






Visually Assessing 1-D Orderings of Contiguous Spatial Polygons

Julius Rauscher , Frederik L. Dennig , Udo Schlegel , Daniel A. Keim , and Johannes Fuchs 

University of Konstanz, Germany

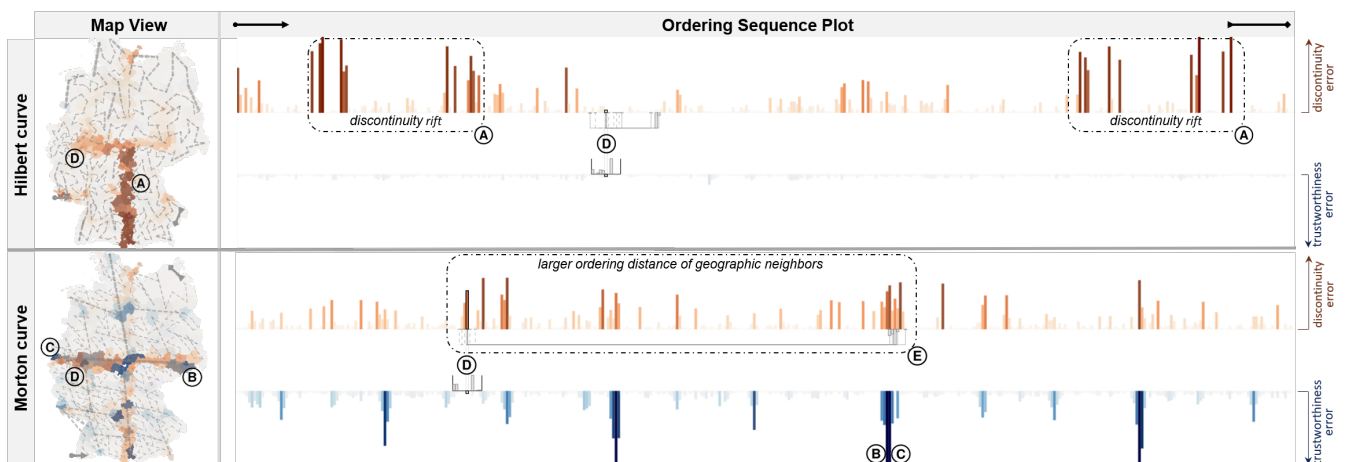


Figure 1: Comparing a Hilbert and Morton curve on the 400 districts of Germany (GER). The Hilbert curve has fewer trustworthiness errors but exhibits a discontinuity rift where geographical neighbors are distant in the ordering (A). The Morton curve performs slightly better regarding discontinuity at the cost of larger jumps in the ordering whenever the diagonal quadrant is switched (B) (C). The geographical neighborhood of a selected polygon (D) is better preserved in the Hilbert layout, indicated by shorter connecting lines of geographical neighbors (E).

Abstract

One-dimensional orderings of spatial entities have been researched in many contexts, e.g. spatial indexing structures or visualizations for spatiotemporal trend analysis. While plenty of studies have been conducted to evaluate orderings of point-based data, polygonal shapes, despite their different topological properties, have received less attention. Existing measures to quantify errors in projections or orderings suffer from generic neighborhood definitions and over-simplification of distances when applied to polygonal data. In this work, we address these shortcomings by introducing measures that adapt to a varying neighborhood size depending on the number of contiguous neighbors and thus, address the limitations of existing measures for polygonal shapes. To guide experts in determining a suitable ordering, we propose a user-steerable visual analytics prototype capable of locally and globally inspecting ordering errors, investigating the impact of geographic obstacles, and comparing ordering strategies using our measures. We demonstrate the effectiveness of our approach through a use case and conducted an expert study with 8 data scientists as a qualitative evaluation of our approach. Our results show that users are capable of identifying ordering errors, comparing ordering strategies on a global and local scale, as well as assessing the impact of semantically relevant geographic obstacles.

CCS Concepts

• **Human-centered computing** → **Geographic visualization; Visual analytics;**

1. Introduction

The mapping of spatial entities to a one-dimensional (1-D) ordering is a valuable technique that enhances spatial data analysis by improving efficiency in spatial indexing structures [Gut84, Sam84] or revealing

spatiotemporal patterns in visualizations [BJC*19, FMKK21]. Nevertheless, an ordering of spatial entities inevitably produces errors, i.e., spatially close objects are placed far apart in the ordering sequence [BP92]. Finding the most suitable ordering heavily depends on the intended use case and becomes a non-trivial task, considering

that the number of possible orderings grows factorially with the number of elements ($N!$). While the exploration of all solutions is computationally unfeasible, designated algorithms such as simulated annealing [GS96], agglomerative hierarchical clustering [GG06], spatial indexing structures [BJC*19], or space-filling curves [FMKK21] aim to approximate a near-optimal result heuristically.

To globally evaluate the quality of a 1-D ordering of point-based geometries, various measures have been introduced [MD86, DM94, BP92, GS96, VK01, GG06]. However, many spatio-temporal trends, such as the development of socio-economic indicators or the spreading of epidemical cases, are captured with polygons as their underlying spatial representation [TJMM20, JSS*22]. While polygonal shapes can be simplified to points using the centroid, this abstraction results in an oversimplified representation since the contiguity between polygons is neglected. Previous studies emphasize that precise measures of topography and similarity are crucial for obtaining reliable evaluation results [GS96], hence we argue for a distance computation based on shared border percentages to account for the topological property of contiguity. As seen in Fig. 2, centroid-based distance measures struggle to accurately capture spatial coherence of polygons, as they assign different nearest neighbor ranks compared to a contiguity-based approach, especially in the presence of natural obstacles such as rivers. The neighborhood preservation also plays an important role when defining such measures, where existing approaches employ the k -nearest neighbor strategy [BP92, GS96, VK01, GG06]. Since real-world polygon-based datasets embody a complex configuration of polygons, we argue that a constant neighborhood size k for all polygons results in an oversimplification of irregular neighborhoods (see Fig. 2). To mitigate this, we present an approach where the neighborhood is defined by all contiguous, neighboring polygons.

Global error measures provide an overall assessment of an ordering, but they often fail to capture local nuances and individual discrepancies crucial for accurate error tracking. Since mapping errors are inevitable, it is essential to identify their geographical locations and their impact within the 1-D ordering. Depending on the use case, certain errors may be tolerable in some regions, but in other areas, even minor discrepancies can lead to misleading interpretations. Visual analytics can play a pivotal role in this context to effectively identify and explore mapping errors on a local scale.

While visual methods to investigate projection-based errors have been proposed [SvLB10, MCMT14, MMT15], they typically assume a 2-D target space, e.g. 2-D scatterplots. In contrast, a 1-D order mapping introduces distinct challenges due to the strict linear, sequential arrangement of only one single dimension. This limits the capability to preserve spatial clusters or neighborhood structures, thus necessitating a customized approach for this specific scenario. To address the limitations of the existing approaches, we contribute the following:

- Novel measures that apply a distance measure based on shared border percentages and utilize varying neighborhood sizes.
- A visual analytics interface designed to support geographic visualization experts in comparing various ordering strategies and identifying mapping errors within them.
- A qualitative expert study and an exemplary use case to evaluate the effectiveness of our approach.
- A comparison of established ordering strategies across spatially diverse datasets from which we derive design considerations.

2. Related Work

In this chapter, we explore aspects essential to understanding the landscape of 1-D spatial orderings: generation methods, quantitative and qualitative evaluation techniques, and their applications.

Generating 1-D Orderings – Spatial indexing structures such as R-Tree [Gut84] or Quad Trees [Sam84] yield an ordering when using a depth-first search. Space-filling curves (SFC) such as Hilbert or Morton curves create an ordering by traversing all points in a given space in a continuous sequential path, where context-based [DCOM00] or data-driven [ZJW21, WGDS22] variants offer a more refined solution. Agglomerative hierarchical clustering (AHC) methods can also be leveraged to generate a 1-D ordering by traversing the leaf nodes of the dendrogram [GG06]. Various dimensionality reduction methods [EMK*21] are also capable of transforming spatial entities into a 1-D encoding. Aside from strictly spatial input topologies, the Traveling Salesman Problem (TSP) [LK75] provides a 1-D ordering by connecting spatial objects through the shortest path and removing the longest link to create a Hamiltonian path. When transforming the underlying spatial topology into a sparse adjacency matrix, the Cuthill-McKee algorithm [CM69] can be applied to minimize its bandwidth and hence produce an ordering where connected nodes are placed closer together in the sequence. We do not directly contribute to generating 1-D orderings, but leverage these techniques to create candidates for validation and refinement by our approach.

