EuroVis Workshop on Visual Analytics (2023)
M. Angelini and M. El-Assady (Editors)

Nonparametric Dimensionality Reduction Quality Assessment based on Sortedness of Unrestricted Neighborhood

Davi Pereira-Santos[1], Tácito T. A. T. Neves[2], Andre C.P.L.F. de Carvalho[1], and F. V. Paulovich[3]

[1]Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, Brazil
[2]NCEx – Federal University of Alagoas (UFAL), Arapiraca, Brazil
[3]Eindhoven University of Technology: Eindhoven, Noord-Brabant, NL

Abstract

High-dimensional data are known to be challenging to explore visually. Dimensionality Reduction (DR) techniques are good options for making high-dimensional data sets more interpretable and computationally tractable. An inherent question regarding their use is how much relevant information is lost during the layout generation process. In this study, we aim to provide means to quantify the quality of a DR layout according to the intuitive notion of sortedness of the data points. For such, we propose a straightforward measure with Kendall τ at its core to provide values in a standard and meaningful interval. We present sortedness and pairwise sortedness as suitable replacements over, respectively, trustworthiness and stress when assessing projection quality. The formulation, its rationale and scope, and experimental results show their strength compared to the state-of-the-art.

CCS Concepts

• ***Human-centered computing*** → ***Visualization design and evaluation methods; Visual analytics;*** • ***Computing methodologies*** → ***Dimensionality reduction and manifold learning;***

1. Introduction

Visualization techniques provide useful tools for data exploration and decision-making. If a data set has up to two (or three) attributes (or dimensions), scatter plots can be good alternatives for exploring intuitive notions, such as spatial distribution, distance, and neighborhood. However, some sort of dimensionality reduction is usually necessary for data sets with more dimensions. In the realm of multidimensional visualization strategies, *Dimensionality Reduction (DR)* techniques have proven their usefulness [NA19]. Despite the advances in the field, it is inevitable that the mapping process, from a high-dimensional to a visually low space, incurs in loss of information due to the high intrinsic dimensionality [FO71] of many data sets. Therefore, DR techniques are often subject to a trade-off between interpretability and reliability, raising the question of how to effectively assess the impact of the implicit loss of information from applying a DR technique.

In the current literature, different quality measures have been proposed Section 2. Despite their popularity, the existing measures have limitations, especially the adoption of arbitrary/empirical thresholds to interpret the quality of a layout and the strong dependence on parametrization. In this work, we address these problems by proposing the concept of *sortedness*, able to provide a value within a meaningful interval of how well-preserved is a data set structure after a mapping process. Sortedness is the level of agreement between a data set and its DR layout regarding the order of the points as measured by the Kendall τ correlation index in relation to a given reference point [Vig15]. Adopting a correlation index provides more meaningful values, i.e. from a standard interval, and consistency to the measure by making it commensurable across different tasks and domains. We also present variants.

Our experimental results show evidence that sortedness is better suited to assess DR layout quality in general, and also in some specific scenarios, than the current practice of adopting metric/non-metric Kruskal *stress*, *trustworthiness*, and other related measures [KNO*03, Kru64b, Kru64a]. Additionally, sortedness is able to indirectly assess the preservation of features usually relied upon by the human interpretation bias during data visualization, and provide a measure for machine learning pipelines.

2. Related Work

In this section, we present DR techniques as they are a frequently used type of data transformation where quality assessment is of interest, and relevant evaluation measures from the literature.

Dimensionality Reduction techniques aim to find a low-dimensional representation of high-dimensional data [EMK*21, NA19]. Existing methods map data points into graphical elements to preserve pairwise distances or neighborhoods. We can classify the techniques as *global* or *local* according to the type of distance intended to be preserved, respectively: all points; and, small neighborhoods. Both aspects are considered in our proposal.

