

Re-spatialization of Time Series Plots

Journal Title
XX(X):1–15
©The Author(s) 2015
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/



A. Abdul-Rahman¹, S. Walton¹, K. Bemis², J. Gonnering Lein³, K. Coles³, D. Silver⁴, and M. Chen¹

Abstract

Visualizing time series is a ubiquitous aspect of many applications from science to business. The conventional designs of line graphs and temporal parallel coordinates plots focus on the depiction of the temporal context, but typically fail to convey the spatial information associated with data values at each time step. In this paper, we present a study of the design space for reintroducing spatial context to time series plots. We consider a range of options, from color encoding to glyph encoding, from abstract illustration to realistic representation, and from interactive brushing to animated 3D viewing. We evaluate these design options in conjunction with two case studies, (i) observing events in temporal sensory data of seafloor hydrothermal plumes, and (ii) observing phonemic dynamics in poems. We report domain experts' comparative assessment, which provides an insight into the effectiveness of different re-spatialization methods.

Keywords

Re-spatialization, visual mapping, time series visualization, parallel coordinates, spatiotemporal data.

Introduction

In many applications, datasets feature attribute values in both spatial and temporal contexts. The visual designs for depicting such data primarily fall into three categories:

- A. *Time-centric designs*^{1–6} focus on the variation of attribute values in the corresponding temporal context. Typically time is mapped to one of the spatial dimensions of the display. Line graphs and temporal parallel coordinates exemplify such a design strategy.
- B. *Space-centric designs*^{7–11} focus on the variation of attribute values in the corresponding spatial context. Typically time is depicted using animation rather than any spatial bandwidth of the display. For instance, several designs for flow visualization (e.g., heatmap, hedgehogs, line integral convolution) depict attribute values and vector properties in each time step of a simulation.
- C. *Spatiotemporal designs*^{12–17} allow for the depiction of both spatial and temporal contexts on a display without using animation. Typically, the three or four spatiotemporal dimensions associated with the data are projected onto the two spatial dimensions of the display, while attribute values are depicted using localized visual channels such as multivariate glyphs. Visualizing

object tracks, usually based on an orthogonal projection along the temporal dimension, exemplifies this design strategy.

The choice of a design strategy is largely application- and/or data-dependent. In some cases, the critical requirement for visualizing temporal variations may override the occasional desire for perceiving the spatial context. In other cases, space-centric or spatiotemporal designs may not be effective in separating different time series because of limited spatial resolution and movement. *Multiple Coordinated Views* (MCV)¹⁸ can provide a linked representation of space and time through two or more views, but can sometimes suffer from higher cognitive demand on users due to context switching when switching focus between views¹⁹. We wanted to investigate complementary designs for re-introducing the spatial context in-place, such that the

¹Oxford e-Research Centre, University of Oxford, UK

²Department of Marine and Coastal Sciences, Rutgers University, US

³Department of English, University of Utah, US

⁴Department of Electrical and Computer Engineering, Rutgers University, US

Corresponding author:

Alfie Abdul-Rahman, Oxford e-Research Centre, University of Oxford,
7 Keble Road, Oxford, OX1 3QG, UK.

Email: alfie.abdulrahman@oerc.ox.ac.uk

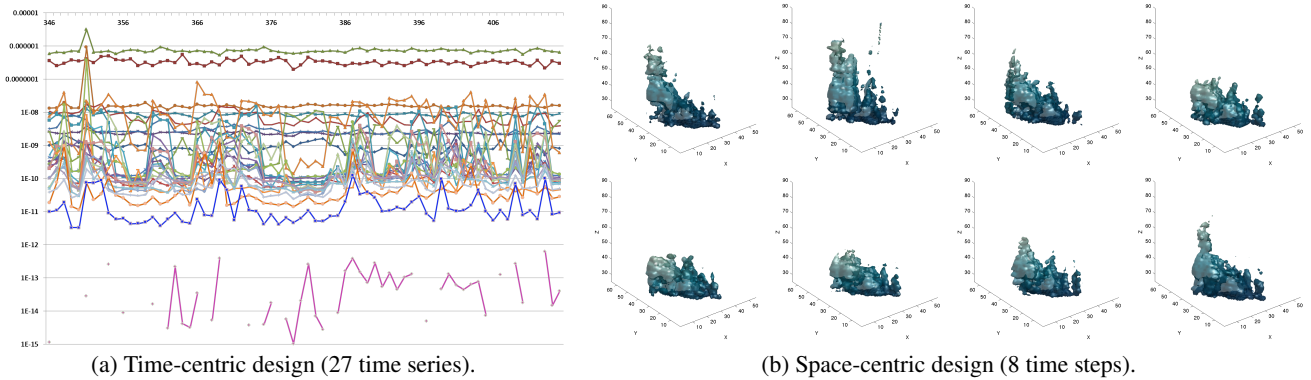


Figure 1. (a) The activities of a hydrothermal plume system are represented by 27 time series of normalized scattering strength on a $3 \times 3 \times 3$ grid. (Note: zero-values appear missing due to the use of a log axis.) (b) Snapshots of the isosurfaces of the same plume system at specific time steps.




Figure 2. Spatial information can be added into a time series plot using different methods. The visualization shows 27 spatially-related time series in the same context as Figure 1. While a temporal parallel coordinates plot allows users to observe scalar and temporal variations of different time series over time, its original form cannot convey the spatial context of each time series effectively. Five different methods for re-introducing such spatial contexts are examined in this paper. They are (a) color encoding, (b) glyph encoding, (c) selection-based exploration, (d) abstract spatial context, and (e) realistic spatial context. (Note: The jigsaw icons indicate missing time-steps in the original data, for example in time-steps 18, 41, and 47.)

user can reason about both space and time in the same plot. This paper focuses on design strategy (A), i.e., time-centric designs, and presents a study of the design space for re-introducing spatial context to time series plots, such as line graphs and temporal parallel coordinates plots. We refer to the collection of methods for such a purpose as *re-spatialization*, distinguishing it from the term *spatialization* that also encompasses the spatial mapping of non-spatial attributes.

In order to strike a balance between applicability and generality, we conducted this design study in conjunction with two different case studies. One is a classical scientific application, where oceanographers need to observe patterns captured in time-varying volume datasets. The other is a humanities application, where poetry scholars need to observe phonetic flow patterns in poems. The datasets in these two case studies exhibit very different spatial characteristics, and thereby lead to different requirements for re-spatialization (see Section ‘Definitions and Requirements’). In both cases, the above-mentioned

design strategies (B) and (C) have been explored, but found ineffective (see Sections ‘Case Study I: Visualizing Events in Temporal Plume Data’ and ‘Case Study II: Visualizing Temporal Phonemic Patterns in Poems’). While initial attempts to deploy strategy (A) showed promising signs of observable temporal patterns, users indicated a noticeable amount of difficulty in understanding the corresponding spatial context at different time steps.

Figure 1(a) shows 27 spatially associated time series summarizing a hydrothermal plume system with normalized scattering strength at $3 \times 3 \times 3$ grid points and are computed from the corresponding time-varying volume dataset. Although one can easily observe both that different time series vary differently over time and that there are diverse conjoint patterns, gaining a sense of spatial context at each time step requires a mental mapping from the identifier of each time series to its spatial position in the $3 \times 3 \times 3$ grid. Animating isosurface images (i.e., design strategy (B)), as illustrated by snapshots in Figure 1(b), has been found ineffective for analytical tasks requiring

comparisons of temporally distant frames or observation of quantitative distinctions^{20;21}. Figure 2 illustrates five different re-spatialization methods, namely *color encoding*, *glyph encoding*, *selection-based exploration*, *abstract spatial context*, and *realistic spatial context*, which are to be detailed in Section ‘Re-spatialization Methods’. We report domain experts’ evaluation in Section ‘Case Study I: Visualizing Events in Temporal Plume Data’.