Quantitative Evaluation of 1-D Orderings – The minimal path length, inspired by the TSP, measures the total length of the line that is formed in the input space when all entities are connected in sequential order of the output sequence [MD86]. Minimal wiring is a concept that aims to minimize the number of connections required to form a 1-D sequence from a 2-D grid [DM94]. Bauer and Pawelzik define the topographic product to quantify neighborhood preservation in self-organizing maps [BP92], which penalizes large-scale order violations between input and output while neglecting ranking differences within the neighborhood. Goodhill and Sejnowski introduce the C measure as the product of the input and output distance summed over all objects [GS96]. The authors show that many quadratic assignment problems are expressible by the C measure, and use it to evaluate and compare different solutions for the square-to-line problem. The above-mentioned approaches only measure the discontinuity of a projection, penalizing objects that are close in the topographical input but distant in the output space while overlooking objects that are distant in the original space but close to each other in the ordering space.

Venna and Kaski introduce the concept of trustworthiness in the context of neighborhood preservation [VK01], which evaluates if objects that are close in the output space are also close in the topographical input space. The authors define two rank-based measures M_1 and M_2 to determine the trustworthiness and discontinuity, respectively. Guo and Gahegan follow this notion and provide Key Similarity (KS) and Spatial Similarity (SS) measures [GG06], which can be defined as either distance-based or rank-based, depending on the specific use case. Numerous evaluation measures exist for assessing the quality of dimensionality reduction techniques that, for example, incorporate high- and low-dimensional properties such as neighborhoods. For a detailed overview of such methods, we defer to an exhaustive survey of Thrun et al. [TMS23].

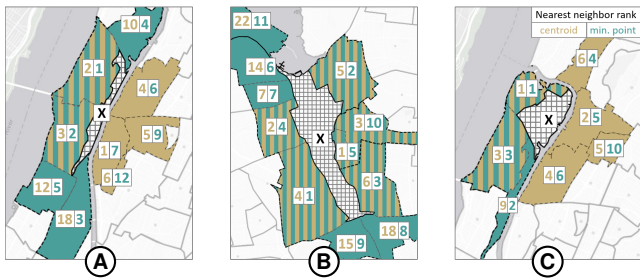


Figure 2: Neighborhoods and nearest neighbor ranks of polygons (x) determined by contiguous borders (■) and fixed k ($=6$) nearest neighbors (■). ■ consider only neighbors on the same side of obstacles like rivers (A) & (C) and prevent underestimation (B) or overestimation (C) when the number of contiguous neighbors differs from k .

All of the presented approaches are designed for point-based data. To the best of our knowledge, no evaluation measure designed for 1-D orderings exists that incorporates the topological properties of polygons into the distance calculation. While topological attributes have been leveraged to arrange layouts of small multiple maps [MDS*17], they have not been incorporated into measures assessing the quality of 1-D orderings. Furthermore, most existing methods rely on a fixed parameter (k) that defines the neighborhood with a fixed size in terms of its k -nearest neighbors and do not consider the possibility of a varying neighborhood size among the objects in the dataset.

Qualitative Evaluation of 1-D Ordering Errors – The existing literature indicates that errors in the ordering introduce artifacts that can stipulate significant misinterpretations, labeled as phantom splits [BJC*19], artificial splits [WBM*21], or referenced as inconsistent visual patterns [FMKK21]. In the domain of dimensionality reduction various layout enrichment techniques have been proposed to visually investigate errors or distortions in the projected layout. These include space-filling background encodings of uncertainty [WM15], neighborhood preservation [MCMT14, MMT15], projection precision [SvLB10], or colored Voronoi cells to visualize distortions [LA11]. Additionally, individual data points can be enriched by color encoding uncertainty [WM15], halos representing errors [SDMT16], or connecting graphs indicating high-dimensional neighborhoods [CVLD19]. However, all of these enrichment techniques are designed for 2-D target space, and a visual framework that facilitates the comparison of mapping errors within 1-D orderings does not exist to the best of our knowledge.

Applications of 1-D Spatial Orderings – Speckmann et al. present Necklace Maps, where spatial regions are projected on a 1-D curve, essentially generating a circular 1-D ordering [SV10]. Kriskograms also order regions to visualize migration flows, where the ordering is obtained either by geographic orientation (i.e. east-west), or demographic criteria such as the population rank [XC09]. In an example application on analyzing traffic accidents, Guo and Gahegan conceptually highlight the potential of spatial orderings and encodings to either sort visual displays or to engineer a spatial similarity feature variable [GG06]. Buchmüller et al. use an ordering strategy based on the Hilbert curve to create MotionRugs, a space-efficient technique for visualizing moving groups of entities [BJC*19]. The ordering is recomputed for each timeframe, possibly introducing undesired




Figure 3: Number of contiguous neighbors for different datasets. The number remains constant (excl. edges) for regular grids (A), whereas real-world examples (B: NYC, C: GER) show a higher variance.

visual jumps of entities in the visualization. To mitigate this issue, Wulms et al. introduce Stable Principal Components as a tradeoff between the spatial quality and temporal stability in computing such orderings [WBM*21]. Franke et al. investigate spatio-temporal phenomena through a matrix-like visualization, where the columns represent the spatially ordered entities, and each row corresponds to a time step [FMKK21]. Valdrighi et al. employ 1-D orderings to analyze moving regions and generate a storyline-based visual summary useful in detecting intersections in object tracing or hurricanes [VFP24]. As these applications inherently suffer from the previously mentioned mapping errors, we aim to contribute an interactive tool to help determine suitable orderings for various application cases.

3. Quantifying Errors in 1-D Polygon Orderings

Based on existing literature, we derive design goals for visually determining the quality of a 1-D ordering of contiguous spatial polygons.

[G1] Contiguous Neighborhood Preservation – Previous studies about quantifying mapping errors emphasize the importance of preserving neighborhoods [BP92, GS96, VK01]. In this context, elements within a neighborhood are classified as either: (1) *True* (consistent in both spaces), (2) *Missing* (absent in output but present in input space), or (3) *False* (present in output but absent in input space) [MCMT14, MMT15].

Existing measures [BP92, VK01, GG06] define the neighborhoods by the k -nearest neighbors, based on a fixed number k , in both input and output space. However, we argue that this is an oversimplification in the context of contiguous spatial polygons and argue for a locally varying neighborhood size defined by the number of contiguous polygons. In tessellated grids, the neighborhood size will be constant (except at the edges), but in real-world examples, the impact of a more complex and variable neighborhood configuration becomes apparent (see Fig. 3). In the 1-D output space, each object typically has two directly adjacent neighbors, positioned at an equal distance. Hence, the neighborhood size shall be even to ensure directional balance.

[G2] Error Locality Exploration – According to Martins et al. [MMT15], projection errors can be captured on different levels of detail, ranging from *global* measures that compute an aggregate value [Ken38, Tor52, Kru64] to measures that capture the errors *locally*, i.e., for every single object [GS96, VK01, GG06]. While global scores are relevant for overall comparison tasks, a local

investigation provides more detailed insights into the quality of the mapping. Previous studies [SvLB10, MCMT14, WM15] have highlighted the importance of both global entropy error quantification and local error identification on an entity level. Depending on the use case, certain geographical subregions can be of particular interest and shall contain only minor mapping errors, reinforcing the need for a local inspection of mapping errors.

[G3] Polygon Distance Assessment – Any error quantification of topographic mapping algorithms is heavily influenced by the employed similarity measure [GS96]. Distances between points are straightforward to calculate in Euclidean space or with the Haversine formula, but polygonal shapes introduce additional complexity that existing measures designed for point-based data cannot handle. Using the centroid as a reference point simplifies calculations but introduces biases, as centroids can lie outside irregular polygons and larger polygons have greater centroid-border distances, even if they are contiguous [MT20]. Minimal point distances between polygons mitigate these issues but treat all contiguous polygons as equally distant. For such cases, the length of the shared border can serve as a measure of contiguity and contribute to the distance assessment. Further polygonal properties such as the area or perimeter could also function as distinction criteria. However, they do not contribute to spatial coherence as they do not capture relative positioning or connectivity between polygons.

