





Visual Analysis of Spatio-temporal Phenomena with 1D Projections

M. Franke¹ , H. Martin^{2,3} , S. Koch¹ , and K. Kurzhals^{1,2} 

¹University of Stuttgart, Germany ²ETH Zurich, Switzerland ³Institute of Advanced Research in Artificial Intelligence, Vienna, Austria

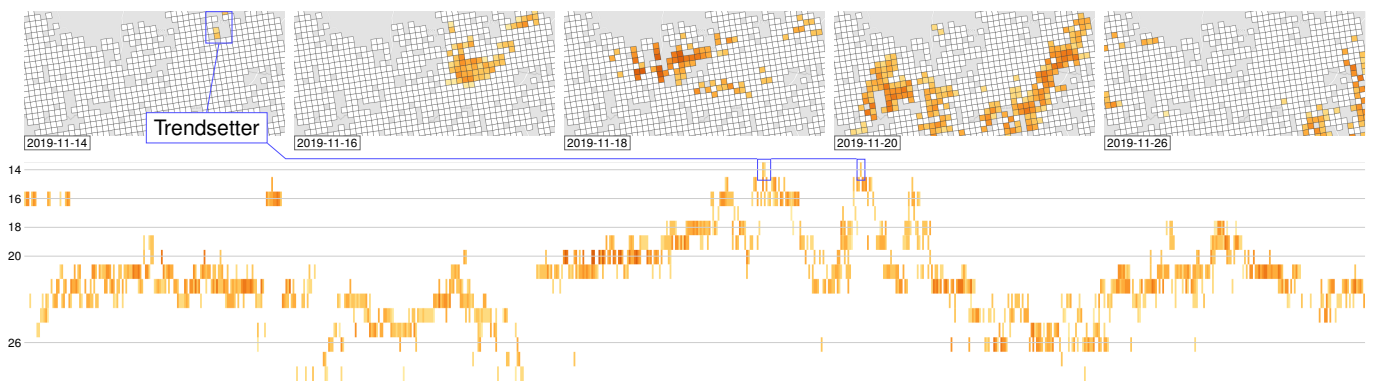


Figure 1: Geospatial entities (here square regions of Australia with side length 1 km) are projected to discrete 1D positions while attempting to conserve local neighborhoods. The resulting ordering of entities is used to generate a spatio-temporal matrix visualization (bottom), where the rows are points in time, and the columns are geospatial entities. The sequence shows a swath of wildfire spreading outwards as a ring over the course of two weeks (top, left to right), and the resulting pattern in the projection (bottom, time increases downwards).

Abstract

It is crucial to visually extrapolate the characteristics of their evolution to understand critical spatio-temporal events such as earthquakes, fires, or the spreading of a disease. Animations embedded in the spatial context can be helpful for understanding details, but have proven to be less effective for overview and comparison tasks. We present an interactive approach for the exploration of spatio-temporal data, based on a set of neighborhood-preserving 1D projections which help identify patterns and support the comparison of numerous time steps and multivariate data. An important objective of the proposed approach is the visual description of local neighborhoods in the 1D projection to reveal patterns of similarity and propagation. As this locality cannot generally be guaranteed, we provide a selection of different projection techniques, as well as a hierarchical approach, to support the analysis of different data characteristics. In addition, we offer an interactive exploration technique to reorganize and improve the mapping locally to users' foci of interest. We demonstrate the usefulness of our approach with different real-world application scenarios and discuss the feedback we received from domain and visualization experts.

CCS Concepts

• **Human-centered computing** → **Visual analytics**; **Geographic visualization**; **Information visualization**; **Visualization systems and tools**;

1. Introduction

For the analysis of spatio-temporal effects, such as patterns of propagation or localized similarities of temporal characteristics, both the spatial and the temporal context are essential to answer questions like: *Where and when did something happen? How will spatio-temporal phenomena spread/propagate? How did an event influence areas over time? Are there temporal coherences between events at different locations?* Answering such questions plays an

important role in many applications when analyzing, for example, the evolution of disastrous events, or for social media analysis. As such, time-to-time mapping visualizations (i.e., animation) are problematic because viewers must remember previous states of the visualization to retain the temporal context, and are subject to change blindness [HE11; TMB02]. Another possibility is the introduction of complex glyphs which visualize the temporal progression for one location. Such approaches, however, have limited

scalability both in the level of detail shown in the glyphs and the number of locations that can be shown [BKC*13], although aggregation can help in the latter case. To this point, many existing techniques focus more on either the temporal or the spatial aspects of spatio-temporal data, or are limited in their applicability for long time spans and large spatial context.

Recent approaches [BJC*18; WFG*18; BSC*20; ZJW21] have explored the possibility of projecting spatial data into a one-dimensional space using their position on a space-filling curve. As a result, locations which are close in 2D or 3D space are also depicted close to each other in 1D representations. Using this ordering as one axis, a temporal progression can be shown for each location on a second axis, which can reveal regional temporal patterns and propagation patterns. We extend this concept in multiple ways to create a visual analytics approach that supports the top-down and bottom-up analysis of spatio-temporal phenomena. We offer different projection methods from the 2D space to the 1D ordering and support analysts in comparing the results visually and with respective metrics to select the most sensible option for a specific dataset interactively; for instance, a more structured projection such as a Hilbert or Morton curve can be appropriate for spatially dense data that was collected at fixed positions along a grid, such as satellite measurements. Heterogeneously distributed data—for example, data attributed to separate cities—might reveal more patterns if ordered along a distance-based clustering. We further extend the approach by hierarchical aggregation and ordering, such that subtrees in themselves are continuously arranged in the 1D ordering. This extension pairs well with the interactive selection of projection methods, where different methods can be used on different aggregation levels. Overall, we provide a comprehensive approach to explore and interpret spatial and temporal aspects in geo-located data with coordinated views.

Our contributions in this work are the following: We present a new approach to explore geospatial data in space and time with neighborhood-preserving 1D timelines linked with a traditional map view. The visualization helps compare different 1D projections to detect and verify visual patterns in the data. With this setup, we support top-down and bottom-up analysis in versatile application scenarios on different scales. We demonstrate this for two real-world scenarios. We discuss the shortcomings and strengths of our approach based on domain expert feedback and generated test distribution patterns. We further consider additional scenarios where the presented approach could be applied.

2. Related Work

Spatio-temporal analysis, especially in the context of geographical data, is a rich research field. It comprises the analysis of trajectories from moving entities, but also the analysis of changes in a static spatial context [DWL08]. Andrienko et al. [AAB*11; AA13] provide an overview of visual analytics of movement, which includes the analysis of dynamic changes in a spatial context. Li et al. [LCZ*18] discuss tasks for the analysis of spatio-temporal co-occurrence. In this work, we visualize dynamic changes such as the geographical propagation of phenomena, which is crucial, for example, for the analysis of natural disasters, geopolitical developments, and social media trends. In such cases, different patterns

can be identified, for example, for the characteristics of pandemic outbreaks. These characteristics include synchrony, wave patterns, and spatial hierarchies [VBS*06]. Furthermore, we apply a spatial ordering to display temporal changes while preserving spatial proximity. Hence, related work focuses on visualization approaches for comparable analysis scenarios.

Spatio-temporal Visualization. The visualization of temporal changes inside a spatial context has been approached with different techniques to support diverse analysis goals. Peña-Araya et al. conducted a user study to compare animations, small multiples, and glyphs for the depiction of correlations [PBBH19] and temporal changes [PBP20] on maps. The authors tested for tasks about propagation distance, direction, speed, and duration of a fictional disease. While small multiples and glyphs generally outperformed animation, animation was rated best in terms of subjective task confidence and for the estimation of directions. Similarly, Boyandin et al. [BBL12] and Griffin et al. [GMH*06] compared animations and small multiples in this context. Animation is problematic for tasks such as comparison [TMB02], but allows to interpret short-term effects in the spatial context without using abstract mappings. We see the results of Peña-Araya et al. as an indicator for the importance of our linked-view approach where analysts can also investigate animations in detail (e.g., animated geo-temporal clusters [CRO14]). Phoenixmap [ZLG*19] proposes an alternative to animation by superimposing aggregated density snapshots on a map.

Visualizations that consider spatio-temporal phenomena, such as propagation effects in geographical data, can be found in numerous publications. Wang et al. [WLY*13] presented a visualization of trajectory data from road segments plotted over time that is visually similar to our projected timelines. Chen et al. [CAA*19] also applied this type of visualization to analyze movement events such as harsh braking. Liang et al. [LAC*16] and Maciejewski et al. [MLR*11] visualized the spread of diseases with small multiples and line graphs. Deng et al. [DWC*19] developed an approach for the analysis of air pollution propagation based on small multiples. MapTriX [YDGM16] combines geographical maps with a matrix visualization. The `corona-data.ch` [Pro20] dashboard also combines a map with a space-time matrix visualization, but uses case count instead of spatial relations for the ordering of geospatial entities in the matrix. Boyandin et al. [BBBL11] presented Flowstrates, a matrix-based visualization that depicts flow values over time between origin and destination locations. As we will discuss in Section 3.4, those techniques have different strengths that we adapted and combined in our visual analytics approach.

