

Watch This: A Taxonomy for Dynamic Data Visualization

Joseph A. Cottam*
Indiana University

Andrew Lumsdaine†
Indiana University

Chris Weaver‡
University of Oklahoma

ABSTRACT

Visualizations embody design choices about data access, data transformation, visual representation, and interaction. To interpret a static visualization, a person must identify the correspondences between the visual representation and the underlying data. These correspondences become moving targets when a visualization is dynamic. Dynamics may be introduced in a visualization at any point in the analysis and visualization process. For example, the data itself may be streaming, shifting subsets may be selected, visual representations may be animated, and interaction may modify presentation. In this paper, we focus on the impact of dynamic data. We present a taxonomy and conceptual framework for understanding how data changes influence the interpretability of visual representations. Visualization techniques are organized into categories at various levels of abstraction. The salient characteristics of each category and task suitability are discussed through examples from the scientific literature and popular practices. Examining the implications of dynamically updating visualizations warrants attention because it directly impacts the interpretability (and thus utility) of visualizations. The taxonomy presented provides a reference point for further exploration of dynamic data visualization techniques.

Keywords: Dynamic Data, Interpretation.

Index Terms: H.5.1 [Information Systems]: Information Interfaces and Presentation—Multimedia Information Systems; I.3.6 [Computing Methodologies]: Computer Graphics—Methodologies and Techniques

1 INTRODUCTION

Dynamic visualizations—visualizations that change over time—are increasingly common. The way that a visualization actually changes impacts how it can be interpreted, both immediately and over time. Simple decisions, such as choosing to modify a color scale in response to updated data, drastically change what the visualization intuitively reveals. Such decisions can be the difference between building insights and misleading. Understanding and controlling visual dynamics requires an appreciation of the origin and manifestation of underlying data dynamics [22]. This paper presents a taxonomy of dynamic visualization techniques.

Simply put, “dynamics” are changes over time. In the pipeline of the information visualization reference model [12], dynamics can arise at any stage (see Figure 1). User interaction for navigation or to build dynamic queries are two of the most common sources of dynamics in visualization tools. However, the underlying data itself may change as well. These changes may be in response to an external event, such as a sensor update, resulting in streaming updates. Changes may also be in response to modified requirements of the visualization with entirely new data sources seen as relevant.

When a visualization changes, regardless of the reason, that change must be considered in its interpretation. Choosing how the

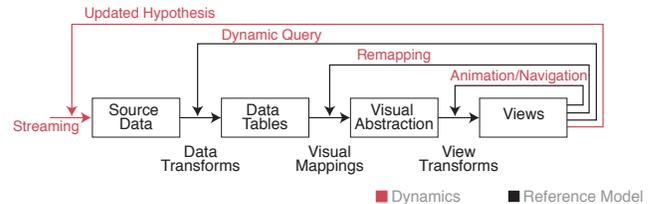


Figure 1: Information visualization reference model [12] and sources of dynamics at each stage.

visualization changes (i.e., in what way a visualization is dynamic) is an important part of the overall visualization design.

Visualization dynamics is too broad a topic to address as a single unit. Important differences are found between user-driven and data-driven changes. This paper focuses on dynamic data, specifically the effects of streaming data on a visualization. Streaming data is important (1) because it is growing in prevalence, (2) because data changes are part of the InfoVis Reference Model, and (3) because it affects techniques further down the pipeline.

Bertin’s description of visualization techniques decomposes the visual space into two groups: Spatial and Retinal variables [9]. Spatial variables determine positions, while retinal variables cover the other visual aspects. ‘Spatial and retinal variables are the basic building blocks of data visualization. Data is encoded using these variables. Dynamic visualizations change over time by modifying the values of these variables.

Our taxonomy organizes the dynamics possible in a streaming-data visualization into four retinal and four spatial categories (Section 2). This organization is the basis for the presented taxonomy of dynamic visualizations (Section 3). A useful description of the role of element identity directly results from the retinal and spatial discussion (Section 4). Guidelines for selecting a dynamic representation (Section 5) and discussion of several existing visualizations (Section 6) are also provided.

2 TAXONOMY DIMENSIONS

Prior on classifying static visualizations provides a foundation for classifying time-varying, dynamic visualizations. This taxonomy decomposes visualizations using distinctions made by Bertin: spatial and retinal variables [9]. The spatial variables are the coordinates in space: X and Y for two-dimensional visualizations. Retinal variables are the size, value, orientation, texture, hue and shape of a visual element. Unlike Bertin, the taxonomy presented in this paper is not concerned with the individual properties in either dimension. Instead, it examines how these variations can be presented within each group and the implications of different variation styles.

2.1 Spatial Variable Treatments

Spatial dimensions describe the position of an element. Position and quantity are intrinsically related because all visual elements must have a position. The fundamental positional information consists of the X and Y position; Z position is a direct extension that can be applied as needed (e.g., for layering). Positional information is considered relative to a canonical “registration point,” not regard-

*jcottam@indiana.edu

†lums@osl.iu.edu

‡weaver@cs.ou.edu

ing the entire extent of the element (which is further influenced by size and rotation, see Section 2.4).