DOI: 10.2312/eurova.20231093

Global techniques fail to preserve neighborhood relationships, especially when sparse high-dimensional spaces are considered [PNML08]. Kruskal [Kru64a] presented the Multidimensional Scaling (MDS) technique to map points from a high-dimensional space to a low-dimensional space by optimization. It minimizes the quadratic difference between the dissimilarities established in the original space and the calculated distances in the transformed space. This quantity is known as *stress*, adopted as the reference in this study. Least Squares Projection (LSP) [PNML08] introduces a different bias into the data transformation which first builds a neighborhood graph between the points. Next, it selects a subset of points to project. Remaining points are handled through interpolation by solving a system of linear equations. Many other global techniques exist. Each approach introduces its specific bias into the data transformation.

Local techniques are intended to preserve the neighborhood relationships. They help to identify groups and define their boundaries, especially for high-dimensional data sets [FFDP15]. The Piecewise Laplacian-based Projection (PLP) [PEP*11] addresses the problem of sensitivity to the positioning of control points, found in techniques such as LSP, by splitting the data set into smaller subsets. The Local Affine Multidimensional Projection (LAMP) LAMP [JCC*11] allows a user-controlled redefinition of the mapping matrix based on a first mapping of control points. This kind of application would benefit directly from our proposed intuitive measure to guide the user through the interactive process. Conversely, some techniques have probabilistic bias. An example is Stochastic Neighbor Embedding (SNE) [HR02]. Each mapped point is positioned next to a group of its original neighbors with a given size. Probability distributions representing the chance of each point choosing another as a neighbor are defined, with higher probabilities assigned to closer points. Finally, local approaches have an inherent type of bias which affects the quality of the DR layout in a very different way than global techniques. Our proposed measures address such diversity by considering the unrestricted neighborhood ordering which encompasses many notions that are intuitive to the human bias, differently from the stress measure which was created having MDS in mind, and from measures limited to a certain number of neighbors.

Evaluation measures help to assess the quality of maps created by DR techniques. Most of them aggregate quality measures based on distance, neighborhood, or cluster segregation measures. However, the values returned by them are not as interpretable as those provided by our proposed measure, sortedness. A popular group of measures proposed to quantify the preservation of distances after the DR process is the set of stress functions. One of the most known is the Kruskal [Kru64a] stress function, represented in this text by σ_1, which measures the difference between the distances calculated in the original space and those calculated in the projected space. The stress value ranges from 0 to 1, when normalized. The smaller the value, the higher the quality. When the computed dissimilarities are not metric, e.g., based on ranking positions, the non-metric Kruskal [Kru64b] stress function, represented by σ_*, can be used. Unreasonably good stress values can be observed for very disordered DR results [PNML08]. Furthermore, similar stress functions can lead to different perceptions of quality. When the target characteristic to be preserved is related to the neighborhood, one can adopt the Neighborhood Preservation (NP) [PM08]. This measure evaluates how many nearest neighbors established in the original space remain as nearest neighbors in the projected space. Its range is also in the interval $[0,1]$, with values closer to 1 representing better neighborhood preservation. While this measure is interpretable, it lacks the standard meaning of our proposed interval based on a correlation index. Other measures are dependent on a parameter. *Trustworthiness* [KNO*03], represented in this text by T_k, depends on the number of neighbors k. It measures the precision of the low-dimensional neighborhoods regarding false positives. *Continuity* [KNO*03] measures the recall of the low-dimensional neighborhoods regarding false negatives. Overall, it is very similar to T_k.

3. Proposed Method

In this section, we introduce the concept of *sortedness* and its respective numeric representation. Sortedness means how well-preserved is the data set structure after a transformation regarding the order of the points. The value is calculated by a function that returns the level of agreement between two sets of points as measured by the Kendall τ correlation index in relation to a reference point [Ken38]. This is the most commonly used statistics to understand the correlation between two different scores for the same set of items [Vig15]. We present in the next subsections the function to evaluate the local sortedness (i.e., for a given point), which can potentially replace trustworthiness, and its reciprocal version, which considers neighborhood in the same perspective as adopted in the *hubness* concept [TRMI13] (Section 3.1). The measures do not depend on a strong parameter like trustworthiness and can be considered non-parametric. We also present a function to evaluate the global (pairwise) sortedness which is sensitive to distortions beyond neighborhood ordering changes (Section 3.2) while still ignoring irrelevant perturbations, which is a potential issue with the Kruskal stress formula I - see Figure 1 in Section 4. Additionally, the pairwise sortedness is generalized to accept a weighting scheme that turns it into a local measure. The other proposed local variants differ by depending on a ranking of neighbors according to their distance to a reference point. The global variant, on the other hand, depends on a ranking of all pairwise distances. The weighting scheme is based on a generalization of the Knight algorithm which has complexity $O(n \log n)$ provided the distances are already ranked [Kni66]. We adopted the Euclidean distance in this work. All measures are provided as an open-source Python package [PSN23].