Poetry scholars are interested in observing the “flow of sound” among different poetic lines in relation to the geometric features of each phoneme as illustrated in Figure 3(a). Glyph-based visualization (i.e., design strategy (C)) as shown in Figure 3(b) has enabled scholars to examine a single poetic line or compare 2-3 lines, but becomes ineffective when there are more than three lines. As shown in Figure 3(c), one may deploy a temporal parallel coordinates plot, where a path traces the vowels across each poetic line in a poem. However, the lack of spatial context in such a plot makes it difficult to perceive “flow of sound” in relation to the mouth geometry. In this study, we applied five re-spatialization methods, analogically equivalent to those in Figure 2. We report domain experts’ evaluation in Section ‘Case Study II: Visualizing Temporal Phonemic Patterns in Poems’.

Spatialization is a generic methodology in visualization for mapping non-spatial attributes (e.g., weight, accuracy, and risk) to spatial variables (e.g., length, size, and location). *Re-spatialization* is different in two ways. (i) There is usually not enough screen space for displaying the original spatial information in its native form, because of, for instance, the use of a spatial dimension for time. (ii) The original spatial information in its native form can serve the ground truth reference and expectation, while some re-spatialization methods may incur extra effort of familiarizing. In the literature, there is not yet much previous work on this topic. The main contribution of this study is to provide a systematic comparison of five different methods of re-spatialization through two distinct application case studies. Our design and implementation can provide future research on this topic with a set of baseline options, while the evaluation by the domain experts involved can offer insight as to the effectiveness of each method.

Related Work

In this section, we give a brief overview of the related work according to the three design strategies described in Section ‘Introduction’. There are numerous visual designs of spatiotemporal visualization, and understandably this section can only cover a small number of these designs. Readers are encouraged to consult several surveys^{23–28} and web sites (e.g., <http://survey.timeviz.net/> and <http://spacetimecubevis.com/>) for a comprehensive coverage.

As most visual designs for depicting spatiotemporal data involve “overloading” the display space, they naturally fall into different categories of *visual multiplexing*²⁹.

Time-centric designs. With this design strategy, the temporal dimension is given the principal attention in the visualization, while the spatial dimensions play a secondary role. It typically employs a conventional time series visualization (e.g., line plots and calendar view) as the base representation augmented with spatial information in abstract forms or available on demand.

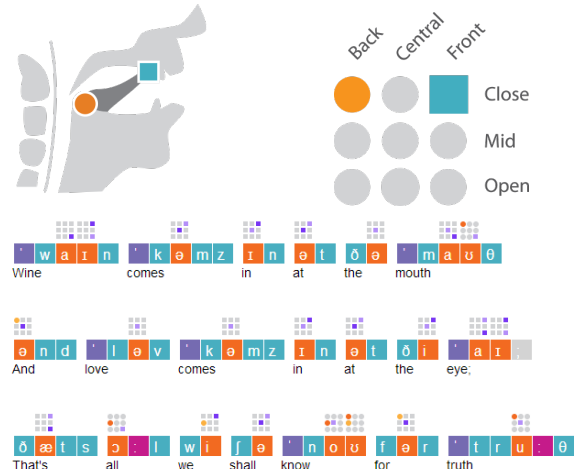
For example, when one visualizes time-varying variables (e.g., GDP, population, or volume of traffic) associated with different locations (e.g., countries or cities), one may simply display the variables using polyline representations of time series data, and the locations in a textual or iconic form (e.g., names or flags). E. J. Marey’s 1880 train schedule visualization, featured on the dust jacket of Tufte’s book³⁰, exemplifies such an approach.

This approach of mapping spatial information to a non-spatial visual channel or a display dimension that has not been utilized for conveying temporal variation usually incurs a loss of spatial information, and often introduces difficulties in spatial awareness and reasoning. On the contrary, there are many examples of spatializing non-spatial data in conjunction with a time-centric visualization, such as documents and baby names in ThemeRiver^{1;2}, and glyph-based event annotation along a time line³. This suggests that spatialization is in general beneficial, while the absence of intrinsic spatial mapping of spatial data can be problematic.

Boyandin *et al.*⁴ annotated a pixel-based time series with spatial information at the beginning and the end. In time-varying volume visualization, it is common to use a time-centric representation as an interface for users to select a specific time step and to visualize the corresponding volumetric spatial data. For example, Lu and Shen⁵ used a temporal storyboard, and Lu *et al.*⁶ used statistics graphics as such an interface for exploring time-varying volume data.

In this paper, we study the time-centric strategy in a structured manner by exploring five different re-spatialization methods, and two applications. Unlike the aforementioned works that tended to focus on one method and one application, we hope to obtain a more objective comparison, while maintaining sufficient relevance to applications.

Space-centric designs. With this design strategy, the spatial dimensions of the data are given the principal attention, while the temporal dimension of data has to utilize the display bandwidth spared by visual objects for depicting spatial information. C. J. Minard’s visualization of the fate of Napoleon’s army in Russia³⁰ is one of the most notable examples of this strategy.



(b) Glyph-based visualization of three lines of a poem.




Figure 3. (a) The vowel chart shows three variables: *frontness* – (from left to right) *back*, *central*, and *front*; *height* – (from top to bottom) *close*, *mid*, and *open*; and *roundedness* – all rounded vowels are shown in circles. (b) Glyph-based visualization of the first three lines of “A Drinking Song” by W. B. Yeats depicts the geometric features of each phoneme, with two images on the top illustrating the glyph representation of a vowel transition from close back rounded to close front unrounded. (c) Two temporal parallel coordinates plots show the frontness and vowel height attributes of the first six lines of “A Drinking Song”.

Daniel and Chen¹⁵ distorted the temporal dimension of a video-cube, while displaying only spatial data featuring noticeable changes. Botchen *et al.*¹⁶ distorted a selected sequence of video frames as a spatial context, while depicting critical spatial data along a time line. Drocourt *et al.*¹⁷ introduced an abstraction by mapping 2D spatial dimensions to 1D while mapping the temporal dimensions to concentric rings.

Spatiotemporal designs. This design strategy attempts to give a similar level of attention to both space and time. It typically compromises the intrinsic depiction of both types of information with some abstraction, omission, or distortion. Hägerstrand's space-time cube¹² exemplifies

Prepared using sagej.cls

the notation, we assume that values in \mathbf{R}_t are organized using a flat structure, and all data records are defined based on the same set of variables. This simplification does not undermine the generality since data records defined with a hierarchical structure can be flattened and those based on different variable sets can be unified using the union of all variable sets. \mathbf{R}_t can thus be written as $\langle v_1, v_2, \dots, v_i, \dots, v_m \rangle$.

In many applications, some values in \mathbf{R}_t are spatial positions, which are typically difficult to convey in traditional time series plots if there is more than one spatial dimension. To differentiate such spatial positions from other attribute values, we separate them into two subsets \mathbf{p}_t and \mathbf{a}_t , i.e.,

$$\mathbf{R}_t = \langle \mathbf{p}_t, \mathbf{a}_t \rangle = \langle v_1, v_2, \dots, v_i, \dots, v_m \rangle.$$

This paper focuses on conveying \mathbf{p}_t for different time series displayed in a shared line graph or parallel coordinates plot. Given the above notation, there are two major classes of multivariate time series, depending on whether or not \mathbf{p}_t changes within the same \mathcal{T}_{id} . For example, in a typical time-varying volume dataset, each voxel defines a time series, for which \mathbf{p}_t does not change. If one displays values of different voxels in a line graph or a parallel coordinates plot, each voxel identifier implicitly defines its spatial position. Nevertheless, as shown in Figure 1(a), it is still difficult for users to “visualize” where each voxel is located based on the voxel identifier. On the other hand, if one tracks a feature (e.g., the center of a vortex) inside the volume, \mathbf{p}_t changes over the time, and has to be explicitly defined in each \mathbf{R}_t . When one visualizes multiple features (e.g., plotting velocity vs. time), it is often more challenging to convey \mathbf{p}_t in such a plot. The problem of visualizing phonemic features in poems as shown in Figure 3(a) falls into this latter class, because when phonemes change along a poetic line, the corresponding positional features in the mouth vary accordingly.