Spatial dimension may be dynamic in two ways: changing existing values and changing the number of visual elements (adding or deleting). These two general change types can be expressed as four general categories:

Fixed: The spatial dimensions do not change. The number of visual elements is fixed at the start of the visualization. Visual elements do not move.

Mutable: The number of elements remains fixed throughout the visualization process, but the location of elements may change over time. (Destructive updates are referred to as “mutation” of that entity.)

Create: New elements may be created in response to incoming data. Existing element positions may be mutated.

Create & Delete: Elements may be created or deleted in response to incoming data. Mutability is implicit in this category because CREATE & DELETE can be used to simulate mutation to retinal attributes.

Mutation of spatial position generally falls into two styles. Both styles of mutation can cause significant spatial changes over time but are not significant at the level of abstraction of this taxonomy. However, a detailed investigation of any given visualization may benefit from noting the degree to which each case applies because they represent different types of relationships. The first style is a change to reflect a changing context. This style is common when the spatial information is not directly tied to underlying data, such as graph layout. When the context changes, visual elements move but that does not imply a change to the related data, only the context. The second style of spatial mutation occurs when spatial positions are directly tied to underlying data. In this case, entity movement directly reflects data changes tied to the given element. Most statistical graphics employ spatial layouts of the second type, and thus have movement dynamics in this second category.

2.2 Retinal Variable Treatments

Bertin’s retinal dimensions are size, value (light vs. dark), orientation (or rotation), texture (or patterning), hue (or color), and shape. Retinal dimensions do not affect the number of individual elements (as position does), but they can affect how elements are perceived to group. Retinal values change over time as the attribute they are associated with do. The user of a visualization employing retinal encoding must interpret the degree of change as a corresponding value change. Depending on how much change is permitted, the complexity of this interpretation changes. The categories of retinal dimension dynamics, in increasing order of complexity, are:

Immutable: Retinal variables are set and left unchanged.

Known Scale: The scale is fixed, but future data may change the retinal presentation of an existing elements. The known scale implies that all values that may be presented are covered by the scale’s regular divisions, which are established when the visualization is initialized. For categorical encodings, known scale means that the number and order of categories is known in advance. For continuous encodings, the range of values is provided in advance and divided into regular intervals.

Extreme Bin: An extreme bins scale is a known scale with sentinel categories (usually at the endpoints). These sentinel categories may be due to an unknown actual extent of values, special/sentinel categories, or to provide details in a specific sub-range. Values outside of the regular range are assigned to

a top or bottom catch-all “bin.” Scales like this are typically include labels such as “100+”, “less than 0”, “No Data” or “Error”. These open-ended categories are distinguished and outside of the normal divisions of the scale.

Mutable Scale: Updates may change the representation of an element and the mapping function itself. In this category, scales may grow or shrink dynamically to accommodate the data being visualized. For example, this category may apply if the number of levels in a categorical variable is not known in advance. When modifying scales, a single new data point may cause existing visual elements to change representations, even though the underlying data did not change.

Retinal dimensions are closely related to gestalt principals [44], and thus do not include symbolic meaning. This distinction indicates that letters, words, geographic outlines of countries, etc. are not visual elements in-and-of themselves. Interpreting words or elements “from the real world” depends strongly on cultural context. In contrast, the ability to recognize variation in retinal variables is less influenced by cultural context [44]. Retinal dimensions have inherent bandwidth. Appropriate limitations are assumed to remain constant in the face of incremental updates and assumed to be respected in the visualization.

2.3 Identity

Comparison is a fundamental task of visualization interpretation [1, 9,41]. However, a visualization that includes dynamics complicates the task of comparison by inviting comparisons to be made across time. This taxonomy examines how dynamics influence the task of comparison over time. Identity can be constructed in spatial or retinal variables or in combinations of the two.

Comparisons in dynamic visualizations can be divided into two rough categories: identity based and nonidentity based. Identity-based comparisons rely on the ability to recognize an element as representing the same real-world entity at multiple times. For example, in a social network, identity based comparisons require the ability to recognize that a given node represents the same person as the visualization changes. Identities may be low-level (e.g., recognizing individual people) or high-level (e.g., recognizing large groups of people that live in the same state). High-level identity is referred to as “group identity.”

Dynamics in a visualization make identity potentially fragile. The more an element changes, and the more that change happens in characteristics with strong visual salience, the more care must be taken if identity must be preserved. Identity preservation can occur at different time scales. In general, identity preservation enables detailed comparisons both within and between different time steps. Long-term identity preservation enables comparisons at arbitrary time differences. However, short-term identity preservation can be used to communicate incremental changes.

Non-identity based comparisons do not enable the fine-grained observations that identity-based comparisons do. However, they do allow observations of distributions (within time steps) and observation of distributional changes (between time steps). Non-identity based comparisons have fewer time-scale issues [42], so small and large time changes affect interpretation less.