3.1. Local Neighborhood Sortedness

Let $\rho : \mathbb{R}^d \to \mathbb{R}^2$ be a transformation function (e.g., a dimensionality reduction), and $X \subset \mathbb{R}^d$ a d-dimensional data set. The *sortedness* $\lambda : \mathbb{R}^d \to [-1;1]$ of a given point $\boldsymbol{x} \in X$ projected as $\hat{\boldsymbol{x}}$ onto a resulting set $\hat{X} \subset \mathbb{R}^2$ by ρ is defined by Equation (1).

$$\lambda_{\tau_w}(\boldsymbol{x}) = \tau_w[r_X(\boldsymbol{x}), r_{\hat{X}}(\hat{\boldsymbol{x}})] \quad (1)$$

where $\tau_w(\boldsymbol{a}, \boldsymbol{b})$ is the *weighted Kendall τ correlation index* between rankings $\boldsymbol{a}$ and $\boldsymbol{b}$ [Vig15]. Such variant of the Kendall τ index is weighted by the function $w(i) = (i+1)^{-1}$ defined for each rank $0 \leq i < |X|$; and, $r_X(\boldsymbol{x})$ is the ranking of all other points in X according to their proximity to $\boldsymbol{x}$. In the unlikely case when there are one or more distance ties, the lexicographical order of $r_X(\boldsymbol{x})$ followed by $r_{\hat{X}}(\hat{\boldsymbol{x}})$ is adopted to break them when defining the weight of each point. When the weight function is constant, e.g., $w(i) = 1$, the measure, represented by λ_{τ_1}, has the same behavior of the non-metric stress function σ_* (Section 2). Despite the similarity, λ_{τ_1} provides more meaningful values than σ_*: the intuitive correlation index interval, and a companion p-value for the null hypothesis H_0 (H_0 implies no correlation). For completeness, a complementary measure is presented in Equation (2), the *reciprocal sortedness*.

$$\overline{\lambda}_{\tau_w}(\boldsymbol{x}) = \tau_w[\overline{r}_X(\boldsymbol{x}), \overline{r}_{\hat{X}}(\hat{\boldsymbol{x}})] \tag{2}$$

where $\overline{r}_X(\boldsymbol{x})$ represents the reciprocal neighborhood ranking, i.e., it contains the rank positions of $\boldsymbol{x}$ within each *list of neighbors*. A list of neighbors is provided by each point $\boldsymbol{u} \in X \setminus \{\boldsymbol{x}\}$. In the case of position ties, the same previously mentioned breaking rule applies. In both equations, all rankings assume the points are listed in the same order.

The mean sortedness $\mu = |X|^{-1} \sum_{\boldsymbol{x} \in X} \lambda_{\tau_w}(\boldsymbol{x})$ can be used as a measure of global sortedness of DR layouts as a balance between local and global evaluation.