It is helpful to distinguish these two classes as the requirements for re-spatialization are likely to be different. Given a multivariate time series, \mathcal{T}_{id} , if \mathbf{p}_t is mapped to an equivalent visual representation (e.g., a screen location) for all t , this is said to be a *stable spatial mapping*. Otherwise, it is an *unstable spatial mapping*.

For the class of *stable spatial mapping*, the users are usually aware of the correspondence between the spatial location and the identifier of each \mathcal{T}_a , but may find remembering such correspondence cognitively demanding. Hence the typical requirements for some types of data analysis may include the following questions:

- Remind me where this time series \mathcal{T}_{id} is;
- Remind me how \mathcal{T}_a and \mathcal{T}_b are spatially related, as they exhibit similarity (or deviation) in attribute values;

- Show me an overview of the attribute values at t in the corresponding spatial context.

On the other hand, for the class of *unstable spatial mapping*, the users do not know where each \mathcal{T}_{id} is at a specific time t , not to mention remembering the spatial context at t . In addition to the above requirements for stable spatial mappings, the typical requirements may also include answers to the following questions:

- Where is time series \mathcal{T}_{id} now at t ?
- How has \mathcal{T}_{id} changed spatially between t_1 and t_2 ?
- Are \mathcal{T}_a and \mathcal{T}_b spatially co-located at t , since they exhibit some similarity (or deviation) in their attribute values at t ?
- Does the spatiotemporal pattern in \mathcal{T}_a between t_1 and t_2 also occur in \mathcal{T}_b , \mathcal{T}_c , etc.?
- What is the spatial distribution of different time series at t , and are there any interesting patterns (e.g., clusters, symmetry)?

Re-spatialization Methods

The goal of *re-spatialization* is to provide visual support to the requirements discussed in the previous section. The main constraint is that we do not have a dedicated screen dimension (e.g., x or y -axis) for encoding spatial data \mathbf{p}_t in different time series \mathcal{T}_{id} . Naturally, there are many design options, such as mapping spatial data to non-spatial visual channels, utilizing unused space, or superimposing a pop-up window for conventional spatial mapping. In this work, we examine five different re-spatialization techniques for helping users observe or recall the spatial context in a time series plot. They are:

- Color encoding:** mapping positions \mathbf{p}_t to colors, and associate them to a \mathcal{T}_{id} ;
- Glyph encoding:** mapping positions \mathbf{p}_t to specially-designed glyphs, and associate them to a \mathcal{T}_{id} ;
- Selection-Based:** organizing different positions \mathbf{p}_t into an ordered list (e.g., using an axis), and using interaction (e.g., brushing) to explore different positions in the list, and to highlight those time series or parts of a time series corresponding to the selected positions;
- Abstract Spatial Context (ASC):** depicting a spatial context using an abstract representation of a set of positions (and associated data);
- Realistic Spatial Context (RSC):** depicting a spatial context using a more realistic representation of a set of positions (and associated data);




Figure 4. Comparative illustration of the basic time series plot (a) and three re-spatialization methods (b-d).

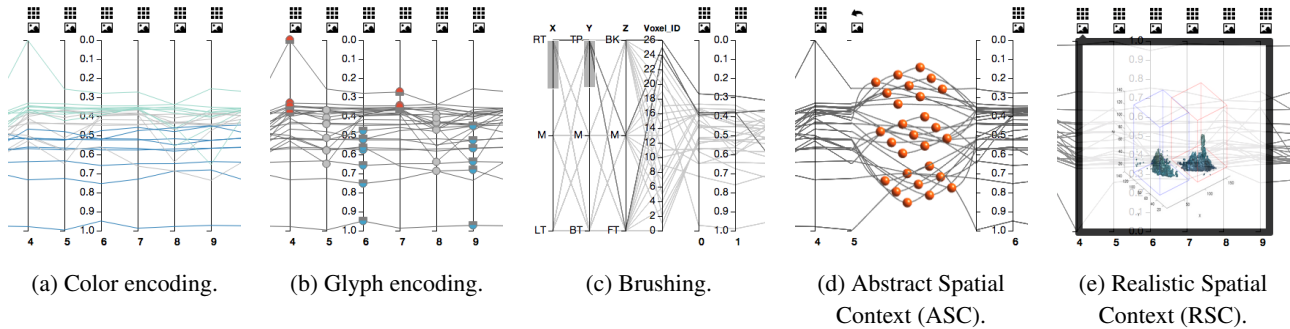


Figure 5. The five visual designs that have been implemented in the Plume Explorer software for Case Study I.

The most common approach to re-spatialization is to display a corresponding space-centric visualization at a selected time t ³¹. This is referred to as *Realistic Spatial Context* (RSC).

Time-centric visualizations (e.g., line plots) often feature several or tens of time series. In such a situation, the re-spatialization can focus on an abstract spatial context at a lower spatial resolution. The two examples as shown in Figures 1 and 3 exhibit this characteristic.

In Figures 1(a,b), the original volumetric resolution is simply too fine and too noisy for observing and discovering signature patterns for different events. Hence a $3 \times 3 \times 3$ low resolution volume was derived from the original volume data set at each time t , yielding 27 time series. The choice of size $3 \times 3 \times 3$ in this application was determined by the domain expert involved. This resolution is expected to show most common types of phenomena being analyzed in the plume data (e.g., the plume bending). The general applicability of this technique to other types of volume data would depend upon the task required by the domain expert. Meanwhile, phonemic positions obtained from text are constrained by the IPA standard translation. In general, the level of resolution that is to be displayed is application-dependent, and small variations of resolution are trivial to implement. A low resolution spatial context offers the ease to distinguish between different locations and to remember them, while its encoding requires a small amount of bandwidth of a visual channel. Hence we refer to such low resolution spatial context as *Abstract Spatial Context* (ASC).

Given a small set of spatial positions, such as the $3 \times 3 \times 3$ grid points of a low resolution volume, we can potentially encode them using a non-spatial visual channel. In this work, we considered using *Color Encoding* and *Glyph Encoding* to represent a small set of spatial positions.

Last but not least, interaction facilitates spatial exploration. If one can select a specific spatial range to make an observation, and select another range for another observation, gradually one can build a spatial context. We refer to this method as *Selection-based Exploration*.

Figures 4(b-d) illustrate three re-spatialization methods: Color Encoding, Glyph Encoding, and Selection-based Exploration. Each method can be used either on its own, or in conjunction with others in an integrated manner. Questions that naturally arise may include: Can users interpret those colors in Figure 4(b), or glyphs in Figure 4(c), as spatial positions? How easy is it for users to build a spatial context from partial views, e.g., Figure 4(d), selected interactively?

Figure 5 and Figures 6(a-e) show two sets of examples of applying these five methods in our case studies, which will be detailed in Sections ‘Case Study I: Visualizing Events in Temporal Plume Data’ and ‘Case Study II: Visualizing Temporal Phonemic Patterns in Poems’, where we will also answer these questions.