2.4 Special Considerations

Because the taxonomy is focused on the influence of dynamic data, data-derived visual attributes are the principal consideration. Non-data derived attributes are assumed to be constant in this taxonomy. Since any change to the representation can introduce interpretation issues comparable to those driven by data, it is important to note that this assumption is not strictly true. Two major circumstances generally introduce changes to non-data derived attributes: graphic design considerations and user interaction. Changes predicated on

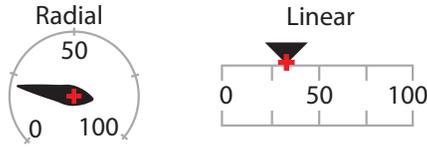


Figure 2: Linear and radial meters with registrations points marked.

graphic design considerations, such as resizing text to avoid overlapping labels, can introduce interpretation issues comparable to those driven by data. However, such events are relatively rare and thus not considered significant in this taxonomy. If the taxonomy focused on interaction or remapping (see Figure 1), such changes would be more significant.

Registration and transparency introduce ambiguities with respect to quantity and position. Treating visual elements as if they were on exact points is an abstraction. Visual elements have spatial extent and therefore contain more than one X/Y position. The canonical position of the visual element is called its “registration point.” Rotation and size changes may induce a motion-like effect without changing the registration point. The difference between radial and linear meters illustrates this difference (see Figure 2). The pointer on a radial meter maintains a fixed registration point, though value changes drive rotation and thus extent changes. In contrast, value changes shown on a linear meter drive changes to the registration point itself. Our taxonomy only considers changes to the registration point as spatial changes. Extent changes are treated as retinal variable changes.

Treatment of transparency depends on the eventual visibility. If the element is always visible, transparency change is treated as a retinal variable change. However, if the visual element is made fully transparent, then it is treated as element deletion.

User interaction is a complex issue and generally deferred to future work (see Section 8). The core difficulty with user interaction is that changes in response to user actions are typically semantically distinct from those in response to source data. On the surface, a user can be treated as any other dynamic data source. User induced changes are no more nor less extensive than data-induced changes. However, the user induced changes imply a shifting mental state in advance of the visualization and an intentional connection to the change. These are the opposite conditions to those found from other data source dynamics. The distinct connection between user interaction and interpretation means that this taxonomy is less applicable to such changes than to data-driven dynamics.

Different retinal dimensions have different properties, but this taxonomy aggregates all retinal dimensions. This simplification focuses the taxonomy on issues above the level of individual retinal dimension variation. It is considered safe because variation between retinal dimensions is generally of degree and not of kind. For example, eight colors is a practical maximum to use in many contexts, despite a theoretical maximum of two million [11]. Similar maxima exist for other retinal dimensions [44]. Additionally, we only consider the retinal dimension with the greatest variation if multiple levels of variation are used in a visualization.

Many visualizations exhibit different behaviors at different phases. Initialization commonly exhibits distinct behaviors. When categorizing visualizations, we focus on the behavior after initialization but do not make further phase distinctions.

3 TAXONOMY

Visualizations can be organized according to the dynamics present in spatial and retinal dimensions. Four divisions of variability for both spatial and retinal variability were identified and described in detail in Section 2. In brief, the spatial properties of a visual element may be (1) fixed at the start (FIXED), (2) mutable over the

		Retinal Categories			
		Immutable	Known Scale	Extreme Bins	Mutable Scale
Quantity/Spatial Categories	Fixed	Identity Preserving			
	Mutate		Transitional		
	Create		Transitional		
	Create & Delete	Immediate			

(a) Technique Matrix

Spatial × Retinal	Description
F × I	No dynamics present
F × KS	Elements are statically positioned with dynamic retinal encodings.
F × EB	Retinal encodings with details in a limited range of the data.
F × MS	Retinal encodings where the value and the value’s scale convey dynamic information.
M × I	Mobile elements with constant appearance.
M × KS	Elements may move and their appearance may change in <i>a priori</i> known ways.
M × EB	Moving elements with emphasized retinal range.
M × MS	A fixed number of elements, but retinal encodings and scales are dynamic.
C × I	Elements may be created with fixed appearance.
C × KS	Created elements are always retained, but their appearance may change over time.
C × EB	Element creation with emphasized retinal range.
C × MS	Number of elements can increase, their appearance (including the scales) change with data.
CD × I	The set of elements can grow or shrink, but appearance is fixed.
CD × KS	Elements appear and disappear; existing elements can change appearance.
CD × EB	Elements appear and disappear, with representation providing details within a range.
CD × MS	Quantity, retinal encoding, and scales can change.

(b) Technique Descriptions

Figure 3: Retinal vs. Spatial Mutability matrix and brief description of technique categories. Descriptions are labeled with abbreviated spatial and retinal categories.

course of the visualization (MUTABLE), (3) new visual elements may be added (CREATE) or (4) existing spatial elements may be deleted (CREATE & DELETE). Spatial modification in options 3 or 4 also changes the quantity of elements in a visualization. Four divisions of retinal variability were also introduced: (1) set at creation time (IMMUTABLE), (2) mutable on a fully-defined and regular scale (KNOWN SCALE), (3) mutable on a defined scale with distinguished ‘bins’ (EXTREME BINS) and (4) scales may update in response to new data (MUTABLE SCALE).

