



Innovations in Visualization
<http://innovis.cpsc.ucalgary.ca/>



iLab
Interactions Lab
<http://ilab.cpsc.ucalgary.ca>

Department of Computer Science
University of Calgary

To Cite this paper:

Bach, Benjamin, Dragicevic, Pierre, Archambault, Daniel, Hurter, Christophe and Carpendale, Sheelagh.

A Descriptive Framework for Temporal Data Visualizations Based on Generalized Space-Time Cubes.

In *Computer Graphics Forum*, 36(6) Pages 36-61. Wiley Online Library.

PUBLICATION DATE
2017

DOI
10.1111/cgf.12804

This is the authors' version of this work



A Descriptive Framework for Temporal Data Visualizations Based on Generalized Space-Time Cubes

B. Bach¹, P. Dragicevic², D. Archambault³, C. Hurter⁴ and S. Carpendale⁵

¹Microsoft Research—Inria Joint Centre, Saclay, France

benjamin.bach@inria.fr

²Inria, Saclay, France

pierre.dragice@gmail.com

³Swansea University, Swansea, UK

d.w.archambault@swansea.ac.uk

⁴ENAC, Toulouse, France

christophe.hurter@enac.fr

⁵University of Calgary, Calgary, AB, Canada

sheelagh@cpsc.ucalgary.ca

Abstract

We present the generalized space-time cube, a descriptive model for visualizations of temporal data. Visualizations are described as operations on the cube, which transform the cube's 3D shape into readable 2D visualizations. Operations include extracting subparts of the cube, flattening it across space or time or transforming the cubes geometry and content. We introduce a taxonomy of elementary space-time cube operations and explain how these operations can be combined and parameterized. The generalized space-time cube has two properties: (1) it is purely conceptual without the need to be implemented, and (2) it applies to all datasets that can be represented in two dimensions plus time (e.g. geo-spatial, videos, networks, multivariate data). The proper choice of space-time cube operations depends on many factors, for example, density or sparsity of a cube. Hence, we propose a characterization of structures within space-time cubes, which allows us to discuss strengths and limitations of operations. We finally review interactive systems that support multiple operations, allowing a user to customize his view on the data. With this framework, we hope to facilitate the description, criticism and comparison of temporal data visualizations, as well as encourage the exploration of new techniques and systems. This paper is an extension of Bach et al.'s (2014) work.

Keywords: space-time Cubes, temporal data, information visualization, video visualization, network visualization

ACM CCS: H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

1. Introduction

Temporal datasets are ubiquitous but notoriously hard to visualize, especially rich datasets that involve more than one dimension in addition to time, such as videos, dynamic networks, multivariate and geo-spatial data.

Previous work on novel visualizations for temporal data has dramatically advanced the field of information visualization (infovis). However, there are so many different techniques today that it has become hard for both researchers and designers to get a clear picture of (i) what has been done, (ii) how to compare and chose visualizations and (iii) how much of the design space of temporal data

visualizations remains to be explored. For similar reasons, teaching this research topic to students is challenging. Therefore, there is a clear need to structure and organize previous work in the area of temporal data visualization.

Part of the problem is that infovis researchers have mostly focused on *nomenclature*. Most familiar charts have an agreed-upon name, for example, *small multiples* or *scatter plots*, and this tradition has been continued in infovis, where each newly published visualization technique is given a different name. Many textbooks and surveys list existing techniques by their name, both for general visualizations [Har99] and for temporal visualizations [AMST11].

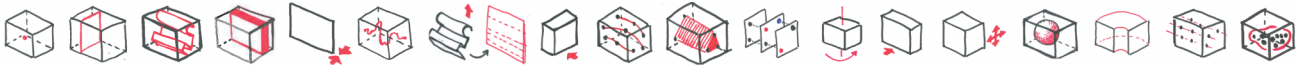


Figure 1: Examples of space-time cube operations, performed on a generalized space-time cube to describe visualizations.

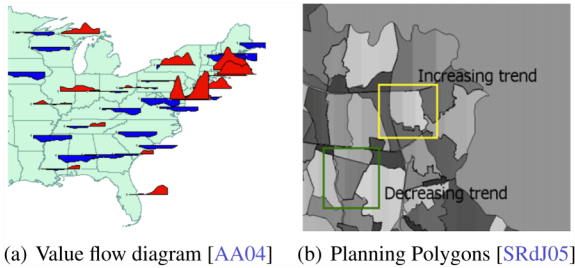


Figure 2: Two conceptually similar temporal visualization techniques showing: (a) the evolution of crime statistics per state; (b) the evolution of high school population per district.

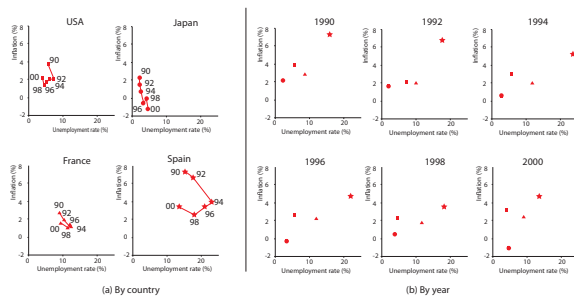


Figure 3: Two visualizations using small multiples to show the same indicator data for four countries over 6 years, but which are conceptually different.

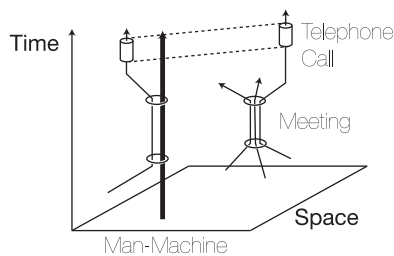


Figure 4: A space-time cube based on an illustration by Hägerstrand [H70] in 1970, showing social interactions across space and time.

While names are essential for indexing, retrieval and communication purposes, they are a poor thinking tool. Because there is no convention for naming techniques, names rarely reflect the essential concepts behind a technique. For example, names such as *Value Flow Maps* [AA04] and *Planning Polygons* [SRdJ05] say little about the possible conceptual similarities between the two techniques (see Figure 2). Names can also be ambiguous. For example,

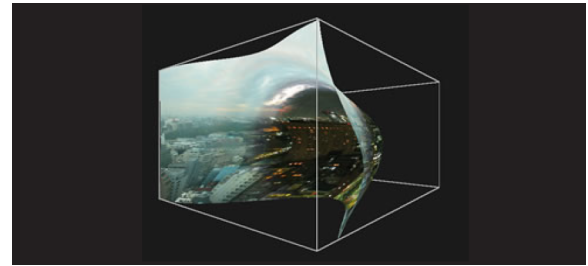


Figure 5: Khronos projector [CI05] lets users dig into video cubes: here, ascene transitioning from day to night.

the term *small multiples* is commonly used to refer to a specific type of temporal data visualization [Tuf86]. But Figure 3 shows that two visualizations can be based on small multiples despite being different with respect to which aspect of the data they show.

Recent effort at proposing taxonomies, conceptual models and design spaces for temporal visualizations mainly categorize visualizations based on analytical tasks or data types (e.g. object movement data [AAH11, AAB*11, AA12], video data [BCD*12] or datasets with different temporal and spatial structures [AMM*07]). While descriptions of visualization techniques by tasks and data types can help choosing a visualization technique for a specific problem, we aim to describe visualizations in their own visualization space, independent from data and task. We contend that it is limitative to reason in a task- and data-agnostic manner, and that this is uncommon in visualization (which emphasizes point designs, point studies and cookbook-type recommendations), but because it is rare enough, it is worth trying.

Our goal is to understand (1) how those visualizations relate conceptually, (2) to provide a clear terminology and (3) a way of organizing these visualizations that can serve as a basis for prescriptive and evaluative work. In other words, our framework is mostly *descriptive*, but, as we show, it can help better understand the strengths and weaknesses of different visualizations with respect to tasks and data, and better formulate research questions in empirical studies.

The benefit of a clear and detailed descriptive framework is that it helps connect techniques that are similar and distinguish techniques that are dissimilar. For example, the two techniques from Figure 2 have different names but are the result of a similar (compound) operation on a *generalized* space-time cube and which we call *repeated drilling*. Figure 3 shows two techniques with the same name ('small multiples') but resulting from very different space-time cube operations; Figure 3(a) involves operations such as *filtering*, *time flattening* and *space shifting*, while Figure 3(b) is the result of a compound operation we call *time juxtaposing*. By discussing such relations, we hope to improve our understanding of which

visualizations are complementary and which ones are possible, but have not yet been realized in visualization. Eventually, we discuss how multiple operations can be combined into an interactive and consistent interface. Our paper is targeted towards designers and researchers, in particular students in both domains.

The term *space-time cube* originates from social geography, where it refers to a three-dimensional (3D) Euclidean space, consisting of a 2D geographical space plus time. Under the term ‘time geography’, Torsten Hägerstrand in 1970 described a ‘*space-time model which could help us to develop a kind of socio economic web model*’ [H70], p. 10]. His intention was to analyse people’s behaviour and interactions across space and time (Figure 4); a moving person on a 2D map becomes a static 3D trajectory visualized in a space-time cube. Since then, the space-time cube has been used as a visualization metaphor in a number of interactive visualization systems in geo-visualization [CCT*99, Kra03], video visualization [FLM00], networks [BPF14b], and art [CI05]. However, it has never been used as a conceptual model for reflecting on visualizing temporal data in general.

With respect to the traditional understanding of a space-time cube—a 3D representation of geographic space plus time—here, we define a *generalized* space-time cube, inspired by Furnas’ notion of generalized fisheye views [Fur86]. The generalized space-time cube has the following two main properties:

(1) A generalized space-time cube is a conceptual representation that helps to think about temporal data visualization techniques in general. It does *not* imply an explicit 3D visualization or a perspective projection thereof. The cube does *not* need to be implemented in the system, used to generate this visualization. For example, the visualizations in Figures 2 and 3 do not show any space-time cube; they are purely 2D visualizations. For most of the examples in this paper, we do not know the exact format of the underlying data (1D, 3D, tables, multivariate, network). However, we imagine a *possible* conceptual cube which the visualization can be derived from. This does not imply how the visualization is implemented or whether the designer was aware of any conceptual cube and generalized space-time cube operations. In Section 6, we discuss systems that explicitly implement the space-time cube and its operations.

(2) A generalized space-time cube does not need to involve geo-spatial data. Many visualizations (e.g. scatterplots or node-link diagrams) convey abstract, non-spatial data. Those (non-temporal) visualizations all embed data into an abstract and 2D *visualization space*. When data change over time, such as in GapMinder’s animated 2D scatterplots [Ros06], each animation frame can be conceptually thought of as a slice of a space-time cube. In the model of the generalized space-time cube, ‘space’ therefore refers to an abstract 2D substrate that is used to visualize (non-temporal) data at a specific time. We call this space *base plane*. Thus, generalized space-time cubes include traditional geo-spatial space-time cubes, video cubes, matrix cubes, as well as a subset of Online Analytical Processing (OLAP) cubes as we discuss later.

It is important to stress that *this paper is not about space-time cube visualizations*, and that 3D space-time cube representations, like the one in Figure 4, represent only a very small subset of the

visualizations we aim to cover. In the remainder of this paper, the term *space-time cube* refers to our model, if not indicated otherwise.

Finally, our framework does not consider how space-time cubes are built for a given dataset. A general heuristic is to decide on a 2D visual representation of the data for a given time slice, and then extrude it over time. For example, networks can be shown as node-link diagrams or matrices, multivariate data as scatterplots or barcharts. Such decisions are independent from the cube. We admittedly do not offer a detailed framework that clarifies the different ways concrete datasets can be turned into concrete space time cubes. Instead, we assume that a conceptual 3D space-time cube is already given and focus on how this cube can be transformed to accommodate 2D media like computer displays and paper while remaining legible. We show how such operations are enough to capture most known techniques for visualizing temporal datasets. We mainly consider datasets that involve two dimensions plus time, although we later discuss how our model can be extended to higher dimensionalities. Datasets with a single data dimension plus time may *not* require a space-time cube and can be visualized with timelines and visualizations for time series.

In this paper, we first review common temporal data visualization techniques, and explain how they can be all seen as operations on a generalized space-time cube. We then describe our framework in more detail by providing definitions of key concepts, as well as a taxonomy of elementary space-time cube operations and how they can be combined into compound operations. Then, we explain how choosing effective space-time cube operations can be motivated by the characteristics of the dataset, mostly because different datasets yield different ‘shapes’ or *inner structures* for space-time cubes. We provide a way to characterize such inner structures, based on characteristics of data. We then informally discuss the relative benefits and drawbacks of different space-time cube operations and briefly review empirical evidence from user studies. We finally discuss interactive visual exploration systems that support a range of space-time cube operations, discuss limitations of our framework and suggest avenues for future work. A catalogue of temporal visualizations classified according to our framework can be found at www.spacetimecubevis.com.

2. Static Visualizations as Space-Time Cube Operations

In this section, we illustrate examples of generalized space-time cube operations and how they can be used to describe a range of common *static* visualization techniques for temporal data, all meant for screen or paper media. We focus on a small but representative selection and describe operations informally, often using analogies from photography and art.

While a generalized space-time cube may not be the only way to describe these visualizations, and the designer may not have thought of it in this way, we construct conceptual space-time cubes for our examples, then apply the respective operation(s) that describe the actual visualization. Note that, while in Hägerstrand’s original illustration the time axis is vertical, in our illustrations time passes from left to right.

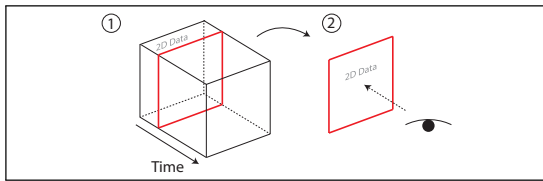


Figure 6: The time cutting operation.

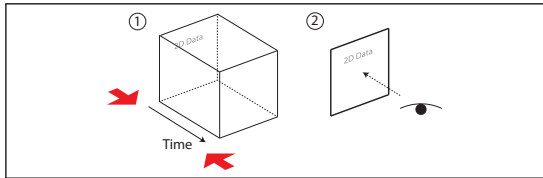


Figure 7: The time flattening operation.



Figure 8: A famous example of time flattening: Napoleon's march to Moscow by Joseph Minard [Tuf86].

2.1. Time cutting

A *time cutting* operation consists in extracting a particular temporal snapshot from the cube to be presented to the viewer. Figure 6 illustrates this operation: the left part (1) shows the initial space-time cube and the temporal snapshot that is being extracted, while the right part (2) shows the resulting image that is presented to the viewer.

For example, consider a photographer who captures a particular instant of a moving scene. If the scene being viewed is represented as a space-time cube (i.e. all possible pictures are piled up to form a cube), then taking a photograph is equivalent to applying a time cutting operation on this cube.

In infovis, an image produced by time cutting is typically called a *time slice*. But a temporal visualization rarely consists in a single time slice. As we will see in Section 3, time cutting is typically either performed multiple times and used in combination with other operations, or it is used in combination with animation and interaction.

2.2. Time flattening

Time flattening collapses the space-time cube along its time axis, by merging all time slices into a single 2D image (Figure 7). An analogy is *long exposure photography*, which collapses several seconds, minutes or even hours of a natural scene into a single image.



Figure 9: Other examples of time flattening: (a) Detail of the map of the cholera outbreak in London 1854, by Dr. John Snow. Piled bars mark the number of death per house. (b) Connected scatterplot showing the relationship between inflation rate and unemployment rate in Spain from 1990 to 2000.