Designing an Abstract Spatial Context (ASC)

The methods of ASC and RSC are designed to be activated for a specific time step t_i upon request by the user. Our goal is to bring the spatial context into focus and integrate this additional visual information into the time series at t_i , in

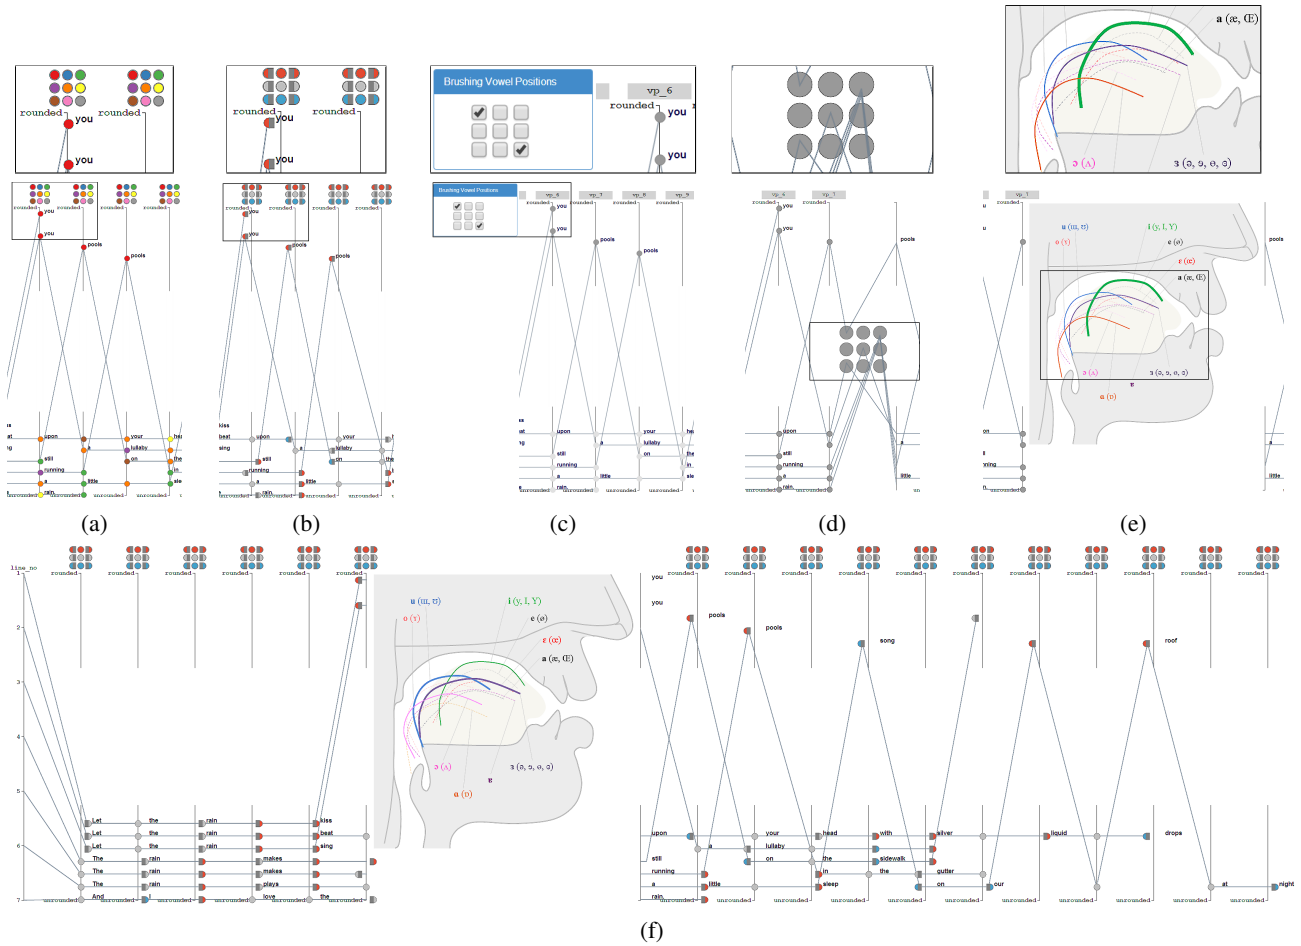


Figure 6. The five different re-spatialization designs considered in the humanities application are (above: zoomed in view and below: standard view) (a) color encoding, (b) glyph encoding, (c) selection-based exploration, (d) abstract spatial context (ASC), and (e) realistic spatial context (RSC). (f) A visualization of a short poem “April Rain Song” by Langston Hughes using an integrated design space of color encoding, glyph encoding, and RSC. A path (from left to right along the x-axis) represents a line in a poem. The first vertical axis on the left is used to mark the line index. The remainder of the vertical axes represent the roundedness attribute (i.e., the shape of the lips during the production of a vowel), where each of these axes is divided into two parts: *rounded* at the upper half and *unrounded* at the lower half.

order to convey the spatial context. We considered many designs for such an abstract representation, and Figure 7 illustrates the six different designs that were considered. Four of them are 2D designs with a 3×3 grid and two are 3D designs with a $3 \times 3 \times 3$ grid.

When an axis is to be spatially enriched, it expands to a rectangle with the same axis duplicated at both sides. Lines of different time series enter the *re-spatialization rectangle* from the left, connect with their associated spatial positions in the rectangle, and leave the rectangle on the right to become the normal time series again. The approach is similar to that of Yuan *et al.*³² who use multidimensional scaling on two or more axes to display connected scatterplots between axes, except that

our approach favours augmenting the user’s ability to recall spatial position and relate it to the axis values.

Spatial Relation for Categorical Axes

In time series visualization, when two or more lines share the same attribute value at a time step, tracking them can be a problem (as shown in Figure 8(a)). The problem becomes worse with cluttering and occlusion, e.g., when lines have to be annotated by glyphs or have to merge into a point within a re-spatialization rectangle for depicting abstract or realistic spatial context. When the axes concerned are categorical axes with a small number of category values, we can relax the path of the time series when they pass through these axes. For each value on a categorical axis, we define a “relaxation” region, and spread the passing lines within the

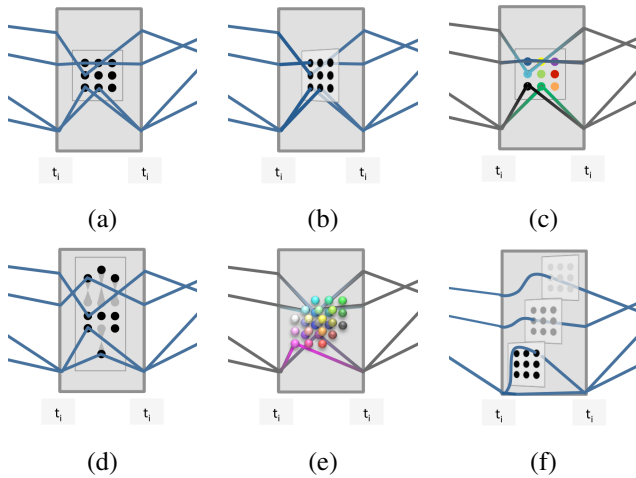


Figure 7. Six ASC design options to be integrated into a time series. (a)-(d) 2D designs with a 3×3 grid. (e)-(f) 3D designs with a $3 \times 3 \times 3$ grid.

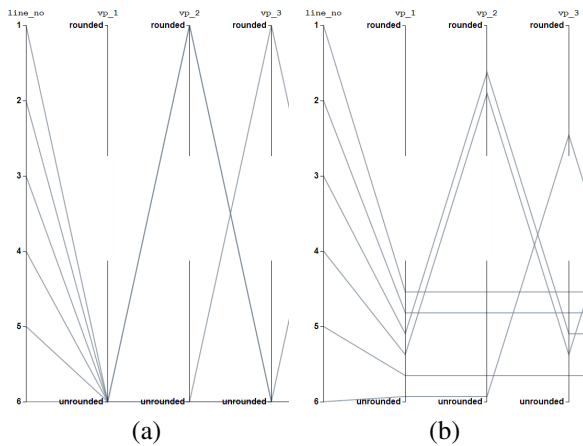


Figure 8. (a) Intersecting points with no relaxation (b) Intersecting points with controlled relaxation.

region. To maintain some form of ordering, for each passing line, the original position p is moved to a new position that is a function of p and its line index as shown in Figure 8(b).

Case Study I: Visualizing Events in Temporal Plume Data

Our first case study represents a scientific application where time series plots are used to visualize 4D hydrothermal plume data. A hydrothermal plume is primarily characterized by the rapid vertical discharge of high-temperature fluid through particle-rich sulfide chimneys along the axis of mid-oceanic ridges^{33–35}. The process is driven by volcanic activity heating groundwater flow under the surface. Hydrothermal plumes exhibit complex temporal behavior, rising through surrounding

water until they become denser than the water, at which point the plume begins to expand laterally. These fluxes and interactions with the ocean are studied in order to further scientific understanding of how the plumes may be distorted by different types of currents and geologic events.