Since spatial and retinal variables are orthogonal to each other, the degrees of variability in each are also orthogonal. Combined, they define a 4-by-4 matrix of technique categories (see Figure 3). Visualizations within a technique category share identity-related properties, and neighboring categories generally represent incremental changes to those properties. This section describes each of the technique categories, including examples in existing visualizations and a discussion of identity properties.

3.1 Dynamic Technique Categories

The taxonomy matrix includes sixteen technique categories. This section provides a more detailed description of each category.

FIXED×IMMUTABLE (F×I)

Visualizations that are fixed in spatial dimensions and immutable in retinal dimensions are static. This is the degenerate case of dynamic visualization. From the standpoint of identity, visualizations of this class have strong identity characteristics and thus enable arbitrary comparisons, but really only ever show one time step.

FIXED×KNOWN SCALE (F×KS)

A fixed spatial representation can present a stable reference system, on which retinal dimensions indicate changing information as an overlay. A road atlas is a static version of such a system, with road positions as the reference and road type (divided highway, highway, unpaved road, etc.) as retinal annotations. A dynamic traffic map [36] extends this idea to dynamic data: as the road congestion changes, so does the coloring (i.e., retinal encoding) of the road segment. In such maps, the road positions do not change, so identity is tied directly to position. Overlay of symbolic data in a heads-up display works in a similar fashion, but the reference system is the real world [3]. A “map of science” can form a metaphorical substrate for other data to be encoded on the retinal variables [10]. Staying within a known scale for retinal variation can be achieved in post-hoc displays of historical information or using scales in a historically or analytically derived safe range.

FIXED×EXTREME BINS (F×EB)

Combining extreme bins with a fixed reference system provides a stable presentation of potentially complex data. A temperature map is a common form in this category. The map coloring corresponds to temperature. However, when the color enters the highest or lowest bins, the only safe statement is that temperature exceeded the normal scale. FIXED×EXTREME BINS visualizations are useful when the entities are known (such as frequency bins in a music equalizer or physical locations on a map) and a subset of possible range is interesting or common. Such visualizations may also provide a focus+context effect [37]. The focus area is a sub-range of the input values, but the context is provided by the extreme bins.

FIXED×MUTABLE SCALE (F×MS)

At any give time, visualizations in the FIXED×MUTABLE SCALE category appear like those in other FIXED categories. However, in this category, the retinal scale itself changes to accommodate the data. For example, if color used to communicate a range of values, the color of an element indicates the relative position of the value in a range *and* the scale communicates what that range actually is. In general, visualizations using a mutable scale must be interpreted as a composite of the value displayed and the scale it is displayed on. If value-vs-extrema (e.g., min or max) is important, this category is appropriate. However, if absolute values are important, visualizations in this category may be more difficult to interpret.

MUTATE×IMMUTABLE (M×I)

Visualizations employing spatial change have objects which may move around but otherwise have fixed appearance. Visualizations in this category are effective when the most important aspect of comparison naturally maps to a coordinate space, such as spatial tracking of specific entities. When used for tracking, identity is established by the retinal characteristics. This places a practical limit on the number of distinguishable elements or element categories (as was done in [16]). Visualization in this category can be used to view distributions of large quantities of elements as well (see NY Taxi cabs [34] and the Sexperience 1000 [35]).

For visualizations where spatial mutation is possible but entity deletion is not, a common technique is to use a distinguished location to indicate “out of selection” or “other” (used by *We Feel Fine* [21] and Sexperience 1000 [35]). Having a distinguished location may introduce spatial non-uniformity, (e.g., if the spatial scales are otherwise linear). This can lead to mis-interpretation if the distinguished location is not properly labeled.

MUTATE×KNOWN SCALE (M×KS)

Visualizations in this category can move elements in space and modify their properties as they do so. Such visualizations generally preserve element identities in the short-term (e.g., if spatial movement is done slowly). However, because position and retinal values can change over time, long-term identity preservation is not guaranteed. Despite the ability to change any attribute, changes are bounded to known ranges. Spatially, the quantity of elements is fixed. Retinal presentation may only vary over pre-determined ranges. The *Zugmonitor* [39], as it presents the position and expected delay of trains, fits in this category.

MUTATE×EXTREME BINS (M×EB)

Visualizations in this category are similar to those in MUTATE×KNOWN SCALE in that variation is bounded. However, in MUTATE×EXTREME BINS only a subset of the dynamics underlying the retinal variables can be represented. In some ways, visualizations in this category focus on an interesting range whereas those in MUTATE×KNOWN SCALE use the entire retinal space more uniformly. MUTATE×EXTREME BINS visualizations provide details in context, and can do so with both spatial and retinal values. However, comparisons across long time spans are difficult since no element of the visualization is necessarily stable.

MUTATE×MUTABLE SCALE (M×MS)

In this category, visualizations present a fixed number of entities, but may modify the retinal encoding in arbitrary ways. The tree-map stock market representation at Smart Money fits this category [8, 45]. The number of tree cells in the maps of the market is fixed, each corresponding to an individual stock. The color of the cells represents percent value change. However, the color scale changes as new data is acquired, always accommodating the largest percentage change currently in the map. Therefore, the same high saturation color in a cell in any given time-step may represent substantially different values. MUTATE×MUTABLE SCALE visualizations show detail for the current state when only a fixed number of visual elements are required.