3.2. Pairwise Sortedness

The *pairwise sortedness* Λ_{τ_1} is defined by Equation (3).

$$\Lambda_{\tau_1}(X, \hat{X}) = \tau(R_X, R_{\hat{X}}) \tag{3}$$

where R_X is the ranking of distances between all combinations of pairs in X. Here, we adopt the original Kendall τ function, i.e., $w(i) = 1$. An advantage of the measure is that it detects a more subtle type of distortion than just changes in the ordering of neighbors: it is sensitive to any change in the relative proximity between points. To illustrate how it works, given three points $a, b, c \in X$, any change in the relationship between original and projected distances like $d(a,b) > d(b,c)$ versus $d[\rho(a), \rho(b)] < d[\rho(b), \rho(c)]$ is penalized by the measure, even if the ordering across a, b, and c is preserved. Such sensitivity is closer to that of the stress function, while still ignoring irrelevant distortions as shown in Figure 1, and with the benefit of providing a p-value. It is a tentative hierarchy of projective distortions, where the higher the position in the stack, the less sensitive the measure. Different from metric stress, none of the variants are differentiable.

The pairwise sortedness Λ_{τ_1} can be generalized to its weighted form Λ_{τ_w} when an independent importance criterion is defined. The proximity of each pair of points to a reference point is a convenient fit for such criterion as it enables the local calculation of the pairwise sortedness. The formula is presented in Equation (4)

$$\Lambda_{\tau_w}(\boldsymbol{x}) = \tau_w[R_X, R_{\hat{X}}, W_X(\boldsymbol{x})] \tag{4}$$

where $W_X(\boldsymbol{x})$ is a ranking that indicates the importance of each pair according to how small the mean of the distances of its two points to $\boldsymbol{x}$ is. The rankings R_X and $R_{\hat{X}}$ are the same for every $\boldsymbol{x} \in X$. Therefore, their values can be calculated only once to assess all data set points.

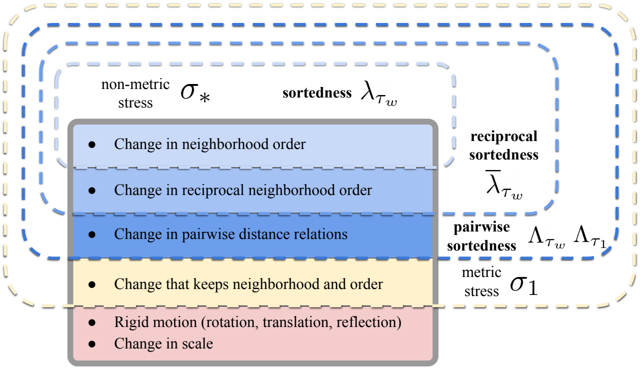

Figure 1: *Overview of measures: each measure is sensitive to narrower subsets of transformations/DR artifacts. σ_1 and λ_{τ_w} are sensitive to strong changes in the data, e.g., neighborhood order. σ_1 is sensitive to small, including irrelevant, changes.*

3.3. Interpretation

All three functions benefit from the meaningful values provided by the Kendall τ rank correlation index. Notable values are 1.0, 0.5, 0.0, and -1.0, meaning respectively: *perfect correspondence*; *half of the sortedness is preserved*; *meaningless transformation* (equivalent to random projection); and, *worst possible transformation* (total unsortedness) which should only happen if it is actively/accidentally engineered to provide the exact opposite of the expected result. These meaningful values contrast with the current practice, which relies upon arbitrary values, e.g., adopting $\sigma_* > 0.200$ as an indication of a low-quality DR, among other arbitrary values for increasing quality (0.100, 0.050, 0.025, and 0.000) [Kru64a]. Therefore, our measures are superior from the interpretability perspective. Each measure is sensitive to a different set of distortion types. Relevant types for this text are illustrated by Figure 2. The distortion degree varies from minor disturbances to major changes in the ordering. The more distorted the DR layout, the more misleading it will be to human interpretation. Rigid motion changes the point of view keeping the data topology.

4. Experiments

The following subsections present an experimental comparison of the proposed and literature measures to illustrate their properties. We applied global and local changes to a set of randomly generated two-dimensional points to simulate projection artifacts. The set is intentionally small to accentuate the effect of such changes. Therefore, both input and output data sets are two-dimensional, allowing to isolate the effects of interest from possible projection biases.