Data Collection and Analysis

In addition to scientific discovery, analyzing the data informs researchers of the data quality and correctness. In this study, we utilize data from the COVIS (Cabled Observatory Vent Imaging System) project^{36,37} in which data is collected from a sonar platform placed near targeted plumes. The multi-beam sonar measures returning backscatter resulting from both the turbulent fluctuations and the concentration of metallic sulfide particles within the plume. Sonar scans are performed at regular intervals of around three hours.

For each dataset representing a time step, post-processing is performed in Matlab to form a volumetric dataset which can be further processed or directly volume rendered. The former methods allow for temporal analysis of the plume's behavior, such as its bending properties which have been the subject of much scientific interest³³; whilst the latter are useful for judging the plume's overall shape³⁸. Assessing the quality and correctness of the datasets is difficult due to its spatiotemporal nature and the uncertainty of the capture method: the sonar platform often produces incomplete or corrupted scans due to network errors, disk space issues, or instrument failure. In addition, the analysis of longer time periods is becoming more commonplace.

Data Visualization and Exploration

In this case study, the temporal volume data belongs to the *stable mapping* class, where each voxel defines a time series in which the voxel position does not change. We wish to provide a tool that will allow researchers to explore the changing properties of the plume's spatial distribution over time, whilst re-introducing the spatial context back into a more traditional time series visualization. As with Figure 1(a), each volume dataset is resampled into a $3 \times 3 \times 3$ grid by taking the mean of the plume's normalized backscatter values within each block. We favored this approach over more advanced aggregation methods to focus on the effectiveness of the visualization's ability to represent bending artefacts caused by changes in density at spatial positions. This aggregation allows not only for the plume's large-scale spatial distribution over time to be analyzed for interesting properties (such as bending), but also provides an effective summarization of the plume's data quality.


We begin by plotting each of the plume's positions in a line plot, as in Figure 1(a). Providing an effective visual mapping for 27 voxel positions in a line plot is

difficult as one quickly runs out of visual channels (color, stipple, markers, etc.) in which to distinguish each line. In addition, we require a visual mapping that (i) does not cause occlusion problems; (ii) does not cause problems for those with color perception issues; and (iii) is *memorable* and/or analogous to a spatial position in 3^3 space such that associating each line with its corresponding voxel is not cognitively demanding.

Our five re-spatialization methods discussed in Section ‘Re-spatialization Methods’ are mapped to five visual design options in an integrated web-based application called *Plume Explorer* (see Figure 2). This time series plot has unique features which we now discuss in relation to the re-spatialization methods.

Using Line Color & Glyphs. Our first two approaches, as seen in Figure 5(a,b), map 3×3 slices through the y-axis to three distinct visual line properties. Such a mapping effectively separates the bottom, middle, and top sections of the plume for independent analysis of each section’s unique characteristics. We provide two visual mappings: line color and line glyph. For each voxel p belonging to a block of 3×3 voxels, the line color approach maps p to a color representing that block, with a user-customizable color scheme. In comparison, the glyph approach uses specially-designed glyphs to denote the bottom, middle, and top sections (shown in Figure 5(b)), which are placed where each p intersects a time axis. To reduce occlusion, we constrained each type of glyph to every third axis.

Axis Brushing for Spatial Axes. Brushing is a standard and widely-used technique in parallel coordinates plots for visually filtering data matching a particular range or intersection of ranges. By introducing X , Y , and Z spatial axes into our time series plot (always shown as the first three axes), we allow the user to brush location as well as data value (see Figure 5(c)). Selecting data value is achieved by brushing on the t -axes, and is useful for identifying spatial areas of the plume that are within a certain value range or exhibit particular temporal patterns of flow. Being able to brush on spatial location achieves the opposite: allowing the user to inspect particular spatial regions of the data along the time axis by brushing on each of the X , Y , and Z axes. Their intersection represents a region in space.

Abstract Spatial Context. The ASC approach allows the user to reintroduce the context of each p at any point along the time series. In the *Plume Explorer* application, this is achieved by clicking the  icon above the axis. Upon re-spatializing an axis, the axis is split into two and laterally expands to show the abstract representation of the $3 \times 3 \times 3$ space. This ASC is visually represented by a 3^3 grid of small spheres (see Figure 5(d)). We chose to render lit spheres to provide a valuable visual cue of the three-dimensional nature of the scene. In a re-spatialized axis, the data lines emerge from the left, connect to their associated





Figure 9. Brushing on a re-spatialized axis.

spatial location by intersecting a sphere, and join back to the remainder of the plot on the right. This technique brings back to the visualization the spatial context of the data surrounding this time step.

The ASC was carefully designed to enhance its effectiveness. We use Catmull-Rom splines to perform the connection between axes and spheres as we found that the gradual nature of the curves improves the users’ ability to follow the lines. Each sphere is Phong shaded to provide visual cue as to its shape. The most apparent feature however is that the ASC performs an animated sinusoidal “wiggle” through its y-axis to take advantage of the motion parallax effect (structure-from-motion) to provide the user with a sense of depth. As well as the additional information that is brought by individual animation frames, the sense of depth provides a “visual decluttering” effect³⁹, allowing the user to locate visually each line’s trajectory and sphere intersection in 3D space. The speed and rotation degree of animation are user-customizable.

During development, our domain expert expressed difficulty in using the single X , Y , and Z axes to reason about 3D areas of the data. Being able to brush directly on the ASC is likely to be more intuitive for those unused to reasoning about data where the dimensions are separated. *Plume Explorer* fully supports brushing operations on the ASC, as shown in Figure 9. The user first hovers over the ASC with the mouse to “freeze” the axis. Clicking an individual voxel sphere toggles its selection in the brush state, and the brushed sphere gains a yellow halo to indicate that it has been selected. A selection can also be dragged over the spheres to select groups at a time. The ASC and all other axes dynamically update to reflect the brushing operations as they occur.

Realistic Spatial Context. The RSC approach provides the user with a realistic representation of the data context. By hovering over the  icon, the user interface provides a Matlab pre-rendered isosurfaces of the plume at that point in time. The pre-rendered isosurfaces are generated nightly from the data directory and are uploaded into the online tool. Currently, typical representations of the plume data use multiple isosurfaces or volume renderings^{37;38}. The main imagery provided in serving the data utilizes multiple isosurfaces combined with another surface representing the

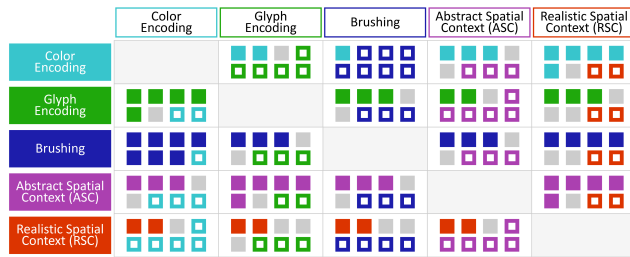


Figure 10. Eight scientists compared pairs of designs. Each color square represents one person preferring one design over another. A filled square corresponds to a row, and hollow to a column. Grey squares indicate *don't know* answers.

seafloor^{37:40}. Although only using a single isosurface, the RSC provides a readily-available means of enabling the user to access such a spatial context.

Expert Evaluation

The evaluation was carried out in a qualitative manner, structured around a questionnaire, a copy of which is included in the supplementary materials. Qualitative expert feedback⁴¹ was used in our evaluation as observing seafloor hydrothermal plumes requires extensive semantic knowledge of 4D hydrothermal plume data that arbitrary participants are unlikely to possess. The questionnaire consists of 11 main questions. Among them, ten focus on pairwise comparison of the five designs, and one on the integrated design as shown in Figure 2. The questions are then followed by a freeform discussion. The dataset that was used in the evaluation was from the COVIS project^{34:37}. The data was chosen as it possessed bending properties and contained gaps between the time-steps due to corrupted packet data.