CREATE×IMMUTABLE (C×I)

Introducing new elements to a visualization introduces an additional level of complexity. Depending on how mutability is used, visualizations in this category may or may not preserve identity. Generally, if identity is created by position then it can be preserved over short time periods, but large position changes (even cumulative) make long-term identity preservation unlikely. However, if identity is preserved via the retinal variables, it is preserved indefinitely. Dynamic timelines belong in this category of visualizations, such as Yanni Loukissas’s Apollo 11 visualization [25]. As time progresses, new elements from a variety of data sources are added to the left-hand side of the display.

CREATE×KNOWN SCALE (C×KS)

Visualizations in this category can employ element creation, positional mutation and mutable appearance along a known scale. Without fixed retinal representation, identity establishment is further complicated over CREATE×IMMUTABLE. As with all categories in the CREATE row, judicious use of retinal and spatial mutation counteracts these effects. Visualization in this category can be used

to track an increasing number of elements. For example, in The New York Talk Exchange [32], positional mutation is not used because location is based on physical geography. In contrast, dynamic graph layout algorithms that do not preserve context use positional mutation extensively as new nodes are added [4, 5].

CREATE×EXTREME BINS (C×EB)

As with all extreme bins visualizations, visualizations in this category enable more detailed presentation of specific ranges of values. Adding the ability to create new elements introduces a new dynamic over just spatial mutability. For example, new information can be retinally emphasized but eventually presented more consistently (as is done in by Nathan Yau to show growth of the retail stores Target, Walmart, and Sam’s Club [48]). Another common pattern in this category of visualizations is to subdue superseded or excluded values without fully removing them. Subdued values remain in position and provide context, but belong to a retinal ‘bin’ that de-emphasizes them. Visualizations in this category can be used to represent deleted elements [5].

CREATE×MUTABLE SCALE (C×MS)

In this category, visualizations tend to use a stable spatial substrate, but accumulate elements and constantly re-encode retinal variables. The Digg Arc [15] visualization belongs in this category. New elements are accumulated into individual arcs, but color re-encoded as percentages shift to different parts of the circle. When used with a slow-changing reference system (such as Digg Arc uses), this visualization can retain identity over short time periods. However, this identity preservation is not generally guaranteed. The mutable scale and ability to add an arbitrary number of elements over time make long-distance comparisons difficult. As a MUTABLE SCALE visualization, the current state shows maximal detail on retinal variables.

CREATE & DELETE×IMMUTABLE (CD×I)

Visualizations that allow elements to be created and deleted (as well as mutated) enable the tracking new elements over time without retaining old elements. Retinal dimensions may also encode additional data. Identity can be established with the retinal encoding. Despite the immutability of the retinal encoding, identity cannot be guaranteed over long periods since elements may be removed. Visualizations in this category are suited to tracking elements of transient interest, like HDPV does with memory objects [40].

Strictly speaking, any visualization that allows creation and deletion of elements can *simulate* mutability. An element can be removed and a new element instantiated with slightly modified retinal characteristics. The distinction between this category and CREATE & DELETE×KNOWN SCALE relies not on capability, but rather observed behavior. If visual elements maintain their retinal characteristics throughout their time in the visualization, they belong in the IMMUTABLE column, even though modification can be simulated without violating the behavioral characteristics of this category.

CREATE & DELETE×KNOWN SCALE (CD×KS)

Visualizations in this category behave much like those in CREATE & DELETE×IMMUTABLE. However, in this category retinal representation of individual elements can also change. Such changes decrease the ability to make identity-based comparisons, but facilitate greater flexibility to represent shifting categorizations of values. Techniques in this category can include fading values out over time [7] and radar-like tracking [6].

CREATE & DELETE×EXTREME BINS (CD×EB)

In this category, elements can be created and deleted, with transient information displayed in retinal variation. Extreme bins allow a specific data range to be highlighted. These techniques are used in Wattenburg’s wind-speed visualization [46], and NASA’s ocean

current diagram [28]. Related techniques are used to represent the stock market via Boids [27]. In all of these cases, the flow is an emergent property of many small motions and supplemental information is overlaid in the flow. Comparisons in short time-steps focus on group trends in position, overlay values or both. Visualizations in this category have very weak identity properties for individual elements, though group identity may be present.

CREATE & DELETE×MUTABLE SCALE (CD×MS)

This final category of visualization techniques can present the current data state with the highest degree of fidelity. However, it is the weakest for comparisons across time. Much like FIXED×IMMUTABLE visualizations, comparisons are restricted to what is presented at any given instant. However, in CREATE & DELETE×MUTABLE SCALE visualizations, the visualization changes over time, inviting comparison. Such comparisons must be made carefully, as changing scales reduce the ability to recognizably encode group identities. Despite this restriction, visualization in this category can communicate step-by-step changes in a complex space. Short term comparisons can be made when changes are made slowly and vary a small number of elements at once (this is further helped by following animation guidelines [18]).