Other publications place a special focus on multivariate data analysis: Guo et al. [GCML06] and Andrienko et al. [AAB*10] visualized multivariate spatio-temporal patterns using self-organizing maps. Turkay et al. [TSH*14] displayed small multiples of multivariate geographical data as attribute signatures for comparisons in a matrix. Kim et al. [KMM*13] introduced Bristle maps for multivariate geo-visualization. Livingston et al. [LD11] discuss multivariate texture evaluation. Tominski et al. [TSAA12] used 3D stacking for multivariate trajectory data. Liu et al. [LXR18] provide multiple linked views for the analysis of spatio-temporal data, with an overview for data partitioning and glyph-based representations on the map. We apply similar concepts, but focus on a bidirectional




Figure 2: The data is hierarchically aggregated based on its geospatial property; for example, by clustering, or grouping by geopolitical entities. Each subtree is re-ordered based on a projection method. A subtree—initially, the hierarchy root’s children—is visualized in the projected timeline. Finally, on selection of an element in the projected timeline, its geospatial neighborhood is visualized in the detail timeline. Projection is performed independently for all subtrees, shown exemplary with a Morton projection on the right.

approach that provides an overview of temporal sequences, as well as their spatial distribution.

Spatial Ordering. Spatial context is typically either two- or three-dimensional. To display temporal changes on one axis of the visualization, the spatial context has to be reduced into one dimension. Popular approaches to achieve such a dimension reduction are space-filling curves [But71] or orders derived from hierarchical clustering [GG06]. Buchmüller et al. [BJC*18] applied space-filling curves to create a dense representation of spatio-temporal movement of multiple objects. In contrast, we focus on the temporal changes in spatial regions, and extend their static overview of the data with interactive means for exploring data instances in detail to support more flexible, in-depth analyses of visual patterns in their spatial context. Space-filling curves have also been applied for the exploration of volume data [WFG*18; ZJW21] on grid-based data layouts. Other approaches include creating layouts for large graphs [MM08] and jigsaw maps [Wat05]. Multivariate data can also be represented with similar techniques. Wong et al. [WFM*11] used a space-filling approach for multivariate graph data and displayed multiple parameters by small multiples. Huang et al. [HHN*17] created multiple stripes for multivariate data visualization. Pavlopoulos et al. [PKS*13] applied space-filling curves to explore the structural variation of genome data. We combine multiple spatial ordering strategies with a hierarchical overview and a linked map view. This way, we can combine the advantages of popular dashboard visualizations (e.g., the COVID-19 dashboard from ESRI [ESR20]) and visualizations that reveal spatio-temporal patterns. The advantage of our approach is that bi-directional analysis from a spatial and a temporal perspective can be performed as needed in one coherent visual analytics approach.

With respect to the quality of spatial orderings, multiple aspects can be considered, such as spatial and temporal coherence [WBM*19]. Dafner et al. [DCM00] introduced context-based space-filling curves which are sensitive to coherent regions in images. Zhou et al. [ZJW21] extend this concept by traversing on a Hamiltonian path based on a space-filling curve. Ngo and Linsen [NL20] combine dimensionality reduction to 2D with interactive 2D to 1D mapping. Guo et al. [GPG02; GG06; Guo07] suggested that complete-linkage hierarchical clustering to derive spatial orderings from geographical data provided better results than space-filling curves. We argue that depending on the task, the anal-

ysis might require different views on the data, making it difficult to determine one optimal spatial order. Hence, we included a dynamic reordering with established space-filling curves, as well as orderings based on hierarchical clustering. This approach is flexible to be extended with future techniques, for example, to emphasize temporal coherence.

3. Approach

The main purpose of our approach is the analysis of spatio-temporal phenomena in a geographical context. Hence, we first define relevant analysis tasks in this context and outline our design rationale before we present the implemented framework.

3.1. Analysis Tasks

We orient ourselves on the task taxonomy presented by Andrienko et al. [AAB*11]. While their work focuses on *movement*, the authors emphasize that the *objects* in motion can be very abstract and do not need to correspond to physical objects. As our data is time-series data for non-moving geospatial entities, we see these entities as *spatial objects*, as defined by Andrienko et al. Within these spatial objects, we further see contiguous sub-sequences of a time series as *temporal objects*. Furthermore, we base our definition of movement patterns on the taxonomy by Dodge et al. [DWL08].

- T_1 **Hotspots:** Identification and exploration of *hotspots*, that is, similar temporal patterns in a close geospatial context. This corresponds to *full* or *lagged co-incidence in space and time* of a *temporal object* in the taxonomy of Dodge et al.
- T_2 **Synchronization:** Identification and exploration of patterns of *synchronization*, either *full* or *lagged*. This would correspond to temporal patterns that reoccur or co-occur [LCZ*18] *outside* of a geospatial context; for instance, a synchronous spike in some measurement in multiple distant locations.
- T_3 **Trendsetters:** Identification and exploration of *trendsetters* for lagged variants of T_1 and T_2 ; that is, spatial objects within a co-incidence or synchronization that are the very first, or the first within the temporal context, to exhibit the relevant progression.
- T_4 **Details:** In-depth *analysis* of a spatio-temporal region of interest with fine control over the temporal progression; for example, via step-wise animation of the temporal data.

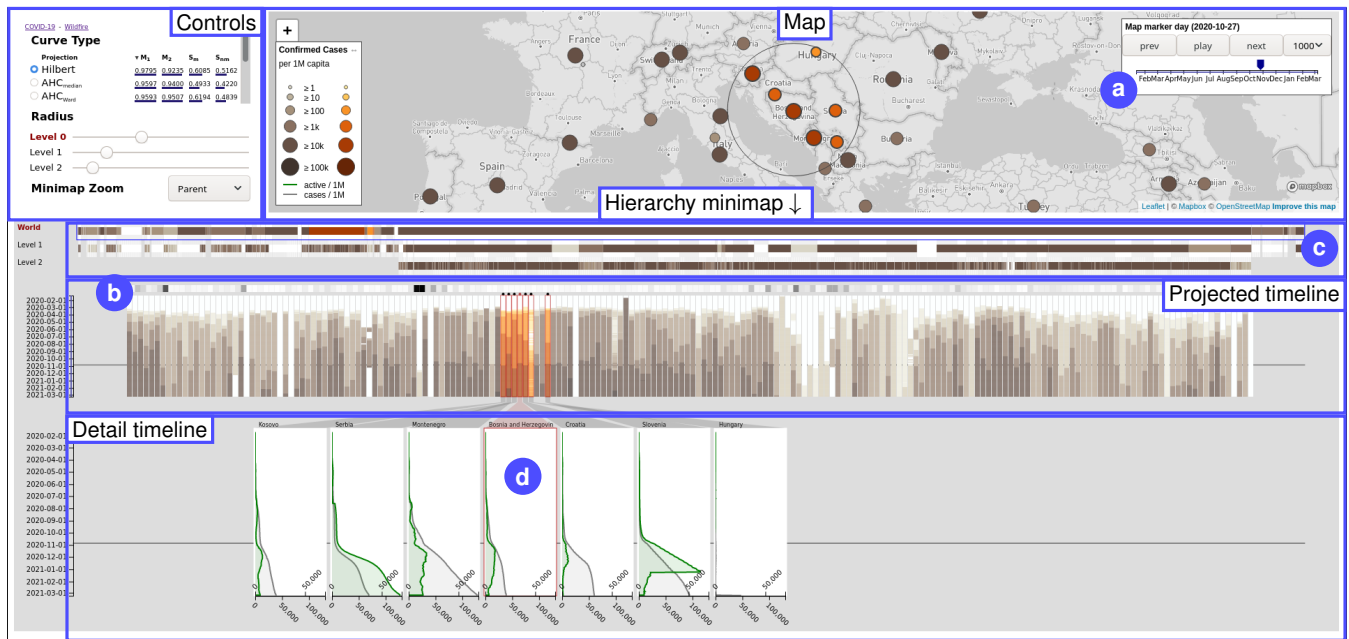


Figure 3: The visualization consists of two main components: the map and the temporal view. The map has a time control panel (a) for adjusting the currently displayed point in time and for animation control. The timeline shows a minimap of the hierarchy, the projected timeline with a bar of pairwise distance indicators (b), and the detail timeline. The subtree of the hierarchy visualized in the projected timeline is indicated in the hierarchy minimap by a rectangle (c). Clicking on an element in the projected timeline selects it as the nexus (d), which is visualized together with a neighborhood in the detail timeline. Here, we see data from the COVID-19 application scenario (Section 5.1).

T₅ Relations: Crosswise *relation* between spatial and temporal objects; for instance, identifying spatially close data, or relating the spatial context of one datum to its temporal progression.

We define generic tasks as the application scenarios and types of data that can be analyzed with our approach are very diverse. However, succinct formulations can be found for these tasks when applied to a specific scenario and dataset; for instance, for a dataset of confirmed cases of infection in different cities over time, **T₃** could be formulated as: “*Identification and exploration of regional and global infection sources.*” We will further discuss these tasks in the application scenarios (Section 5). **T₄** and **T₅** are general tasks that are typically found with interactive visual analysis approaches.