We evaluated the five re-spatialization designs, together with an integrated design, by consulting nine domain experts (one of whom is a co-author of the paper). All of the participants have good knowledge about seafloor hydrothermal plumes, and appreciate the need for and difficulties in observing temporal volume data captured by observatories such as Neptune Canada. As scientists working with hydrothermal plume have diverse research interests and focuses, we did expect that their evaluation would be heavily influenced by personal knowledge and experience. All of the participants have prior data analysis experience, and are familiar with Matlab pre-rendered isosurface of the plume since this is their primary method of qualitative analysis at present. Each of the evaluation sessions lasted approximately an hour.

The evaluation began with a brief introduction. The introduction included an explanation on how the data is being represented at each time step, followed by a brief description of the five re-spatialization designs. After the

introduction, each of the participants was then asked to fill in a questionnaire consisting of 11 main questions. Questions 1 – 10 represent a pairwise comparison of the five re-spatialization designs. The pairwise comparison was then followed by an integrated design question where the participants were asked to interact with the integrated web-based application *Plume Explorer*.

Eight domain experts completed the questionnaire. The results of the ten questions for pairwise comparison are summarized in Figure 10. Their comments on the relative merits of each design are as shown in Figure 11. One domain expert preferred to provide feedback without using the questionnaire, and rated brushing and RSC positively.

While considering Questions 1–10, several domain experts suggested using different designs in an integrated manner. In answering Question 11, all experts agreed that integrated design was better than individual designs. One expert explained: “*The realistic one connects with my prior understanding of what is happening, the abstract and the colored lines help me tease out the temporal pattern and see what parts are changing and what parts are not.*”

ASC was considered good for conveying local spatial context, while color and glyph encoding could convey spatial changes in time. As one of the experts described: “*Having the spatial context (cube) and then selecting individual lines by clicking on the spheres was the easiest way for me to explore individual voxels and understand the spatial context. Being able to select the “Red, Green, Blue” colour scheme was also helpful - this made the lines stand out a lot better than the “underwater” colour scheme.*” One expert found that opening two ASCs at the same time was helpful. Experts also found that wiggling helped spatial perception. They suggested implementing similar capabilities for RSC. They also suggested conducting further research into the glyph design to explore its potential, introducing a legend window for lines, and showing how lines related to the isosurfaces.

Case Study II: Visualizing Temporal Phonemic Patterns in Poems

This case study was built on early work of poem visualization in support of close reading⁴². Poetry scholars who worked with Poem Viewer – a web-based poem visualization tool – have found that visualizing locations and degrees of sonic turbulence helps them begin to capture poetic qualities and characteristics (such as relative intensity and dynamism), which contribute to literary experience and inform critical interpretation. This glyph-based visualization has helped scholars observe where similar mouth actions generate distinct but related phonemic sounds; still, these scholars wish to have more

	A: Color Encoding	B: Glyph Encoding	C: Brushing	D: Abstract spatial context	E: Realistic spatial context
Spatial context	Spatial context is more obvious than <u>B</u> , <u>C</u> , <u>E</u> .	Spatial context is more obvious than <u>A</u> , <u>C</u> , <u>D</u> , <u>E</u> .	Spatial context is more obvious than <u>A</u> , <u>B</u> , <u>E</u> .	Spatial context is more obvious than <u>A</u> , <u>B</u> , <u>C</u> , <u>E</u> .	Spatial context is more obvious than <u>A</u> , <u>B</u> , <u>C</u> .
Learning & remembering		Easier to remember than <u>A</u> , <u>C</u> , <u>D</u> .	Easier to learn than <u>A</u> .	Easier to remember than <u>A</u> , <u>B</u> , <u>C</u> . Easier to learn than <u>C</u> , <u>E</u> .	
Other reasons / comments	Tracking connecting lines is difficult.	Glyphs are intuitive. Tracking connecting lines is difficult. Glyphs are confusing.	Conveys temporal change effectively. <u>C</u> alone is not of much use.	More cluttering than <u>A</u> , <u>C</u> . Need to learn to use <u>D</u> . Tracking connecting lines is difficult.	More information conveyed than <u>A</u> , <u>B</u> , <u>C</u> , <u>D</u> . Interactive on-demand is good. It gives a good qualitative feel. Connects with prior experience. Does not add much to the quantitative analysis.

Figure 11. The comments from the eight scientists on the relative merits of the five re-spatialization techniques for visualizing plume data. Those comparative comments agreed by all are underlined, otherwise in italic.

effective means to observe the presence and patterns of a poem’s unique ‘sound flow’.

For example, scholars wish to look not only at where and how specific sounds appear in a poem but also at whether phonemic positions in a given poem remain in a specific region of the mouth, evolve slowly from one region to another, or shift more often and abruptly. Since such patterns of sound flow are spatiotemporal, and the phonemic positions may change over time, the problem belongs to the class of *unstable spatial mapping* (see Section ‘Definitions and Requirements’). It is hence particularly challenging for scholars to build mental models for such patterns by observing a space-centric visualization (e.g., Figure 3(b)), or a time-centric visualization (e.g., Figure 3(c)).

Phoneticians have utilized a variety of spatial representations to study the differences between vocalic systems, ranging from tongue position illustrations to x-ray videos depicting muscular motion⁴³. In this work, we used the International Phonetic Alphabet (IPA)²², to cluster 30 vowel phonemes into 3×3 groups. As depicted by the boxes in Figure 3(a), the 9 vowel position groups are: (left to right) back, central, and front, and (from top to bottom) close, mid, and open. The frontness or backness dimension describes which part of the tongue is raised, while the openness and closedness dimension (commonly referred to as vowel height) describes how far the tongue is raised. The third attribute, vowel roundedness or lip rounding, describes the shape of the lips when a vowel is produced. In Figure 3(a) rounded vowels are indicated by a circle surrounding the corresponding phonetic characters. For studying sound flow in poems, meaningful insights can be derived in such a 2D representation. In this paper, we focus

on the visualization of vowel phonemes, though our re-spatialization methods are easily adaptable to consonants. With the increasing use of medical imaging modalities for recording the movements of speech organs, the re-spatialization methods in Case Study I in Section ‘Case Study I: Visualizing Events in Temporal Plume Data’ can potentially be deployed for such data.

Visual Designs

Figure 6 shows the application of five re-spatialization methods to a temporal parallel coordinates plot of a poem. The poem in the example is “April Rain Song” by Langston Hughes. Each line of the poem is shown as a path, from left to right, intersecting with vertical axes. The first vertical axis on the left indicates the different lines in a poem. The remainder of the vertical axes represent the roundedness attribute. Each of these axes is divided into two parts, rounded vowels at the upper half and unrounded at the lower half. Spatial relaxation is used in order to reduce cluttering and allow for spatial variations at junction points. To make the two roundedness categories more distinguishable, each axis is split into two sections. Although we also have spatiotemporal visual designs that map one spatial dimension (i.e., either frontness or vowel height) to the vertical axes, and use re-spatialization for the other, the five designs discussed below allow us to focus on evaluating re-spatialization methods for time-centric visual designs.

Color Encoding. Figure 6(a) shows a partial view of this re-spatialization method. Each of the nine vowel position groups is encoded with a different color (selected using ColorBrewer⁴⁴). As a quick aid to learning and memorization, a small legend is displayed at the top of

all vertical axes except the first axis for line number. A possible advantage of this design may be its simplicity. One possible disadvantage may be the difficulty in learning and remembering the mapping between spatial positions and colors.

Glyph Encoding. In this design, each of the nine vowel position groups is represented by a glyph, as shown in Figure 6(b). The glyphs utilize two visual channels to encode the different positions. The vowel height is mapped to a color channel, i.e., red for close, gray for mid, and blue for open. The frontness is mapped to a combined shape-orientation channel, i.e., back is shown as a semicircle on the left, central a circle on its own, and front a semicircle on the right. Similar to the color encoding, a small legend is shown on top of each axis. To make the frontness dimension more observable, we shifted the glyph for a back position slightly leftward and that for a front position slightly rightward. The possible advantages of this design may include the ease to learn and remember, to separate the two dimensions, and to distinguish the three values in each dimension.