4 HIGHER-LEVEL CATEGORY: IDENTITY GROUPS

Identity properties form three groups from the techniques presented in Section 3. The three identity groups, in decreasing order of identity preservation, are: Preserving, Transitional and Immediate. These groups are formed from the interplay between spatial and retinal categories. The identity groups are an emergent structure in this taxonomy. They have implications for the pervasive task of comparison through time [5, 18, 42].

Preserving

Preserving techniques maintain the association between visual elements and underlying data across arbitrary time scales. This preservation of identity is achieved by holding some part of the representation of an element constant. Fixed locations preserve identity in four of the techniques groups in this category. Position is of high visual salience, and thus the identity association is strong and easy to decode. In contrast, visualizations in the MUTATE×IMMUTABLE category keep the retinal encoding constant, while permitting position to change. This allows greater expressiveness in the spatial dimension, but increases the cognitive load for fine-grained comparisons (though they remain possible). Preserving techniques do not require additional effort for identity preservation; it is inherent in the bounds on dynamics. Preserving techniques can represent a great deal of dynamic information. They are most easily achieved when the range of dynamic inputs are known. Preserving techniques are essential when comparisons need to be made over long periods of time.

Transitional

Transitional techniques often retain identity associations, but the association is not as strong as for Preserving techniques. Transitional techniques retain identity associations across time by limiting changes to only known values. From the standpoint of quantity, new elements may be added or changed (but not both if long-term identity is needed). Retinal encodings are also bounded. Transitional techniques preserve identity association over short time scales without difficulty. However, each technique category has at least one way that identity associations can be destroyed by design decisions. Transitional techniques balance encoding flexibility with detailed comparison across time. Fine-grained comparison is easily supported over short time spans, but degrades over longer spans. They provide greater degrees of freedom for presenting and contextualizing current information than Preserving techniques.

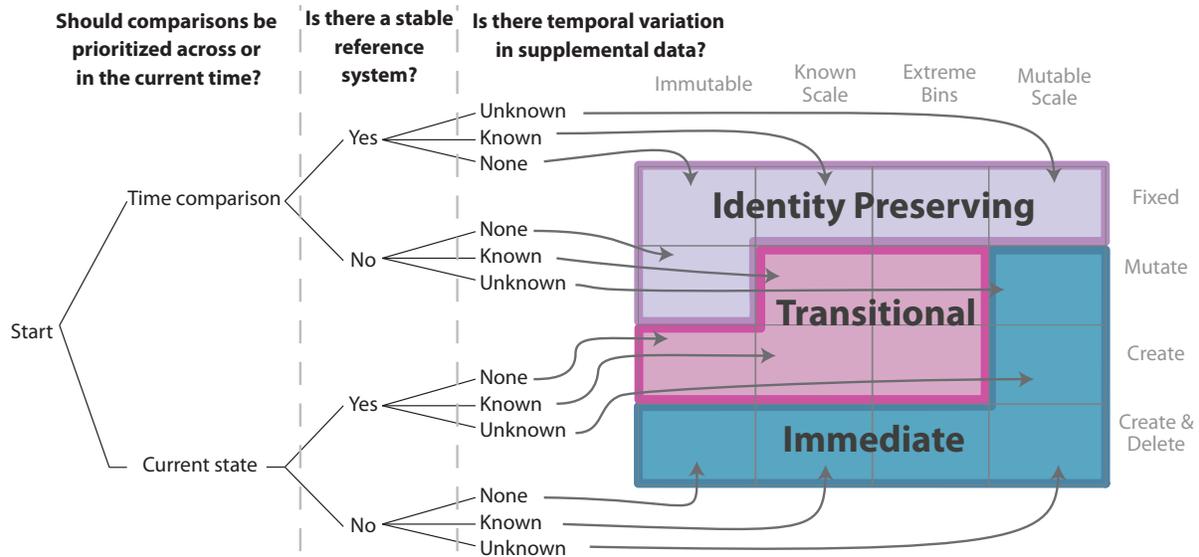


Figure 4: Decision tree for mapping from specific tasks to a techniques through the taxonomy.

The CREATE×IMMUTABLE technique category is the only IMMUTABLE Transitional category. The inclusion of spatial mutation in CREATE visualizations makes this category transitional. Unless the retinal characteristics are specifically used to encode identity, identity is difficult to establish across time. New elements and moving old elements quickly degrade identity characteristics.

Immediate

Immediate techniques *may* retain identity associations, but only through careful design choices. Those design choices also fall outside the scope of this taxonomy (such as varying one visual dimension widely but fixing another, see Section 2.4). By default, Immediate techniques do not preserve the identity association. In the case of categories in the CREATE & DELETE row, identity is destroyed whenever an element is removed from the visualization. For visualizations in categories of the MUTABLE SCALE column, identity is destroyed whenever a retinal scale is remapped. This identity association destruction may prevent comparisons over short time-scales as well as long time scales. (Removed items cannot be found and remapped color scales cannot be compared in subsequent time steps, even if there is only one time step between the two images).