[G4] Geographic Boundary Interaction – Geographic boundaries, whether topographic (e.g., rivers, mountain ranges) or anthropogenic (e.g., train tracks, major roads) can impact distance calculations between spatial entities [Wan20] and, therefore, affect the quality assessment of a 1-D ordering. Such properties have been handled by using a visibility graph [THH01] or Delaunay triangulation [ECL04]. However, implementing these approaches requires prior knowledge of obstacle details, and the impact of barriers such as rivers, mountain ranges, cliffs, or forests varies depending on the use case. In this context, Visual Analytics provides an interactive approach to examining obstacles, allowing users to flexibly interpret spatial barriers based on their specific application requirements.

4. Error Measures

In this section, we review existing measures that have been employed to quantify errors in 1-D orderings and assess their alignment with our design goals established in Sec. 3 (see Eq. 1a-d in the supplementary material for formulas). In light of their shortcomings, we introduce novel measures. For formal notations, we define $i, j \in P$ as polygons in geographical space and the ordering respectively.

4.1. Existing Measures

M_1 & M_2 – Venna and Kaski propose the M_1 and M_2 measures to capture the trustworthiness and discontinuity of a projection, respectively [VK01], which have been applied by Franke et al. [FMKK21]. Dependent on a fixed neighborhood of size $k \in \mathbb{N}^+$, these rank-based measures determine the *false* and *missing* neighbors for every object. The distance in ranks is computed for each of these neighbors, and the sum of all rank differences yields the final score. The formulas permit a deconstruction to the local entity level and hence fulfill [G2]. However, the globally *fixed* neighborhood size k does not properly

Name	Description
$knn_{geo}(i)$	The k -nearest neighbors of i in geographical space.
$knn_{order}(i)$	The k -nearest neighbors of i in the ordering.
$d_{geo_centr}(i, j)$	The geographic centroid distance between i & j .
$d_{geo_min_pt}(i, j)$	The geographic minimum point distance between i & j .
$d_{order}(i, j)$	The distance between i & j in the ordering.
$r_{geo}(i, j)$	Nearest-neighbor rank of j in relation to i in the geography.
$r_{order}(i, j)$	Nearest-neighbor rank of j in relation to i in the ordering.

Table 1: Relevant definitions and notations for the error measures.

capture the *varying* neighborhoods in real-world datasets and for an odd k may introduce a directional bias, i.e., $\lfloor k/2 \rfloor + 1$ neighbors on the left and $\lfloor k/2 \rfloor$ neighbors on the right, while the polygon might be placed on either side in 1-D orderings, thus violating [G1]. The authors define neighborhoods simply by the *closest* objects which is ambiguous in terms of spatial polygons [G3]. Furthermore, all neighbors are weighted equally, meaning that the error for the most distant neighbor is just as important as that for the closest neighbor.

KS & SS – Guo and Gahegan provide several Key Similarity (KS) and Spatial Similarity (SS) measures and employ them to assess the quality of various 1-D ordering strategies [GG06]. The largest difference to M_1 and M_2 is a weighting term $w(i, j)$ that can either be defined by geographical distances (denoted by KS_{nd} & SS_{dn}) or neighbor ranks (denoted as KS_{nn} & SS_{nn}). While these measures may diminish the impact of more distant neighbors through a user-defined weighting factor, the issue with a globally constant neighborhood size k [G1] remains. All four measures are designed for point-based data, so they lack a proper assessment of polygon distances [G3].

4.2. Error Measures for Varying Neighborhood Sizes

As none of the existing measures properly fulfill all design goals we established in Sec. 3, we adapt the formulas of Venna and Kaski [VK01, VK06] and define a set of novel measures. To satisfy [G1], we determine an *even* neighborhood size k for every polygon based on its contiguous neighbors from the set of all polygons P :

$$k(i) := \lceil |\{j \in P \setminus \{i\} : d_{geo_min_pt}(i, j) = 0\}| \cdot 0.5 \rceil \cdot 2 \quad (1)$$

To determine distances between polygons [G3], we rely on the minimum point distance instead of the centroid. We introduce an additional term based on the shared border percentage to retain distance variations between contiguous polygons, where the minimum point distance will be zero:

$$d_{geo}(i, j) := \frac{d_{geo_min_pt}(i, j)}{\max_{x, y \in P} d_{geo_min_pt}(x, y)} + \left(1 - \frac{border_{length}(i, j)}{border_{sum}(i)}\right) \quad (2)$$

In accordance with Venna and Kaski [VK01] we determine the k -nearest neighbors in the geographical space using Eq. 2 and ordering space as $knn_{geo}(i)$ and $knn_{order}(i)$, respectively. We exclude *true* neighbors present in both neighborhoods, reducing the objects in $knn_{geo}(i)$ to *missing* neighbors and objects in $knn_{order}(i)$ to *false* neighbors. The M_1 and M_2 measures consider all neighborhood violations of these objects as equally important. We argue that the closest neighbors shall have a higher influence, especially in larger neighborhoods, and account for that by introducing a weighting factor. A geographical weighting can be obtained by calculating the border percentage (Eq. 3). This ensures that non-contiguous polygons (in the case of an odd number of contiguous polygons, the closest non-contiguous polygon will be

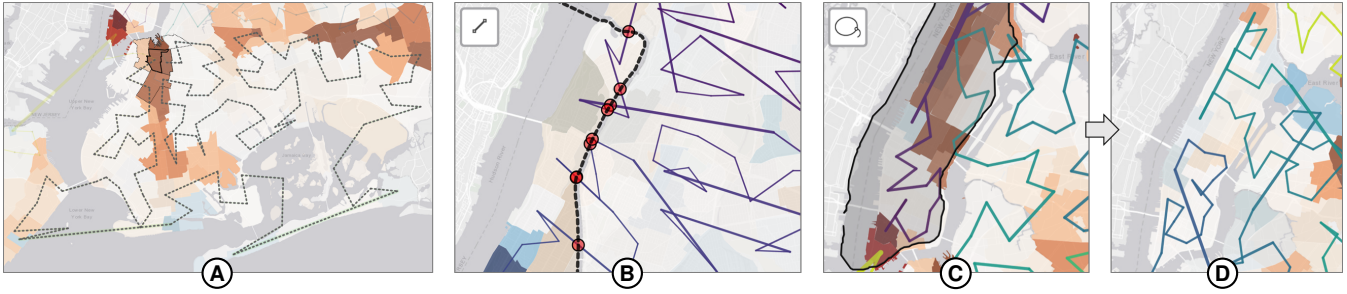


Figure 4: (A) Highlighting a long ordering path between 2 contiguous polygons. (B) Drawing a polyline along the Harlem River and determining the crossings with the current ordering. (C & D) Using the lasso to determine a better local ordering for a subregion.

part of the neighborhood to ensure it is *even*) will not contribute to the error score, and the sum of all weights results in 1. To achieve the same behavior for the order weighting, we define Eq. 4 as a term that decays symmetrically with increasing order distance (d_{order}).

$$weight_{geo}(i,j) := \frac{border_{length}(i,j)}{border_{sum}(i)} \quad (3)$$

$$weight_{order}(i,j) := \frac{(1 + \frac{k(i)}{2}) - d_{order}(i,j)}{(1 + \frac{k(i)}{2}) \cdot \frac{k(i)}{2}} \quad (4)$$

Similar to Guo and Gahegan [GG06], we multiply the weighting with the rank distance in the geographical or ordering space. However, we subtract the neighborhood size k from the rank distance to account for larger neighborhoods, essentially computing the distance from the neighborhood bound of i to j . The discontinuities d_{poly} of the geographical neighborhood for a given polygon i is hence defined as:

$$d_{poly}(i) := \sum_{j \in missing(i)} weight_{geo}(i,j) \cdot ((2 \cdot d_{order}(i,j) - 1) - k(i)) \quad (5)$$

As rank distances in an ordering cannot be uniquely resolved (there are always 2 objects with the same distance and rank), we use twice the distance in the ordering space (d_{order}). The trustworthiness error t_{poly} of the corresponding ordering is defined as:

$$t_{poly}(i) := \sum_{j \in false(i)} weight_{order}(i,j) \cdot (r_{geo}(i,j) - k(i)) \quad (6)$$

These measures allow for a detailed inspection of mapping errors, enabling traceability down to individual polygons and further relationships between two polygons [G2]. Since these measures include asymmetric terms, they do not satisfy the properties of a metric space. Hence, they are unsuitable for algorithms that require symmetric distance functions such as AHC, TSP, and certain dimensionality reduction methods.