3.2. Dimensionality Reduction

Our visual analytics approach is based on spatio-temporal data which is processed in multiple stages before visualization. Figure 2 shows the data processing pipeline. Initially, we hierarchize the data based on its spatial properties, for instance by applying spatial clustering. The hierarchization step can be skipped if the structure is already given, as it would be the case with geopolitical entities (counties, countries, continents). A hierarchical processing of the data is essential to provide scalability for spatial content, but is not strictly necessary for data with few geospatial entities. The hierarchization takes place in the geospatial domain. Our approach then orders each subtree into a 1D sequence based on one from a set of projection methods. The subtrees are projected independently, as shown in Figure 2, and the hierarchization is not time-dependent.

In our prototype, we offer two types of space-filling curves—Hilbert and Morton—as a basis for the projection. A space-filling curve is constructed in the geospatial domain such that each location or entity is mapped to a unique discrete location on the curve. Space-filling curves can be constructed on uniform regular grids. For our purposes, even adaptive regular grids can be used, and we utilize this fact to construct the curves based on the leaf nodes of a quadtree, where each leaf node contains at most one entity.

Besides the space-filling curves, we offer projections based on the traversal of binary trees. These binary trees are the result of agglomerative hierarchical clustering (AHC) [Sib73]. AHC can use different linkage criteria; such as *single, average, complete, median, centroid*, or *Ward* [War63] linkage; and our prototype offers those variants where applicable. We use Great Circle distance between two geospatial locations as distance function. For the COVID-19 application scenario (Section 5.1), we also utilize a dataset of air traffic volume as an inverse distance metric between two countries [HWG*13]. We also consider dynamic time warping (DTW) [SC78] as distance metric between the time series data of two entities, using the `tslearn` [TFV*20] toolkit, as well as a *first-peak* projection which orders entities based on the first point in time their time series exceed a threshold.

3.3. Visualization Framework

We visualize this processed data in multiple coordinated views: a map view, a temporal view with spatial cues, and a control panel. In the following, we explain the different parts of our web-based

prototype (Figure 3), which we built using JavaScript, Leaflet, and D3.js [BOH11]. We have made the prototype available in the supplemental material and published the source code [FMKK21].

Map View. The *map* visualizes the geospatial aspect of the data. All data from the same hierarchy level as the currently visualized subtree are shown. A *legend* in the bottom left shows two versions for the symbols: the regular version, and the version that indicates a selection. We do this because the combination of a color scale and a selection highlight that utilizes saturation lead to confusion with darker colors. The top right of the map also contains a *time/animation control* panel (Figure 3a), which mainly allows to configure the point in time used for the map markers; for instance, a day of a dataset on day-by-day temporal granularity can be selected, and the map markers then encode the data of that day. The time/animation control panel also has buttons for stepping through time steps (T_4), and for starting and stopping animation. Animation shows each time step for a configurable duration.

Temporal View. The temporal view consists of the *hierarchy minimap* and *projected* and *detail timelines*. The *hierarchy minimap* shows an icicle plot [KL83] of the data hierarchy, which can be zoomed to only show the current subtree. A rectangle (Figure 3c) indicates the currently visualized subtree of the dataset. To increase visibility of the indicator, only the vertically central band of the icicle plot's cells encodes value by color. The cell's backgrounds are filled white and light gray in alternation to emphasize the hierarchical structure. Navigation through the hierarchy is realized by scrolling through the *projected timeline*. It visualizes the elements of the current subtree and their time series data as a matrix. In the horizontal direction, the order of elements depends on the applied projection method. The vertical direction displays the temporal extent of the dataset, where time increases downward. Adjacency in the projected timeline does not necessarily imply geospatial closeness of the respective data. To indicate geospatial distance of neighboring elements in the projected timeline, the *distance bar* (Figure 3b) maps pairwise great-circle distance to a gray-scale value, where darker cells indicate larger distances. In addition, the *detail timeline* shows a selected set of elements with a higher level of detail (T_4). These elements are generally larger, which allows for more flexibility in how to visualize the temporal data; for instance, in the COVID-19 scenario (Section 5), we visualize the temporal data using vertical line charts. This concept is generally extendable in the future by techniques such as stacked bar charts, stacked graphs, Gantt charts, or others, depending on the type of data.

Control Panel A *control panel* provides different projection method options. Selecting a different method re-orders the data hierarchically, in-place. The quality metrics for the current subtree are listed for each projection, and the projections can be sorted by any of them. The control panel also contains sliders for setting the selection cutoff radius individually for each level of the hierarchy.

Interactive Exploration. Data is linked throughout the views by using another frame color (T_5). Selecting an element in the map or the projected timeline sets that element as the *nexus* (Figure 3d). The detail timeline visualizes the nexus, as well as all elements from the same hierarchy level within a configurable *cutoff radius*.

The circle around the current selection depicts the radius. In the detail timeline, elements are sorted based on their geospatial distance to the nexus, but grouped by whether they appear before or after the nexus in the projected timeline; that is, elements left of the nexus are sorted in descending order, elements on the right in ascending order. This reduces edge crossings in the links between the two timelines while indicating *geospatial* closeness. Users can further select horizontal and vertical ranges, where the vertical range restricts the temporal extent of the detail timeline.

3.4. Design Rationale

The aforementioned tasks (Section 3.1) are important for spatio-temporal data in a geographical context. For such data, numerous solutions have been proposed [AMST11]. After reviewing existing methods, we identified animations, small multiples, glyphs, space-time cubes, and timelines as potential candidates to depict the data. We chose timelines for representing spatial context in a 1D ordering because of their good representation of temporal properties and their scalability. In particular, we aimed at producing clear spatio-temporal patterns by using timelines in combination with a neighborhood-preserving 1D projection, which we were able to confirm with test data (Figure 5) as well as real datasets (Figure 6). For the map, we selected glyphs for their representation of spatial properties, as well as animation on demand. This way, we follow the information seeking mantra [Shn96] by providing an overview with the timelines, details in the map glyphs, and a drill-down to the animated changes at locations.

We include a map view that displays animations of the data changes and uses glyphs to depict multivariate aspects. Different timeline visualizations can be displayed compactly—down to a single pixel per value—without aggregation [Kei00], providing a good scalability over time. For the timelines, we provide a compact overview with a hierarchical navigation structure and spatially ordered 1D projections. Our approach supports bottom-up analysis, for example, by selecting a specific region on the map and investigating temporal changes in this region and its neighborhood. Top-down analysis is supported by the timelines; for example, an analyst can investigate temporal developments in an overview and select regions and time frames which are then shown in detail and linked on the map supported by animation.

To represent the 2D spatial data in a discrete 1D order, we apply different techniques for this dimensionality reduction. It is essential to provide the analyst with visual support about how the results of dimensionality reduction techniques differ. Patterns might not be visible with all techniques, so guidance for comparing and selecting the best projection for the current task is necessary.

3.5. Visual Support for Comparing 1D Projections

Visual patterns that help solving the tasks (Section 3.1) are not always consistent between different projection methods. Hence, the visualization has to support the analysis by providing comparability between projections. This way, users can identify if a pattern is constant between projections and if another projection would be better for the current area under investigation.

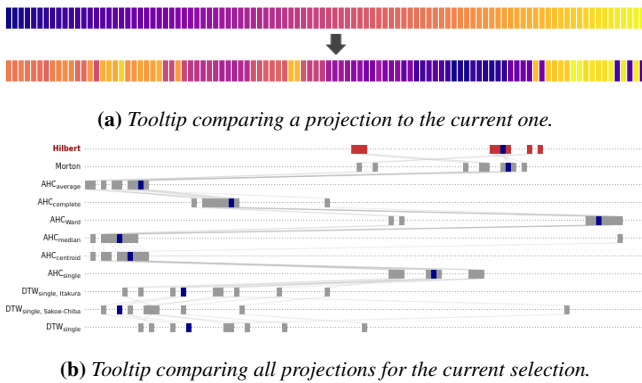


Figure 4: When hovering over a projection in the control panel, differences between the current and the hovered projection are (a) shown by applying a consistent color scheme to entities. When hovering over the selection in the timeline overview or map view, the selection is (b) visualized in all projections as a parallel coordinate plot. The current projection is red, the nexus blue.

To compare two projections on a global scale, we provide a tooltip which is visible when hovering over a projection in the *control panel*. This tooltip visualizes all entities of the currently visualized subtree as rectangles, which are ordered by the current projection order and colored using a continuous color scheme. Below that, the entities are shown again, in the order of the hovered projection, but with the same color mapping. In the example in Figure 4a, blocks that stay contiguous in the hovered projection show up clearly as blocks of similar color. We also add a path connecting all entities in the order of the hovered projection to the map, which can be used to understand the geospatial layout of the projection.

We create a second popup when hovering over a selection, either in the projected timeline or in the map. This popup visualizes the selection in the form of a vertical parallel coordinate plot. Each row corresponds to one projection, the data cases are the selected entities, and the horizontal position within each row is the index of the entity in the order of that projection. The selection's entities are colored red for the current projection, and the *nexus* is highlighted in blue in all projections. The popup provides a way to compare the effect of different projections on a selection; for instance, in the example shown in Figure 4b, we see that the selected geospatial neighborhood is not contiguous in the current Hilbert projection (red), but mostly contiguous in some AHC-based ones.