4.1. Subset Randomization

The lesser the correspondence between the data set and the DR layout, the lower its quality for visual interpretation. To simulate this, we uniformly sampled 1000 random points bounded by a 100x100 square. We randomized the location of an increasingly larger subset of the points to show how this progressive change in the ordering of the points affects the measures. Figure 3 shows meaningful near

Proceedings published by Eurographics - The European Association for Computer Graphics.

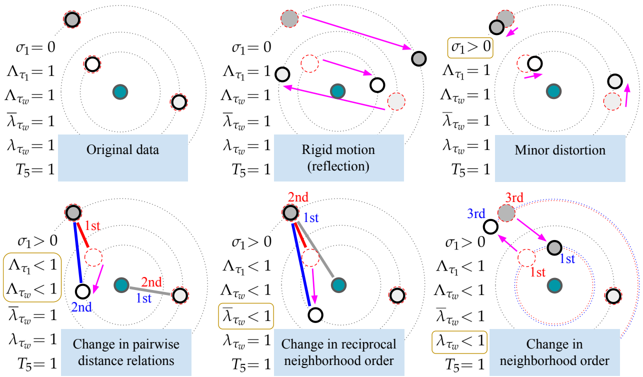

Figure 2: *Examples of transformations/DR artifacts related to a central point from most trivial to most misleading, and respective values for each measure (left to right, starting from the top). Dashed red circles indicate the expected positions for an exact transformation. Bold lines highlight relevant distances for comparison. Red: original. Blue: new. Gray: reference. Refer to Figure 1 to see how artifacts might relate to each other.*

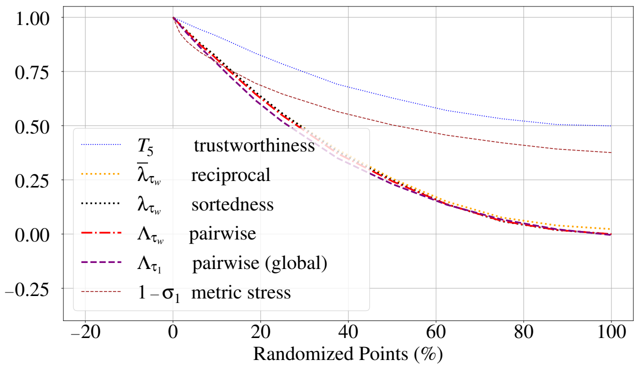

Figure 3: *Effect of increasing the randomized points subset size. The proposed measures match the intuition that the absence of randomization has value 1, while total randomization has value 0.*

zero end values for all proposed variants, while T_5 and σ_1 end values are 0.5 and 0.4, respectively. A more detailed comparison of the local variants is provided by boxplots in Figure 4 where each point in the plot represents the local measure for a data point. Non-global curves are calculated by averaging local values. The boxplot illustrates the stability of the measures across all points. Pairwise sortedness has better overall stability when the whiskers and outliers are considered and a shorter box when higher shuffling levels are applied ($\Lambda_{\tau_w} > 12.5\%$). This can be explained by the calculation of the measure being the same for all points except for the weighting of pairs, i.e., only the weights are relative to the point under measurement. T_5 has the overall tallest boxes, except for level 50% which affects mostly λ_{τ_w} and $\overline{\lambda}_{\tau_w}$.