To obtain a global measure for a given ordering strategy, the sum of all objects is taken, where the multiplication with the neighborhood size ($k(i)$) acts as a scaling factor to ensure global measures between different datasets are comparable. Dividing by N results in the average quality of a polygon for the given ordering:

$$d_{sum} := \frac{1}{N} \sum_{i=1}^N k(i) \cdot d_{poly}(i) \quad (7) \quad t_{sum} := \frac{1}{N} \sum_{i=1}^N k(i) \cdot t_{poly}(i) \quad (8)$$

5. Visual Analysis Prototype

In accordance with the design goals from Sec. 3, we propose a visual analytics framework consisting of multiple linked views

to investigate errors in ordering strategies. The prototype can be accessed under: assess.1d-poly-order.dbvis.

Orderings – In Sec. 2, we outlined algorithmic methods to generate 1-D orderings from spatial datasets. We focus on AHC-based methods and SFCs, as they have provided promising results in previous studies [GG06, FMKK21]. Different orderings can be generated by rotating the initial orientation of SFCs (see Fig. 5). Hence, we provide four orderings for these curves as different initializations may alter the discontinuity and trustworthiness error results. The TSP relates to our problem of generating a 1-D ordering of spatial entities, as a Hamiltonian path can be obtained from the cycle by removing the longest connection. We further include an ordering obtained from the Cuthill-McKee algorithm, initially designed to permute a sparse matrix to reduce its bandwidth. From a set of polygons, we create a binary distance matrix that captures if polygons are contiguous or not. Then, we apply the Cuthill-McKee algorithm to obtain our ordering sequence.




Figure 5: Depending on the orientation, different SFC orderings with locally different properties can be generated.

5.1. Geographical Map View

To visually assess mapping errors within the geography, we provide a map component consisting of several layers that can be interactively toggled. To investigate the distribution of discontinuity d_{poly} and trustworthiness t_{poly} errors across the geography, we employ a choropleth map where mapping errors of every polygon are encoded by color (see Fig. 4). To visualize both errors simultaneously, we opted for a bivariate colorscale and followed the guidelines of Brewer [Bre94], using complementary hues of orange and blue (resulting in a grayscale diagonal). Polygons with no error are thus colored in white and thus less conspicuous than problematic ones. When hovering over a polygon, its geographical neighborhood is indicated by a dotted line [G1]. A click on a polygon selects this region and anchors the neighborhood visualization, and a subsequent hover on another polygon indicates the path in the ordering between these objects (see Fig. 4 A).




Figure 6: Schematic of Order Sequence Plot. Upper bars encode discontinuity, lower bars trustworthiness error. In between, the geographic and ordering neighborhood of one polygon (x) is depicted.

Ordering Path – The entire geographical path of the ordering can be visualized as a line, resembling the *minimal path length* [MD86] (see Fig. 5). The width of each line segment is scaled according to its length, highlighting undesired long geographic distances between ordering neighbors. The color is determined by the ordering index and mapped to the Viridis color scale, where darker colors indicate segments at the beginning of the order, and lighter colors segments at the end of the order. This allows the user to easily depict large distances in the ordering between spatially proximate polygons.

Interactions – While a certain ordering strategy might produce the smallest global errors, it still contains local errors that might make it unsuitable depending on the use case. **Using a lasso**, the user can make a geographical selection to obtain local error scores d_{local} and t_{local} for the selected subregion. These scores then support finding an optimized ordering for this subregion (see Fig. 4 © & ④) [G2]. Geographical barriers such as rivers or mountain ranges impose a natural obstacle that an optimal ordering shall circumvent in certain cases [G4]. To explore such barriers, the user can **draw a polyline** across the map, for which the number of crossings n_{cross} with the ordering path is computed and depicted on the map (see Fig. 4 ⑤).

5.2. Ordering Sequence Plot

Visualizations of 1-D sequences are found in Kriskograms [XC09] or arc diagrams [Wat02], where nodes are ordered along a single axis and rcs indicate links between entities. We draw inspiration from these concepts and visualize all objects along a horizontal axis, ordered according to the currently selected ordering strategy. The design space of the Ordering Sequence Plot (Fig. 6) has three parts: the discontinuity, trustworthiness, and the neighborhood of an object.

Discontinuity & Trustworthiness – Given a specific ordering strategy, we can obtain discontinuity and trustworthiness error scores for every polygon according to Eq. 5 and Eq. 6 respectively. We employ two bar charts where the values of the error scores are encoded for every object in the sequence of the currently selected ordering [G2]. The bars are colored using shades of orange and blue, aligning with the bivariate colorscale used for the choropleth map. This encoding directs attention to objects with higher error scores, while perfectly projected regions become imperceptible.

Neighborhood – Between the two barcharts, we allocate space for a local inspection of a specific neighborhood [G1]. Similar to Kriskograms or arc diagrams, we model the results from our error measures pairwise as links between different objects. Instead of arcs, we chose rectangular connecting lines and only show the links of one single object on hover to reduce visual clutter. The geographical neighborhood

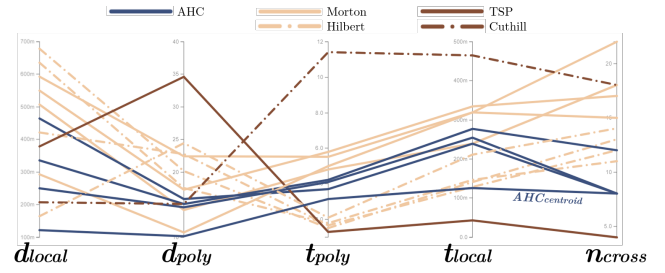


Figure 7: PCP plot showing the global (d_{poly} & t_{poly}), as well as local scores (d_{local} , t_{local} & n_{cross}), obtained using the interactions visible in Figure 4 (⑤ & ④)) for different orderings for the NYC dataset.

of a selected object (x) is visualized adjacent to the discontinuity errors, where every neighborhood member is indicated by connecting lines. The longer this line, the larger the distance of these geographical neighbors in the ordering. The error value corresponding to a given link is encoded on the height of a gray rectangle, allowing the user to identify the most critical relationships. Since *true* neighbors do not produce an error, no rectangle will be visible and only the connecting line is drawn. The ordering neighborhood is depicted adjacent to the trustworthiness error bars and is visually enclosed to indicate its bounds. Again, *true* neighbors do not produce an error and are hence invisible. Ideally, both neighborhoods overlap and encompass the same objects, indicating that only *true* neighbors are present. Fig. 6 shows the neighborhood of object x , which is unbalanced as all geographical neighbors lie to the right and thus have a higher ordering index. The ordering neighborhood therefore features three *false* neighbors to the left, which are indicated by gray rectangles.

5.3. Strategy Comparison Plot

To compare ordering strategies on a global level, we employ a parallel coordinate plot, where every axis represents a measure, and every line represents an ordering strategy. Clicking on a line will select the respective ordering strategy and update the Map and the Ordering Sequence Plot. Aside from the measures we defined in Eq. 7 and Eq. 8, the local subregion error scores (d_{local} , t_{local}) and number of crossings (n_{cross}) are also available as axes [G3]. The different strategies are colored according to their algorithm family (AHC, SFC, Network & Projection-based) and can be toggled with a button click.

6. Evaluation

We quantitatively compare the measures established in Sec. 4, and provide a case study to outline the workflow of the prototype introduced in Sec. 5. We further performed an expert user study.

6.1. Correlation Analysis

To compare our measures against the existing literature, we conduct a correlation analysis on synthetic and real-world datasets to establish in which scenarios our measures provide different results.

Methodology – Correlations between the measures vary depending on the applied ordering strategy. Orderings that provide good results

Dataset	k	σ	Discontinuity (d_{poly})			Trustworthiness (t_{poly})		
			KS _{nd}	KS _{nn}	M ₂	SS _{dn}	SS _{nn}	M ₁
HEX	6	1.02	.9789	.8428	.9500	.9698	.9808	.8772
	8	1.35	.7589	.8270	.5986	.9650	.9753	.8967
SQR	4	1.35	.9845	.9126	.9717	.9033	.9110	.9229
	4	1.60	.8488	.8489	.7273	.9311	.9664	.9130
AUS	6	1.86	.8716	.9336	.7845	.9351	.9524	.8837
CHI	4	1.87	.7879	.7999	.7345	.8634	.9489	.8971
NYC	4	1.95	.7868	.7954	.7025	.9200	.9465	.8410
GER	6	2.23	.6355	.8159	.6839	.9294	.9413	.8222
WOR	2	2.70	.5580	.6539	.6562	.9339	.9348	.9351

Table 2: Correlation of our scores against existing literature obtained from 1000 random orderings on synthetic (top) and real-world (bottom) datasets. The number after the k indicates the standard deviation σ of the number of contiguous neighbors.

produce fewer errors and thus naturally exhibit a higher correlation. To alleviate this, we compute random orderings, which are expected to perform poorly and produce high error scores. To ensure a stable basis for comparing these correlations, we iterate this procedure and compute 1000 orderings for every dataset, from which we take the average. As comparative measures we select the measures introduced in Sec. 4: M_2 and M_1 [VK01], as well as KS_{nd} , KS_{nn} , SS_{dn} and SS_{nn} [GG06]. Since these measures expect a parameter k to indicate the size of the neighborhood, we estimate this by taking the mean of contiguous polygons and rounding it to the nearest even number.