3.6. Quantitative Support for Comparing 1D Projections

We also provide quantitative support for the quality of the different projections. We use two different types of quality metrics: two neighborhood preservation measures defined by Venna and Kaski [VK01] in the context of self-organizing map dimensionality reduction; and as well-known references we use two stress measures derived from the loss function of multi-dimensional scaling (MDS) [Tor52; Kru64; GS96]. The stress measures how well the projected data represents the original data globally. In this work, the focus lies on the preservation of local neighborhoods. There-

fore, we focus more on the M_1 and M_2 scores. We provide more details on the calculations in the supplemental material.


M_1 and M_2 Scores. Venna and Kaski [VK01] proposed the M_1 and M_2 scores specifically for the application of dimensionality reduction methods for the visualization of high-dimensional data. For this use case, the authors define two possible errors with respect to the preservation of the k -neighborhood of a data point: Either a point is wrongly projected into a new neighborhood, or a point is projected away from its original neighborhood. They define the M_1 score to measure the first error, which they label the *trustworthiness* of the projection; and the M_2 score to measure the second error, which they label the *preservation* of high-dimensional neighborhoods. Both values are defined for neighborhoods $k \leq \frac{N}{2}$ and have values between 0 and 1, where higher values indicate better scores. A low M_1 score indicates that many data points are projected into new neighborhoods, which means that new cluster could emerge, or that existing clusters are exaggerated. The M_1 score can therefore be interpreted as a measure for the trustworthiness of the visualized neighborhoods. The M_2 score measures how many points were projected away from their original neighborhood and, thus, indicates how well the original neighborhoods were preserved in the projection. We use the equations as they were proposed by Venna and Kaski [VK01]. We choose the k depending on the nature of the data: For the grid-based data in the Wildfire scenario (Section 5.2), we choose $k = 8$, such that all immediate neighbor cells on the grid are selected. As the spatial layout in the COVID-19 scenario (Section 5.1) is based on countries, and most countries have at most 5 neighbors [CIA20, Field 281: “Land boundaries”], we choose $k = 5$ as a suitable heuristic. We cap k to $k_{\max} = \lfloor \frac{N}{2} \rfloor$ for a subtree of size N to keep the metrics well-defined for smaller subtrees.

Metric and Non-metric Stress. We calculate two different stress metrics. *Metric stress* (S_m), which is the loss function of classical MDS [Tor52; GS96], and *non-metric stress* (S_{nm}), which is the loss function of non-metric MDS [Kru64; GS96]. S_m is zero if the elements of the original and projected distance matrices are equal, S_{nm} applies a rank-preserving transformation beforehand and is zero if their rank order is equal. Before calculating S_m , we normalize the distance matrix by its maximum value. For both measures, a low value indicates a good projection, and both are within $[0, 1]$. For the stress measures we followed the implementation guidelines Galbraith et al. [GMBS02] provide and otherwise use the formulas as introduced in the original papers. Because we do a dimensionality reduction from 2D to 1D, some stress is unavoidable, and we argue that the M_1 and M_2 scores are better for assessing projection quality, especially for locality preservation. We still include the stresses because analysts might be more familiar with their interpretation.

4. Comparing Projections

Different projections from the geospatial domain to the discrete 1D ordering are most suited for different tasks and data. In this section, we demonstrate the effect of the projections on different test datasets we generated, and discuss the implications. For the test datasets, we use simple mathematical models without the addition of noise to generate a ground truth for patterns. The first dataset is

(a) Spreading Ring



(b) Graph-based Spreading

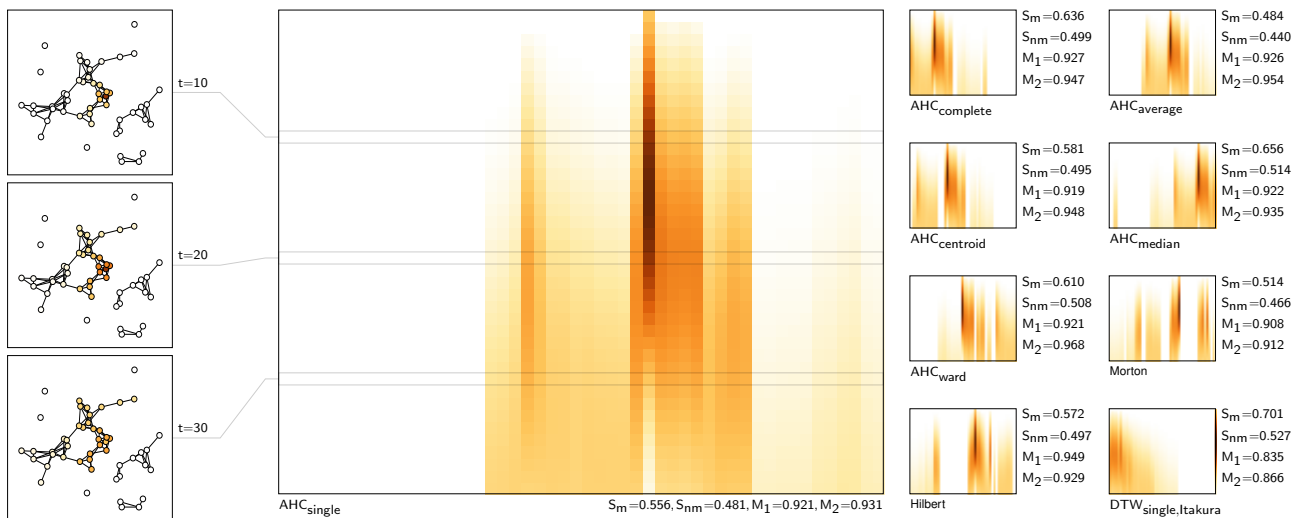


Figure 5: Two artificial spatio-temporal patterns; (a) a spreading ring of higher values, and (b) diffusion of a higher value through a graph network; are projected using our approach. For the grid-based data in (a), the space-filling curve projections perform better, whereas for the irregular entity positions in (b), the AHC-based projections perform better. We refer to the supplemental material for more examples.

generated on a regular grid, resembling the data in the Wildfire scenario (Section 5.2), whereas the second dataset is generated with an irregular geospatial layout, resembling the data in the COVID-19 scenario (Section 5.1). We describe the dataset generation in more detail and explore different variants in the supplemental material.

Spreading Ring. The first test dataset we generated was of a hotspot (T_1) that started as a single location (*trendsetter*; T_3) and then spread outwards as a ring over time, with the values falling off to zero again in the center. We observed such patterns, for example, in the Wildfire dataset, where fires would start in one place, spread outward, and run out of fuel in the center (Figure 6c). The resulting projections (Figure 5a) show the cells mapped to the horizontal axis, and the time steps mapped to the vertical axis. The increase of cells with high values over time can be seen in all projections, but Hilbert and Morton projection seem to perform best, revealing

a chevron-like pattern with a clear origin. This is not surprising, as the data is grid-based, just like these space-filling curves. The AHC-based projections perform worse, which we attribute to the fact that all geospatial entities are the same distance apart. The latter affects the hierarchical clustering negatively. The criteria which consider distances between clusters, not between cluster members, seem to perform better, with larger bits of contiguity. With these linkage criteria, the result is close to a balanced binary tree, and the projection very close to what we get with a Morton curve. For the Hilbert projection, we can also see that the contiguity of the projected pattern is very good initially, up to the point where the ring extends into other quarters of the spatial domain. At that point, the projection introduces discontinuities.

Graph-based Spreading. We generate one dataset with entities in an irregular geospatial layout. We connect the entities by links

within a cutoff-radius and simulate the spreading of a disease via the links, using a simple model. In this model, one entity's values over time are fixed, and the other entities' values at time step $t + 1$ are derived from its own value at time t , as well as those of the linked entities. This results in the values from one entity slowly propagating through the graph, as can be seen in Figure 5b. The fixed entity is visible because of its higher values, and the spreading pattern shows up in all projections. In this case, the data is not on a regular grid, and the AHC-based projections perform far better. With the single linkage criterion, the connected parts of the graph are even projected to a continuous range.


Both datasets reveal patterns, and the resulting projections are comprehensible. The shapes differ between projections, and will do so as well with different spatial layout of the data, but we can observe general similarities in the generated patterns which could provide starting points for interactive analyses in real-world application scenarios. We generate those patterns without the noise, incompleteness, multiplicity, and distortions present in real-world data as a ground truth. We were also able to find them in real-world data, such as in Figures 1, 6c and 7a. As size, rate of change, and starting position affect the clarity of patterns, we also explore different parametrizations, as well as an additional pattern, in the supplemental material. We also notice that the stress measures never go below 0.4, but increase for the DTW-based projection, which does not consider spatial layout. We conclude that these stress values are sensible for a dimensionality reduction from 2D to 1D. For reference, classical MDS on the grid positions in Figure 5a yielded $S_m = 0.4805$, $S_{nm} = 0.4424$, and non-metric MDS yielded $S_m = 0.4711$, $S_{nm} = 0.4364$. The M_1 and M_2 scores are also very good for all projections but the DTW-based one.