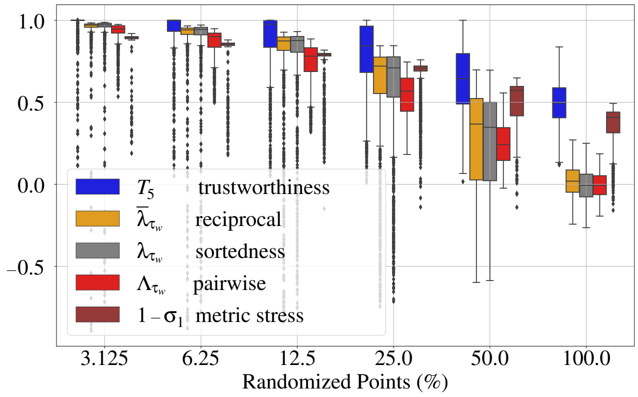

Figure 4: *Effect of increasing the randomized points subset size.*

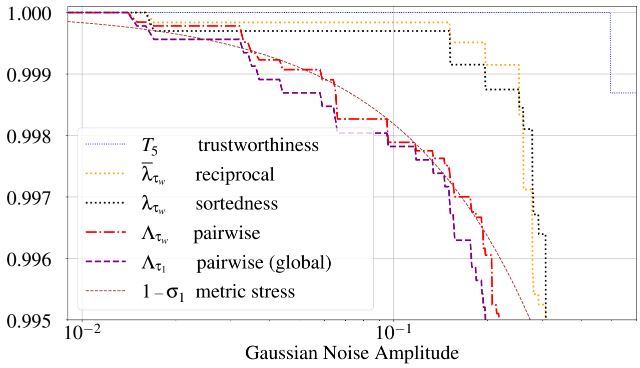

Figure 5: *Distortion level affects σ_1 earlier, and T_5 later.*[‡]

4.2. Gaussian Distortion

The application of Gaussian noise to each point with increasing amplitude is a way to highlight which measure is affected earlier or later as a result of its degree of sensitivity. We sampled 17 points bounded by a 100x100 square, and translated each point in a fixed random direction. The first affected measure is σ_1, followed by the pairwise variants (Λ_τ), and finally, the λ_{τ_w} variants as shown in Figure 5. The last affected measure is T_5, which means it was the least sensitive to Gaussian noise. This shows how sensitive is σ_1 to even the smallest change in the data while T_k is the least sensitive due to its hit-or-miss nature which is limited to k neighbors.

4.3. Single Point Translation

A user trying to interpret a point far from its original neighbors would inevitably draw wrong conclusions. This experiment illustrates how each measure is affected when the point of interest is moved away from its original position. Here, such translated point is part of a set of uniformly sampled 25 collinear points within the interval $[0;50]$ along the x axis. Each natural number in the axis corresponds to an offset in Figure 6. On one hand, it shows that Λ_{τ_1} values are penalized by less than 0.2 in the complete translation (step=50) due to the measure global nature. On the other hand, local measures are directly related to the translated point which makes

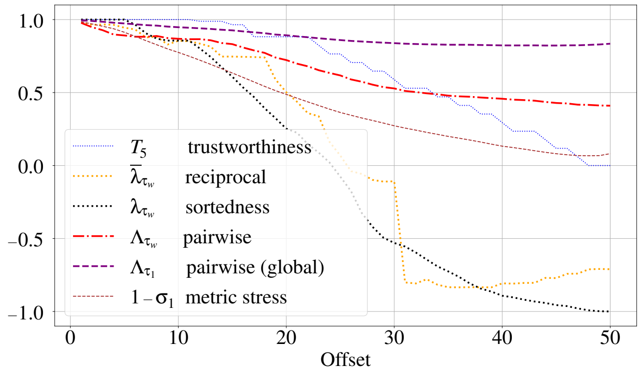

Figure 6: *Effect of single point translation in a linear data set. Non-global curves represent only the values for the translated point.*

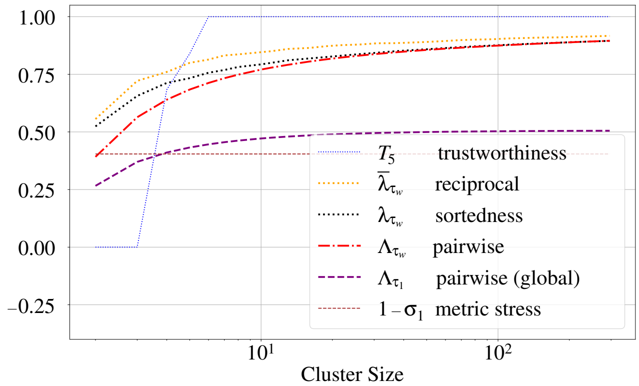

Figure 7: *Effect of swapping two out of three non-overlapping clusters. Curve T_5 was set to zero when undefined.*