Datasets – Franke et al. [FMKK21] provided two application scenarios to showcase their prototype: *World countries* (WOR) and a tessellated grid of squares spanning over southeastern *Australia* (AUS). We include both of these datasets in our analysis and further selected geographical regions with different characteristics. The districts of *Germany* (GER) consist of 400 polygons, whereas larger cities manifest their own district. These city districts are often fully encapsulated by their surrounding neighbor. The country of *Italy* (ITA) resembles a boot and further includes the two islands Sicily and Sardinia. This dataset was chosen for its irregular, non-squared shape comprising 107 polygons representing the provinces. *Chile's* (CHI) narrow and elongated country shape also largely deviates from a square, where we selected the communal level, consisting of 341 regions. *New York City* (NYC) features a complex topography segmented by bays and rivers. We selected the 262 Neighborhood Tabulation Areas that form the geographical basis for the American Community Survey [Uni24].

Since polygons are often the result of aggregation tasks in geospatial data mining [BZP22, SE21, WD23], we further include two synthetic, tessellated datasets similar to Fig. 3 (A): a 16x16 square grid (SQR) and a 187-tile hexagon grid (HEX).

Results – Tab. 2 provides the correlations of our measures against existing methods for the above-mentioned datasets. Regarding the discontinuity, we can observe the strongest correlation for the HEX dataset, which is expected since the neighborhood size remains constant throughout the dataset except for the borders of the grid. While this is also the case for the SQR dataset, the correlation here is amongst the lowest when we consider $k = 8$. Since our weighting factor depends on the shared border percentage, the four diagonal (bishop) neighbors will not produce an error as they are merely tangent. We obtain the expected correlation when we change to $k = 4$ (rook neigh-

bors). Since the AUS dataset contains holes in the tessellated grid, the estimated k reduces to 6, which decreases the correlation factor. For every dataset except WOR, the correlation to the KS measures is higher, although our d_{poly} measure originates from the M_2 formula. While the latter does not perform any weighting of the objects in the neighborhood, the former measures incorporate a weighting by dividing the nearest neighbor rank in the ordering by the geographical distance. The KS_{nn} measure exhibits the highest correlation to our d_{poly} measure, which we attribute to the fact that both are rank-based and include a geographical weighting factor. The tessellated regular grid datasets (HEX, SQR) are an exception to this case, as rankings between objects within the neighborhood are ambiguous, as they all share the same distance. Real-world datasets generally show lower correlation values, attributed to variable neighborhood sizes, particularly in datasets with many enclaves or islands (GER, WOR). This indicates that our method for determining neighborhood size significantly impacts discontinuity errors.

For trustworthiness errors, Tab. 2 reveals substantially higher correlations overall, with the gap between grid and real-world datasets remaining, though less pronounced. Similar to the discontinuity measures, the correlations against the SS scores tend to be higher than the M_1 , which does not feature a weighting factor for the objects in the neighborhood. This effect is more pronounced for larger neighborhood sizes (AUS, GER, 6) and understandably not as apparent for smaller neighborhoods (WOR, 2).

Consequently, using a varying neighborhood size for every polygon does not have as considerable an impact on trustworthiness error measures as it does on discontinuity. For the discontinuity, the geographical neighborhood is of importance, where a semantic meaning of direct neighbors can be established with the number of contiguous polygons. Inversely, the ordering neighborhood determines the trustworthiness errors, where the number of direct neighbors is always 2, and larger neighborhood sizes merely contribute to smoothly distributing the error across the neighboring objects.

6.2. Case Study

To provide insights about household income development throughout New York City, a data scientist is tasked to develop a matrix-based dense-pixel visualization that captures trends in local neighborhoods. The growth over time shall be depicted along the columns, leaving one row for every Neighborhood Tabulation Area. The ordering of these rows shall reflect geographical proximity to capture spatio-temporal trends within the dense-pixel visualization. Initially, the trustworthiness property is determined as a more important property than the discontinuity since the ordering neighbors in the resulting visualization shall be interpretable with respect to geographical proximity. The Strategy Comparison Plot (Fig. 7) reveals that the TSP ordering accumulates the lowest global trustworthiness error (t_{poly}). The Ordering Sequence Plot (Fig. 9 (A)) confirms this, but further exhibits grave discontinuity errors. Fig. 9 (1) shows the neighborhood for the polygon with the highest error, and the connecting lines indicate that 5 geographic neighbors have very large ordering distances.

The four Hilbert curve orderings are the next best strategies concerning trustworthiness errors, but the analyst wants to make a more locally based decision, focusing on the subregion of Manhattan. For




Figure 8: Map view of the TSP ordering on the NYC dataset.




Figure 9: Ordering Sequence Plot of the TSP on NYC.

this, he draws a polyline along the Harlem River to inspect crossings with the ordering line (see Fig. 4 (B)) and selects the entire island using the lasso (see Fig. 4 (C)). From the Strategy Comparison Plot (see Fig. 7) it can be observed that most Hilbert curve orderings score among the worst in terms of local discontinuity (d_{local}). The $AHC_{centroid}$ ordering, although worse globally (t_{poly}), is on par with the Hilbert curve orderings in terms of trustworthiness errors (t_{local}), and outperforms them with regard to n_{cross} . Therefore, the strategy is selected for a more detailed inspection.

Fig. 11 (1) shows the neighborhood for the polygon with the highest error, where the connecting lines are much shorter compared to the TSP ordering. This comes at the cost of multiple trustworthiness errors, indicated by blue spikes (A)(B)(C). These jumps can be clearly identified when looking at the ordering path on the map (see Fig. 10) and mostly occur when the ordering is transitioning from one borough (district) to another. Hence, the analyst is satisfied with the tradeoff between discontinuity and trustworthiness errors and selects the $AHC_{centroid}$ ordering for his visualization.

6.3. Expert User Study

To qualitatively evaluate the effectiveness of our prototype, we conducted pair analytics sessions [AHKGF11] with potential endusers.

Participants – 8 participants (E1 – E8) aged between 20-35, partook in the study, most being PhD students in data science who had previous experience in developing spatiotemporal data visualizations.

Procedure – The study began with a video of the prototype introducing all components and interactions. All tasks were performed on a synthetic dataset first, serving as a test condition. Given the HEX dataset and the $AHC_{complete}$ ordering, participants had to locate the polygons with the highest discontinuity and trustworthiness error and identify all their missing and false neighbors respectively as a first task [G1]. Next, they had to compare the best and worst-performing global strategies for both measures and comment on any visual patterns they encountered. Then, the participants had to assess the quality of all orderings within a subregion spanning the left half




Figure 10: Map view of the $AHC_{centroid}$ ordering on the NYC dataset.




Figure 11: Ordering Sequence Plot of the $AHC_{centroid}$ on NYC.

of the hexagon grid and compare the results with the global discontinuity and trustworthiness error scores [G2]. As a last task, the participants had to find the ordering strategy with the least number of crossings along a vertical line through the center of the layout [G4]. All tasks were repeated on a real-world dataset (NYC), where the subregion was defined as Manhattan and the crossing line corresponded to the Harlem River (akin to Fig. 4). The study concluded with an interview in which the participants shared their feedback regarding usability, faced challenges, and suggestions for improvement.