5. Application Scenarios


To show the versatility of our approach, we applied it to two datasets with different data properties and spatio-temporal resolution. The first scenario visualizes the current spreading of COVID-19 virus infections on a world-wide scale with daily updates. The second scenario shows the development of wilderness fires in Australia on a smaller spatial scale. Our approach can also be applied to other datasets of time series data attributed to geospatial entities.

5.1. COVID-19 Cases

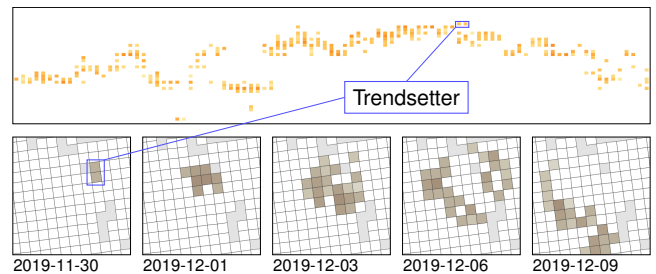
In the first scenario, we investigate the spatio-temporal spreading of diseases. During an outbreak, decision makers need to understand the spreading behavior of a disease, its dynamics, and its recurring spatio-temporal patterns. We chose data from the COVID-19 pandemic to assess the ability of our approach to help identify these patterns. The public domain data is obtained from *COVID Atlases* [Cov20a; Cov20b]. We generate a hierarchical dataset of confirmed and active cases for each day, starting from January 12, 2020 and spanning about 500 days at the time of publication. We normalize both measures to the count per 1 M capita. On the highest level, the data is aggregated on a per-country basis. For most larger countries, state or province data is included on a second level, and for the United States, a third level of hierarchy contains data on county level. Figure 3 shows a screenshot of the prototype with the data.



(a) Spatio-temporal hotspots: Several central European countries experience a strong increase in COVID-19 cases at the same time (l). A contiguous region of western Australia experiences a period of wildfire activity (r).



(b) Synchronization of an early and drastic increase in COVID-19 infection numbers in New Jersey and New York (below label). Infection numbers in Rhode Island; Washington, D.C.; and Connecticut follow a few days later.



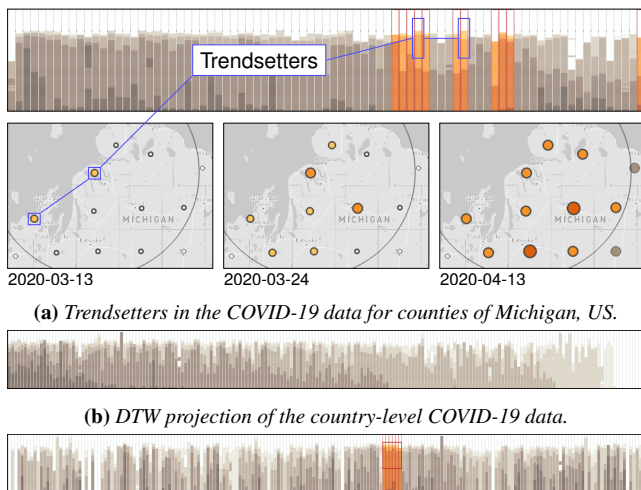
(c) Ring-like wildfire spreading pattern with trendsetters. The wildfires run out of fuel or get extinguished in the center, but propagate outwards over the course of several weeks.

Figure 6: We observe several spatio-temporal patterns of interest in the application scenarios: Hotspots (T_1) are contiguous spatio-temporal regions of high values, synchronizations (T_2) are similar temporal progressions outside of the geospatial context, and trendsetters (T_3) are the starting point of a local pattern.

We use a logarithmic threshold scale to categorize the confirmed case counts, to account for the exponential growth of the pandemic. In the projected timeline, each day in the data is represented by one rectangle for each element. The progression of confirmed cases over time for one location is visualized as a line chart in the detail timeline, as well as the number of active cases, if available.

Figure 6 shows examples for analysis targets that match the domain-specific tasks (Section 3.1). One important task is to identify hotspots (T_1), which indicate regions with a surge of cases. Such hotspots might trigger allocating more resources towards that region for disaster mitigation. Figure 6a (left) shows one such hotspot in central Europe, where some countries' infection numbers rose faster than in surrounding countries.

Another task is to identify locations outside of the geospatial context with similar progression, synchronization (T_2). In Figure 6b, New York and New Jersey (as well as RI, DC, and CT) can be identified as synchronizing, which might indicate paths of infection or similar strategies in dealing with the infection. For this



(a) Trendsetters in the COVID-19 data for counties of Michigan, US.
 (b) DTW projection of the country-level COVID-19 data.
 (c) Projection of the country-level COVID-19 data using AHC with the centroid linkage criterion, using air traffic volume between countries as an inverse distance measure. Switzerland, the Netherlands, France, Spain, the United Kingdom, and Ireland form a synchronization pattern (T_2) that could be due to high passenger exchange between the countries.

Figure 7: Projected timeline patterns for the COVID-19 scenario.

regional analysis, we could either identify these early, simultaneous peaks directly in the timeline and interactively retrieve more information about the entities; or we could find them via animation using the map and the animation pane. We could also find trans-regional synchronization, for example, by selecting contiguous similarities in the DTW projection (Figure 7b). Finding synchronization, as well as regions with different progressions, might also help decision makers to quickly identify regions with a synchronized spreading behavior to take local measures for regions with a very dynamic spreading behavior.

Finally, *trendsetters* (T_3) might indicate regional outbreak origins of the pandemic (Figure 7a). Identifying trendsetters reveals trans-regional paths of infection and helps identify what movement to restrict. One goal of decision makers during the COVID-19 outbreak 2019/2020 was to contain the spreading of the virus to protect the population, while keeping the economic impact minimal. Identifying the most problematic sources of disease spreading would allow other, vital, infrastructure to remain operating. Trendsetters can be found manually by searching in the projected timeline. However, our approach also offers two projections which facilitate this task; the DTW projection groups timelines with similar progression, and the *first-peak* projection.

5.2. Australian Wildfires

The second scenario is a forest fire management scenario. Better planning and organization of forest fire prevention and forest fire suppression are necessary to counter their increasing impact. A key role to improve forest fire management performance is “*how fire managers assimilate various types of information into decision processes*” [DCT17, p. 551]. We chose the 2019/2020 Australian wildfire season to showcase the ability of our approach to help iden-

tify spreading patterns. The data has a spatial resolution of 1 km^2 and a temporal resolution of approximately one day, depending on the overpass times of the satellites. We use data from September 2019 until February 2020 and aggregate it into three hierarchy levels based on squared grid cells with a side length of 1 km, 10 km, and 100 km, respectively. As main attribute for the fire intensity we use the fire radiative power (FRP) [GNL20]. We aggregate by summation, following the provided coarsening procedure [GSHJ15].

The identification of a fire and its spatio-temporal extent is the first task in any wildfire analysis. These *hotspots* (T_1) are directly visible in the projected timeline overview. As shown in Figure 6a (right), regions with fire activity can be identified as clusters in the overview. By selecting a rectangular area around a hotspot, the analyst sees the affected area highlighted in the map above, and more information in the detail timeline (T_4 and T_5). This allows them to further explore the data. In the context of the wildfire scenario, a *trendsetter* (T_3) describes a region where a fire started and propagated to its neighboring regions. A trendsetter can be identified in the projected timeline overview as a time-lagged staggered pattern, like for example in Figures 1 and 6c. Starting from this region, an analyst can identify the time-lagged pattern in the regions around the trendsetter region.

For an exemplary usage scenario, an analyst is interested in wildfires in the end of November, beginning of December, 2019. Using the time control (Figure 3a), they select November 30, and notice a hotspot in the map and the timeline (Figure 6a right). They identify a central cell of interest in the hotspot and select it by clicking on it in the map. They then fine-tune the cutoff radius of the neighborhood to include only the neighboring 8 cells. As the analyst is interested in a projection that preserves neighborhoods well, they order the projections in descending order of their M_2 metric (Section 3.6). By hovering over the selection in the timeline, the analyst can compare the contiguity of the selection in different projections (Figure 4b), and identifies AHC with complete linkage criterion as a good choice that has a high M_2 rating as well as good contiguity. Hovering over AHC_{complete} in the projection list, they can also confirm that the projection overall retains large blocks of the current Hilbert projection (Figure 4a). The analyst now drills down into the selected entity by scrolling in the timeline, and then again into an individual entity of high FRP value. On the lowest level of the hierarchy, they now see two entities with high values, *trendsetters* (T_3). By stepping through the following days, they can see the wildfire spreading in a ring-like pattern, which they can also identify in the detail timeline by its chevron shape when selecting a four-week time frame in the overview timeline (Figure 6c). Such analyses can be sped up by preparing analysts with generated patterns that reveal the typical shape of projected spatio-temporal phenomena, such as the *spreading-ring* pattern discussed in Section 4.