One of the earliest uses of time flattening is Minard's illustration of Napoleon's march towards Moscow (Figure 8). The illustration shows on a single image the state of Napoleon's army (position, size, key events) at different points in time during the Russian campaign in 1812 [Tuf86]. Another early example is Dr. John Snow's map showing where deaths from cholera occurred in London in 1854 (Figure 9(a)). The map shows events from several days aggregated over time.

Many maps that show temporal data can be seen as time-flattened geo-spatial space-time cubes. But the time flattening technique is not limited to geographical data and has been employed in a large variety of infovis systems as well as in static data graphics. Figure 9(b) for example, shows the evolution of inflation and unemployment rate in Spain from 1990 to 2000. This diagram can be seen as time-flattened version of a space-time cube created from a scatterplot that contains a single data point evolving over time (*Spain*).

2.3. Discrete time flattening

Discrete time flattening is similar to time flattening, but instead of merging all time slices, a selection of meaningful time slices is made before combining them (Figure 10).

An analogy for discrete time flattening is *multiple exposure photography*, where several photos are taken at different times and blended

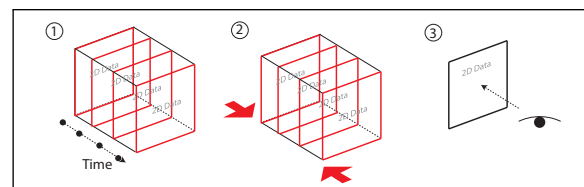


Figure 10: The discrete time flattening operation.

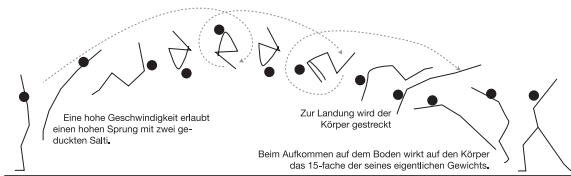


Figure 11: An example of discrete time flattening. For a better infographic by Megan Jaegerman, see [Tuf].

into a single image. Etienne-Jules Marey pioneered this technique in 1882 with an instrument (the chronophotographic gun) that records 12 photos per second on the same film. He used it to visualize human and animal motion [Mar78]. Modern art has also employed a similar technique to convey movement, for example, Marcel Duchamp's 'Nude Descending a Staircase, No. 2'.

Tufte [Tuf86] comments on several examples of infographics that employ discrete time flattening. He calls them *sequences*. One of his famous examples is the life cycle of the Japanese beetle [Tuf86]. Figure 11 is a sequence showing a dancer's move. Discrete time flattening has also been used for summarizing videos [BDH04].

2.4. Coloured time flattening

The *coloured time flattening* operation is similar to the time flattening operation, but time slices are assigned a colour before being combined (Figure 12). Although this operation does not map to

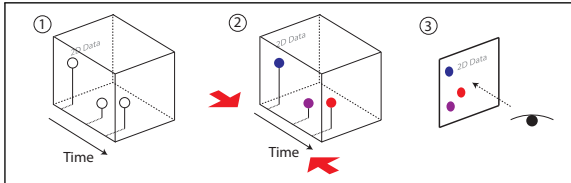


Figure 12: The coloured time flattening operation.

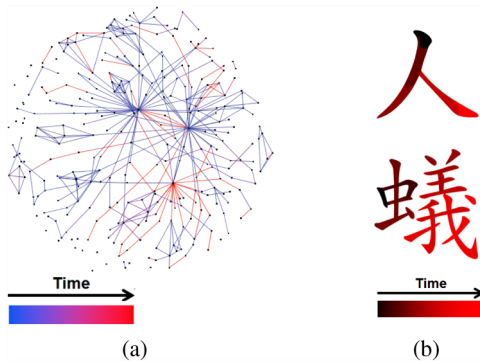


Figure 13: Two visualizations using coloured time flattening. (a) Illustration of a dynamic graph visualization as used in Gevol [CKN03]. (b) Stroke order in Chinese characters [Wik13]; the colour legends have been added.

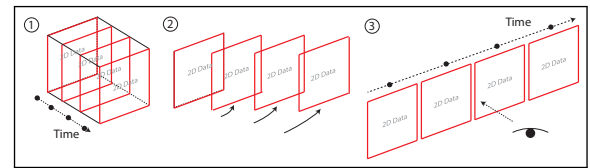


Figure 14: The time juxtaposing operation.

any photography technique we know of, similar results could in principle be obtained by rapidly switching colour filters during a long-exposure photography.

Two examples of visualizations obtained by coloured time flattening are shown in Figure 13: (a) a dynamic graph where old links (in red) are distinguished from new links (in blue) [CKN*03]; (b) Chinese characters where first strokes (in black) are distinguished from later strokes (in red) [Wik13]. Minard's map (Figure 8) also makes use of a simplified form of coloured time flattening, since the army's forward march and return are distinguished using two different colours.

2.5. Time juxtaposing

Time juxtaposing consists in extracting multiple time slices and placing them side-by-side or on a grid (Figure 14).

An analogy is Eadweard Muybridge's *multiple camera chronophotography* [Muy87]. In contrast with Marey, Muybridge used multiple cameras that recorded snapshots on *different* locations on the film. He used it for the scientific study of, for example, horse gaits, and his pictures famously settled the question as to whether horses have all four feet off the ground while trotting. Time juxtaposing is also the base for many forms of sequential art, from ancient Egyptian murals and Greek vase paintings to today's comics [McC94, BKH*16].

Time juxtaposing is often used in infovis to show temporal data such as time-evolving maps, trajectories in space [TBC13] and dynamic graphs [LNS11, BBL12, RM13, BPF14a, BHRD*15, BKH*16]. Figure 15 shows forest harvest data over 11 years. In infovis, time juxtaposing is usually referred to as *small multiples* [CKN*03], although small multiples are not necessarily built from time slices (see Figure 3(a)). Time juxtaposing has been also widely used for video summarization [TV07].

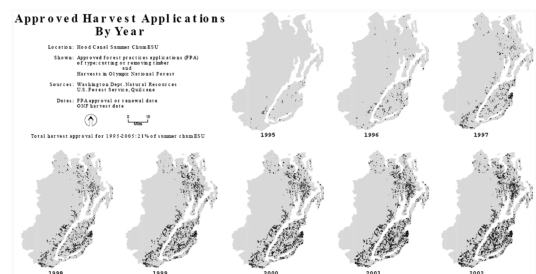


Figure 15: Time juxtaposing approved forest harvest applications across 10 years [Gre11].

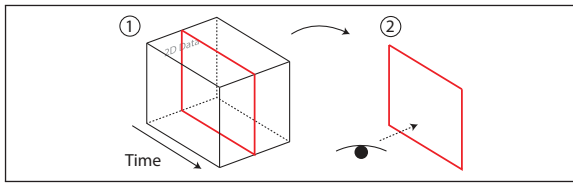


Figure 16: The space cutting operation.

2.6. Space cutting

Space cutting consists in extracting a planar cut in a direction perpendicular to the data plane (Figure 16). An analogy is *slit-scan photography*, a process where a plate into which a slit has been cut is inserted in front of a camera and then moved while the film is being exposed [TGF08]. Slit-scan photography has been used to create special effects in movies, artwork and photo finishes in sports.

Space cutting has also been employed for visualizing temporal data. In the 19th century, Marey created a visualization of train connections between major French cities (Figure 17). This visualization can be described as space cutting on a geographical space-time cube, along the tracks connecting cities while diagonal lines indicate positions of trains at any time [Tuf86, Mar78].

More recently, space cutting was shown to be useful for analysing video logs [TGF08]: Figure 18 shows a space cut (called ‘tear’ in the original work) extracted from a video space-time cube, and revealing traffic activity (car count, speed and direction) on a road. The time slice at t_1 is shown to the left, together with the position of the segment extracted. The system is also able to show multiple longitudinal slices on top of each other (i.e. *space juxtaposing*).

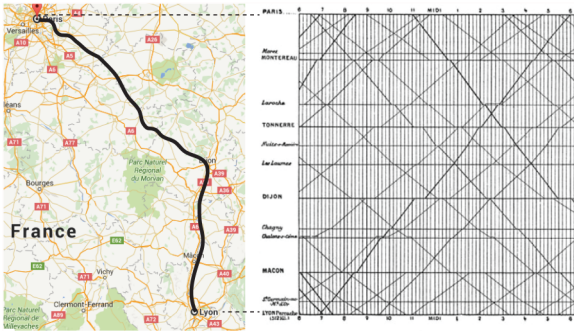


Figure 17: Example of space cutting (right): horizontal lines indicate train stops, vertical lines indicate times, and diagonal lines indicate moving trains [Mar78].



Figure 18: Space cutting used to show road traffic [TGF08].

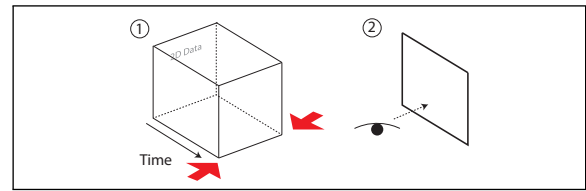


Figure 19: The space flattening operation.

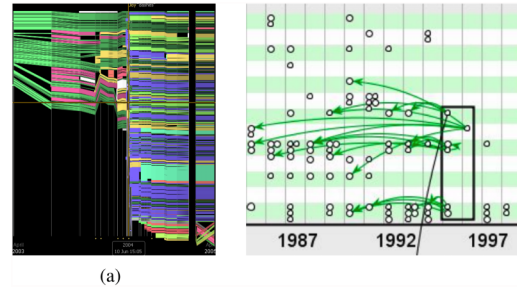


Figure 20: Examples for space-flattening. (a) Edit history of a Wikipedia article [VWD04].; (b) article citations over time [SA06, AS].

2.7. Space flattening

Space flattening is similar to space cutting, but involves flattening the cube along a particular direction on the data plane instead of extracting a cut (see Figure 19).

An example of use of space flattening in infovis is the History Flow technique for visualizing document histories [VWD04], illustrated in Figure 20(a). A space-time cube would consist of the different document version piled over time. The visualization shows the history of the paper, built by collapsing each article revision into a one-pixel column, and then displaying all columns side-by-side. These operations are equivalent to flattening the article’s space-time cube along the x data axis. Colours correspond to contributors.

Space flattening has also been used for visualizing dynamic networks, based on node-link representations on the base plane [FBS06, SA06, BVB*11]. For example, Figure 20(b) shows a screenshot from Semantic Substrates [SA06] where the y-axis is a 1D graph layout, and the x-axis shows when connections are established.

2.8. Repeated drilling

Repeated drilling is a more complex operation that consists in extracting drilling cores from a space time cube at several locations on the visualization plane, then rotating those cores in-place so they face the viewer (Figure 21).

Two examples of drilling are mentioned in this paper’s introduction (Figure 2). The top one shows the evolution of crime statistics in every US state [AA04] obtained from a geographical space-time cube containing a series of values for each state. The bottom one shows the evolution of high school population in several districts across 3 years [SRdJ05]. Although additional operations are

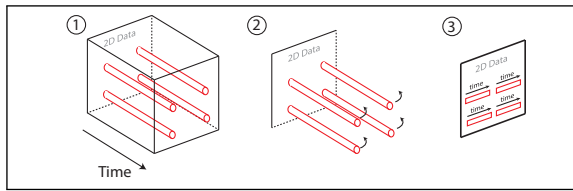


Figure 21: The repeated drilling operation.

involved (e.g. using silhouette graphs to encode values), both examples are conceptually based on repeated drilling. Repeated drilling has also been used in dynamic network visualization, for conveying changes in edge weight [BN11] and in attribute values [HSCW13].

2.9. 3D rendering

3D rendering consists in showing a space-time cube the way 3D objects are typically displayed on 2D media, that is, by projecting it onto a 2D plane (Figure 22).

3D rendering is essentially a flattening operation but in contrast with time flattening and space flattening, it is (i) typically done on a plane not orthogonal to the cube's principal axes; (ii) can involve a non-orthographic projection (e.g. perspective projection); (iii) can involve 3D shading, that is, the addition of light reflections and shadows.

In geography, 3D rendering has been used to visualize events such as earthquakes (Figure 23(a)), attributes over time [TSWS05, TR09] or the movement of objects [Kra03, GAA04]. 3D rendering has also been used to visualize dynamic networks based on node-link diagrams; nodes become columns and links become bridges between columns [DG04, BC03] (Figure 23(b)). When the layout of the dynamic network also changes, nodes become worms [DE02, GHW09]. Worms also result from 2D scatterplots extruded into time [MR97].

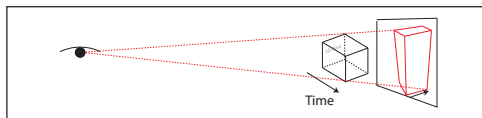


Figure 22: The 3D rendering operation.

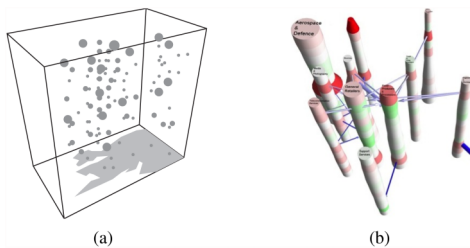


Figure 23: Two examples of 3D rendering. (a) Occurrence of earthquakes (authors' illustration after [GAA04]), and (b) a dynamic Network [DG04]

3. The Design Space of Space-Time Cube Operations

The previous section reviewed several common operations that turn a conceptual time-space cube into a final 2D visualization. Examples were selected for illustration, and the list is not meant to be exhaustive. Some of the operations were rather simple (e.g. *time cutting*), while others were more complex (e.g. *repeated drilling*) and could be described as a composition of several lower level operations. Therefore, we provide in this section a more systematic description of the design space of space-time cube operations.

3.1. Basic terminology

A *space-time cube operation* takes a space-time object and produces another space-time object. A *space-time object* is a geometrical object within a space-time coordinate system (i.e. two dimensions for the visualization space and one for time). Possible space-time objects include (i) *space-time volumes* (of which a complete space-time cube is an example), (ii) *space-time surfaces* (planar and non-planar), (iii) *space-time curves*, (iv) *points* as well as (v) sets of disconnected volumes, surfaces, curves and points.

The ultimate goal of space-time cube operations is to transform a space-time cube into a space-time object whose shape is compatible with the shape of the media employed to convey the information. By media we mean the visualization's *physical presentation*, which is the physical object or apparatus that makes a visualization observable to the viewer [JD13]. In the vast majority of cases (i.e. computer displays and paper) the media has a planar shape.


















For a given media, a space-time cube operation is *complete* if it takes space-time volumes as input and produces space-time objects whose shape match the media's shape. Otherwise, the operation is *incomplete*: it cannot be used to produce a valid visualization from a space-time cube. Several elementary space-time cube operations can be chained, in which case they form *compound operations*. A compound operation is complete if the first operation takes space-time volumes as input, and the last operation produces space-time objects whose shape is compatible with the media.

3.2. A taxonomy of elementary space-time cube operations

A taxonomy of elementary space-time cube operations is shown in Figure 33 on the next page. The taxonomy breaks down space-time cube operations into four main classes:

- **Extraction** consists in selecting a subset of a space-time object (e.g. extracting a line or cut from a volume),
- **Flattening** consists in aggregating a space-time object into a lower dimensional space-time object (e.g. projecting a volume onto a surface),
- **Geometry transformation** consists in transforming a space-time object spatially without change of content,
- **Content transformation** consists in changing the content of a space-time object without affecting its geometry.