Results – E1, E6, and E8 detected recursive patterns of SFC orderings, which were more apparent and predictable in the HEX dataset compared to NYC. The lasso and crossing-line interactions were deemed as useful, E8 stated they integrate geographic semantics into the analysis process. E3, E5, E7 occasionally faced semantic confusion, mixing up the terms for discontinuity and trustworthiness. E5 labeled a geographic jump in the ordering as a discontinuity in the ordering path. E1 – E8 found that the bivariate colorscale was beneficial for identifying poorly mapped polygons and could distinguish between the mapping errors, E3 and E8 suggested highlighting the hovered polygon on the colorscale legend. E2 and E4 raised concerns that with higher error scores and saturation levels, distinguishing between errors becomes more difficult, and suggested using texture shading as a mitigation. Many participants relied on the Ordering path on the map to identify ordering patterns. E4 stated that the path based on centroids is ambiguous as contiguous polygons will also generate a line even though they share a border. E1, E3, and E7 reported that the map component can suffer from overplotting; however, they further stated that this can be effectively addressed by deactivating certain layers. As additional features, E1, E4, E7 desired the possibility to slightly modify the orderings to relocate critical mapping errors. E5 suggested an axis reordering feature for the Strategy Comparison Plot. E8 proposed to determine subregions and crossing-line candidates algorithmically instead of manually drawing them on the map.

7. Discussion and Design Considerations

In this section, we discuss adjustments made to the prototype in response to the feedback from the study participants and reflect on

Dataset	AHC _{complete}		AHC _{ward}		AHC _{centroid}		AHC _{single}		Morton		Hilbert		TSP		Cuthill-McKee		MDS	
HEX	3.14	0.39	3.44	0.65	3.26	0.44	2.28	1.70	2.16	0.75	2.75	0.16	5.40	0.41	2.32	1.93	11.57	5.71
SQR	2.46	0.49	2.32	0.45	2.85	0.42	1.62	1.71	1.57	0.65	2.02	0.12	4.79	0.44	1.79	2.69	22.01	7.56
ITA	0.95	0.37	1.26	0.44	1.06	0.38	4.31	0.77	1.05	0.58	1.54	0.32	1.56	0.27	0.93	1.29	12.78	4.28
AUS	1.20	0.36	1.45	0.33	1.06	0.34	2.22	0.85	0.70	0.44	1.32	0.15	3.17	0.48	1.40	1.74	1.01	1.28
CHI	1.84	0.53	1.55	0.40	1.58	0.48	7.60	1.52	4.29	1.76	7.66	1.61	8.98	0.27	3.22	1.54	22.92	4.06
NYC	3.01	0.84	2.91	0.82	2.01	0.62	7.64	2.11	2.14	1.00	4.01	0.30	6.89	0.26	3.01	2.27	20.15	6.70
GER	4.98	0.51	5.10	0.59	4.81	0.70	12.77	2.73	3.90	1.14	4.88	0.26	9.76	0.28	4.92	3.97	39.03	8.41
WOR	4.67	0.73	3.65	0.85	5.94	1.15	9.42	1.82	6.42	2.04	7.19	1.16	8.52	0.46	11.52	3.59	56.79	21.71

Table 3: Global results for our evaluation measures for different ordering strategies applied to various datasets. The first number corresponds to the discontinuity error d_{poly} , and the second to the trustworthiness error t_{poly} . The best strategy for every dataset is highlighted in bold.

the effects of our novel evaluation measures for discontinuity and trustworthiness errors. Furthermore, we contribute a numerical comparison of various ordering strategies on the datasets introduced in Sec. 6.1, derive design considerations for their usage, and outline limitations as well as future research opportunities.

7.1. Study Implications

The expert study results show that our prototype can effectively support analysts in identifying mapping errors and comparing ordering strategies on a local and global scale. Based on the qualitative feedback from the study participants, we made minor adaptations to the prototype (see Fig. 1 in the supplementary material). To mitigate overplotting issues on the map, we reduced the opacity and dash the line of true neighbors in the ordering path. Although the path is still drawn through the centroids, only links with larger geographic distances are accentuated, where the difference to the minimum point distance is negligible. Furthermore, the subregion as well as the crossings are also displayed in the Ordering Sequence Plot. This supports the visual investigation of the fracturedness of a subregion as well as determining at which ordering locations the crossings occur.

7.2. Measures

Sec. 6.1 provides a comprehensive comparison between our novel measures and the existing literature. We manifested that our measures are capable of capturing characteristics of complex polygon datasets with highly varying neighborhood sizes, while established methods tend to overgeneralize these circumstances. These effects are more dominant when measuring discontinuity errors, where a precise definition of the geographical neighborhood is imperative to obtain reliable results. Since the trustworthiness errors are computed by investigating neighbors in the ordering, using the number of contiguous polygons to define the neighborhood size is less semantically grounded. We decided to adhere to the concept of varying neighborhood sizes to maintain consistency between the measures and to ensure that the neighborhood sizes for a given polygon are equal for both formulas. However, the results of our correlation analysis indicate that a constant, even number will yield comparable results.

The differences between our measures and existing methods vanish for polygon datasets arranged in a regular, tessellated grid structure. Such aggregation techniques are commonly used in the analysis of geospatial point data [BZP22, SE21, WD23]. We argue, however, that for real-world datasets with complex polygonal arrangements, our measures provide a significant contribution to capturing the neighborhood preservation of the 1-D ordering more accurately.

Computational Complexity – Our measures are computationally more expensive than existing approaches due to the use of $d_{geo_min_pt}$, which operates in $O(n \cdot \log n)$ compared to $O(n)$ for centroid distances [TB83], where n represents the number of points per polygon. This can pose scalability challenges for large datasets with detailed polygons, as polygon simplification is not a viable solution due to the risk of compromising contiguity. A practical solution is precomputing an asymmetric distance matrix (complexity $O(N^2 \cdot n \cdot \log n)$) once for efficient lookups. Since k is not a fixed parameter, the worst-case of neighborhood sizes can be $N - 1$. However, Tab. 2 indicates that typical average values range between 2 and 6.

Choosing the right measure – A promising application area of 1-D spatial orderings is dense-pixel visualizations that are able to visually expose trends over space and time. In previous works [BJC*19, WBM*21], orderings were selected with respect to exhibiting low trustworthiness errors. These strategies are preferable for visualizations that provide a global overview of the entire dataset, as a large geographical gap in the ordering can result in visual patterns that are not reflected in the underlying geography. In contrast, for visualizations that focus on local details within a geographical subregion, orderings with a lower discontinuity score are preferable, as they better preserve the local original neighborhood. This facilitates the identification of dynamic phenomena, such as trendsetters and spreading events, as visual patterns within the visualization.

7.3. Ordering Strategies

Tab. 3 provides the results of our global comparison scores defined in Eq. 7 and 8 for various established ordering strategies on the datasets introduced in Sec. 6.1. From these values, we can confirm the statement from Guo and Gahegan [GG06, p. 252] that “very often, an ordering or encoding method works well on one group of measures but performs poorly against the other.”, since no ordering strategy excels in both measures for a given dataset.

SFCs excel in squared layouts – For tessellated, grid-based layouts (HEX, SQR, AUS), SFCs provide the most optimal results, where the Morton curve is more beneficial for generating orderings with minimal discontinuities and the Hilbert curve produces a more trustworthy result. They also perform well on layouts that can be approximated with a square (GER) but encounter problems with more elongated shapes (CHI). The recursive nature of these algorithms inherently introduces discontinuity rifts: lines where geographic neighbors exhibit large ordering distances (see Fig. 1). Due to their recursive, predictive structure, given regions will consistently occupy similar ordering positions (i.e. southwestern regions at the start and southeastern regions

at the end of the ordering). Depending on the initial orientation of the curve, these ordering positions and the location of rifts can be shifted.

AHC accomplish a tradeoff – AHC-based methods provide mostly balanced results, with the exception of the single-linkage method. For regular, tessellated grid layouts, single linkage provides better results concerning discontinuity but performs significantly worse otherwise. The remaining AHC-based methods achieve a reasonable tradeoff between discontinuity and trustworthiness errors, and all exhibit very similar results with only minor variations. Compared to SFCs, these methods tend to align better with a complex, irregular layout. However, they are similarly susceptible to discontinuity rifts with the disadvantage that the location of these rifts is not as deterministic and can vary when the input data changes slightly. Hence, if stability of the ordering is desired, SFCs are a more viable option.

TSP provides most trustworthy orderings – Regarding real-world datasets, the TSP ordering strategy consistently yields among the lowest trustworthiness error scores and is able to align with complex layouts as it outperforms the Hilbert curve on irregularly shaped datasets (*ITA*, *CHI*, *NYC*, *WOR*). This comes at the cost of high discontinuity scores, as the preservation of the geographic neighborhood is not considered. A noteworthy drawback is the geographic proximity of the first and the last object in the ordering sequence (see Fig. 8) since the TSP essentially computes a cycle where every object is visited exactly once. The Hamiltonian Path is then generated by removing the longest connection from the cycle. While undesired for 1-D orderings, this might be beneficial for circular layouts [SV10].