them more sensitive to its translation. Despite its focus on weighting pairwise distance relationships in the neighborhood, Λ_{τ_w} has a noticeable decrease in value from 1.0 to around 0.4. This shows the point displacement still penalizes the measure as one would expect from a local measure, instead of exclusively taking into account the remaining perfect ordering of the neighbors among themselves. The penalization is fully present in the case of the λ_{τ_w} variants due to their use of rankings always relative to the point of interest. This is illustrated by the extreme case of $\lambda_{\tau_w} = -1$ in the plot. The value $\lambda_{\tau_w} = 0$ is consistent with the translation having completed half of the linear data set length: half of the points still keep their relative position; the other half is ordered in the exact opposite direction of the expected.

4.4. Global Distortion

We also investigated the effect of changing the overall data structure, while keeping most local structures preserved. This represents changes that affect, e.g., the visual interpretation of how large structures are related inside the DR layout. We created a data set with three non-overlapping clusters with 100 random points each. The global distortion consisted in swapping the position of two clusters. Figure 7 shows the effect of such change for increasing cluster sizes for different measures. Notice that the curve of the global variant (Λ_{τ_1}) is the most penalized by the cluster swap, while the local variants are mainly focused on the fact that the local neighborhood is preserved for the majority of the points. On the other hand, trustworthiness (T_k) is completely unable to detect changes beyond the scope of k neighbors. Λ_{τ_1} behaves similarly to σ_1, with the advantage of providing a meaningful value. Interestingly, the curve stays around 0.5, meaning half of the ordering is lost. Notice that 0 represents a total absence of correlation, and 1 represents a perfect correlation. This suggests that, for the global variant, cluster ordering is as important as local ordering in this experiment. Conversely, the local variants (mean values) are penalized by a decrease between 0.1 and 0.15 (right half of the chart) due to the global change. Among the local variants, Λ_{τ_w} and $\overline{\lambda}_{\tau_w}$ are respectively the most and the less sensitive to global changes, specially for smaller cluster sizes.

5. Conclusion

We presented the concept of sortedness as a measure of how ordered a DR layout is when compared to the original data points.

The simplest variant quantifies the correlation between the original and projected neighborhoods of a given point. When unweighted, sortedness provides an interval of values more meaningful than non-metric stress while still presenting the same behavior, i.e., one increases monotonically with the other. When weighted, sortedness is a suitable non-parametric replacement for trustworthiness. In practice, we consider it non-parametric as the weighting function does not need to be changed. The weighting function recommended in the literature already has the desired behavior of decreasing the importance of each neighbor according to its proximity [Vig15]. Additionally, we presented the reciprocal counterpart of sortedness as a slightly more sensitive measure that considers neighborhoods in the same way adopted in the concept of hubness.

Pairwise sortedness is the most sensitive variant. It quantifies the correlation between the rankings of pairwise distances of a DR layout when compared to the original data. This measure can be a replacement for Kruskal stress formula I. The former is less sensitive to irrelevant changes in the points location than the latter, and, like all proposed variants, the interval of values provided by Kendall τ is more interpretable. When unweighted, it provides a global measure. When weighted, pairwise sortedness is a local (point-wise) measure that can be averaged across all points if a balance between locality and globality is desired.

6. Acknowledgment

This work was supported by Wellcome Leap 1kD Program; São Paulo Research Foundation (FAPESP) - grant 2020/09835-1; Canadian Institute for Health Research (CIHR) Canadian Research Chairs (CRC) stipend [award number 1024586]; Canadian Foundation for Innovation (CFI) John R. Evans Leaders Fund (JELF) [grant number 38835]; Dalhousie Medical Research Fund (DMRF) COVID-19 Research Grant [grant number 603082]; and the Canadian Institute for Health Research (CIHR) Project Grant [award number 177968].