The table in Figure 24 shows how general operations break down into more specific operations. On each of the two columns, general operations are on the left while more specific operations are on

Operations			Time	Space
Extraction	Point			
	Curve	Planar Drilling	 Time Drilling	 Space Drilling
				
		Planar Curvilinear Drilling		
		Non-Planar Drilling		
	Surface	Planar Cutting	 Time Cutting	 Linear Space Cutting
				
		Non-Planar Cutting		
			Other	
	Volume	Planar Chopping	 Time Chopping	 Linear Space Chopping
				
		Non-Planar Chopping		
			Other	
	Flattening	Planar Flattening	 Time Flattening	 Space Flattening
				
		Non-Planar Flattening		






















Operations		Time	Space	
Geometry Transformation	Translation	 Time Shifting	 Space Shifting	
		Rotation	Yaw 	Pitch 
	Roll 			
	Scaling		 Time Scaling	 Space Scaling
	Bending			
	Unfolding			
Content Transformation	Filling	Orthogonal Interpolation	 Time Interpolation	 Space Interpolation
		Volume Interpolation 		
	Encoding	Recoloring	 Time Coloring	 Space Coloring
			 Difference Coloring	
		Others		
	Labeling		 Time Labeling	
	Repositioning	Stabilizing		
		Bundling		
	Shading			
	Filtering			
Aggregation				

Figure 24: Taxonomy of elementary space-time cube operations with schematic illustrations. Grey shading indicates non-leaves. The Time column regroups operations that are applied according to the time axis, while the Space column regroups operations that are applied according to the base plane.

the right. Operations that are the most specialized (i.e. leaves on the taxonomy tree) are shown on a white background. Operations written in bold are those which produce planar surfaces, that is, can be used as final operations on screen-based and paper-based media.

We quickly review the most specialized operations (white background), going from top to bottom on the left column, then on the right column. We also describe the parameters necessary to specify each space-time cube operation. Most of the operations have already been used in infovis, others have been added to fill gaps in the taxonomy.

- **Extraction:**

- **Point extraction** consists in selecting a specific point inside a space-time volume. This operation is defined by a 2D position on the data plane and a time value.
- **Time drilling** consists in extracting a line parallel with the time axis. It is uniquely specified by a 2D position on the data plane. For example, *repeated drilling* (Section 2.8) uses several drilling operations.
- **Space drilling** extracts a line perpendicular with the time axis. It is specified by a 2D line and a time value.
- **Oblique drilling** consists in extracting an arbitrarily oriented straight line from within a space-time volume.
- **Planar curvilinear drilling** consists in extracting a planar 3D curve from a space-time volume. This operation, as well as all operations above, is complete for 2D media.
- **Non-planar curvilinear drilling** consists in extracting an arbitrary 3D curve from a space-time volume. It is incomplete, and hence needs to be combined with other operations like *flattening* or *unfolding*. This operation can be used to extract object trajectories [KW04, RFF*08].
- **Time cutting** consists in extracting a planar cut from a space-time volume in a direction orthogonal to the time axis (see Section 2.1). It takes as parameter a time value that defines the cut position on the time axis. It is a complete operation for 2D media.
- **Linear space cutting** consists in extracting a planar cut from a space-time volume in a direction orthogonal to the data plane (see Section 2.6). It is also complete, and takes as parameter a line or a segment parallel to the data plane that once extruded over time defines the cutting surface.
- **Oblique cutting** consists in extracting a planar cut from a space-time volume that is neither orthogonal to the time axis, nor orthogonal to the data plane (e.g. [FLM00]). It takes as parameter a 3D cutting plane.
- **Curvilinear space cutting** is similar to linear space cutting except the cutting surface is produced by extruding a curve parallel to the data plane that is neither a line nor a segment. This operation produces non-planar space-time surfaces that further need to be flattened (e.g. using 3D rendering [TS12]) or unfolded (as in Figure 17).
- **Time chopping** is similar to time cutting but slices have a thickness instead of being infinitely thin. Since it produces volumes it is not complete for 2D media, and thus needs to be complemented with additional operations. It takes as parameter a time segment that defines the two cutting slabs (a slab is the infinite region between two planes).

- **Linear space chopping, oblique chopping and curvilinear space chopping** are similar to the previous cutting operations, with the difference that they produce volumes with a certain thickness instead of infinitely thin surfaces.

- **Flattening:**

- **Time flattening** aggregates a space-time volume into a plane orthogonal to the time axis (see Section 2.2). This operation takes as parameters a time value, a projection function and an aggregation function. The *projection function* maps 3D points to points on the plane. Examples include orthographic projection and perspective projection. The *aggregation function* describes how point values are combined. If values are defined in an RGBA colour space, the function maps vectors of RGBA colours to a single RGBA colour. Examples of such functions include alpha-blending (e.g. averaging all colours) and overplotting (i.e. only keeping the last colour).
- **Space flattening, oblique flattening and non-planar flattening** are similar operations, but the surface on which the volume is projected is different (see previous cutting operations as well as Sections 2.7 and 2.9 for more details).

- **Geometry transformation:**

- **Space shifting, time shifting, yaw, roll and pitch** consist in moving or rotating space-time objects. They can be used, for example, for placing multiple cuts side-by-side or for rotating an entire space time cube rendered in three dimensions (e.g. [KW04, CCT*99, BPF14b]).
- **Time scaling and space scaling** rescale space-time objects along their principal axes. They take as parameters one and two scalar values, respectively, that define the scaling factor.
- **Bending** deforms space-time objects. For example, a space-time volume can be bent such that the time axis follows an arc instead of a line [DC03]. This operation takes as parameter a deformation function that maps 3D locations to 3D locations.
- **Unfolding** transforms a non-planar space-time surface into a planar space-time surface. An analogy is a map projection function that transforms a sphere or portion of sphere into a plane. An example of space-time unfolding is Maray's train schedule (Figure 19), which can be seen as an unfolded curvilinear space cut performed on a time-evolving 2D map.

- **Content transformation:**

- **Time colouring** consists in altering the colours of each time slice according to time. Examples include colouring each time slice uniformly according to a linear colour scale (Figure 13), changing the hue of each time slice, or dividing the time axis in different regions and applying a discrete colour scale (Figure 8).
- **Space colouring** alters the colour of points in a space-time volume depending on their 2D position on the data plane.
- **Difference colouring** consists in altering the colours of each time slice according to the difference between time slices. One example is highlighting appearing nodes and disappearing nodes in a dynamic graph [RM13, BPF14a].

- **Time labelling** consists in adding time labels to each time slice or to objects inside a space-time volume (Figure 9(b)).
- **Stabilizing** consists in repositioning objects on each data plane so that their trajectories are as parallel as possible to the time axis. Examples include computing stable layouts for dynamic networks [AP12a], dynamic maps [HKV12] and stabilizing videos [BGPS07].
- **Bundling** consists in repositioning objects on each data plane in order to bring their trajectories closer to each other. One example is bundling air plane routes [HEF*13].
- **Shading** consists in altering the colour of a space-time volume's content by simulating light propagation mechanisms (e.g. diffusion, specular reflection, drop shadows).
- **Filtering** consists in removing parts of a space-time volume's content. One example is removing all points of a certain colour or value [CCT*99, DC03, BPF14b].
- **Aggregation** replaces multiple space-time objects by a single, larger space-time object. Different methods exist. For example, 3D kernel density estimation transforms a set of space-time points or space-time curves into 3D volumes or 2D (iso) surfaces [DV10].

3.3. Adaptive and semantic operations

So far, we mostly described operations that are agnostic to the data and the content of the cube. *Adaptive operations* take into account the shape or content of the particular space-time objects they operate on. For example, an *adaptive time cutting* operation can be used to cut cubes according to regions with large changes instead of cutting them into regularly spaced slices. This technique is used, for example, in adaptive video fast-forward [PJH05] and to highlight states and repetition in dynamic networks, using a technique called 'Small MultiPiles' [BHRD*15] (*time chopping* + *space shifting*). Similarly, an unfolding operation that works on any surface (as opposed to, e.g. only spheres), would be an *adaptive unfolding* operation.

Semantic operations take into account the data semantics of the space-time objects they operate on. One example would be a *semantic volume interpolation* operation that connects discrete sets of moving objects with lines or tubes (see Figures 9(b) and 23(b), as well as [Ros06, BPF14a]). This type of operation is semantic because it needs to know the identity of the objects to be able to match them on successive time slices. *Time labelling* operations such as the one used in Figure 9(b) are also semantic, because they need to know the location of datapoints of interest to place the labels appropriately. *Filtering*, *aggregation* and *bundling* can also be seen as *semantic operations* as well as all *difference colouring* and *space colouring* [VWD04, RFF*08, BPF14b]. Finally, semantic operations can also be used to cut or chop cubes according in regular distances (e.g. hours, days, weeks). Flattening such time chops changes the temporal granularity of the cube.

3.4. Compound operations

We previously defined *compound operations* as several operations applied in sequence. According to our taxonomy from Figure 24,

Table 1: Example compound operations decomposed.

Compound Operation	Elementary Operations
Discrete time flattening	time cutting* + time flattening
Colored time flattening	time coloring + time flattening
Time juxtaposing	(time cutting + space scaling + space shifting)* + time flattening
Marey's schedule	curvilinear space cutting + yaw + unfolding
Slit tears	(linear/curvilinear space cutting + yaw + [unfolding] + space shifting)*
Repeated drilling	(time drilling + time scaling + yaw)*
3D rendering	[shading] + oblique flattening

some of the operations we introduced in Section 2 are elementary, namely *time cutting*, *time flattening*, *space flattening*. Others are compound and can be broken down as indicated in Table 1. In our notation, the symbol + refers to a composition, the symbol * refers to operations applied multiple times and the symbols [] refer to optional operations.

A compound operation inherits the parameters of its suboperations. For example, a *discrete time flattening* operation is specified by a sequence of time values, as well as a projection function and an aggregation function. But in practice, most compound operations enforce constraints between their parameters. For example, all *space scalings* from a *time juxtaposing* operation are typically the same. Possibly those parameters could even result from a function. For example distance between individual time cuts in a *discrete time-flattening* could follow a logarithmic distribution, to give more importance to recent time steps.

Many elaborate temporal data visualization techniques can be described as compound operations. For example, the 'Visits' technique (Figure 25(a)) shows the history of a person's travels using *semantic (time chopping + space chopping + time flattening + space shifting)**. The conceptual space-time cube captures the location of the person in the entire world over a period of time. Each circle in Figure 25(a) is a *time-flattened* segment from this space-time cube. Each segment has also been *space-chopped* to show only a portion of the entire world. Time-flattened segments are then *shifted* on the data plane to distinguish them visually and to reflect the temporal order in which the regions have been visited.

A further example of a compound operation is shown in Figure 25(b). The technique is similar to that of Figure 3(a) and shows one connected scatterplot per country, visualizing infant mortality versus life expectancy over time. While the original article [RFF*08] calls this technique 'Small multiples', our framework would describe it as the following compound operation: (*filtering + interpolation + time flattening + space shifting*)*. That means, the initial space-time cube consists of a 2D scatterplot containing all countries, and time as third dimension. For each country, the technique extracts all the corresponding points, interpolates between them, flattens them across time and finally juxtaposes the resulting space-time surfaces.

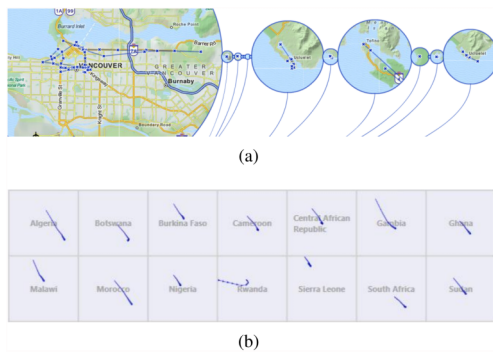


Figure 25: Examples for compound operations: (a) in Visits [TBC13] (b) in a study by Robertson et al. [RFF08]

3.5. Dynamic operations

So far we only considered operations (both elementary and compound) that transform a space-time cube into a static visual representation. On computer displays, operations can also be applied in a dynamic manner. Dynamic operations can involve either animation or interaction.

3.5.1. Animation

We refer to *animation* as the process of applying different operations on a space-time cube over time, or similarly, varying the parameters of an operation over time.

The most common form of animation consists in changing the position of a cutting operation over time, that is, *animated time cutting*. This results in the space-time cube content being ‘played back’. For example, if the space-time cube represents a visual scene like video surveillance data, synchronizing the motion of the slice with a clock will result in a real-time playback of the original scene. When significant data are skipped during playback, the animation is closer to a *discrete time juxtaposing* operation, except slices are shown in sequence instead of being laid out side-by-side.

An animated time cutting operation can be preceded by a *filling* operation in order to produce smooth animated transitions. Many examples exist in the literature, for example when animating dynamic networks [ATMS*11, RM13, BPF14a] or scatterplots [Ros06, RFF*08]. Most of these examples can be described as *semantic volume interpolation* + *animated time cutting* operations. Animated time cutting can also be combined with other space-time cube operations such as *time flattening*. For example, Gapminder can combine scatterplot animations with static trails for points of interest (a *filtering* + *time flattening* operation) [Ros06].

While many animation techniques can be described as animated time cutting on *static* space-time cubes, more elaborate techniques require operations to be applied in real time. For example, Hurter et al.’s system [HEF*13] uses *animated time chopping* to animate a network over time while preserving temporal context information. At every animation frame, a *time flattening* is applied that produces coloured trails and a dynamic *bundling* operation is applied that

guarantees a continuous animation without jumping bundles [HET12].

Although *animated time cutting* and its many variants are the most common forms of animation, other animated operations exist. For example, *animated 3D rendering* can explain a transition between two space-time cube operations to a user by smoothly rotating a space-time cube representation [BPF14b]. This technique will be discussed in Section 6, where we review space-time cube visualization systems.

3.5.2. Interaction

Interaction is similar to animation, except the changes in the space-time cube operations are under the user’s control.

Consider *animated time cutting*: if the position of the cutting plane is controlled by the user (e.g. by dragging a slider) instead of being automatically moved, then the operation becomes *interactive time cutting*. A common implementation of interactive time cutting is the seeker bar on a video player. As with animations, any operation can be made interactive. Examples of interactive operations abound, and we will review some of them in the next section.

Another example of interaction is *interactive curvilinear space-cutting* + *oblique flattening*. Tominski and Schulz [TS12] allow user to draw a curve onto a map, then show the curvilinear space cut in a perspective projection.

4. The Inner Structure of Space-Time Cubes

So far, we mostly considered the generalized space-time cube as monolithic entity and operations were defined on the entire cube. However, two space-time cubes can look quite different (see Figures 23, 29, 30, 32, 33), depending on the characteristics of the data being visualized; migration of animals, earthquakes, changes in vegetation, dynamic networks, surveillance videos or scatterplots evolving over time. Such differences define how the space-time cube ‘looks’, and which we refer to as its *inner structure*.

On the other hand, different datasets can result in similar inner structures. For example networks in the form of node-link diagrams, and multivariate data visualized as scatterplots over time, both yield structures similar to the first row in Figure 26. The inner structure of a space-time cube is an important factor when choosing between space-time cube operations and is independent of the specific type of data. This allows for a general discussion about the advantages and disadvantages of an operation, as well as to transfer visualization techniques between domains.