Immediate techniques are valuable techniques, despite limited identity preservation. They provide the detailed information in the most flexible format *right now*. They are poor for making comparisons across time steps.

5 TASK TO TECHNIQUE MATCHING

The Preserving, Transitional and Immediate identity groups presented in Section 4 are central to effectively applying this taxonomy to new visualization problems. A common task is to have a dynamic-data visualization problem, and to need a technique to fit the problem. Since not all techniques enable the same kinds of comparisons, selection is a process of matching project priorities to technique capabilities. Once a suitable technique category is identified, neighboring techniques can also be considered because neighbors in the taxonomy hold similar characteristics. Changes from one technique category to a neighboring technique category are incremental but cumulative (therefore, diagonal neighbors are two significant changes away). The process of selection can be pursued in three questions: (1) “Should comparisons be prioritized across

time or in the current time?”, (2) “Is there a stable reference system?” and (3) “Is there temporal variation in supplemental data?” Answers to these questions select an identity group, a spatial technique category, and a retinal row. This mixture of task and data-type questions is not uncommon in visualization taxonomies (for example [1, 37, 47]). Practically, the questions may be answered in any order convenient. These questions, and the suggestions offered in this section based on their responses, are derived from Gestalt principals [44] and supported by pre-existing evaluations of techniques presented by others. Figure 4 illustrates the process.

The initial question of the decision tree is: “Should comparisons be prioritized across time or in the current time?” Selecting the desired characteristics of comparison determines identity preservation properties. If comparisons across a large amount of time are required, then (identity) Preserving techniques are preferable. Preserving techniques retain a strong sense of context over time. Context construction and preservation improves interpretation of analysis for questions that rely on it [4]. In contrast, if only the immediate context is required, visualizations from the Immediate categories present this context with increased fidelity. Transitional techniques may be adapted to either circumstance, but are so categorized because they preserve short-term context. Short-term context can be used for analysis tasks where long-term context is not required [7].

“Is there a stable reference system?” is the second question. A reference system is composed of the quantity of elements combined with some task-significant relationship between them. Geography and slow/unchanging social relationships can be used as stable reference systems [1, 10]. If a reference system exists and it is relevant to the task, then representing it improves analysis performance [5]. Visualizations commonly represent reference systems via spatial encoding [10]. Spatial techniques in the upper taxonomy rows retain spatial reference systems more than those lower in the matrix. If a reference system is not applicable, the lower parts of the matrix yield more representational flexibility.

The final question is: “Is there temporal variation in supplemental data?” This question clarifies the role of data supplementing which supplements the data used in spatial encoding. The supplemental data is encoded in retinal variables. “None”, “Known” and, “Unknown” are the possible answers to this question. These responses map directly on to columns in the matrix.

In general, using a technique from a neighboring category will yield similar results because encoding and identity properties are similar. For example, the EXTREME BINS column is not directly indicated by any of the responses to the third question. It represents an intermediate category between the “Known” and “Unknown” cases, and may be effective in many circumstances for either. Whether to choose the indicated column or shift to EXTREME BINS is task-dependent. Providing more detail in a specific range is a reason to shift from KNOWN SCALE; improving long-term interpretability is a reason to shift from MUTABLE SCALE. In general, deciding to make such a shift may be predicated on conventions of the target audience, available mapping functions, additional knowledge about the data set or detailed task requirements.

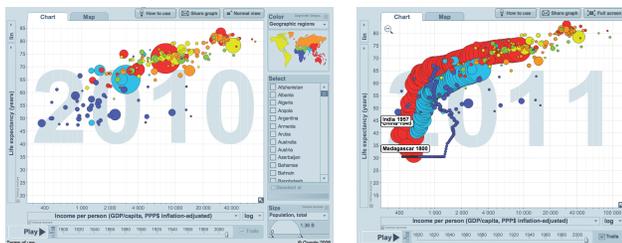
An example application will clarify the search process. Starting with a group of people in a social network and a set of elements they may “up-vote”, design a visualization to address the question, “Do people who like the same things tend to provide new, shared up-votes around the same time?” This question is a precursor to treating social behavior as an “infectious” phenomenon [20]. Since the time-scale of comparisons is not known, we will assert (in response to question 1) that long-distance comparisons are important. The social network is not a stable substrate—membership changes, as may member roles—providing a negative response to question two. Finally, “up-voting” an element is a binary response, so variation is “known” in supplemental data. Preferring long-distance comparisons over an unstable substrate with known variability in supplemental information leads to the CREATE×KNOWN SCALE category. CREATE×IMMUTABLE may be used if up-vote quantity does not matter; asserting that it is only significant that an item received some up-vote. If the number of items that may be voted on is fixed, a MUTATE×IMMUTABLE may be used instead, but this changes the visualization’s focus from people to items. If the number of people is known in advance (changing the dynamics assumed in question two), a MUTATE×IMMUTABLE technique may also be employed.