6. Expert Feedback

The presented framework has been developed by applying an iterative design process where we presented an intermediate version of the prototype to two external visualization experts and four domain experts. Each of them tested the prototype with the presented datasets. Overall, the participants mentioned the projected timeline as the key advantage of our approach compared to their usual set

of tools. One participant noted that “*traditional GIS-based visualizations mostly resort to animation or temporal filtering to let the user interact with the temporal features of the data.*” They argued that change blindness made it hard to draw conclusions over larger time spans, and that existing 1D embeddings suffer from projection errors that are not properly communicated to the user, whereas the distance bar of our approach gave them more confidence in interpreting the data.

Based on their feedback, we added the tooltips for comparing and understanding different projections. In addition to the geographically driven projections, we also added projections based on DTW [SC78]. This addition allows to identify geospatial entities with similar time series progressions (*synchronization*, T_2) outside of the geographical context. While this approach is currently limited to global time series similarity, it already yields interesting results, as shown in Figure 7b. Finally, one expert suggested including air traffic as a topology for the COVID-19 data. We did so, considering the number of passengers flying between two countries as an inverse distance metric. For pairs of countries with no flight data, and for geopolitical subdivisions, we fall back to geospatial distances. Using the air traffic distance metric reveals some patterns (Figure 7c), but we do not know if those patterns are meaningful, especially considering the age of the flight data. We consider this projection to be a proof of concept, where real-world analyses would require acquisition of more recent and more detailed travel data, including on sub-country levels.

7. Discussion

Based on the expert feedback and by applying the approach to different datasets, we identified some aspects worthwhile discussing.

Guidance on Projection Methods. The choice of projection method affects the visual outcome of the projected timeline, and thereby the start of an analysis. To support analysts, we have provided the quality metrics, and specifically the M_1 and M_2 metrics that consider locality preservation. While these metrics indicate potential candidates, our approach supports versatile applications and datasets with different data characteristics and requirements. Therefore, we provide facilities to interactively compare projections, both globally and for data selections, allowing analysts to interactively select a fitting projection for the data and task they are currently facing. Nevertheless, we have observed some general guidelines that supplement the quality metrics and interactive comparison: For grid-based data such as our wildfire scenario, the space-filling curve projections are more suited. Additionally, while geospatial context is an important part of our approach, for time-oriented tasks such as synchronization (T_2) or trendsetters (T_3), switching to our time-focused projection options can help. For the identification of spatio-temporal regions of interest, knowledge about typical patterns that emerge, and how they manifest in the projected timeline, is helpful. We identified two such patterns in the application scenarios and generated isolated variants in Section 4. In cases where patterns of interest are already known, this process can help get the analysis of new application scenarios started faster.

Target Audience. The presented approach addresses domain experts to support decision making. We can see our approach being

applied, for example, by disaster managers and prevention analysts, who need to identify regions in particular need of support. Depending on domain-specific tasks, individual parts of the visualizations (e.g., glyphs) could be modified. Based on the feedback we received by the experts (Section 6), we argue that our approach can be greatly beneficial especially when context data is incorporated into the visualization to see the larger picture.

Scalability. The presented datasets cover different scales of temporal and spatial resolution (Section 5). Our hierarchical aggregation approach, combined with timeline visualization, supports a high visual scalability in temporal as well as spatial dimension. By also hierarchizing the temporal data, the approach’s scalability in the temporal dimension could be improved even further. Our web-based prototype has been able to run smoothly on larger datasets (> 50,000 entities), such as one for the wildfire (Section 5.2) covering Australia entirely. Further optimizations, such as using WebGL for rendering, could be applied in the future.

Generalizability. Our approach is applicable to various scenarios that involve temporal changes within a spatial context. This includes single-value, as well as multivariate data. Furthermore, the concept is extendable by replacing glyphs or the visualizations in the detailed selection by alternatives that fit to a specific application scenario. While we focus on spatio-temporal analyses (Section 3.1), we want to point out that both the spatial and the temporal aspect of our approach could be substituted; for instance, the temporal aspect could be replaced by a different, continuous variable such as depth, visualizing nutrient density or temperature measurements in the oceans to support marine biologists.

8. Conclusion

We introduced a new technique for the analysis of spatio-temporal phenomena in a geographical context. By combining animation, glyphs, and spatially ordered 1D projections in multiple coordinated views, we support top-down and bottom-up analysis scenarios. The visualization supports the choice of appropriate projections for the task at hand and the analyst can compare different projections while investigating. We demonstrated the versatility of our approach in the application scenarios with feedback from domain experts. Additionally, we explored and discussed the effect of our projections on simple, artificial spatio-temporal patterns modeled after observed real-world patterns, generating a ground truth for the detection of such patterns in real-world data. Future work might include exploring time series similarity projections based on local similarity of an interactively selected time frame, thereby extending the current, global, time series similarity projection.

Acknowledgments

This work has been funded and supported by the Volkswagen-Stiftung as part of the Mixed Methods project “Dhimmis & Muslims”, by the DFG as part of the VGI priority program project “VA4VGI” and Germany’s Excellence Strategy EXC 2120/1 (project #390831618), and by a Mobility Initiative grant funded through the ETH Zurich Foundation (“Mobility of the Future” Application MI-01-19). We would also like to thank the experts for

their participation, and the reviewers for their valuable feedback. Open access funding enabled and organized by Projekt DEAL. [Correction added on 05 November 2021, after first online publication: Projekt Deal funding statement has been added.]