4.1. Decomposing space-time cubes

We assume that any space-time cube can be subdivided into lower level space-time objects (see Section 3.1). A common example is a set of 3D curves or tubes generated by moving objects such as cars, persons, nodes in a network or points in a scatterplot (see Figure 26 ①). These curves can yield different patterns depending on the data ②③④. Another typical structure is obtained from dense matrix data such as videos or dynamic heatmap visualizations.

Density:	Dense		Sparse	
Variability:	Low	High	Low	High
Moving Objects	①	②	③	④
Matrix Data	⑤	⑥	⑦	⑧

Figure 26: Different inner structures for a space-time cube.

For these datasets, the object of interest (e.g. a pixel) is fixed in space but its colour changes over time. The space-time objects are therefore multi-coloured lines parallel with the time axis and with no intervening space ⑤. Again, different patterns of variation are possible ⑥.

A space-time cube can be subdivided in different ways. For example, pixels with an alpha channel can yield higher level space-time objects with intervening spaces: ⑦ could be obtained by filtering a pixel-based landscape visualization according to the type of vegetation while ⑧ could be obtained by isolating bush fires [CCT*99]. Alternatively, optical flow can be extracted from a video, which would yield a structure more similar to ①. Conversely, any animated visualization could be turned into a video and yield a structure similar to ⑤. These are all different ways of looking at a space-time cube. Ideally, however, the focus should not be on the specific decomposition used, but rather on the overall properties of the inner structure.

4.2. Dimensions of inner structures

Figure 26 only shows eight possible examples of inner structures, but many more structures are possible. For example, moving objects can also change colour or shape over time, or disappear and reappear, and this would yield different structures than the ones shown in the figure. Four dimensions are however enough to capture the most important properties of inner structures, namely: *density*, *variability in positions*, *variability in visual attributes* and *object lifespan*.

Density: The density of a space-time cube refers to how much non-empty space it contains. The inner structure of a space-time cube can be *dense* (Figure 26 ①②⑤⑥) or *sparse* (Figure 26 ③④⑦⑧). Datasets like videos (Figure 26 ⑤⑥) produce inner structures that are maximally dense, and which will be referred to as *complete*. Furthermore, a space-time cube can be *spatially dense* but *temporally sparse* (e.g. a few photo snapshots), or vice versa (e.g. a GPS log over several days). Density is a key property of inner structures, and it is the only dimension that is independent from the decomposition used.

Variability in positions: This dimension describes the amount of motion that 2D objects undergo in the space-time cube, which in turn affects the shape of the 3D space-time objects that make up its inner structure. For example, case ③ in Figure 26 (which could be moving people) involves rather straight curves, while case ④ (which could be animal movements) involves spaghetti-like curves. Matrix data (cases ⑤ to ⑧) yield zero variability in positions, since pixels are fixed and space-time objects are lines parallel with

the time axis. Using an optical flow decomposition, case ⑥ would exhibit more variability in positions than case ⑤.

Variability in visual attributes: This dimension refers to the amount of changes (colour and shape) that 2D objects undergo in the space-time cube. For example the pixels in ⑤ and ⑥ undergo changes in colour, with a higher variability in ⑥. In contrast, the moving objects in ① to ④ have a fixed shape and colour, and therefore the cube's inner structure exhibits no variability in visual attributes. In a multivariate network [KPW13] setting, this variation could exist, for example, the objects were moving taxis and their colour or size (e.g. radius) encoded the number of customers they carry. Case ⑧, which involves blobs moving over time (e.g. a spreading bush fire), yields high variability while case ⑦ does not.

Object lifespan: Object lifespan refers to how long 2D space-time objects persist on average, which has an influence on the portion of the time axis covered by the 3D space-time objects. The lifespan is maximum when all objects span the entire time axis (Figure 26 ① to ④). Regular matrix data also yield a maximum lifespan (Figure 26 ⑤ and ⑥). In contrast, objects can reflect events that have a start and an end: ⑧ involves short object lifespans whereas ⑦ involves relatively long object lifespans. Events without duration produce the shortest possible lifespans, and yield inner structures consisting of points. Since objects do not persist, there is no variability in positions or visual attributes.

The inner structure of any space-time cube can be described along the four previous dimensions. For example, the earthquake cube in Figure 23(a) is relatively sparse and has minimal object lifespan. The dynamic network cube in Figure 23(b) is also relatively sparse, but has various lifespans (long for nodes and minimal for links). Also, nodes are columns that exhibit no variability in position but a high variability in visual attributes (colour).

Thereby, it does not matter whether the initial data are discrete or continuous in time or space. Discrete and continuous are two complementary aspects, both encompassed by the generalized space-time cube. For example, *repeated (time chopping + time flattening)* can transform a continuous time dimension into a discrete time dimension. On the other hand, *time interpolation* can yield a continuous impression of otherwise discrete data.

In the next section, we will see why a dataset's characteristics matter, and how the inner structure of a space-time cube can be altered with space-time operations.

5. Which Operations Work Best?

So far we essentially focused on the descriptive function of our framework: the different types of operations, how they can be combined or made dynamic and the inner structure of space-time cubes. We now have all the elements to discuss practical strengths and weaknesses of space-time cube operations. This aims to support the choice of operations when designing visualizations, to researching novel visualization techniques to leverage the strengths and overcome the drawbacks of space-time cube operations.

Strengths and weaknesses of operations are directly related to user tasks, since certain operations yield visualizations that may better

suit a certain task. So far, only few user studies compare space-time cube operations but will be used to ground our discussion. However, only a small subset of all possible operations have been covered by user studies; many of the existing findings are hard to generalize beyond the domain, type of data, and operation parameters used in the stimuli. Thanks to the ability of our framework to reduce complex techniques into elementary operations, useful advice can often be provided based on more general knowledge on visual perception and on common sense. With respect to data type, the inner structure of a space-time cube can in many cases inform the suitability of certain operations. For example, if a dataset produces maximally dense space-time cubes, a 3D representation will always be problematic unless other operations (filtering, cutting) are used.

5.1. 3D rendering

3D rendering ([*shading*] + *oblique flattening*) is the operation most commonly associated with the space-time cube concept, so it seems necessary to start with a cautionary note. 3D visualizations are known to pose major difficulties to users and are overall not recommended [Shn03, Mun15]. Difficulties include (i) occlusion, (ii) depth ambiguity, (iii) perspective distortion, (iv) colour distortion and (v) navigation difficulties, further discussed below. *3D rendering* can however be useful for providing a general overview of a space-time cube's inner structure (density, variability and object lifespan). It is also the most effective approach when the goal is to have observers explicitly think in terms of a space-time cube, as it is the case for most diagrams in this paper.

Occlusion issues are due to the use of a *flattening* operation together with an *overplotting* aggregation function (i.e. near objects are drawn on top of far objects). The amount of occlusion depends on the *density* of the inner structure (see Figure 27). Therefore, *3D rendering* is mostly adapted to *sparse* structures, as in Figure 23. Additional space-time cube operations can however be employed for visualizing *dense* (and sometimes *complete*) structures. An effective technique involves sparsifying the cube's structure using *interactive filtering* [CFC*96, BPF14b]. A related approach involves the use of *translucency filtering* or more generally, translucent space-time objects (Figure 26). However, translucency produces visual clutter without fully addressing the problem of occlusion, and we found few instances of this technique yielding legible results. Also, occlusion is an important depth cue [WFG92] that is mostly lost when translu-

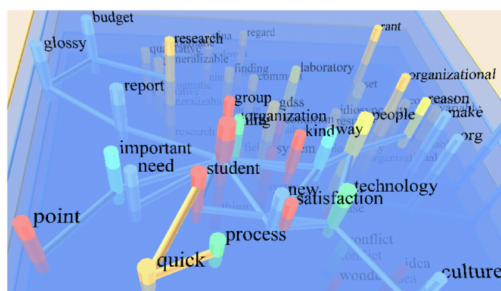


Figure 27: A time-evolving network rendered in 3D with translucency [BC03].

ency is employed. For similar reasons, we recommend against making the whole space-time cube translucent—or equivalently, using an *alpha blending* aggregation function. Elaborate interactive operations exist that can decompose a space-time cube using *extraction* and *rigid transformation* operations, and allow users to look inside [CCF96].

In cases where space-time objects are densely clustered around certain areas with the space-time cube, *aggregation* can help reducing visual complexity and occlusion. In order to avoid rigid objects inside the cube, clusters of space-time cube objects can be estimated through transparency [DV10].

The *flattening* operation is further responsible for causing depth ambiguity. In Figure 23(a) for example, it is impossible to extract the spatio-temporal coordinates of the points. In the real world, most depth ambiguities are resolved through stereoscopic vision and structure from motion [Tod04, WFG92]. Structure from motion is as powerful as stereoscopic vision [Tod04], and this is why *3D rendering* is the most effective when the space-time cube is allowed to rotate through *animated* or *interactive 3D rendering*. Incidentally, being able to ‘look behind’ space-time objects also mitigates occlusion problems. When interaction cannot be used (e.g. on paper media), static depth cues are crucial. One approach for structures with no or very low *variability in positions* is to have space-time objects “touch” the ground. This is already the case for structures with maximum *object lifespans* such as in Figure 23(b). For short *object lifespans* like in Figure 23(a), it is advised to either add lines connecting objects to the ground [KW04] or add drop shadows [WFG92, TKB07] by applying a *shading* operation with perpendicular lighting.

Perspective distortion is due to the use of a *perspective* projection function: far objects appear smaller than near objects, making it difficult to accurately compare sizes and lengths. This issue can be eliminated with the use of an *orthographic* projection function instead. However, doing so will eliminate an important depth cue and may produce depth reversals analogous to the necker cube illusion [WFG92]. Again, this type of ambiguity can be alleviated by the use of other depth cues such as occlusion, contact and *shading*.

Colour distortion is a side-effect of the *shading* operation, since it alters the colour of the space-time objects. This is not a problem when colour is not encoding information or when simple colour scales are used as in Figure 23(b), but could be an issue when numerical data are encoded through brightness.

Navigation difficulties with 3D visualizations are due to the *oblique flattening* operation having too many degrees of freedom. One way to facilitate navigation is to add constraints, that is, only let users manipulate *3D renderings* in ways that are essential to their task [SSS*01, Shn03].

Generally speaking, space-time cubes are a convenient conceptual tool for thinking about temporal visualizations, but explicitly showing them using *3D rendering* poses a number of problems and should be done only when necessary (e.g. for showing an overview or for explanatory purposes). Computer screens and paper are 2D media, and are optimized for 2D visualizations [Shn03]. As we have already seen, many 2D solutions exist to show temporal data.

5.2. Time flattening

Time flattening (Figure 7) can be seen as a 2D technique, or as a specific form of *3D rendering* where ambiguities are purposeful and animation or interaction are not necessary to make sense of the data. *Time flattening* essentially turns a space-time cube into a spatial visualization: each point on the resulting image shows the history of this point aggregated over time. As a result, spatial relationships and trends are preserved and are clearly visible, such as the path of Napoleon's army in Figure 8 or the distribution of data values in Figure 9(b). Any aggregation function can be used, for example, alpha-blending for showing averages, or more elaborate techniques for showing cumulative statistics (Figure 9(a)).

The major drawback of *time flattening* is that it reduces the time axis to a point (Figure 7), so most temporal information is lost, including the ordering of events, temporal intervals and absolute timestamps. However, part of this information can sometimes be cognitively reconstructed from object trajectories. For example, when reading Napoleon's march (Figure 8), one knows that an army does not jump between random locations but progresses smoothly, so not all orderings of events are plausible. Knowing that an army generally shrinks in size is then enough to infer its direction of motion: eastward then westward. The assumption of motion smoothness also helps to read the infographics of Figure 11: the ice skater either progresses from left to right, or from right to left. Basic knowledge about how a human body moves then suffices to entirely reconstruct the direction of motion and the ordering of events.

For flattened trajectories to be legible and informative, the space-time cube needs to be *spatially sparse*, with long *object lifespans*, and exhibit a reasonable amount of *variability in positions* and in *visual attributes*. A space-time cube can sometimes be processed to meet these criteria. For example, *time colouring* (Figure 13) is a way of adding *variability in visual attributes*. This operation has the additional advantage of explicitly encoding time. However, colour is a poor visual attribute since the human visual system cannot discriminate between many colours [CAS05], and no colour scale exists that naturally maps to time. For example, red in Figure 13(a) represents old events, while it represents new events in Figure 13(b). If time intervals (or object speeds) are more important than absolute times, a more effective approach is to change the colour or shape of objects at regular time intervals.

Additional operations can be applied to a space-time cube in order to reintroduce missing time information. For example, the Phillips curve in Figure 9(b) uses *semantic volume interpolation* to connect points and show their ordering. Both Figure 9(b) and Figure 8 (for the army's return trip) employ *time labelling* to explicitly convey event timestamps.

Discrete time flattening (*time cutting** + *time flattening*) is a variant that can also clearly expose time intervals and object speeds. Since this technique can result in loss of information, it works best for structures that are *sparse in time*. The way time cuts are selected is important. Equidistant cuts generally expose relative times, durations and rates of change. When not all cuts are equally informative, non-equidistant cuts may be preferred, but at the cost of losing most of this temporal information. If cuts are chosen so

that objects do not overlap, these objects become clearly visible. This technique is used in Figure 11 to expose the skater's posture, an information that helps to read the visualization.

In summary, *time flattening* is an operation that produces easy-to-interpret visualizations: space simply maps to space. It is useful for showing data patterns aggregated over time and where space is more important than time. But since it collapses the time axis into a point, it is most powerful when it can leverage humans' ability to reconstruct time from trajectories. This often requires the use of additional space-time cube operations and a very careful visual design.

5.3. Space flattening

Space flattening (Figure 19) is another form of *orthogonal flattening* that avoids some of the drawbacks of regular *3D rendering*. The orthogonal projection function creates a meaningful spatial encoding, since one axis maps to time while the other axis maps to space. The resulting visualization can be read as a regular timeline.

Space flattening is effective when both time and space are important, but one dimension of space is more informative than the other. This is the case for edit patterns in Wikipedia articles (Figure 20(a)). This technique is also useful for focusing on one data dimension at a time, when the two data axes have clearly different roles (e.g. in a 2D scatterplot). However, it produces visualizations that are generally less easy to interpret if the base plane does not have a natural coordinates system (e.g. in a node-link diagram or on a map).

Other cases require a linear ordering of data elements along the single remaining spatial dimension, in order to avoid occlusion. Such a linear ordering can be obtained by *unfolding* as in Marey's train schedule (Figure 17). In networks, a linear ordering of nodes can help reduce visual clutter caused by edge crossings. An ordering can be defined based on different criteria and strongly influence the overall readability of the visualization, as well as emphasize particular patterns [vdEHBvW13].

5.4. Time cutting

So far we covered different types of *flattening* techniques. These techniques essentially aggregate the data in a space-time cube and are helpful for overview tasks, but fail at providing details. We now turn to *extraction* techniques. Among these techniques, *time cutting* is by far the most common. A single time cut provides a clean and fully detailed snapshot of a particular time stamp. It can potentially fill the entire screen and show many details, and is furthermore free of using any additional visual variable for time (e.g. colour, brightness). A simple *time labelling* of the time cut can be enough.

However, since the goal of a temporal visualization is to show information across time, more than one time cut must be shown. We already mentioned *discrete time flattening* as a possible approach, although it still suffers from the drawbacks of aggregation. Two aggregation-free alternatives are *time juxtaposing* and *animated time cutting*. We review them separately, then discuss empirical research comparing both.

