithm**. For example, some methods prioritize **linear relationships** (e.g., PCA), others aim to capture **manifolds of varying densities** (e.g., t-SNE, PHATE), **local neighborhood continuity** (e.g., UMAP), or even **diffusion components**.

There are several existing **benchmark datasets** and methodologies that could be used to evaluate these structural aspects more rigorously. I suggest incorporating some of these into your evaluation framework to better contextualize the **Saturn coefficient** within the landscape of structure-preserving methods.

I made sure, for example using your PIP tutorial, that the original space and the low-dimensional space are of the same type (class).

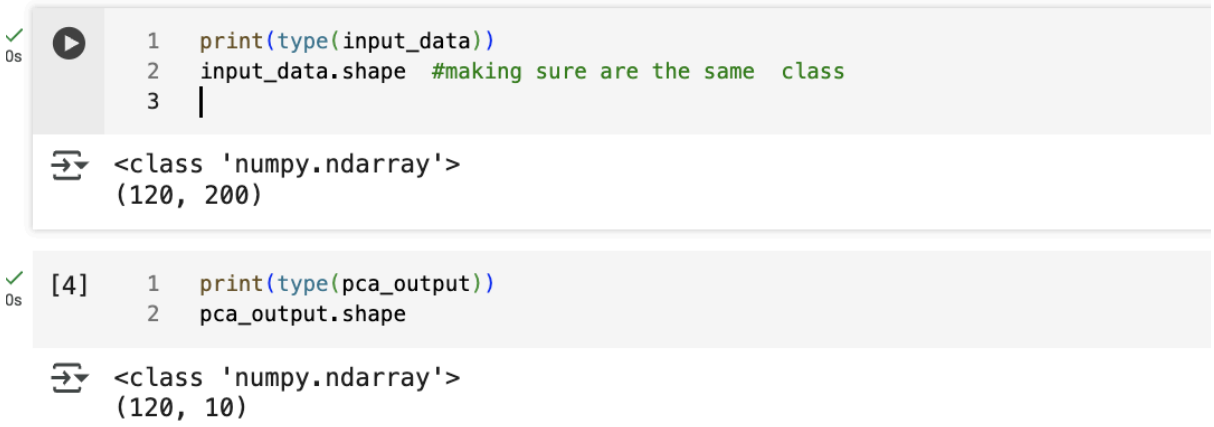

```
print(type(input_data))
input_data.shape  #making sure are the same  class
```

```
<class 'numpy.ndarray'>
(120, 200)
```

```
[4]
print(type(pca_output))
pca_output.shape
```

```
<class 'numpy.ndarray'>
(120, 10)
```

I made sure, for example using your PIP tutorial, that both the original space and the low-dimensional space were of the same type (class). I applied the same procedure to MNIST, but the result returned NaN. Why does this happen?

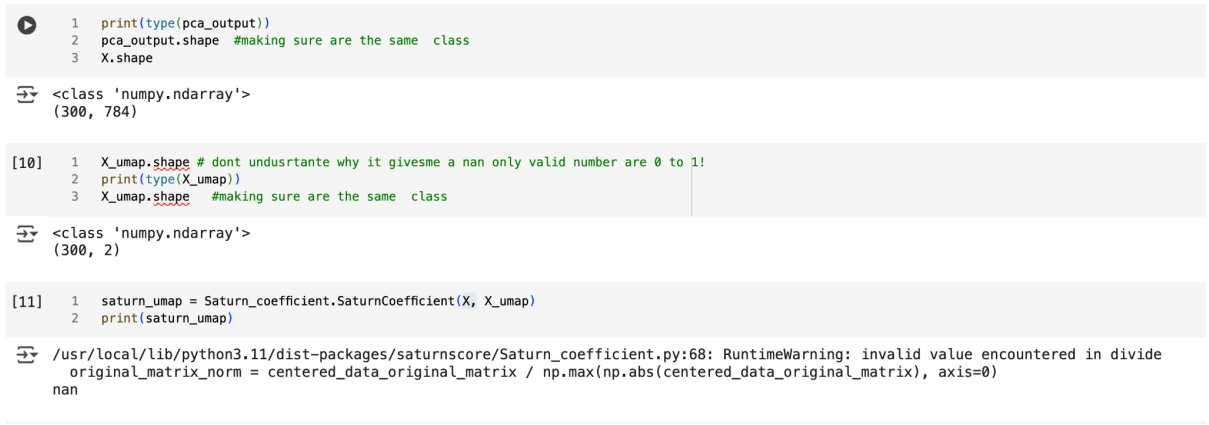

```
print(type(pca_output))
pca_output.shape  #making sure are the same  class
X.shape
```

```
<class 'numpy.ndarray'>
(300, 784)
```

```
[10]
X_umap.shape # dont undusrtante why it givesme a nan only valid number are 0 to 1!
print(type(X_umap))
X_umap.shape   #making sure are the same  class
```

```
<class 'numpy.ndarray'>
(300, 2)
```

```
[11]
saturn_umap = Saturn_coefficient.SaturnCoefficient(X, X_umap)
print(saturn_umap)
```

```
/usr/local/lib/python3.11/dist-packages/saturnscore/Saturn_coefficient.py:68: RuntimeWarning: invalid value encountered in divide
  original_matrix_norm = centered_data_original_matrix / np.max(np.abs(centered_data_original_matrix), axis=0)
nan
```

Final note for eticals:

I would like to disclose that I am the author of *sc-PHENIX*, a published article that discusses several of the conceptual issues raised in this manuscript. In particular, *sc-PHENIX* highlights that PCA generally preserves global structure, while UMAP more effectively captures local neighborhood continuity. The metric proposed in the current manuscript appears to favor global structure preservation, rather than offering a balanced evaluation of both global and local structures. *sc-PHENIX* also emphasizes how the concentration of distances in high-dimensional spaces can distort neighborhood relationships, making it difficult to meaningfully separate distinct clusters. In such cases, PCA may struggle to resolve local structures, whereas UMAP is better suited to uncovering cluster separation in

non-linear manifolds. That said, in high-dimensional settings, multidimensional PCA embeddings can still approximate the geometry of the original space more faithfully than aggressive low-dimensional projections, highlighting the trade-offs between interpretability and structural fidelity. This conceptual overlap is central to the feedback provided in my review.

I have explicitly given the authors the option to cite other references that address the ambiguity in defining “structure preservation,” since this issue is not exclusive to my own work. However, these aspects are not adequately addressed in the current version of the manuscript. I believe it is essential that the authors acknowledge this ambiguity and contextualize it within the broader literature.

Additionally, I have independently tested the implementation provided by the authors and encountered several issues—such as inconsistent behavior and NaN outputs—which I describe in detail in my review. These technical limitations further support the need for clarification, validation, and potentially revisions of the method presented.