Selection-based Exploration. In this design, users can select the different vowel positions that they wish to view using a set of 9 tick-boxes. The positions that have been selected are encoded using dark gray and those unselected with light gray. Figure 6(c) shows an example of using this design where only the closed back and open front vowel positions have been selected. By interacting with the different vowel positions, users may be able to build a mental model of the spatial context and explore the different spatial relationships between different lines.

Abstract Spatial Context. In this design, each vertical axis (except the first one) can be expanded to display the abstract representation of the nine position groups. As illustrated in Figure 6(d) each path is connected to the corresponding vowel position in the expanded illustration, where lines are also relaxed to maintain traceability. The possible advantage of this design may be due to its on-demand nature as a user could easily examine where concentrations of phonemic positions occur in a poem.

Realistic Spatial Context. In this design, each vertical axis (except the first one) can be expanded to display a more detailed representation of the tongue geometries for all the vowel sounds on the axis. The illustrations depict cardinal tongue positions, and are drawn based on x-ray images in the literature⁴³. The different curves, representing the top of the tongue, also implicitly suggest how different sounds travel above these curves through the mouth. The acoustic structures of the vowels are determined by the position of the tongue in the mouth⁴³. We used line thickness to indicate the frequencies of the vowel positions, i.e., the thicker the line is the more times the vowel positions are repeated at the corresponding temporal step. A




Figure 12. Five humanities scholars compared pairs of designs. Each color square represents one person preferring one design over another. A filled square corresponds to a row, and hollow to a column. Grey squares indicate *don't know* answers.

possible advantage with this design is that it provides extra information that is not available in the ASC.

Expert Evaluation

Similar to the study in Section ‘Expert Evaluation’ above, this evaluation was also carried out in a qualitative manner, structured around a questionnaire featuring 11 main questions and freeform discussions (see the supplementary materials). Among the questions, ten of the questions were focused on pairwise comparisons and one on an implementation where four designs are integrated together (excluding color encoding as it is mutually exclusive with glyph encoding in this case).

Observing sound flow along different lines in poems is a highly specialized visualization task. We thus conducted our evaluation of the five re-spatialization designs, together with an integrated design, by consulting seven humanities scholars (two of whom are co-authors). All have studied poetry, and have experience in close reading. The poem that was used in the evaluation was “A Drinking Song” by W. B. Yeats. Each of the evaluation sessions lasted approximately an hour.

The evaluation session began with a brief introduction. The introduction included an explanation of the IPA vowel chart that was used in the study and its representation, followed by a brief description of the five re-spatialization designs. After the introduction, the participants were then asked to fill in the questionnaire consisting of 11 main questions. Similar to the questionnaire in Section ‘Expert Evaluation’ above, Questions 1–10 are designed as a pairwise comparison of the five re-spatialization designs and Question 11 is an integrated design question where the participants were asked to explore the integrated design using an implemented web-based prototype.

Five domain experts completed the questionnaire. The results of the 10 questions for pairwise comparison are summarized in Figure 12. Their comments on the relative merits of each design are as shown in Figure 13. Two domain experts (the 6th and 7th) gave feedback without using the questionnaire. One indicated a preference towards

	A: Color Encoding	B: Glyph Encoding	C: Brushing	D: Abstract spatial context	E: Realistic spatial context
Spatial context	Spatial context is more obvious than <u>B</u> , <u>C</u> , <u>D</u> , <u>E</u> .	Spatial context is more obvious than <u>C</u> , <u>D</u> , <u>E</u> .	Spatial context is more obvious than <u>A</u> , <u>E</u> .	Spatial context is more obvious than <u>C</u> , <u>E</u> .	
Learning & remembering		Easier to learn and remember than <u>A</u> .			
Other reasons / comments	Easier to discern patterns.	Better for overall patterns.	More flexible and adaptable for exploration.	Closer to the observation point.	Easier to relate to the actual structure of the mouth.
	Difficult to remember colors.		Easier to remember.	Available on demand.	The picture is too large.
			The interface is unfriendly.	Difficult to track connection lines.	Difficult to relate it to the words in the poem.

Figure 13. The comments from the five humanities scholars on the relative merits of the five re-spatialization techniques in poem visualization. Those comparative comments agreed by all are underlined, otherwise in italic.

color and glyph encoding, and another preferred glyph encoding.

While considering Questions 1–10, several domain experts suggested different ways of combining visual designs. In answering Question 11, all experts agreed that the integrated design was better than individual designs, and liked the flexibility in choosing different designs for different tasks. One suggested that brushing should always be available. Another expert thought the integrated design would provide both an orienting overall view of the poem (through glyph encoding) and the flexibility to focus explorations (via brushing). A number of suggestions for improving the visual designs were made, ranging from using a color wheel metaphor, to changing the contrast of lines during brushing.

Discussions and Conclusions

In this work, we studied a challenging problem in spatiotemporal data visualization of how to reintroduce spatial context back into time-centric designs. We examined five design options in conjunction with two application case studies. Based on the evaluation by domain experts who understand the requirements of visualization tasks, we can make the following observations:

- The integrated design is much better than individual designs. It provides different levels of support to different tasks, e.g., discovering the local spatial context at t vs. observing spatial changes over a period $[t_1, t_2]$. We did not observe any user difficulties in using more than one re-spatialization mechanism.
- Domain experts suggested that providing further information to accompany the Realistic Spatial Context (RSC) may be required as it would allow them to relate to the data using their existing

knowledge. However, doing so may not necessarily provide support for quantitative analysis.

- Domain experts are able to relate more to the Abstract Spatial Context (ASC) in 2D than in 3D. The possible reasons may include that 3D spatial perception is more difficult, and glyph encodings are less effective in 3D. The wiggle motion, which received positive feedback, can help to improve 3D perception.
- Brushing (selection-based exploration) is considered important by both sets of domain experts, though the humanities scholars appreciated it more. This may be due in part to poetry scholars' greater desire for flexibility in exploration – a desire facilitated by the comparably smaller data size of poems, which allows for more time to interact with the visualizations. In comparison with RSC, the humanities scholars felt that the spatial context is more clearly related to the poetic context.
- Glyph encoding is much appreciated in 2D but not at all in 3D. When combined with brushing, it could potentially offer a good practical solution to many 2D applications. For example in Case Study II, Glyph encoding was able to convey a sense of gradation and relationship to nearby spatial positions, i.e., horizontally, vertically, and diagonally.
- Color encoding received fairly positive reviews in 2D in terms of spatial context, but was scored lower than glyph encoding due to the difficulties in learning and memorization.

In addition to the above observations the domain experts also found that for the plume study, the ASC provided the best quantitative analysis support since it helped them to recall the spatial context with more ease. They however found that RSC provided the best qualitative analysis support since it provided a level of familiarity with their existing knowledge of isosurface renderings of the data.

For the poem data, the humanities scholars found that both glyphs and color encoding helped them to make additional observations as it enabled them to relate to the poem's phonetic positions spatially. Using RSC on the other hand, they found it difficult to relate to the individual words in the poems. For both the plume and the poem data, brushing was appreciated as it allowed for filtering operations that are essential for analysis.

Like most expert-based evaluations in application case studies, these observations may be specific to the applications, tasks and users concerned. Nevertheless, they do offer a useful template for visualization designers to map out various design options and to anticipate users' requirements and feedback.

This work demonstrates that the re-spatialization of time series plots is feasible under the condition that the spatial context is of a relatively low resolution. The scalability with regards to the number of time steps is additionally a factor in its usability, which needs to weigh the analytical and interactive functions afforded by our plot against the simplicity of a basic time series plot. As we can see from the two application case studies, low resolution is often desirable either because high-resolution spatial context is too complex and noisy to comprehend (e.g., Section 'Case Study I: Visualizing Events in Temporal Plume Data'), or because data for high-resolution spatial contexts are not available (e.g., Section 'Case Study II: Visualizing Temporal Phonemic Patterns in Poems'). Of course, one should not generalize the findings of this work to that in arbitrarily high resolutions or large datasets, which is a much more difficult challenge. In future work, we wish to explore the design spaces for the other two design strategies mentioned in Section 'Introduction' as well as their suitability at various data sizes.