5.5. Time juxtaposing

Time juxtaposing ($\{time\ cutting + space\ scaling + space\ shifting\}^* + time\ flattening$) allows to show different time slices simultaneously. Due to our exposure to sequential art, it produces familiar visualizations, that are easy to interpret. Usually, *space shifting* follows the occidental convention of laying out time slices either linearly from left to right, or on 2D from left to right then top to bottom [McC94].

The main advantage of *time juxtaposing* over *flattening* is the legibility of individual time slices due to the absence of occlusion. However, the *space scaling* operation also has the effect of shrinking time slices: the more slices are shown, the less space—and therefore resolution—every slice has. The *time juxtaposing* operation therefore requires finding the best trade-off between spatial and temporal resolution.

A potentially serious drawback of *time juxtaposing* is that time slices use inconsistent spatial coordinate systems, due to the *space shifting* operation. This can make it difficult to relate moving objects across time slices and reconstructing their trajectories. However, if objects on a time slice can be visually distinguished and have a consistent appearance across time (i.e. low *variability in visual attributes* and *in positions*), relating time slices becomes easier. With few objects and pre-attentive visual encodings [War12], relating slices can be effortless. If there are many objects in each time slice and/or high variability in positions and visual attributes, tasks may still require visual search and split attention. However, this effort can be alleviated through interaction by adding brushing and linking support [BC87].

Time juxtaposing is analogous to *discrete time flattening*, except time slices are laid out next to each other instead of being blended into a single image. As a result, the technique yields similar benefits compared to a regular *time flattening*. One is that it is able to expose the temporal information that is otherwise lost by aggregation, including event ordering, time intervals, speeds and absolute time stamps. But again, this assumes that objects can be easily related across slices, which can be difficult when slices do not share the same coordinates system.

Like with *discrete time flattening*, the proper choice of time slices is crucial with *time juxtaposing*, with the difference that granularity also affects spatial resolution. Again, equidistant cuts are better at preserving temporal information. But when not all cuts are equally informative, non-equidistant cuts (i.e. *semantic time cutting*) may be preferred. For example, when showing changes in territory, historical atlases and newspapers often employ this technique to emphasize different periods and historical events (Figure 28). Individual dates



Figure 28: Territory of later Yugoslavia (red line).

typically need to be added (i.e. *time labelling*) to reintroduce the missing time information.

5.6. Animated time cutting

A popular way of exploiting *time cutting* is through *animated time cutting*. This approach inherits the benefits of *time cutting* compared to *flattening* in that there is no aggregation. Therefore, individual time slices can be seen clearly, without occlusion. In addition, *animated time cutting* is probably the most ‘natural’ and easiest to interpret among all possible techniques, since space maps to space and time maps to time. Animations can also be visually appealing and users can find them engaging.

In contrast with *time juxtaposing*, *animated time cutting* does not suffer from the reduced spatial resolution due to *space scaling*: each time slice can span the entire screen. Also, many time slices can be shown in a short period of time (e.g. 60 in only 1 s on a 60-Hz display). However, fast animations are not necessarily easy to comprehend. If a space-time cube’s structure is *temporally dense* (e.g. a long surveillance video), an animation can take time to be perceived and understood [TMB02]. As with previous techniques, *semantic time cutting* can be used to speed up the playback when not all time cuts are equally informative [PJH05]. Nevertheless, like *time juxtaposing*, the technique may not be suitable for conveying information with extremely high temporal density, and aggregation (e.g. in the form of *time flattening*) may be preferable.

One major drawback of animated time cutting is that time slices must be presented in a pre-defined (typically chronological) order. As a result, the user must rely on memory in order to fully understand the data [Mun15]. Since time slices are not shown simultaneously, they cannot be freely accessed [TMB02]. This makes it difficult to carry out many types of exploratory tasks, such as comparing two non-successive slices, or examining a slice in context with other slices. Annotations are also impossible. These are serious issues that can be partly alleviated by using *interactive time slicing* instead of *animated time slicing* and allow to animate between non-adjacent time cuts [BPF14a]. But these techniques do not preclude the use of memorization. Also, both *animated* and *interactive time slicing* usually requires dynamic displays (though one can think of physical flipbooks), whereas all previous methods are compatible with actual static media such as newspapers.

Nevertheless, *animated time cutting* can be very effective at revealing changes between successive slices, as our visual system is equipped for pre-attentively detecting in-place changes [Bar97]. It can reveal subtle visual changes that are hard or impossible to detect otherwise [WB05, GMH*06], as well as large and drastic visual changes. Our perceptual ability to detect and interpret changes, however, requires the absence of visual disturbances (e.g. flicker) between time slices of interest [Ren02], and ideally, that successive time slices are similar enough to be visually integrated and perceived as a motion [FT06]. This means the *variability in positions* should not be too high as to produce large discontinuities. If it is the case, *time interpolation* or *volume interpolation* need to be used to produce smooth animated transitions. Changes can also be highlighted (e.g. through a *difference colouring* operation [APP10, RM13, BPF14a]) in order to facilitate their detection and interpretation.

5.7. Stabilization

The *stabilization* operation can be used to complement several of the previously seen approaches. While *interpolation* allows to produce smooth transitions between pre-defined time slices, *stabilization* allows to reduce *variability in positions* when positions do not encode data. Although it may reduce the legibility of trajectories when using *time flattening*, this technique may facilitate object tracking tasks in *animated time cutting* and object mapping tasks in *time juxtaposing*. It is mostly used in dynamic graph drawing. The goal of dynamic graph drawing is to compute optimal layouts for time-evolving networks by either partially stabilizing nodes, or by computing a fixed global layout (Figure 23(b)). The underlying motivation is to allow users to maintain a spatial ‘mental map’ of the data [ELMS91, MELS95, AP13].

5.8. Empirical evidence

There has been a considerable amount of literature in infovis, cartography and educational psychology about the respective merits and weaknesses of animated versus static representations [TMB02]. The scientific debate essentially opposes *time juxtaposing* with *interactive time cutting*. Other studies, mostly from the domain of geo-spatial visualization, compared *time flattening* with *3D rendering* and sometimes with *animated time cutting*.

Overall, studies suggest that *time juxtaposing* can be effective for many tasks as long as objects of interest can be easily related across time slices or visually indicated through linked highlighting. It has been shown to be faster and more accurate than *interactive time cutting* for understanding temporal trends in dynamic scatterplots [RFF*08]. For dynamic network analysis, it has been consistently shown to be faster for a number of tasks (e.g. reading change in node degree or number), with no measurable difference in error rates [APP11, FQ11]. Thus, for most tasks, time juxtaposing has been shown to be faster than interactive animation with no impact on the error rate [TMB02].

However, in specific circumstances, *interactive time cutting* has been shown to bring advantages. For example, it yields less errors when detecting the simultaneous appearance of nodes or edges in a network visualization, although it is slower [APP11]. *Animated time cutting* is very effective for highlighting subsets of data against a similarly looking background. Ware and Bobrow [WB05] demonstrate this quite convincingly for highlighting subgraphs in a large graph, and Griffen *et al.* find that it helps identify moving clusters against a background with identical colours [GMH*06]. In both cases, animation was both more accurate and faster. There is also some evidence that participants use both spatial and motion information when memorizing animated network sequences, suggesting that animation can help with the memorization of dynamic data [AP12b]. Thus, *interactive time cutting* and *animated time cutting* seem useful for highlighting data subsets (using motion) and for representing a small number of dynamic changes. Short animations, consisting of just two to three time steps, may be useful for clarifying drastic reorientations in the visual representation of the data [TMB02]. There is some evidence to support this finding in the context of dynamic graphs where the stability of the drawing is low and important nodes in the task cannot be highlighted throughout the time series [AP16].

Concerning the efficiency of *3D rendering*, results are also mixed. On a dataset of four people moving through space, Kristensson *et al.* [KDA*09] found that *time flattening* with *space colouring* (each person had an individual colour) performed better for simple tasks than *3D rendering* with *shading*. However, the latter was found to better support more complex overview tasks. Investigating similar techniques but on a much larger dataset (203 trajectories), Kveladze *et al.* [KK12] conducted a participatory design session with experts in geography. Experts reported that *shading* and mapping data attributes to trajectory thickness cluttered the visualization. Instead, they required *space colouring* in order to better distinguish people. *Time flattening* was generally perceived too overloaded. However, experts further reported that *3D rendering* ‘helps in earlier stages of the research to get some hints about the data’.

Willems *et al.* [WvdWvW11] compared *non-interactive animated time cutting* with *3D rendering* and an enhanced version of *time flattening* [WvdWvW09]. Results show that overall, each technique was best for one task. The study discusses many exceptions, depending on data *density* and *variability in position*. *Animated time cutting* was reported to be the least robust to a *high variability in position*. For the number of objects, Willems *et al.* report that *3D rendering* was least scalable, while *animated time cutting* performed best. Comparing *interactive animated time cutting*, *time juxtaposing* and *aggregation (iso-surfaces)* in a space-time cube, Brunsdon *et al.* [BCH07] came to similar results. As movement.

The *stabilization* operation has also been empirically studied. Many studies found no positive effect [PS08, SP08, APP11, AP12b]. However, studies on directed graphs [PHG06] and recent studies on general, undirected graphs [GEY12, AP12a] have demonstrated benefits on mental map preservation. Overall, stabilization can help for a number of tasks, but further study is needed.

Overall, previous empirical studies paint a complex picture, and several controversies such as static versus animated and 2D versus 3D are far from being settled. Many studies emphasize task as a main source of trade-off. But most likely, the respective advantages of different approaches also heavily depend on the inner structure of space-time cubes and on the specific space-time cube operations being used. For example, employing smooth transitions (an *interpolation* operation) and giving subjects control over animations (*interactive time cutting*) can have a huge impact on the effectiveness of animations. Also, different datasets may produce widely different *inner structures*, and dramatically influence the efficacy of techniques. By clarifying these differences and providing a terminology to describe them, we hope that our framework will help design more informative user studies and eventually settle ongoing debates.

5.9. Other operations

We discussed the benefits and drawbacks of space-time cube operations commonly used for visualizing temporal data. As suggested by our framework (the elementary operation taxonomy in Figure 24 and the different ways to combine them), a possibly infinite variety of operations exist and it is impossible to analyse them all. It is likely that among all possible operations—especially the least constrained ones such as *oblique* and *non-planar* operations—the vast majority will be harder to read and to interpret. Popular approaches

are particular in that they mostly involve orthogonal operations, especially along the time axis (i.e. *time flattening* and *time cutting*). One striking exception is *3D rendering*, which combines several complex operations to emulate how we perceive solid objects in the real world (e.g. rotation, perspective projection, shading) and produces familiar (but sometimes illegible) images.

It is however also likely that the immense design space of space-time cube operations contains techniques that need to be learned, but can yield unique benefits for specific temporal data analysis tasks that usefully complement standard approaches. Some of these have already been discovered (see, e.g. Figures 2, 5, 15, 17, 18, 20(a), 31). The exciting possibility that the design space contains a number of other ‘pearls’ to be discovered stresses the importance of research in novel temporal visualization techniques. By discussing the multiple trade-offs involved in known techniques and explaining how these trade-offs originate from the intrinsic properties of elementary operations, we hope that our framework will be able to guide designers and accelerate future discoveries. Another promising line of work consists in supporting multiple space-time cube operations, and is discussed next.

6. Space-Time Cube Systems

Choosing an appropriate space-time cube operation depends on many factors and almost always involves trade-offs. In this section, we review a representative sample of visualization systems that address this issue by supporting multiple space-time cube operations. Those systems almost invariably use *3D rendering* as an explicit representation of the space time cube, both for showing an overview and for explaining how different operations relate. Operations are implemented on the visual model of the cube. We call these systems *space-time cube systems* and the traditional understanding of space-time cube in geo-visualization is in fact a geographical space-time cube system. Because space-time cube systems work by letting people switch between different operations and tune their parameters, *interaction* and *animation* are central features.

6.1. GeoTime

GeoTime is a carefully designed commercial system for analysing spatio-temporal data [geo, KW04]. Events are shown as spheres on a *3D rendering* view that can be freely rotated (Figure 29(a)). This view also uses a reference plane, and a *semantic volume interpolation* operation is applied to indicate event ordering. Users can perform *time chopping* operations by dragging on a timeline widget. GeoTime also supports *time flattening* and *space flattening*. Figure 29(a) shows a *space flattening* view where time runs from top to bottom, and a reference plane is provided that can be rotated. Thin grey lines connect the two views. Finally, pan and zoom is supported through *space chopping* + *space scaling*.

6.2. Tardis

Tardis [CCT*99, CFC*96] is a system for visualizing environmental data using *3D rendering* in combination with advanced space-time cube operations. The voxels in the cube are colour-coded depending

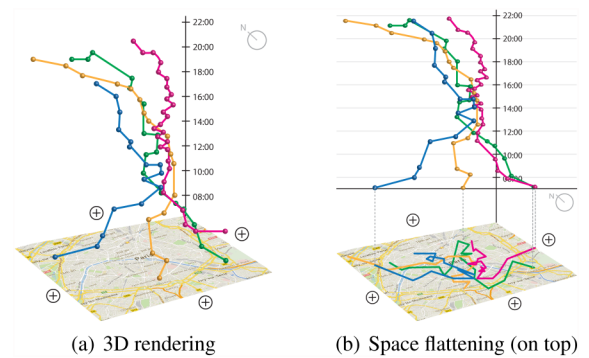


Figure 29: Illustration after GeoTime [geo]

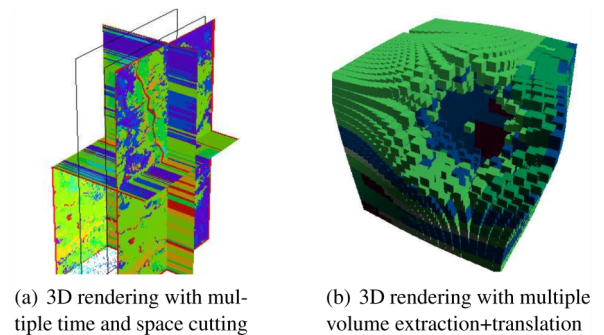


Figure 30: Tardis [CCT99] and visual access [CFC96]

on the type of vegetation, its age, soil characteristics or the presence of bush fires (Figure 30).

Tardis implements *interactive semantic filtering*: users can, for example, select a particular type of vegetation or a range of vegetation ages. In addition, Tardis supports *interactive orthogonal cutting*, but in contrast with our previous examples, cutting is always used in combination with *3D rendering*. Users can define and manipulate multiple orthogonal cutting planes (Figure 30(a)). Further operations include opening the cube like a book (*interactive (volume extraction + rotation)**) or applying a 3D fisheye effect (*interactive (volume extraction + translation)**), pushing away voxels from the cursor (‘Visual Access Distortion’, Figure 30(b)).

6.3. VISUAL-TimePacTS

VISUAL-TimePacTS is a system for analysing activity diaries [VFC10]. It uses non-geographical space-time cubes. The cube’s two data axes can be mapped to data dimensions such as individuals, locations or activities. We focus on the case where one axis maps to individuals while the other axis maps to activities. Activities are also encoded using colour.