References

- [AAI3] ANDRIENKO, NATALIA and ANDRIENKO, GENNADY. “Visual analytics of movement: An overview of methods, tools and procedures”. *Information Visualization* 12.1 (2013), 3–24. DOI: [10.1177/14738716124576012](https://doi.org/10.1177/14738716124576012).
- [AAB*10] ANDRIENKO, GENNADY, ANDRIENKO, NATALIA, BREMM, SEBASTIAN, et al. “Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns”. *Computer Graphics Forum*. Vol. 29. 3. Wiley Online Library, 2010, 913–922. DOI: [10.1111/j.1467-8659.2009.01664.x2](https://doi.org/10.1111/j.1467-8659.2009.01664.x2).
- [AAB*11] ANDRIENKO, GENNADY, ANDRIENKO, NATALIA, BAK, PETER, et al. “A conceptual framework and taxonomy of techniques for analyzing movement”. *Journal of Visual Languages & Computing* 22.3 (2011), 213–232. DOI: [10.1016/j.jvlc.2011.02.00323](https://doi.org/10.1016/j.jvlc.2011.02.00323).
- [AMST11] AIGNER, WOLFGANG, MIKSCH, SILVIA, SCHUMANN, HEIDRUN, and TOMINSKI, CHRISTIAN. *Visualization of time-oriented data*. Springer Science & Business Media, 2011. DOI: [10.1007/978-0-85729-079-35](https://doi.org/10.1007/978-0-85729-079-35).
- [BBBL11] BOYANDIN, ILYA, BERTINI, ENRICO, BAK, PETER, and LALANNE, DENIS. “Flowstrates: An approach for visual exploration of temporal origin-destination data”. *Computer Graphics Forum*. Vol. 30. 3. 2011, 971–980. DOI: [10.1111/j.1467-8659.2011.01946.x2](https://doi.org/10.1111/j.1467-8659.2011.01946.x2).
- [BBL12] BOYANDIN, ILYA, BERTINI, ENRICO, and LALANNE, DENIS. “A qualitative study on the exploration of temporal changes in flow maps with animation and small-multiples”. *Computer Graphics Forum* 31.3pt2 (2012), 1005–1014. DOI: [10.1111/j.1467-8659.2012.03093.x2](https://doi.org/10.1111/j.1467-8659.2012.03093.x2).
- [BJC*18] BUCHMÜLLER, JURI, JACKLE, DOMINIK, CAKMAK, EREN, et al. “MotionRugs: Visualizing collective trends in space and time”. *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2018), 76–86. DOI: [10.1109/TVCG.2018.286504923](https://doi.org/10.1109/TVCG.2018.286504923).
- [BKC*13] BORGIO, RITA, KEHRER, JOHANNES, CHUNG, DAVID H S, et al. “Glyph-based visualization: Foundations, design guidelines, techniques and applications”. *Eurographics STARS*. The Eurographics Association, 2013, 39–63. DOI: [10.2312/conf/EG2013/stars/039-0632](https://doi.org/10.2312/conf/EG2013/stars/039-0632).
- [BOH11] BOSTOCK, MICHAEL, OGIEVETSKY, VADIM, and HEER, JEFFREY. “D³: Data-driven documents”. *Proc. InfoVis* 17.12 (2011), 2301–2309. DOI: [10.1109/TVCG.2011.1855](https://doi.org/10.1109/TVCG.2011.1855).
- [BSC*20] BUCHMÜLLER, JURI F, SCHLEGEL, UDO, CAKMAK, EREN, et al. “SpatialRugs: Enhancing spatial awareness of movement in dense pixel visualizations”. *Proc. EuroVA*. The Eurographics Association, 2020, 1–5. DOI: [10.2312/eurova.202010782](https://doi.org/10.2312/eurova.202010782).
- [But71] BUTZ, ARTHUR R. “Alternative algorithm for Hilbert’s space-filling curve”. *IEEE Transactions on Computers* 100.4 (1971), 424–426. DOI: [10.1109/T-C.1971.2232583](https://doi.org/10.1109/T-C.1971.2232583).
- [CAA*19] CHEN, SIMING, ANDRIENKO, GENNADY, ANDRIENKO, NATALIA, et al. “Contextualized analysis of movement events”. *Proc. EuroVA*. The Eurographics Association, 2019, 1–5. DOI: [10.2312/eurova.201911242](https://doi.org/10.2312/eurova.201911242).
- [CIA20] CENTRAL INTELLIGENCE AGENCY. *The world factbook 2020*. Washington, DC, 2020. URL: <https://www.cia.gov/the-world-factbook/6>.
- [Cov20a] COVID ATLAS. *COVID Atlas: The most comprehensive source of local COVID-19 (Coronavirus) outbreak information*. 2020. URL: <https://covidatlas.com/> (visited on 03/30/2021) 8.
- [Cov20b] COVID ATLAS. *Li: Next-generation serverless crawler for COVID-19 data*. 2020. URL: <https://github.com/covidatlas/li> (visited on 03/30/2021) 8.
- [CRO14] CRAIG, PAUL, ROA-SEÏLER, NÉNA, and OLVERA CERVANTES, ANA DELIA. “Animated geo-temporal clusters for exploratory search in event data document collections”. *Proc. InfoVis*. IEEE, 2014, 157–163. DOI: [10.1109/IV.2014.692](https://doi.org/10.1109/IV.2014.692).
- [DCM00] DAFNER, REVITAL, COHEN-OR, DANIEL, and MATIAS, YOSSI. “Context-based space filling curves”. *Computer Graphics Forum* 19.3 (2000), 209–218. DOI: [10.1111/1467-8659.004133](https://doi.org/10.1111/1467-8659.004133).
- [DCT17] DUNN, CHRISTOPHER J, CALKIN, DAVID E, and THOMPSON, MATTHEW P. “Towards enhanced risk management: planning, decision making and monitoring of US wildfire response”. *International Journal of Wildland Fire* 26.7 (2017), 551–556. DOI: [10.1071/WF170899](https://doi.org/10.1071/WF170899).
- [DWC*19] DENG, ZIKUN, WENG, DI, CHEN, JIAHUI, et al. “AirVis: Visual analytics of air pollution propagation”. *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2019), 800–810. DOI: [10.1109/TVCG.2019.29346702](https://doi.org/10.1109/TVCG.2019.29346702).
- [DWL08] DODGE, SOMAYEH, WEIBEL, ROBERT, and LAUTENSCHÜTZ, ANNA-KATHARINA. “Towards a taxonomy of movement patterns”. *Information Visualization* 7.3-4 (2008), 240–252. DOI: [10.1057/palgrave.ivs.950018223](https://doi.org/10.1057/palgrave.ivs.950018223).
- [ESR20] ESRI’S DISASTER RESPONSE PROGRAM. *ESRI COVID-19 overview*. 2020. URL: <https://coronavirus-resources.esri.com/> (visited on 10/20/2020) 3.
- [FMKK21] FRANKE, MAX, MARTIN, HENRY, KOCH, STEFFEN, and KURZHALS, KUNO. *Visual Analysis of Spatio-temporal Phenomena with 1D Projections: Code Repository*. 2021. URL: <https://github.com/UniStuttgart-VISUS/spatiotemporal1d> (visited on 04/18/2021) 5.
- [GCML06] GUO, DIANSHENG, CHEN, JIN, MACEACHREN, ALAN M, and LIAO, KE. “A visualization system for space-time and multivariate patterns (VIS-STAMP)”. *IEEE Transactions on Visualization and Computer Graphics* 12.6 (2006), 1461–1474. DOI: [10.1109/TVCG.2006.842](https://doi.org/10.1109/TVCG.2006.842).
- [GG06] GUO, DIANSHENG and GAHEGAN, MARK. “Spatial ordering and encoding for geographic data mining and visualization”. *Journal of Intelligent Information Systems* 27.3 (2006), 243–266. DOI: [10.1007/s10844-006-9952-83](https://doi.org/10.1007/s10844-006-9952-83).
- [GMBS02] GALBRAITH, JANE I, MOUSTAKI, IRINI, BARTHOLOMEW, DAVID J, and STEELE, FIONA. *The analysis and interpretation of multivariate data for social scientists*. CRC Press, Feb. 2002. DOI: [10.5860/choice.40-03386](https://doi.org/10.5860/choice.40-03386).
- [GMH*06] GRIFFIN, AMY L, MACEACHREN, ALAN M, HARDISTY, FRANK, et al. “A comparison of animated maps with static small-multiple maps for visually identifying space-time clusters”. *Annals of the Association of American Geographers* 96.4 (2006), 740–753. DOI: [10.1111/j.1467-8306.2006.00514.x2](https://doi.org/10.1111/j.1467-8306.2006.00514.x2).
- [GNL20] GIGLIO, LOUIS, NASA, and LANCE. *MODIS/Aqua+terra thermal anomalies/fire locations 1km FIRMS V006 NRT (Vector data)*. 2020. DOI: [10.5067/FIRMS/MODIS/MCD14DL.NRT.006](https://doi.org/10.5067/FIRMS/MODIS/MCD14DL.NRT.006). URL: <https://earthdata.nasa.gov/earth-observation-data/near-real-time/firms/c6-mcd14dl> (visited on 10/20/2020) 9.
- [GPG02] GUO, DIANSHENG, PEUQUET, DONNA, and GAHEGAN, MARK. “Opening the black box: Interactive hierarchical clustering for multivariate spatial patterns”. *Proc. ACM GIS*. ACM, 2002, 131–136. DOI: [10.1145/585147.5851753](https://doi.org/10.1145/585147.5851753).
- [GS96] GOODHILL, GEOFFREY and SEJNOWSKI, TERRENCE. “Quantifying neighbourhood preservation in topographic mappings”. *Proc. Neural Computation* 6 (Oct. 1996) 6.
- [GSH15] GIGLIO, LOUIS, SCHROEDER, WILFRID, HALL, JOANNE V, and JUSTICE, CHRISTOPHER O. “MODIS collection 6 active fire product user’s guide revision A”. *Department of Geographical Sciences. University of Maryland* (2015) 9.