Cuthill-McKee minimizes variance of discontinuity – Orderings based on the Cuthill-McKee algorithm exhibit low discontinuity values and, therefore, look promising to utilize when geographic neighborhoods shall be depicted. Upon closer inspection, we can observe that the variance between the discontinuity errors is lower compared to other strategies, which is achieved by spreading out the geographic neighbors along the ordering sequence (see Fig. 12). Hence, geographic neighborhoods are rarely perfectly represented, and numerous geographic jumps are introduced resulting in the highest trustworthiness errors aside from the projection-based *MDS* ordering.




Figure 12: Ordering Sequence Plot of a Cuthill-McKee ordering, where the geographic neighborhood is spread along the ordering.

8. Limitations and Future Work

Errors Inside Neighborhood – While our measures incorporate the relative importance of its neighbors to determine an error score, a score is only computed if we consider it either a *false* or *missing* neighbor. Akin to Venna and Kaski [VK01], we do not assign errors for *true* neighbors. This does not account for mapping errors inside the neighborhood, i.e., when the closest geographical neighbor is placed as the third closest in the ordering. This can be mitigated by also computing an error value for *true* neighbors, and only considering the error as zero if the neighbor is placed in the correct position.

Ordering Possibilities – For a dataset with N objects, the number of available 1-D orderings can be calculated by $N!$. We only consider

solutions generated by established algorithms to generate such orderings, resulting in a solution set that is far from exhaustive. While these algorithms are optimized to provide the best possible results, it is not guaranteed that the global minimum for any measure is contained in this set. This limitation is magnified when locally investigating regions, as only the precomputed solutions are available to choose from. Implementing further ordering strategies, such as simulated annealing [KGV83] or data-driven SFC [ZJW21] could prove to be fruitful in determining semantically coherent solutions.

Local Ordering Improvements – As demonstrated in Sec. 6.2, our prototype can expose suboptimal sections within an ordering, but is not capable of resolving such minor local issues. Providing means to make changes to an existing ordering coupled with direct feedback through our measures can be an effective workflow to iteratively improve an ordering locally that was requested by study participants.

Further Validation of Measures – Our correlation analysis showed a weaker alignment of our measures with existing point-based methods for irregular real-world polygon datasets. To further validate our approach, a more in-depth comparison across diverse layouts and orderings could strengthen their validity and usefulness. Additionally, performance evaluations on diverse datasets will help assess the computational trade-offs.

Distance functions – Several algorithms, such as AHC, TSP, and dimensionality reduction methods, can be supplied with a dissimilarity function to obtain a 1-D ordering. Since these algorithms expect a metric dissimilarity function, our proposed distance function is unsuitable as it is neither symmetric nor satisfies the triangle inequality. Hence, we generated these orderings using the Haversine distance between the centroids of the polygons, which is known to introduce biases [MT20]. Since our prototype is capable of comparing different ordering strategies regarding discontinuity and trustworthiness errors, the influence of different dissimilarity functions can be examined in a similar fashion. Future research may explore the effects of alternative dissimilarity measures on the generation of 1-D orderings.

9. Conclusion

This work introduces evaluation measures tailored to capture the neighborhood preservation and trustworthiness of 1-D orderings of spatial polygon datasets by adapting spatial distance functions to polygonal shapes and incorporating varying neighborhood sizes based on contiguous neighbors. We further present a user-steerable visual analytics application to support users in identifying various mapping errors and determining suitable ordering strategies. We demonstrate its capabilities with a use case and further provide a qualitative expert study showing that users are able to assess the effects of local subregions and semantically meaningful obstacles. By means of our measures, we examine and discuss the results of established 1-D ordering strategies for a number of synthetic and real-world datasets and provide design considerations for their usage.

Acknowledgement

This work was funded by the Federal Ministry for Economic Affairs and Climate Action – 03EI1048D, and Deutsche Forschungsgemeinschaft (DFG) – TRR 161 (Project A03) – 251654672. Open Access funding enabled and organized by Projekt DEAL.