VISUAL-TimePacTS supports *linear space flattening* on both data axes. Figure 31(a) shows six individuals (horizontal axis) and their activities (colours) across time (vertical axis). Figure 31(c) shows the evolution of activities aggregated across all people over time. VISUAL-TimePacTS supports a seamless transition between

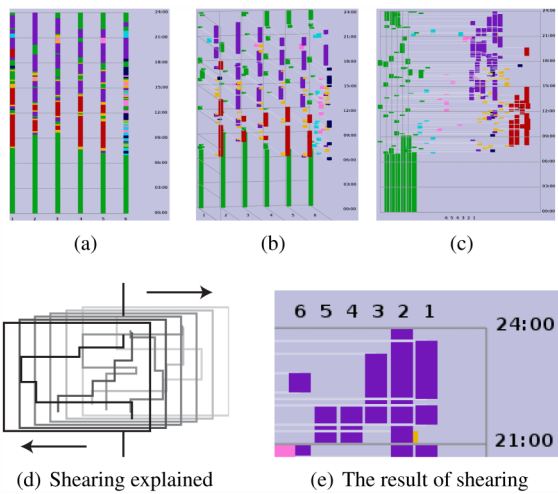


Figure 31: *VISUAL TimePacTS [VFC10]. (a) Space flattening on activities, (b) Oblique flattening, and (c) Sspace flattening on individuals.*

the two operations through *interactive 3D rendering* (Figure 31(c)). Since 3D rendering employs orthographic projection and no shading, it is essentially an *oblique flattening* operation.

VISUAL-TimePacTS supports a more elaborate space-time cube operation that prevents visual marks from overlapping due to *flattening*. In Figure 31(c), for example, individuals are horizontally offset when several of them do the same activity at the same time. This technique is called *shearing* by the authors, and is further explained in Figures 31(d), 31(e). This technique is essentially a *(linear space cutting + space shifting)* + space flattening* operation, and is a hybrid between *space juxtaposing* and *space flattening*.

6.4. Cubix

Cubix is a system for analysing dynamic weighted networks through adjacency matrix representations [BPF14b]. A *3D rendering* provides an overview of the data (Figure 32(a)). Time goes from left to right. Each cell of the cube represents a connection between two nodes at a given time, with size depending on connection weight. Cells can be colour-coded according to time, weight or direction.

Cubix supports a range of space-time cube operations, including *time juxtaposing* (Figure 32(b)), *space juxtaposing* (detail in Figure 32(d)), *animated time cutting*, *animated space cutting*, *time flattening* and *space flattening*. For flattening operations, cells can be made translucent to visually aggregate the history of connections. Cubix also supports *semantic filtering* on connections based on their weight.

Cubix provides a control widget in the form of a stylized cube, and whose different parts can be clicked or dragged to switch between operations. All operation switches are explained using animated transitions through rotations of the *3D rendering* representation, or through staged animations of *extraction* and *rigid transformation* operations.

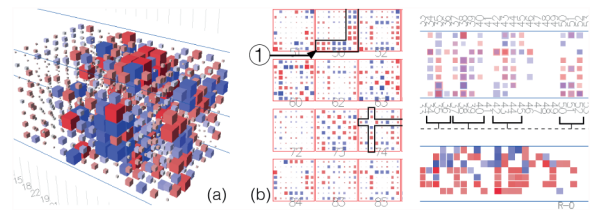


Figure 32: *Different operations applied to a time-evolving adjacency matrix in Cubix [BPF14b].*

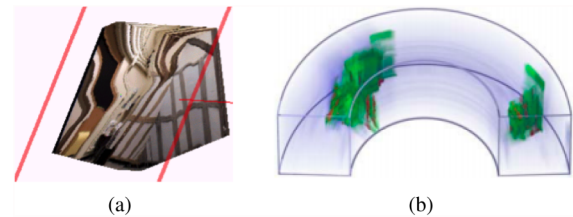


Figure 33: (a) *Video Cubism [FLM00]*; (b) *V3 [DC03]*.

6.5. Video cube systems

Several space-time cube systems have been proposed to support video analysis [MB98, FLM00, DC03, CI05]. Video Cubism [FLM00] uses a *3D rendering* representation together with an *interactive volume extraction* operation that is defined by manipulating a planar cutting plane (Figure 33(a)). Similarly, Khronos projector [CI05] supports manipulation of a non-planar cutting plane using touch or mid-air gestures (Figure 5). *V³* [DC03] (Volume Visualization for Videos) supports different operations, including *time juxtaposing* and a *3D rendering* view that can be combined with a *bending* operation (Figure 33(b)). *V³* also supports *filtering* operations that allow removal of pixels of a certain colour, or pixels that do not change across a given time period.

6.6. Wakame

Wakame [FW10] is a system for visualizing multivariate data by extruding starplots along a time dimension. The base plane of the space-time cube lays out starplots by their geographical location, although any layout could be imagined (Figure 34(a)). Colours correspond to dimension in the data.

Wakame supports three operations: *oblique flattening* (Figure 34(a)), *time flattening* (Figure 31(b)) and *time cutting*. In oblique flattening, extruded starplots can be rendered as traditional timeline charts as shown in Figure 31(b). The system provides an animated transition to the time-flattening where those timeline charts are shown aligned and in 2D. For time-cutting, the user selects a single time point on the complementary timeline and the system moves the camera orthogonal to the base plane.

Besides the space-time cube systems reviewed in this section, there is a wealth of general 3D visualization systems. Commercial and research tools exist in domains such as geovisualization (e.g. Voxler [Vox], ArcGIS [arc]), scientific visualization (e.g.

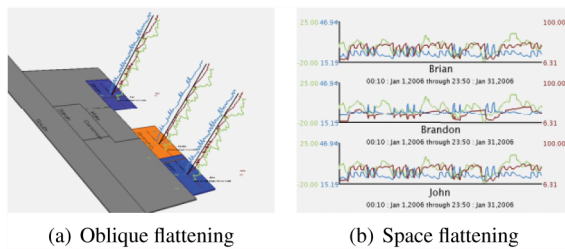


Figure 34: Wakame [FW10].

VTK [SAH00], Matlab [mat] and R [r]) and medical visualization [MTB03]. Although these tools do not treat time as a specific dimension, they can be used to inspire the design of interactive space-time cube systems.

7. General Discussion

We now discuss the limitations of our descriptive framework and consider areas for future research, including unifying our framework with the infovis pipeline model, extending it to other dimensionalities, discussing OLAP-cubes, considering non-planar media (e.g. physical visualizations), as well as discussing user tasks.

7.1. Comparison with the infovis pipeline model

Since our framework builds on the notion of sequential composition of operations, it shares similarities with another common model: the infovis reference model, also called the *infovis pipeline* [CMS99, Chi00, JD13]. The infovis pipeline sees visualization as a data-flow process, that is, a sequence of stages and transformations that turn raw data into a final image. These transformations commonly include *data transformation*, *visual mapping*, *presentation mapping* and *rendering* [JD13]. Interactivity is implemented by having data analysts alter these transformations at different stages.

There is clearly an analogy between transformations on the infovis pipeline and operations on a space-time cube, since both describe steps in creating a visualization. However, the infovis pipeline describes steps for the *implementation* of a visualization, that is, when and in which order data are transformed and mapped onto visual variables in a visualization. Space-time cube operations, on the other hand, are purely conceptual and do not imply any direct implementation; a slice from a cube could be seen as a data object, a graphical object, something in-between or something blending the two. Our model rather turns a 3D conceptual space into a 2D visualization space, regardless of how the visualization is eventually implemented on the infovis pipeline.

In fact, some space-time operations can be performed at different stages of the infovis pipeline, for example, *time flattening* can be done at the *data transformation* stage, by aggregating raw data over time. Alternatively, it could be emulated by explicitly rendering a 3D space-time cube on the screen and using a proper camera placement and projection transformation. In that case, it would be implemented at the *rendering* level. Frame buffers objects are another example as they can store data, while running on the GPU [ME09]. Both

models, the infovis pipeline and the generalized space-time cube, are complementary.

Finally, the infovis pipeline is a *general* model for visualizations, where the sequence of operations is *fixed*, but the operations themselves are rather *abstract*. For example, the infovis pipeline model provides no specific details about what happens in the *visual mapping* transformation. In contrast, our model only captures a *specific* family of visualizations (temporal visualizations), its sequence of operations is *not fixed* and the operations are more *concrete*. The infovis pipeline is more general but too high level to capture the similarities and differences between the visualizations we presented. On the other hand, our model is incomplete in that it does not define how the space-time cube is built.

To fully support a variety of space-time cube operations on the same dataset, the generalized space-time cube must be implemented as a first-class object, as done in many space-time cube systems (Section 6). Thus space-time cube operations are best seen as *presentation mapping* or *rendering* transformations, that is, transformations that turn the abstract visual form into a fully specified 2D image or 3D model [JD13]. In other terms, our space-time cube operations can be used to decompose and refine the *presentation mapping* transformation of the infovis pipeline. We believe that implementing our framework in this way, and potentially based on abstract inner structures of space-time cubes, could dramatically facilitate the exploration of a wide range of temporal datasets.

7.2. Other dimensionalities

This paper focused on temporal visualizations that involve *two dimensions in the base plane* plus time. These two dimensions can be inherently spatial or can result from 2D spatial encodings of non-spatial data. However, temporal visualizations with other dimensionalities exist or may be required for specific problems.

Most notably, a rich variety of temporal visualizations exist that involve a *single* spatial dimension plus time, for example, timelines and time-series visualizations [AMST11]. In principle, our framework still applies if the 3D space-time cube is turned into a 2D space-time plane. Operations analogous to our *geometry transformation* operations would capture techniques such as spiral visualizations, calendar visualizations or cycle plots [AMST11]. However, since a 2D space-time plane already naturally maps to a 2D planar display, and since the richness of time-series visualizations and timelines mostly stem from the visual encodings used, the usefulness of our framework would be less clear in this case.

Other temporal visualization techniques, although less common, show *three* spatial dimensions plus time. We believe most of them can be captured with operations on 4D space-time hypercubes. For example, Tufte explains how small multiples can be used to show the evolution of a 3D storm [BB95]. This approach amounts to applying a *time juxtaposition* operation on a space-time hypercube, where each *time cutting* operation yields a 3D image. Similarly, From-DaDy [HTC09] uses 3D trails to show the trajectories of airplanes in space. This technique amounts to performing a *time flattening* on a space-time hypercube. Extending our framework to higher data dimensionalities is an exciting topic for future research. However, it is less easy to imagine a hypercube than a cube, so the merits of

such a conceptual model still remain to be seen. Another solution to explore high-dimensional temporal data is to use dimensional reduction [BSH*16, vdeHBvW16].

7.3. OLAP cubes

Related to both implementation and higher dimensionalities of the generalized space-time cube are data cubes [GCB*97] in OLAP cubes. OLAP cubes consist of three data dimensions, selected freely from a potentially high dimensional dataset and do not have to involve a single temporal dimension. In this sense, OLAP cubes could be seen as superset to all space-time cubes. However, OLAP cubes are unable to capture ‘moving’ objects such as animals or nodes in a network (Figure 26, ① to ④), hence representing only a subset of all space-time cubes.

OLAP cubes are meant for querying and aggregating multi-dimensional data, rather than for visualization and exploration. They are pure logical models, and we are not aware of any visualization system that explicitly uses OLAP cubes. Some OLAP operations indeed share striking analogies with STC operations, namely *slice* (time and space cutting), *dice* (space and time chopping) and *roll-up* (space and time flattening). Other operations on OLAP cubes, not explicitly supported by space-time cubes, are *drill-down* (zooming in and effectively being the reverse of *flattening*) and *pivot* (reassigning dimensions of the cube). In supporting these operations, the spatial structure of the OLAP cube can vary and is not persistent, as in space-time cubes.

Future work could extend our framework to OLAP cubes. This would include (i) to differentiate between the two dimensions of the base plane, and consequently to (ii) extend the number of existing operations (e.g. enable cutting on each of the three cube dimensions). Finally, operations currently supported by OLAP cubes should be discussed for their general applicability to actual space-time cubes.

7.4. Non-planar media

Throughout this paper, we assumed the presentation medium to be planar. Although these are by far the most common, other display shapes are being explored in HCI, some of which are even deformable [RPPH12, HV08]. In these cases, the conditions for an operation to be *complete* are not the same. This opens up a wide range of possibilities for new visualization designs. For example, one implementation of the Khronos projector (Figure 5) employs back projection on a freely deformable cloth, allowing the use of *non-planar cutting* operations that are *complete*. In addition, physical visualizations make it possible to faithfully display 3D space-time cubes without any additional operations [JDF13, JD13]. Many such visualizations have been already crafted by scientists, artists and designers [DJ13].

Physical temporal visualizations can even be made modular to support interactive space-time cube operations. Figure 35(a) shows two physical representations of a dynamic network [BPF14b] made of laser-cut and laser-engraved acrylic. The left version supports interactive *time cutting* while the right version supports interactive *space cutting*. Cuts can be taken apart and manipulated freely, allowing for *time juxtaposing* and *space juxtaposing* as well as *time*

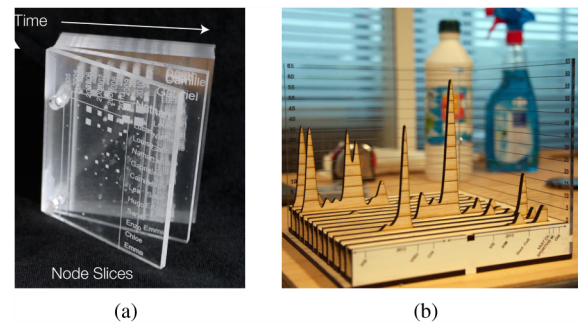


Figure 35: Examples of physical space-time cubes: (a) Mmatrix Ccube for dynamic network [BPF14b] (made by Benjamin Bach), (b) Ssubversion commits on different research projects over time [SSJ14].

flattening and *space flattening*, if viewed from a proper orthogonal angle and distance. Figure 35(b) shows SVN commits over time. Time goes from left to right, and different layers represent different projects. Slices can also be taken apart and reordered. For another example see [STB13].

7.5. Towards a prescriptive model

Our framework is mostly descriptive. A prescriptive model in turn can help us making implications about which operations to use for which data and tasks. Unfortunately, we need more empirical studies to assess the relative effectiveness of different space-time cube operations with respect to specific tasks as well as data types and characteristics. Although many studies have already been conducted (as discussed before), the current body of evidence paints a complex and partly inconclusive picture. This may be due to the lack of a clear and detailed descriptive framework. A clear descriptive framework allows us to better tease apart the effects of different subtle design features and to better control for confounds. We showed how many temporal visualization techniques can be decomposed into elementary operations. These operations can be combined in many ways or made dynamic at different levels, either through animation or interaction. The characterization of the inner structure of space-time cubes may already provide many elements to discuss the practical strengths and weaknesses of space-time cube operations, mostly based on well-established knowledge on perception and HCI, and on common sense.

Running studies for answering specific research questions will naturally remain important, and we hope our descriptive framework will facilitate the design of such studies and the discussion of findings in a more informative manner, leading to a more prescriptive framework in the future.