- [Guo07] GUO, DIANSHENG. “Visual analytics of spatial interaction patterns for pandemic decision support”. *International Journal of Geographical Information Science* 21.8 (2007), 859–877. DOI: [10.1080/136588107013490373](https://doi.org/10.1080/136588107013490373).
- [HE11] HEALEY, CHRISTOPHER and ENNS, JAMES. “Attention and visual memory in visualization and computer graphics”. *IEEE Transactions on Visualization and Computer Graphics* 18.7 (2011), 1170–1188. DOI: [10.1109/TVCG.2011.1271](https://doi.org/10.1109/TVCG.2011.1271).
- [HHN*17] HUANG, TZE-HAW, HUANG, MAO LIN, NGUYEN, QUANG VINH, et al. “A space-filling multidimensional visualization (SFMDVis) for exploratory data analysis”. *Information Sciences* 390 (2017), 32–53. DOI: [10.1145/2636240.26368413](https://doi.org/10.1145/2636240.26368413).
- [HWG*13] HUANG, ZHUOJIE, WU, XIAO, GARCIA, ANDRES J, et al. “An open-access modeled passenger flow matrix for the global air network in 2010”. *PloS one* 8.5 (2013), e64317. DOI: [10.1371/journal.pone.00643174](https://doi.org/10.1371/journal.pone.00643174).
- [Kei00] KEIM, DANIEL A. “Designing pixel-oriented visualization techniques: Theory and applications”. *IEEE Transactions on Visualization and Computer Graphics* 6.1 (2000), 59–78. DOI: [10.1109/2945.8411215](https://doi.org/10.1109/2945.8411215).
- [KL83] KRUSKAL, JOSEPH B and LANDWEHR, JAMES M. “Icicle plots: Better displays for hierarchical clustering”. *The American Statistician* 37.2 (1983), 162–168. DOI: [10.1080/00031305.1983.104827335](https://doi.org/10.1080/00031305.1983.104827335).
- [KMM*13] KIM, SUNGYE, MACIEJEWSKI, ROSS, MALIK, ABISH, et al. “Bristle maps: A multivariate abstraction technique for geovisualization”. *IEEE Transactions on Visualization and Computer Graphics* 19.9 (2013), 1438–1454. DOI: [10.1109/TVCG.2013.662](https://doi.org/10.1109/TVCG.2013.662).
- [Kru64] KRUSKAL, JOSEPH B. “Nonmetric multidimensional scaling: A numerical method”. *Psychometrika* 29.2 (June 1964), 115–129. DOI: [10.1007/BF022896946](https://doi.org/10.1007/BF022896946).
- [LAC*16] LIANG, XING, AGGARWAL, RAJAT, CHERIF, ALHAJI, et al. “Visualizing malaria spread under climate variability”. *Proc. Workshop on Visualisation in Environmental Sciences*. 2016, 29–33. DOI: [10.5555/3057016.30570222](https://doi.org/10.5555/3057016.30570222).
- [LCZ*18] LI, JIE, CHEN, SIMING, ZHANG, KANG, et al. “COPE: Interactive exploration of co-occurrence patterns in spatial time series”. *IEEE Transactions on Visualization and Computer Graphics* 25.8 (2018), 2554–2567. DOI: [10.1109/TVCG.2018.285122723](https://doi.org/10.1109/TVCG.2018.285122723).
- [LD11] LIVINGSTON, MARK and DECKER, JONATHAN. “Evaluation of trend localization with multi-variate visualizations”. *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), 2053–2062. DOI: [10.1109/TVCG.2011.1942](https://doi.org/10.1109/TVCG.2011.1942).
- [LXR18] LIU, DONGYU, XU, PANPAN, and REN, LIU. “TPFlow: Progressive partition and multidimensional pattern extraction for large-scale spatio-temporal data analysis”. *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2018), 1–11. DOI: [10.1109/TVCG.2018.28650182](https://doi.org/10.1109/TVCG.2018.28650182).
- [MLR*11] MACIEJEWSKI, ROSS, LIVENGOOD, PHILIP, RUDOLPH, STEPHEN, et al. “A pandemic influenza modeling and visualization tool”. *Journal of Visual Languages & Computing* 22.4 (2011), 268–278. DOI: [10.1016/j.jvlc.2011.04.0022](https://doi.org/10.1016/j.jvlc.2011.04.0022).
- [MM08] MUELDER, CHRIS and MA, KWAN-LIU. “Rapid graph layout using space filling curves”. *IEEE Transactions on Visualization and Computer Graphics* 14.6 (2008), 1301–1308. DOI: [10.1109/TVCG.2008.1583](https://doi.org/10.1109/TVCG.2008.1583).
- [NL20] NGO, QUYNH QUANG and LINSEN, LARS. “Interactive generation of 1D embeddings from 2D multi-dimensional data projections”. *Vision, Modeling, and Visualization*. The Eurographics Association, 2020. DOI: [10.2312/vmv.202011903](https://doi.org/10.2312/vmv.202011903).
- [PBBH19] PÁLENIK, JURAJ, BYŠKA, JAN, BRUCKNER, STEFAN, and HAUSER, HELWIG. “Scale-space splatting: Reforming spacetime for cross-scale exploration of integral measures in molecular dynamics”. *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2019), 643–653. DOI: [10.1109/TVCG.2019.29342582](https://doi.org/10.1109/TVCG.2019.29342582).
- [PBP20] PEÑA-ARAYA, VANESSA, BEZERIANOS, ANASTASIA, and PIETRIGA, EMMANUEL. “A comparison of geographical propagation visualizations”. *Proc. SIGCHI Conference on Human Factors in Computing Systems*. 2020, 223:1–223:14. DOI: [10.1145/3313831.33763502](https://doi.org/10.1145/3313831.33763502).
- [PKS*13] PAVLOPOULOS, GEORGIOS A, KUMAR, PARVEEN, SIFRIM, ALEJANDRO, et al. “Meander: Visually exploring the structural variome using space-filling curves”. *Nucleic Acids Research* 41.11 (2013), e118. DOI: [10.1093/nar/gkt2543](https://doi.org/10.1093/nar/gkt2543).
- [Pro20] PROBST, DANIEL. “COVID-19 information for Switzerland” dashboard. 2020. URL: <https://www.corona-data.ch/> (visited on 03/31/2021) 2.
- [SC78] SAKOE, HIROAKI and CHIBA, SEIBI. “Dynamic programming algorithm optimization for spoken word recognition”. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26.1 (1978), 43–49. DOI: [10.1109/TASSP.1978.1163055410](https://doi.org/10.1109/TASSP.1978.1163055410).
- [Shn96] SHNEIDERMAN, BEN. “The eyes have it: A task by data type taxonomy for information visualizations”. *Proc. VL IEEE*, 1996, 336–343. DOI: [10.1109/VL.1996.5453075](https://doi.org/10.1109/VL.1996.5453075).
- [Sib73] SIBSON, ROBIN. “SLINK: An optimally efficient algorithm for the single-link cluster method”. *The Computer Journal* 16.1 (Jan. 1973), 30–34. DOI: [10.1093/comjnl/16.1.304](https://doi.org/10.1093/comjnl/16.1.304).
- [TFV*20] TAVENARD, ROMAIN, FAOUZI, JOHANN, VANDEWIELE, GILLES, et al. “tslearn, a machine learning toolkit for time series data”. *Journal of Machine Learning Research* 21.118 (2020), 1–6. URL: <http://jmlr.org/papers/v21/20-091.html4>.
- [TMB02] TVERSKY, BARBARA, MORRISON, JULIE BAUER, and BERTANCOURT, MIREILLE. “Animation: Can it facilitate?”. *International Journal of Human-Computer Studies* 57.4 (2002), 247–262. DOI: [10.1006/ijhc.2002.101712](https://doi.org/10.1006/ijhc.2002.101712).
- [Tor52] TORGERSON, WARREN S. “Multidimensional scaling: I. Theory and method”. *Psychometrika* 17.4 (Dec. 1952), 401–419. DOI: [10.1007/BF022889166](https://doi.org/10.1007/BF022889166).
- [TSAA12] TOMINSKI, CHRISTIAN, SCHUMANN, HEIDRUN, ANDRIENKO, GENNADY, and ANDRIENKO, NATALIA. “Stacking-based visualization of trajectory attribute data”. *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), 2565–2574. DOI: [10.1109/TVCG.2012.2652](https://doi.org/10.1109/TVCG.2012.2652).
- [TSH*14] TURKAY, CAGATAY, SLINGSBY, AIDAN, HAUSER, HELWIG, et al. “Attribute signatures: Dynamic visual summaries for analyzing multivariate geographical data”. *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), 2033–2042. DOI: [10.1109/TVCG.2014.23462652](https://doi.org/10.1109/TVCG.2014.23462652).
- [VBS*06] VIBOUD, CÉCILE, BJØRNSTAD, OTTAR N, SMITH, DAVID L, et al. “Synchrony, waves, and spatial hierarchies in the spread of influenza”. *Science* 312.5772 (2006), 447–451. DOI: [10.1126/science.11252372](https://doi.org/10.1126/science.11252372).
- [VK01] VENNA, JARKKO and KASKI, SAMUEL. “Neighborhood preservation in nonlinear projection methods: An experimental study”. *Artificial Neural Networks — ICANN 2001*. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2001, 485–491. DOI: [10.1007/3-540-44668-0_686](https://doi.org/10.1007/3-540-44668-0_686).
- [War63] WARD, JOE H JR. “Hierarchical grouping to optimize an objective function”. *Journal of the American Statistical Association* 58.301 (1963), 236–244. DOI: [10.1080/01621459.1963.105008454](https://doi.org/10.1080/01621459.1963.105008454).
- [Wat05] WATTENBERG, MARTIN. “A note on space-filling visualizations and space-filling curves”. *Proc. InfoVis*. IEEE, 2005, 181–186. DOI: [10.1109/infvis.2005.15321453](https://doi.org/10.1109/infvis.2005.15321453).
- [WBM*19] WULMS, JULES, BUCHMÜLLER, JURI, MEULEMANS, WOUTER, et al. “Spatially and temporally coherent visual summaries”. *CoRR* abs/1912.00719 (2019). arXiv: [1912.00719 \[cs.HC\]](https://arxiv.org/abs/1912.00719) 3.

- [WFG*18] WEISSENBÖCK, JOHANNES, FRÖHLER, BERNHARD, GRÖLLER, EDUARD, et al. “Dynamic volume lines: Visual comparison of 3D volumes through space-filling curves”. *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2018), 1040–1049. DOI: [10.1109/TVCG.2018.2864510](https://doi.org/10.1109/TVCG.2018.2864510) 2, 3.
- [WFM*11] WONG, PAK CHUNG, FOOTE, HARLAN, MACKEY, PATRICK, et al. “A space-filling visualization technique for multivariate small-world graphs”. *IEEE Transactions on Visualization and Computer Graphics* 18.5 (2011), 797–809. DOI: [10.1109/TVCG.2011.993](https://doi.org/10.1109/TVCG.2011.993).
- [WLY*13] WANG, ZUCHAO, LU, MIN, YUAN, XIAORU, et al. “Visual traffic jam analysis based on trajectory data”. *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), 2159–2168. DOI: [10.1109/TVCG.2013.2282](https://doi.org/10.1109/TVCG.2013.2282).
- [YDGM16] YANG, YALONG, DWYER, TIM, GOODWIN, SARAH, and MARRIOTT, KIM. “Many-to-many geographically-embedded flow visualisation: An evaluation”. *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2016), 411–420. DOI: [10.1109/TVCG.2016.2598885](https://doi.org/10.1109/TVCG.2016.2598885) 2.
- [ZJW21] ZHOU, LIANG, JOHNSON, CHRIS R, and WEISKOPF, DANIEL. “Data-driven space-filling curves”. *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2021), 1591–1600. DOI: [10.1109/TVCG.2020.3030473](https://doi.org/10.1109/TVCG.2020.3030473) 2, 3.
- [ZLG*19] ZHAO, JUNHAN, LIU, XIANG, GUO, CHEN, et al. “Phoenixmap: An abstract approach to visualize 2D spatial distributions”. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–14. DOI: [10.1109/TVCG.2019.2945960](https://doi.org/10.1109/TVCG.2019.2945960) 2.