References

- [AHKGF11] ARIAS-HERNANDEZ R., KAASTRA L. T., GREEN T. M., FISHER B.: Pair Analytics: Capturing Reasoning Processes in Collaborative Visual Analytics. In *Int. Conf. Syst. Sci.* (2011), pp. 1–10. doi:10.1109/HICSS.2011.339. 8
- [BJC*19] BUCHMÜLLER J., JACKLE D., ÇAKMAK E., BRANDES U., KEIM D. A.: MotionRugs: Visualizing Collective Trends in Space and Time. *IEEE Trans. on Vis. and Comp. Graph.* 25, 1 (2019), 76–86. doi:10.1109/TVCG.2018.2865049. 1, 2, 3, 9
- [BP92] BAUER H.-U., PAWELZIK K.: Quantifying the Neighborhood Preservation of Self-Organizing Feature Maps. *IEEE Trans. Neural Networks* 3, 4 (1992), 570–579. doi:10.1109/72.143371. 1, 2, 3
- [Bre94] BREWER C. A.: Color Use Guidelines for Mapping and Visualization. In *Visualization Mod. Cartogr.*, vol. 2 of *Mod. Cartogr. Ser.* Academic Press, 1994, pp. 123–147. doi:10.1016/B978-0-08-042415-6.50014-4. 5
- [BZP22] BURIAN J., ZAPLETAL J., PÁSZTO V.: Disaggregator - a tool for the aggregation and disaggregation of spatial data. *Earth Sci. Inf.* 15, 2 (2022), 1323–1339. doi:10.1007/S12145-021-00737-9. 7, 9
- [CM69] CUTHILL E., MCKEE J.: Reducing the bandwidth of sparse symmetric matrices. In *Proc. 24th Natl. Conf.* (1969), ACM, pp. 157–172. doi:10.1145/800195.805928. 2
- [CVLD19] COLANGE B., VUILLON L., LESPINATS S., DUTYKH D.: Interpreting distortions in dimensionality reduction by superimposing neighbourhood graphs. In *IEEE Visualization Conf. (VIS)* (2019), pp. 211–215. doi:10.1109/VISUAL.2019.8933568. 3
- [DCOM00] DAFNER R., COHEN-OR D., MATIAS Y.: Context-based space filling curves. *Comput. Graph. Forum* 19, 3 (2000), 209–218. doi:10.1111/1467-8659.00413. 2
- [DM94] DURBIN R., MITCHINSON G.: A dimension reduction framework for understanding cortical maps. *Nature* 343 (1994), 644–647. doi:10.1038/343644a0. 2
- [ECL04] ESTIVILL-CASTRO V., LEE I.: Clustering with obstacles for geographical data mining. *ISPRS J. Photogramm. Remote Sens.* 59, 1 (2004), 21–34. doi:10.1016/j.isprsjprs.2003.12.003. 4
- [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 Trans. on Vis. and Comp. Graph.* 27, 3 (2021), 2153–2173. doi:10.1109/TVCG.2019.2944182. 2
- [FMKK21] FRANKE M., MARTIN H., KOCH S., KURZHALS K.: Visual Analysis of Spatio-temporal Phenomena with 1D Projections. *Comput. Graph. Forum* 40, 3 (2021), 335–347. doi:10.1111/CGF.14311. 1, 2, 3, 4, 5, 7
- [GG06] GUO D., GAHEGAN M.: Spatial ordering and encoding for geographic data mining and visualization. *J. Intell. Inf. Syst.* 27, 3 (2006), 243–266. doi:10.1007/S10844-006-9952-8. 2, 3, 4, 5, 7, 9
- [GS96] GOODHILL G. J., SEJNOWSKI T. J.: Quantifying neighbourhood preservation in topographic mappings. In *Proc. 3rd Joint Symp. on Neural Comput.* (1996), vol. 6, pp. 61–82. 2, 3, 4
- [Gut84] GUTTMAN A.: R-trees: a dynamic index structure for spatial searching. *SIGMOD Rec.* 14, 2 (jun 1984), 47–57. doi:10.1145/971697.602266. 1, 2
- [JSS*22] JENTNER W., SPERRLE F., SEEBACHER D., KRAUS M., SEVASTJANOVA R., FISCHER M. T., SCHLEGEL U., STREEB D., MILLER M., SPINNER T., ÇAKMAK E., SHARINGHOUSEN M., MESCHENMOSER P., GÖRTLER J., DEUSSEN O., STOFFEL F., KABITZ H.-J., KEIM D. A., EL-ASSADY M., BUCHMÜLLER J. F.: *Visualisierung der COVID-19-Inzidenzen und Behandlungskapazitäten mit CoronaVis*. Kohlhammer, 10 2022, ch. 2, pp. 176–189. 2
- [Ken38] KENDALL M. G.: A New Measure of Rank Correlation. *Biometrika* 30, 1-2 (06 1938), 81–93. doi:10.1093/biomet/30.1-2.81. 3
- [KGV83] KIRKPATRICK S., GELATT C. D., VECCHI M. P.: Optimization by simulated annealing. *Science* 220, 4598 (1983), 671–680. doi:10.1126/science.220.4598.671. 10
- [Kru64] KRUSKAL J. B.: Nonmetric multidimensional scaling: A numerical method. *Psychometrika* 29 (06 1964), 115–129. doi:10.1007/BF02289694. 3
- [LA11] LESPINATS S., AUPETIT M.: CheckViz: Sanity Check and Topological Clues for Linear and Non-Linear Mappings. *Comput. Graph. Forum* 30, 1 (2011), 113–125. doi:j.1467-8659.2010.01835.x. 3
- [LK75] LENSTRA J. K., KAN A. H. G. R.: Some Simple Applications of the Travelling Salesman Problem. *J. Oper. Res. Soc.* 26, 4 (1975), 717–733. doi:10.1057/jors.1975.151. 2
- [MCMT14] MARTINS R. M., COIMBRA D. B., MINGHIM R., TELEA A.: Visual analysis of dimensionality reduction quality for parameterized projections. *Comput. & Graphics* 41 (2014), 26–42. doi:10.1016/j.cag.2014.01.006. 2, 3, 4
- [MD86] MITCHISON G., DURBIN R.: Optimal Numberings of an N x N Array. *SIAM J. Algebraic Discrete Methods* 7, 4 (1986), 571–582. doi:10.1137/0607063. 2, 6
- [MDS*17] MEULEMANS W., DYKES J., SLINGSBY A., TURKAY C., WOOD J.: Small Multiples with Gaps. *IEEE Trans. on Vis. and Comp. Graph.* 23, 1 (2017), 381–390. doi:10.1109/TVCG.2016.2598542. 3
- [MMT15] MARTINS R. M., MINGHIM R., TELEA A. C.: Explaining Neighborhood Preservation for Multidimensional Projections. In *Comput. Graph. Visual Comput. (CGVC)* (2015), The Eurographics Association. doi:10.2312/cgvc.20151234. 2, 3
- [MT20] MU W., TONG D.: Distance in Spatial Analysis: Measurement, Bias, and Alternatives. *Geogr. Anal.* 52, 4 (2020), 511–536. doi:10.1111/gean.12254. 4, 10
- [Sam84] SAMET H.: The Quadtree and Related Hierarchical Data Structures. *ACM Comput. Surv.* 16, 2 (jun 1984), 187–260. doi:10.1145/356924.356930. 1, 2
- [SDMT16] STAHNKE J., DÖRK M., MÜLLER B., THOM A.: Probing Projections: Interaction Techniques for Interpreting Arrangements and Errors of Dimensionality Reductions. *IEEE Trans. on Vis. and Comp. Graph.* 22, 1 (2016), 629–638. doi:10.1109/TVCG.2015.2467717. 3
- [SE21] SHAITO M., ELMASRI R.: Map Visualization using Spatial and Spatio-Temporal Data: Application to COVID-19 Data. In *14th Pervasive Technol. Relat. to Assistive Environ. Conf.* (2021), ACM, pp. 284–291. doi:10.1145/3453892.3461336. 7, 9
- [SV10] SPECKMANN B., VERBEEK K.: Necklace maps. *IEEE Trans. on Vis. and Comp. Graph.* 16, 6 (2010), 881–889. doi:10.1109/TVCG.2010.180. 3, 10
- [SvLB10] SCHRECK T., VON LANDESBERGER T., BREMM S.: Techniques for Precision-Based Visual Analysis of Projected Data. *Inf. Visualization* 9, 3 (2010), 181–193. doi:10.1057/ivs.2010.2.2.3, 4
- [TB83] TOUSSAINT G. T., BHATTACHARYA B. K.: Optimal algorithms for computing the minimum distance between two finite planar sets. *Pattern Recognit. Lett.* 2, 2 (1983), 79–82. doi:10.1016/0167-8655(83)90041-7. 9
- [THH01] TUNG A., HOU J., HAN J.: Spatial clustering in the presence of obstacles. In *Proc. Int. Conf. Data Eng.* (2001), pp. 359–367. doi:10.1109/ICDE.2001.914848. 4
- [TJMM20] TÖRÖK A., JANTYIK L., MARÓ Z. M., MOIR H. V. J.: Understanding the Real-World Impact of Geographical Indications: A Critical Review of the Empirical Economic Literature. *Sustainability* 12, 22 (2020). doi:10.3390/su12229434. 2
- [TMS23] THRUN M. C., MÄRTE J., STIER Q.: Analyzing Quality Measurements for Dimensionality Reduction. *Mach. Learn. Knowl. Extr.* 5, 3 (2023), 1076–1118. doi:10.3390/make5030056. 2
- [Tor52] TORGERSON W. S.: Multidimensional scaling: I. Theory and method. *Psychometrika* 17 (12 1952), 401–419. doi:10.1007/BF02288916. 3

- [Uni24] UNITED STATES CENSUS BUREAU: American Community Survey (ACS). <https://www.census.gov/programs-surveys/acs>, 2024. Accessed 09.04.2025. 7
- [VFP24] VALDRIGHI G., FERREIRA N., POCO J.: MoReVis: A Visual Summary for Spatiotemporal Moving Regions. *IEEE Trans. on Vis. and Comp. Graph.* 30, 4 (2024), 1927–1941. doi:10.1109/TVCG.2023.3250166. 3
- [VK01] VENNA J., KASKI S.: Neighborhood Preservation in Nonlinear Projection Methods: An Experimental Study. In *Artif. Neural Networks (ICANN)* (2001), pp. 485–491. doi:10.1007/3-540-44668-0_68. 2, 3, 4, 7, 10
- [VK06] VENNA J., KASKI S.: Local multidimensional scaling with Controlled Tradeoff Between Trustworthiness and Continuity. *Neural Networks* 19, 6 (2006), 889–899. doi:10.1016/j.neunet.2006.05.014. 4
- [Wan20] WANG I. J.: Topographic path analysis for modelling dispersal and functional connectivity: Calculating topographic distances using the topoDistance r package. *Methods Ecol. Evol.* 11, 2 (2020), 265–272. doi:10.1111/2041-210X.13317. 4
- [Wat02] WATTENBERG M.: Arc diagrams: visualizing structure in strings. In *IEEE Symp. on Inf. Visualization* (2002), pp. 110–116. doi:10.1109/INFVIS.2002.1173155. 6
- [WBM*21] WULMS J., BUCHMÜLLER J., MEULEMANS W., VERBEEK K., SPECKMANN B.: Stable Visual Summaries for Trajectory Collections. In *IEEE 14th Pac. Visualization Symp. (PacificVis)* (2021), pp. 61–70. doi:10.1109/PacificVis52677.2021.00016. 3, 9
- [WD23] WECKMÜLLER D., DUNKEL A.: An Application-oriented Implementation of Hexagonal on-the-fly Binning Metrics for City-scale Georeferenced Social Media Data. *Int. Arch. the Photogramm., Remote Sens. and Spatial Inf. Sci. XLVIII-4/W7-2023* (2023), 253–260. doi:10.5194/isprs-archives-XLVIII-4-W7-2023-253-2023. 7, 9
- [WGDS22] WANG H., GUPTA K., DAVIS L., SHRIVASTAVA A.: Neural Space-Filling Curves. In *Eur. Conf. Comput. Vision* (2022), pp. 418–434. 2
- [WM15] WANG Y., MA K.-L.: Revealing the Fog-of-War: A visualization-directed, Uncertainty-aware Approach for Exploring High-dimensional Data. In *IEEE Int. Conf. Big Data* (2015), pp. 629–638. doi:10.1109/BigData.2015.7363807. 3, 4
- [XC09] XIAO N., CHUN Y.: Visualizing Migration Flows Using Kriskograms. *Cartogr. Geogr. Inf. Sci.* 36, 2 (2009), 183–191. doi:10.1559/152304009788188763. 3, 6
- [ZJW21] ZHOU L., JOHNSON C. R., WEISKOPF D.: Data-driven space-filling curves. *IEEE Trans. on Vis. and Comp. Graph.* 27, 2 (2021), 1591–1600. doi:10.1109/TVCG.2020.3030473. 2, 10