Acknowledgements

The authors wish to thank all those who participated in the evaluation. They are humanities scholars from University of Oxford: Iain Emsley, Alexander Huber, Laura Ludtke, David Robey, and Natasha Ryan; and COVIS data experts: Ben Biffard, Maia Hoeberechts, Darrell Jackson, Reyna Jenkins, Mike Kenney, Russ Light, Li Liu, Sedat Ozer, Jay Takle, and Guangyu Xu.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

1. Havre S, Whitney P and Nowell L. ThemeRiver: Visualizing thematic changes in large document collections. *IEEE Trans Vis Comput Graph* 2002; 8(1): 9–20.
2. Wattenberg M. Baby names, visualization, and social data analysis. In *Proc. IEEE Symp. Information Vis. 2005*. pp. 1–7.
3. Legg P, Chung D, Parry M et al. MatchPad: Interactive glyph-based visualization for real-time sports performance analysis. *Computer Graphics Forum* 2012; 31(3): 1255–1264.
4. Boyandin I, Bertini E, Bak P et al. Flowstrates: An approach for visual exploration of temporal origin-destination data. *Computer Graphics Forum* 2011; 30(3): 971–980.
5. Lu A and Shen HW. Interactive storyboard for overall time-varying data visualization. In *Proc. IEEE Pacific Visualization 2008*. pp. 143–150.
6. Lu A, Chen W, Ribarsky W et al. Year-long time-varying 3D air quality data visualization. In *Advances in Information and Intelligent Systems 2009*. pp. 289–306.
7. Joshi A and Rheingans P. Illustration-inspired techniques for visualizing time-varying data. In *Proc. IEEE Visualization 2005*. pp. 679–686.
8. Duffy B, Thiyagalingam J, Walton S et al. Glyph-based video visualization for semen analysis. *IEEE Trans Vis Comput Graph* 2015; 21(8): 980–993.
9. Bak P, Mansmann F, Janetzko H et al. Spatiotemporal analysis of sensor logs using growth ring maps. *IEEE Trans Vis Comput Graph* 2009; 15(6): 913–920.
10. Andrienko N and Andrienko G. Interactive visual tools to explore spatio-temporal variation. In *Proc. Working Conference on Advanced Visual Interfaces 2004*. pp. 417–420.
11. Walton S, Berger K, Ebert D et al. Vehicle object retargeting from dynamic traffic videos for real-time visualisation. *The Visual Computer* 2014; 30(5): 493–505.
12. Hägerstrand T. What about people in regional science? *Papers in Regional Science* 1970; 24(1): 6–12.
13. Kapler T and Wright W. GeoTime information visualization. In *Proc. IEEE Symp. Information Visualization 2004*. pp. 25–32.
14. Forlines C and Wittenburg K. Wakame: Sense making of multi-dimensional spatial-temporal data. In *Proc. Int. Conf. Advanced Visual Interfaces 2010*. pp. 33–40.
15. Daniel G and Chen M. Video visualization. In *Proc. IEEE Visualization 2003*. pp. 409–416.
16. Botchen RP, Bachthaler S, Schick F et al. Action-based multifield video visualization. *IEEE Trans Vis Comput Graph* 2008; 14(4): 885–899.
17. Drocourt Y, Borgo R, Scharrer K et al. Temporal visualization of boundary-based geo-information using radial projection. *Computer Graphics Forum* 2011; 30(3): 981–990.
18. Roberts JC. State of the art: Coordinated & multiple views in exploratory visualization. In *Proc. Fifth Int. Conf. Coordinated and Multiple Views in Exploratory Visualization 2007*. pp. 61–71.

19. Wang Baldonado MQ, Woodruff A and Kuchinsky A. Guidelines for using multiple views in information visualization. In *Proc. Working Conference on Advanced Visual Interfaces 2000*. pp. 110–119.
20. Fisher D. *Animation for Visualization: Opportunities and Drawbacks*, chapter 19. Beautiful Visualization, O'Reilly Media, 2010. pp. 329 – 352.
21. Robertson G, Fernandez R, Fisher D et al. Effectiveness of animation in trend visualization. *IEEE Trans Vis Comput Graph* 2008; 14(6): 1325–1332.
22. International Phonetic Association. *Handbook of the International Phonetic Association*. Cambridge University Press, 1999.
23. Andrienko N, Andrienko G and Gatalsky P. Exploratory spatio-temporal visualization: an analytical review. *J Visual Languages & Computing* 2003; 14: 503–541.
24. Aigner W, Miksch S, Muller W et al. Visual methods for analyzing time-oriented data. *IEEE Trans Vis Comput Graph* 2008; 14(1): 47–60.
25. Andrienko G, Andrienko N, Demsar U et al. Space, time and visual analytics. *Int J Geographical Information Science* 2010; 24(10): 1577–1600.
26. Zhong C, Wang T, Zeng W et al. Spatiotemporal visualisation: A survey and outlook. In Arisona SM, Aschwanden G, Halatsch J et al. (eds.) *Digital Urban Modeling and Simulation 2012*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 299–317.
27. Andrienko N and Andrienko G. Visual analytics of movement: An overview of methods, tools and procedures. *Information Visualization* 2013; 12(1): 3–24.
28. Bach B, Dragicevic P, Archambault D et al. A review of temporal data visualizations based on space-time cube operations. In *EuroVis State of the Art Report 2014*. pp. 23–41.
29. Chen M, Walton S, Berger K et al. Visual multiplexing. *Computer Graphics Forum* 2014; 33(3): 241–250.
30. Tufte ER. *The Visual Display of Quantitative Information*. 2nd ed. Graphics Press USA, 2001.
31. Höllt T, Magdy A, Zhan P et al. Ovis: A framework for visual analysis of ocean forecast ensembles. *IEEE Trans Vis Comput Graph* 2014; 20(8): 1114–1126.
32. Yuan X, Guo P, Xiao H et al. Scattering points in parallel coordinates. *IEEE Trans Vis Comput Graph* 2009; 15(6): 1001–1008.
33. Rona PA, Bemis KG, Jones CD et al. Entrainment and bending in a major hydrothermal plume, Main Endeavour Field, Juan de Fuca Ridge. *Geophysical Research Letters* 2006; 33(19): 971–980.
34. Di Iorio D, Lavelle JW, Rona PA et al. Measurements and models of heat flux and plumes from hydrothermal discharges near the deep seafloor. *Oceanography* 2012; 25(1): 168–179.
35. Xu G, Jackson DR, Bemis KG et al. Observations of the volume flux of a seafloor hydrothermal plume using an acoustic imaging sonar. *Geochemistry, Geophysics, Geosystems* 2013; 14(7): 2369–2382.
36. Ocean Networks Canada. Data Archive. University of Victoria, Canada. Downloaded in Mar. 2014.
37. Rona P and Light R. Sonar images hydrothermal vents in seafloor observatory. *Eos, Transactions American Geophysical Union* 2011; 92(20): 169–170.
38. Santilli K, Bemis K, Silver D et al. Generating realistic images from hydrothermal plume data. In *IEEE Visualization 2004*. pp. 91–98.
39. Holten D and van Wijk JJ. Evaluation of Cluster Identification Performance for Different PCP Variants. *Computer Graphics Forum* 2010; 29(3): 793–802.
40. Ocean Networks Canada. Data & Tools. www.oceannetworks.ca/data-tools. (Accessed in Sept. 2014).
41. Isenberg T, Isenberg P, Chen J et al. A systematic review on the practice of evaluating visualization. *IEEE Trans Vis Comput Graph* 2013; 19(12): 2818–2827.
42. Abdul-Rahman A, Lein J, Coles K et al. Rule-based visual mappings – with a case study on poetry visualization. *Computer Graphics Forum* 2013; 32(3): 381–390.
43. Ladefoged P and Disner SF. *Vowels and Consonants*. 3rd ed. Wiley-Blackwell, 2012.
44. Brewer CA, Hatchard GW and Harrower MA. ColorBrewer in print: A catalog of color schemes for maps. *Cartography and Geographic Information Science* 2003; 30(1): 5–32.