7.6. User tasks

In our discussion on the advantages and drawbacks of certain space-time cube operations, we have mostly focused on the inner structure of the space-time cubes as a source of trade-offs. However, the process of creating a successful visualization also has to take into account the user’s tasks. There have been efforts at proposing task

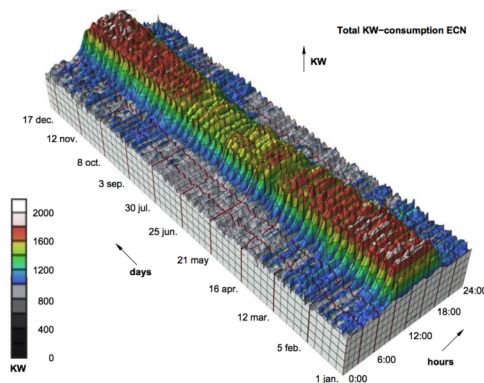


Figure 36: Energy consumption over time [vWvS99]: the horizontal axis is mapped to days while the other one is mapped to hours.

taxonomies that should complement our framework (e.g. [Mac95, AA05, Rot13, APS14, BPF14a]). However, we do not yet see a straightforward mapping between operations and the tasks they support best. As already mentioned, more effort is needed to come up with domain-agnostic task taxonomies.

7.7. Other limitations

Our framework is only meant to be a thinking tool. Like any model, it is necessarily incomplete. First, our taxonomy of elementary operations in Figure 24 is not meant to cover all possible operations since their number is potentially infinite and the taxonomy can easily be extended. Second, our framework does not provide much guidance for interaction design: the design space for interactive operations has only been partially explored in research. Finally, not all techniques for visualizing temporal data can be captured with space-time cubes. For example, temporal data can be visualized using two time axes instead of a single data axis (see Figure 36). Also, space-time cubes may not account for branching times. However, our framework is not meant to restrict creativity but rather to help visualization designers find new solutions, extend or generalize existing ones and think out of the box.

Our framework assumes that the space-time cube already exists. It does not provide guidance for producing the space-time cube itself, mainly because this process involves many decisions and should not restrict the range of possible techniques. For non-spatial data, many visual mappings can be used to produce individual slices, for example, locations on a map can yield values for altitude, temperature, rain, vegetation and soil type. How to visualize all these attributes at any particular point in time is a general problem of infovis, but the choice may also affect the efficiency of later space-time cube operations. We hope that future work can clarify some of these questions.

8. Conclusion

We reviewed various techniques to visualize temporal data, by describing them as sequences of parametric operations applied to a conceptual space-time cube. Our operations are independent from

the underlying data and can be applied across a range of application domains, be they cartography, dynamic network analysis, geopolitics or video analytics. When choosing the appropriate operations, a designer only needs to take into account the inner structure of the space-time cubes generated by the data, and the users' tasks.

Furthermore, by introducing domain-agnostic concepts and a clear terminology, this paper aims at facilitating the comparison of different approaches for visualizing temporal data. Existing visualizations from one data domain can be analysed in terms of elementary operations and then be adapted to other domains and problems. By giving a better vision of the richness of this design space, we hope our model will also motivate the exploration of new approaches. It stresses the importance of developing fully integrated interactive systems and toolkits that can support a range of techniques in a consistent manner.

Our model further aims at facilitating the design of studies and discussing their results in a more informative as well as formal manner. We hope the presented terminology and low-level concepts will help design better experiments that tease out important factors in dynamic data visualization. Controlled studies are needed to understand the trade-offs between different space-time cube operations and how they perform depending on task, data and the people using them.

This work mostly arose out of the need to teach temporal infovis to undergrad students. We therefore hope that it will help other people teach this field effectively, by providing a clear structure and a clear terminology on which to base higher level discussions and analyses.

Acknowledgements

We would like to thank the anonymous reviewers, as well as the presentation audience at EuroVis 2014 for their valuable feedback. Collaboration on this work was initiated by Dagstuhl Seminar 'Putting Data on the Map 12261', June 2012, and supported in part by NSERC, SMART Technologies, AITF, Surfnets and GRAND NCE, as well as by Clique Strategic Research Cluster funded by Science Foundation Ireland; (SFI) Grant No. 08/SRC/I1407.

References

- [AA04] ANDRIENKO N., ANDRIENKO G.: Interactive visual tools to explore spatio-temporal variation. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (New York, NY, USA, 2004), AVI '04, ACM, pp. 417–420.
- [AA05] ANDRIENKO N., ANDRIENKO G.: *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Springer-Verlag, New York, Inc., Secaucus, NJ, 2005.
- [AA12] ANDRIENKO N., ANDRIENKO G.: Visual analytics of movement: An overview of methods, tools and procedures. *Information Visualization* 12, 1 (2012), 3–24.
- [AAB*11] ANDRIENKO G., ANDRIENKO N., BAK P., KEIM D., KISILEVICH S., WROBEL S.: A conceptual framework and taxonomy of

- techniques for analyzing movement. *Journal of Visual Languages and Computing* 22, 3 (June 2011), 213–232.
- [AAH11] ANDRIENKO G., ANDRIENKO N., HEURICH M.: An event-based conceptual model for context-aware movement analysis. *International Journal on Geographic Information Science* 25, 9 (September 2011), 1347–1370.
- [AMM*07] AIGNER W., MIKSCH S., MÜLLER W., SCHUMANN H., TOMINSKI C.: Visualizing time-oriented data: A systematic view. *Computers and Graphics* 31, 3 (June 2007), 401–409.
- [AMST11] AIGNER W., MIKSCH S., SCHUMANN H., TOMINSKI C.: *Visualization of Time-Oriented Data* (1st edition). *Human-Computer Interaction*. Springer Verlag, 2011.
- [AP12a] ARCHAMBAULT D., PURCHASE H. C.: Mental map preservation helps user orientation in dynamic graphs. In *Proceedings of Graph Drawing* (Berlin-Heidelberg, 2012), Springer, GD '12, pp. 475–486.
- [AP12b] ARCHAMBAULT D., PURCHASE H. C.: The mental map and memorability in dynamic graphs. In *Proceedings of Pacific Visualization Symposium* (PacificVis, Songdo, Korea, 2012), IEEE '12, pp. 89–96.
- [AP13] ARCHAMBAULT D., PURCHASE H. C.: The 'Map in the mental map: Experimental results in dynamic graph drawing. *International Journal of Human-Computer Studies* 71, 11 (2013), 1044–1055.
- [AP16] ARCHAMBAULT D., PURCHASE H. C.: Can animation support the visualisation of dynamic graphs? *Information Sciences* 330 (2016), 495–509.
- [APP10] ARCHAMBAULT D., PURCHASE H. C., PINAUD B.: Difference map readability for dynamic graphs. In *Proceedings of Graph Drawing (GD 2010)* (Berlin-Heidelberg, 2010), vol. 6502 of *LNCS*, Springer, pp. 50–61.
- [APP11] ARCHAMBAULT D., PURCHASE H., PINAUD B.: Animation, small multiples, and the effect of mental map preservation in dynamic graphs. *IEEE TVCG* 17, 4 (2011), 539–552.
- [APS14] AHN J.-W., PLAISANT C., SHNEIDERMAN B.: A task taxonomy for network evolution analysis. *IEEE TVCG* 20, 3 (Mar. 2014), 365–376.
- [arc] ArcGIS. <http://www.esri.com/software/arcgis>. (accessed 02 April 2014).
- [AS] ARIS A., SHNEIDERMAN B.: NVSS: Network Visualization by Semantic Substrates. <http://www.cs.umd.edu/hcil/nvss/>. (accessed 02 April 2014).
- [ATMS*11] AHN J.-W., TAIEB-MAIMON M., SOPAN A., PLAISANT C., SHNEIDERMAN B.: Temporal visualization of social network dynamics: prototypes for nation of neighbors. In *Proceedings of International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction* (Berlin, Heidelberg, 2011), SBP'11, Springer-Verlag, pp. 309–316.
- [Bar97] BARTRAM L.: Perceptual and interpretative properties of motion for information visualization. In *Proceedings of the 1997 Workshop on New Paradigms in Information Visualization and Manipulation* (New York, NY, USA, 1997), ACM, pp. 3–7.
- [BB95] BAKER M. P., BUSHELL C.: After the storm: Considerations for information visualization. *IEEE Computer Graphics and Applications* 15, 3 (1995), 12–15.
- [BBL12] BOYANDIN I., BERTINI E., LALANNE D.: 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.
- [BC87] BECKER R. A., CLEVELAND W. S.: Brushing scatterplots. *Technometrics* 29, 2 (1987), 127–142.
- [BC03] BRANDES U., CORMAN S. R.: Visual unrolling of network evolution and the analysis of dynamic discourse. *Information Visualization* 2, 1 (March 2003), 40–50.
- [BCD*12] BORGO R., CHEN M., DAUBNEY B., GRUNDY E., HEIDEMANN G., HÖFERLIN B., HÖFERLIN M., LEITTE H., WEISKOPF D., XIE X.: State of the art report on video-based graphics and video visualization. *Computer Graphics Forum* 31, 8 (December 2012), 2450–2477.
- [BCH07] BRUNSDON C., CORCORAN J., HIGGS G.: Visualising space and time in crime patterns: A comparison of methods. *Computers, Environment and Urban Systems* 31, 1 (2007), 52–75. [Extracting Information from Spatial Datasets.]
- [BDA*14] BACH B., DRAGICEVIC P., ARCHAMBAULT D., HURTER C., CARPENDALE S.: A review of temporal data visualizations based on space-time cube operations]. In *Eurographics Conference on Visualization* (Swansea, Wales, United Kingdom, June 2014).
- [BDH04] BARTOLI A., DALAL N., HORAUD R.: Motion panoramas. *Computer Animation and Virtual Worlds* 15, (2004), 501–517.
- [BGPS07] BATTIATO S., GALLO G., PUGLISI G., SCELLATO S.: SIFT features tracking for video stabilization. In *Proceedings of Conference on Image Analysis and Processing (ICIAP 2007)*, '07, IEEE, pp. 825–830.
- [BHRD*15] BACH B., HENRY-RICHE N., DWYER T., MADHYASTHA T., FEKETE J.-D., GRABOWSKI T.: Small MultiPiles: Piling time to explore temporal patterns in dynamic networks. *Computer Graphics Forum* (2015).
- [BKH*16] BACH B., KERRACHER N., HALL K. W., CARPENDALE S., KENNEDY J., RICHE N. H.: Telling storied about dynamic networks with graph comics. In *Proceedings of ACM Conference on Human Factors in Information Systems* (2016), CHI, ACM, to appear.
- [BN11] BRANDES U., NICK B.: Asymmetric relations in longitudinal social networks. *IEEE TVCG* 17, 12 (December 2011), 2283–2290.

- [BPF14a] BACH B., PIETRIGA E., FEKETE J.-D.: GraphDiaries: Animated transitions and temporal navigation for dynamic networks. *IEEE TVCG* 20, 5 (2014), 740–754.
- [BPF14b] BACH B., PIETRIGA E., FEKETE J.-D.: Visualizing dynamic networks with matrix cubes. In *Proceedings of SIGCHI Conference on Human Factors and Computing Systems* (New York, NY, USA, 2014), CHI '14, ACM, pp. 877–886.
- [BSH*16] BACH B., SHI C., HEULOT N., MADHYASTHA T., GRABOWSKI T., DRAGICEVIC P.: Time curves: Folding time to visualize patterns of temporal evolution in data. *IEEE TVCG* 22, 1 (January 2016), 559–568.
- [BVB*11] BURCH M., VEHLW C., BECK F., DIEHL S., WEISKOPF D.: Parallel edge splatting for scalable dynamic graph visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2344–2353.
- [CAS05] CONTI G., AHAMAD M., STASKO J.: Attacking information visualization system usability overloading and deceiving the human. In *Proceedings of the 2005 Symposium on Usable Privacy and Security* (Pittsburgh, PA, 2005), ACM, pp. 89–100.
- [CCF96] COWPERTHWAIT D. J., CARPENDALE M. S. T., FRACCHIA F. D.: Visual access for 3D data. In *Proceedings on ACM CHI* (1996), pp. 175–176.
- [CCT*99] CARPENDALE S. T., COWPERTHWAIT D. J., TIGGES M., FALL A., FRACCHIA F. D.: The Tardis: A visual exploration environment for landscape dynamics. In *Proceedings of Conference on Visual Data Exploration and Analysis* No. 3643, (San Jose, 1999), pp. 110–119.
- [CFC*96] CARPENDALE S. T., FALL A., COWPERTHWAIT D. J., FALL J., FRACCHIA F. D.: Case study: Visual access for landscape event based temporal data. In *Proceedings of Visualization* (1996), IEEE Press Piscataway, NJ, USA.
- [Chi00] CHI E. H.-h.: A taxonomy of visualization techniques using the data state reference model. In *Proceedings of IEEE Symposium on Information Visualization* (Salt Lake City, UT, 2000), Infovis '00, IEEE, pp. 69–75.
- [CI05] CASSINELLI A., ISHIKAWA M.: Khronos projector. In *Proceedings of SIGGRAPH 2005* (2005).
- [CKN*03] COLLBERG C., KOBOUROV S., NAGRA J., PITTS J., WAMPLER K.: A system for graph-based visualization of the evolution of software. In *Proceedings of ACM Symposium on Software Visualization* (New York, NY, USA, 2003), SoftVis '03, ACM, pp. 77–ff.
- [CMS99] CARD S. K., MACKINLAY J. D., SHNEIDERMAN B. (Eds.): *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.
- [DC03] DANIEL G., CHEN M.: Video visualization. In *Proceedings of IEEE Visualization* (Washington, DC, USA, 2003), VIS '03, IEEE, pp. 409–416.
- [DE02] DWYER T., EADES P.: Visualising a fund manager flow graph with columns and worms. In *Proceedings of International Conference on Information Visualisation* (2002), IV '02, London, UK. IEEE, Los Alamitos, CA, pp. 147–152.
- [DG04] DWYER T., GALLAGHER D. R.: Visualising changes in fund manager holdings in two and a half-dimensions. *Information Visualization* 3, 4 (December 2004), 227–244.
- [DJ13] DRAGICEVIC P., JANSEN Y.: List of physical visualizations. <http://www.tinyURL.com/physvis>, 2013. [Accessed 04 September 2013].
- [DV10] DEMÄJÄR U., VIRRANTAU K.: Space-time density of trajectories: Exploring spatio-temporal patterns in movement data. *International Journal of Geographical Information Science* 24, 10 (2010), 1527–1542.
- [ELMS91] EADES P., LAI W., MISUE K., SUGIYAMA K.: Preserving the mental map of a diagram. In *Proceedings of Computer Graphics* (1991), Tokyo, Japan, Australian Computer Society, Inc. Darlinghurst, Australia, pp. 24–33.
- [FBS06] FALKOWSKI T., BARTELHEIMER B., SPILIOPOULOU M.: Mining and visualizing the evolution of subgroups in social networks. In *Web Intelligence, 2006. WI 2006. IEEE/WIC/ACM International Conference on* (Hong Kong, Los Alamitos, CA, 2006), IEEE, pp. 52–58.
- [FLM00] FELS S., LEE E., MASE K.: Techniques for interactive video cubism (poster session). In *Proceedings of International conference on Multimedia* (New York, NY, USA, 2000), MULTIMEDIA '00, ACM, pp. 368–370.
- [FQ11] FARRUGIA M., QUIGLEY A.: Effective temporal graph layout: A comparative study of animations versus static display methods. *Information Visualization* 10, 1 (2011), 47–64.
- [FT06] FELDMAN J., TREMOULET P. D.: Individuation of visual objects over time. *Cognition* 99, 2 (2006), 131–165.
- [Fur86] FURNAS G. W.: Generalized fisheye views. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 1986), CHI '86, ACM, pp. 16–23.
- [FW10] FORLINES C., WITTENBURG K.: Wakame: Sense making of multi-dimensional spatial-temporal data. In *Proceedings of the International Conference on Advanced Visual Interfaces* (New York, NY, USA, 2010), AVI '10, ACM, pp. 33–40.
- [GAA04] GATASKY P., ANDRIENKO N., ANDRIENKO G.: Interactive analysis of event data using space-time cube. In *Proceedings of the Information Visualisation, Eighth International Conference* (Washington, DC, USA, 2004), IV '04, IEEE, pp. 145–152.
- [GCB*97] GRAY J., CHAUDHURI S., BOSWORTH A., LAYMAN A., REICHAERT D., VENKATRAO M., PELLOW F., PIRAHESH H.: Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-totals. *Data Mining and Knowledge Discovery* 1, 1 (January 1997), 29–53.