References

[EMK*21] ESPADOTO M., MARTINS R. M., KERREN A., HIRATA N. S. T., TELEA A. C.: Toward a quantitative survey of dimension reduction techniques. *IEEE Transactions on Visualization and Computer Graphics 27*, 3 (2021), 2153–2173. doi:10.1109/TVCG.2019.2944182. 1

[FFDP15] FADEL S. G., FATORE F. M., DUARTE F. S. L. G., PAULOVICH F. V.: Loch: A neighborhood-based multidimensional projection technique for high-dimensional sparse spaces. *Neurocomputing 150* (2015), 546–556. doi:10.1016/j.neucom.2014.07.071. 2

[FO71] FUKUNAGA K., OLSEN D. R.: An algorithm for finding intrinsic dimensionality of data. *IEEE Transactions on Computers 100*, 2 (1971), 176–183. doi:10.1109/T-C.1971.223208. 1

[HR02] HINTON G. E., ROWEIS S.: Stochastic neighbor embedding. *Advances in neural information processing systems 15* (2002). doi:10.5555/2968618.2968725. 2

[JCC*11] JOIA P., COIMBRA D., CUMINATO J. A., PAULOVICH F. V., NONATO L. G.: Local affine multidimensional projection. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 2563–2571. doi:10.1109/TVCG.2011.220. 2

[Ken38] KENDALL M. G.: A new measure of rank correlation. *Biometrika 30*, 1/2 (1938), 81–93. doi:10.1093/biomet/30.1-2.81. 2

[Kni66] KNIGHT W. R.: A computer method for calculating kendall's tau with ungrouped data. *Journal of the American Statistical Association 61*, 314 (1966), 436–439. doi:10.1080/01621459.1966.10480879. 2

[KNO*03] KASKI S., NIKKILÄ J., OJA M., VENNA J., TÖRÖNEN P., CASTRÉN E.: Trustworthiness and metrics in visualizing similarity of gene expression. *BMC bioinformatics 4*, 1 (2003), 1–13. doi:10.1186/1471-2105-4-48. 1, 2

[Kru64a] KRUSKAL J. B.: Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika 29*, 1 (1964), 1–27. doi:10.1007/BF02289565. 1, 2, 3

[Kru64b] KRUSKAL J. B.: Nonmetric multidimensional scaling: a numerical method. *Psychometrika 29*, 2 (1964), 115–129. doi:10.1007/BF02289694. 1, 2

[NA19] NONATO L. G., AUPETIT M.: Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE Transactions on Visualization and Computer Graphics 25*, 8 (2019), 2650–2673. doi:10.1109/TVCG.2018.2846735. 1

[PEP*11] PAULOVICH F. V., ELER D. M., POCO J., BOTHA C. P., MINGHIM R., NONATO L. G.: Piece wise laplacian-based projection for interactive data exploration and organization. In *Computer Graphics Forum* (2011), Wiley Online Library, pp. 1091–1100. doi:10.1111/j.1467-8659.2011.01958.x. 2

[PM08] PAULOVICH F. V., MINGHIM R.: Hipp: A novel hierarchical point placement strategy and its application to the exploration of document collections. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (2008), 1229–1236. 2

[PNML08] PAULOVICH F. V., NONATO L. G., MINGHIM R., LEVKOWITZ H.: Least square projection: A fast high-precision multidimensional projection technique and its application to document mapping. *IEEE Transactions on Visualization and Computer Graphics 14*, 3 (2008), 564–575. doi:10.1109/TVCG.2007.70443. 2

[PSN23] PEREIRA-SANTOS D., NEVES T.: sortedness/sortedness: Release v0.230501.10, May 2023. URL: https://pypi.org/project/sortedness, doi:DOI:10.5281/zenodo.7888017. 2

[TRMI13] TOMASEV N., RADOVANOVIC M., MLADENIC D., IVANOVIC M.: The role of hubness in clustering high-dimensional data. *IEEE transactions on knowledge and data engineering 26*, 3 (2013), 739–751. doi:10.1109/TKDE.2013.25. 2

[Vig15] VIGNA S.: A weighted correlation index for rankings with ties. In *Proceedings of the 24th international conference on World Wide Web* (2015), pp. 1166–1176. doi:10.1145/2736277.2741088. 1, 2, 5