- [geo] GeoTime. <http://www.geotime.com>. [Accessed 24 January 2014].
- [GEY12] GHANI S., ELMQVIST N., YI J. S.: *Perception of Animated Node-Link Diagrams for Dynamic Graphs*. The Eurographics Association and Blackwell Publishing Ltd, Eurographics, Aire-la-Ville, Switzerland, 2012.
- [GHW09] GROH G., HANSTEIN H., WOERNL W.: Interactively visualizing dynamic social networks with DySoN. In *Workshop on Visual Interfaces to the Social and the Semantic Web* (New York, NY, USA, February 2009), ACM.
- [GMH*06] GRIFFEN A. L., MACEachREN A. M., HARDISTY F., STEINER E., LI B.: 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.
- [Gre11] GRETCHEN P.: A Cartographer's Toolkit - Small Multiples. <http://www.gretchenpeterson.com/blog/small-multiples>, 2011. (Accessed 24 January 2014).
- [Har99] HARRIS R. L.: *Information Graphics: A Comprehensive Illustrated Reference*. Oxford University Press, New York, NY, 1999.
- [HEF*13] HURTER C., ERSOY O., FABRIKANT S., KLEIN T., TELEA A.: Bundled visualization of dynamic graph and trail data. *IEEE TVCG*, 99 (2013), 1141–1157.
- [HET12] HURTER C., ERSOY O., TELEA A.: Graph bundling by kernel density estimation. *Computer Graphics Forum* 31 (June 2012), 865–874.
- [HKV12] HU Y., KOBouROV S. G., VEERAMONI S.: Embedding, clustering and coloring for dynamic maps. In *Proceedings of Pacific Visualization Symposium* (Songdo, South Korea, 2012), PacificVis '12, IEEE, pp. 33–40.
- [HSCW13] HADLAK S., SCHUMANN H., CAP C. H., WOLLENBERG T.: Supporting the visual analysis of dynamic networks by clustering associated temporal attributes. *IEEE TVCG* 19, 12 (2013), 2267–2276.
- [HTC09] HURTER C., TISSOIRES B., CONVERSY S.: FromDaDy: Spreading aircraft trajectories across views to support iterative queries. *IEEE TVCG* 15, 6 (2009), 1017–1024.
- [HV08] HOLMAN D., VERTEGAAL R.: Organic user interfaces: Designing computers in any way, shape, or form. *Communications of the ACM* 51, 6 (2008), 48–55.
- [Hä70] HÄGERSTRAND T.: What about People in regional science. In *Papers of the Regional Science Association* (Angra do Heroísmo, Azores, 1970), Regional Science Association International, pp. 7–24.
- [JD13] JANSEN Y., DRAGICEVIC P.: An interaction model for visualizations beyond the desktop. *IEEE TVCG* 19, 12 (2013), 2396–2405.
- [JDF13] JANSEN Y., DRAGICEVIC P., FEKETE J.-D.: Evaluating the efficiency of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France, 2013), CHI '13, ACM, pp. 2593–2602.
- [KDA*09] KRISTENSSON P., DAHLBACK N., ANUNDI D., BJORNSTAD M., GILLBERG H., HARALDSSON J., MARTENSSON I., NORDVALL M., STAHL J.: An evaluation of space time cube representation of spatiotemporal patterns. *IEEE TVCG* 15, 4 (July 2009), 696–702.
- [KK12] KVELADZE I., KRAAK M.-J.: What do we know about the space-time cube from cartographic and usability perspective? In *Proceedings of the 2012 AutoCarto International Symposium on Automated Cartography* (Columbus, OH, 2012), AuroCarto'12.
- [KPW13] Kerren A., Purchase H. C., Ward M. O.: (Eds.) *Multivariate Network Visualization, Vol. 8380 of LNCS*, (Berlin–Heidelberg, 2013), Springer.
- [Kra03] KRAAK M. J.: The space-time cube revisited from a geovisualization perspective. *Proceedings of the International Cartographic Conference* (2003), pp. 1988–1996.
- [KW04] KAPLER T., WRIGHT W.: GeoTime Information Visualization. In *Proceedings of IEEE Information Visualization* (Washington, DC, USA, 2004), INFOVIS '04, IEEE, pp. 25–32.
- [LNS11] LIU Z., NAVATHE S., STASKO J.: Network-based visual analysis of tabular data. In *Proceedings of Conference on Visual Analytics Science and Technology* (2011), VAST '11, pp. 41–50.
- [Mac95] MACEachREN A. M.: *How Maps Work*. Guilford Press, 1995.
- [Mar78] MAREY É.-J.: *La méthode graphique dans les sciences expérimentales et particulièrement en physiologie et en médecine*. G. Masson, 1878.
- [mat] MatLab. <http://www.mathworks.fr/products/matlab/>. (Accessed 02 April 2014).
- [MB98] MACKAY W., Beaudouin-Lafon, M.: DIVA: Exploratory Data Analysis with Multimedia Streams. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems* (April 1998), CHI '98, ACM, pp. 416–423.
- [McC94] MCCLOUD S.: *Understanding Comics*, HarperCollins, New York, NY, 1994.
- [ME09] McDONNELL B., ELMQVIST N.: Towards utilizing gpus in information visualization: A model and implementation of image-space operations. *IEEE TVCG (Proc. InfoVis 2009)* 15, 6 (2009), 1105–1112.
- [MELS95] MISUE K., EADES P., LAI W., SUGIYAMA K.: Layout adjustment and the mental map. *Journal of Visual Languages and Computing* 6 (1995), 183–210.
- [MR97] MATTHEWS G., ROZE M.: Worm plots. *Computer Graphics and Applications, IEEE* 17, 6 (November 1997), 17–20.

- [MTB03] MCGUFFIN M. J., TANCAU L., BALAKRISHNAN R.: Using deformations for browsing volumetric data. In *Proceedings of IEEE Visualization* (October 2003), VSI '03, pp. 401–408.
- [Mun15] MUNZNER T.: *Visualization Analysis & Design*. CRC Press, Boca Raton, FL, 2015.
- [Muy87] MUYBRIDGE E.: *Animal Locomotion*. Published under the auspices of the University of Pennsylvania, 1887.
- [PHG06] PURCHASE H. C., HOGGAN E., GÖRG C.: How important is the “Mental Map”? An empirical investigation of a dynamic graph layout algorithm. In *Proceedings of Graph Drawing* (Berlin Heidelberg, Germany, 2006), Springer, vol. 4372 of *GD '06*, pp. 184–195.
- [PJH05] PETROVIC N., JOJIC N., HUANG T. S.: Adaptive video fast forward. *Multimedia Tools and Applications* 26, 3 (2005), 327–344.
- [PS08] PURCHASE H. C., SAMRA A.: Extremes are better: Investigating mental map preservation in dynamic graphs. In *Proceedings of Diagrams* (Berlin Heidelberg, Germany, 2008), Springer, Vol. 5223 of *Diagrams '08*, pp. 60–73.
- [r] R. <http://www.r-project.org/>. (Accessed: 02 April 2014).
- [Ren02] RENSINK R. A.: Change detection. *Annual Review of Psychology* 53, 1 (2002), 245–277.
- [RFF*08] ROBERTSON G., FERNANDEZ R., FISHER D., LEE B., STASKO J.: Effectiveness of animation in trend visualization. *IEEE TVCG* 14, 6 (Nov. 2008), 1325–1332.
- [RM13] RUFIANGE S., MCGUFFIN M. J.: DiffAni: Visualizing dynamic graphs with a hybrid of difference maps and animation. *IEEE TVCG* 19, 12 (2013), 2556–2565.
- [Ros06] ROSLING H.: Stats that reshape your worldview. http://www.ted.com/talks/hans_rosling_the_good_news_of_the_decade.html, (2006). (Accessed: 24 January 2014).
- [Rot13] ROTH R.: An empirically-derived taxonomy of interaction primitives for interactive cartography and geovisualization., *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (December 2013), 2356–2365.
- [RPPH12] RASMUSSEN M. K., PEDERSEN E. W., PETERSEN M. G., HORNBAEK K.: Shape-changing interfaces: a review of the design space and open research questions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, 2012), ACM, pp. 735–744.
- [SA06] SHNEIDERMAN B., ARIS A.: Network visualization by semantic substrates. *IEEE TVCG* 12, 5 (2006), 733–740.
- [SAH00] SCHROEDER W. J., AVILA L. S., HOFFMAN W.: Visualizing with VTK: A tutorial. *Computer Graphics and Applications, IEEE* 20, 5 (2000), 20–27.
- [Shn03] SHNEIDERMAN B.: Why not make interfaces better than 3D reality? *Computer Graphics and Applications, IEEE* 23, 6 (2003), 12–15.
- [SP08] SAFFREY P., PURCHASE H. C.: The “Mental Map” versus “Static Aesthetic” compromise in dynamic graphs: A user study. In *Proceedings of Australasian User Interface* (2008), Vol. 76 of *AUIC '08*, Australian Computer Society, pp. 85–93.
- [SRdJ05] SHANBHAG P., RHEINGANS P., DES JARDINS M.: Temporal visualization of planning polygons for efficient partitioning of geo-spatial data. In *Proceedings of the 2005 IEEE Symposium on Information Visualization* (Los Alamitos, CA, 2005), INFOVIS '05, IEEE, pp. 211–218.
- [SSJ*14] SWAMINATHAN S., SHI C., JANSEN Y., DRAGICEVIC P., OEHLBERG L. A., FEKETE J.-D.: Supporting the design and fabrication of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), CHI '14, ACM, pp. 3845–3854.
- [SSS*01] SMITH G., STUERZLINGER W., SALZMAN T., WATSON B., BUCHANAN J.: 3D scene manipulation with 2D devices and constraints. *Graphics Interface 1* (2001), 135–142.
- [STB13] STUSAK S., TABARD A., BUTZ A.: Can physical visualizations support analytical tasks? *Proceedings of IEEE VIS*, poster abstracts, 2013.
- [TBC13] THUDT A., BAUR D., CARPENDALE S.: Visits: A spatiotemporal visualization of location histories. In *EuroVis: Short Papers* (2013), The Eurographics Association.
- [TGF08] TANG A., GREENBERG S., FELS S.: Exploring video streams using slit-tear visualizations. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (New York, NY, USA, 2008), AVI '08, ACM, pp. 191–198.
- [TKB07] TURDUKULOV U. D., KRAAK M.-J., BLOK C. A.: Designing a visual environment for exploration of time series of remote sensing data: In search for convective clouds. *Computers & Graphics* 31, 3 (2007), 370–379.
- [TMB02] TVERSKY B., MORRISON J., BETRANCOURT M.: Animation: Can it facilitate? *International Journal of Human-Computer Studies* 57, 4 (2002), 247–262.
- [Tod04] TODD J. T.: The Visual Perception of 3D Shape. *Trends in Cognitive Sciences* 8, 3 (2004), 115–121.
- [TR09] THAKUR S., RHYNE T.-M.: Data vases: 2d and 3d plots for visualizing multiple time series. In *Advances in Visual Computing*. G. Bebis, R. Boyle, B. Parvin, D. Koracin, Y. Kuno, J. Wang, R. Pajarola, P. Lindstrom, A. Hinkenjann, M. Encarnação, C. Silva, D. Comin (Eds.), vol. 5876 of *Lecture Notes in Computer Science*. (Berlin, Heidelberg, 2009), Springer, pp. 929–938.
- [TS12] TOMINSKI C., SCHULZ H.-J.: The great wall of space-time. *Vision, Modeling, and Visualization* (Aire-la-Ville, Switzerland, 2012), The Eurographics Association.

- [TWS05] TOMINSKI C., SCHULZE-WOLLGAST P., SCHUMANN H.: 3D information visualization for time dependent data on maps. In *Proc. of IEEE Information Visualisation* (Los Alamitos, CA, July 2005), IEEE, pp. 175–181.
- [Tuf] TUFTE E.: Megan Jaegerman's brilliant news graphics. http://www.edwardtufte.com/bboard/q-and-a-fetch-msg?msg_id=0002w4. (Accessed: 12 April 2014).
- [Tuf86] TUFTE E. R.: *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, USA, 1986.
- [TV07] TRUONG B. T., VENKATESH S.: Video abstraction: A systematic review and classification. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMCCAP)* 3, 1 (2007), 3.
- [vdEHBvW13] VAN DEN ELZEN S., HOLTEN D., BLAAS J., VAN WIJK J.: Dynamic network visualization with extended massive sequence views. *IEEE TVCG PP*, 99 (2013), 1–1.
- [vdEHBvW16] VAN DEN ELZEN S., HOLTEN D., BLAAS J., VAN WIJK J.: Reducing snapshots to points: A visual analytics approach to dynamic network exploration. *Visualization and Computer Graphics, IEEE Transactions on* 22, 1 (2016), 1–10.
- [VFC10] VROTSOU K., FORSELL C., COOPER M.: 2D and 3D representations for feature recognition in time geographical diary data. *Information Visualization* 9, 4 (December 2010), 263–276.
- [Vox] Voxler. <http://www.goldensoftware.com/products/voxler>. (Accessed: 02 April 2014).
- [VWD04] VIÉGAS F. B., WATTENBERG M., DAVE K.: Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2004), CHI '04, ACM, pp. 575–582.
- [vWvS99] VAN WIJK J. J., VAN SELOW E. R.: Cluster and calendar based visualization of time series data. In *Proceedings of IEEE Symposium on Information Visualization* (Washington, DC, USA, 1999), INFOVIS '99, IEEE, pp. 4–.
- [War12] WARE C.: *Information Visualization: Perception for Design* (Orlando, FL, 2012), Elsevier.
- [WB05] WARE C., BOBROW R.: Supporting visual queries on medium sized node-link diagrams. *Journal of Information Visualization* 4, 1 (2005), 49–58.
- [WFG92] WANGER L. R., FERWERDA J., GREENBERG D. P.: Perceiving spatial relationships in computer-generated images. *IEEE Computer Graphics and Applications* 12, 3 (1992), 44–58.
- [Wik13] Wikimedia: Stroke order project. <http://www.tinyurl.com/strokeOrderProject>, 2013. (Accessed 24 January 2014).
- [WvdWvW09] WILLEMS N., VAN DE WETERING H., VAN WIJK J. J.: Visualization of vessel movements. In *Proceedings of the Conference on Visualization* (2009), EuroVis'09, Eurographics Association, pp. 959–966.
- [WvdWvW11] WILLEMS N., VAN DE WETERING H., VAN WIJK J. J.: Evaluation of the visibility of vessel movement features in trajectory visualizations. In *Proceedings of the Conference on Visualization* (2011), EuroVis'11, Eurographics Association, pp. 801–810.