The example question above is closely related to the research questions underlying New York Times Labs: Project Cascade [29]. One Project Cascade visualization belongs to the CREATE×IMMUTABLE category. In Project Cascade, nodes represent events. New events induce new nodes that are given stable locations (placing it in the CREATE row). Each event is categorized at creation time, with retinal variables representing this category. Therefore, Project Cascade essentially presents a dynamic timeline of events. Because events are created, but not generally moved, Project Cascade visualizations have higher identity preservation than most Transitional visualizations.

6 DYNAMIC VISUALIZATION EXAMPLES

The taxonomy presented in Section 3 can be used to describe and critique dynamic data visualizations. In this section, we have selected exemplar visualizations to analyze using it.

6.1 Gapminder World



(a) Default Configuration

(b) Traces Enabled

Figure 5: Gapminder World’s time-dependent visualizations [17].

The *Gapminder World* tool [17] generates visualizations for observing trends over time. In its default configuration, it presents a scatter plot of values as they progress over time (Figure 5a). The scatter plot mutates the position of a fixed number of elements and modifies size on a known scale as data changes. Therefore the presentation mode is in the MUTATE×KNOWN SCALE category. This category is consistent with the goal of enabling comparisons between time steps that focus on the values selected for the axes.

Gapminder World can be configured to provide different types of visualizations. For example, prior node positions can be retained over time (the “traces” option), creating a visualization in the CREATE×KNOWN SCALE category where no mutation is used. The taxonomy structure suggests that performance between traces and non-traces would be similar, given that they belong to adjacent taxonomy categories. Robertson, et al. confirm this hypothesis [33].

6.2 Flow Lines

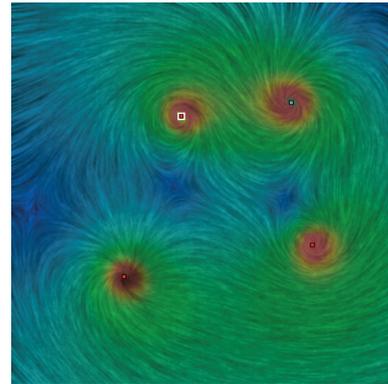


Figure 6: Flow lines example taken from [43].

Flow lines represent dynamically shifting, emergent structures (Figure 6). This visualization is constructed by generating a texture of discrete points in response to an input topology [43]. As the topology changes, so does the texture. The texture may be animated to illustrate flow direction and magnitude. Flow lines present information about the underlying topology through ephemeral entities; therefore, long-term identity preservation is not required and may not even be desirable. In its common form, this visualization belongs to the CREATE & DELETE×KNOWN SCALE category, a member of the Immediate identity group. Textual points, the basic unit of this visualization, may be created and deleted. Retinal variables may be used to encode specific information about points. For example when using “dyes,” color communicates origin. It may also be used to emphasize velocity or direction (as is used in Wat-tenburg’s wind visualization [46]).

6.3 code_swarm

The *code_swarm* visualization [31] shows the shifting relationship between documents and their multiple authors. It uses the position of points (representing documents) around text labels (representing authors) to describe the activity of a software repository (Figure 7). The labels are the only elements that preserve identity; they have a consistent representation and move slowly. The positional information and quantity of elements are revealed dynamically as emergent properties of developer activity. Color is used to indicate both file type and recent activity on an EXTREME BINS scale. This is a more complex encoding than the KNOWN SCALE used to encode size. The net effect of the presentation decisions is to generate a sense of trends that are composed of individual elements in the visualization, but trend membership changes over time. This presentation conforms to the CREATE & DELETE×EXTREME BINS category. Com-



Figure 7: code_swarm visualization of developer activity in a repository. Documents are represented as points that are animated in response to developer activities.

parisons do not involve individual elements, but rather patterns of fluid collections of elements shifting through time.

6.4 Taxis cab Tracking



(a) Taxis of London

(b) Taxi!

Figure 8: Taxi activity in London [24] and New York [34] (contrast adjusted). These visualizations reveal how different treatments of the same dynamics affect interpretation.

The *Taxis of London* visualization [24] is one of many in the British Broadcasting Company “Britain from Above” series (see Figure 8a). It is similar to other usage-based visualizations [23, 30, 39]. The task of this visualization is to represent road utilization in London. The road network forms a fixed reference system and taxi information is presented in the context of that usage. Though the precise mapping function is not indicated, it appears to use a known-scale projection (levels appear consistently spaced, with no catch-all extreme-bin at the top). Therefore, this visualization belongs to the $\text{FIXED} \times \text{KNOWN}$ SCALE. In this category, arbitrarily separated times can be compared directly to discern differences in road usage.

In contrast, *Taxi!* project [34] presents a $\text{CREATE} \& \text{DELETE} \times \text{IMMUTABLE}$ visualization of taxis in New York. In *Taxi!* (Figure 8b), the principal elements are the taxis themselves. The roads appear through the distribution of the taxis over time. Since the quantity of taxis is not constant across time (taxis are added and deleted), this project presents a $\text{CREATE} \& \text{DELETE} \times \text{IMMUTABLE}$ visualization of taxi activity. The *Taxi!* visualization is less suitable for comparison than the *Taxis of London* and it is unclear which comparisons *Taxi!* invites.

6.5 Map of the Market

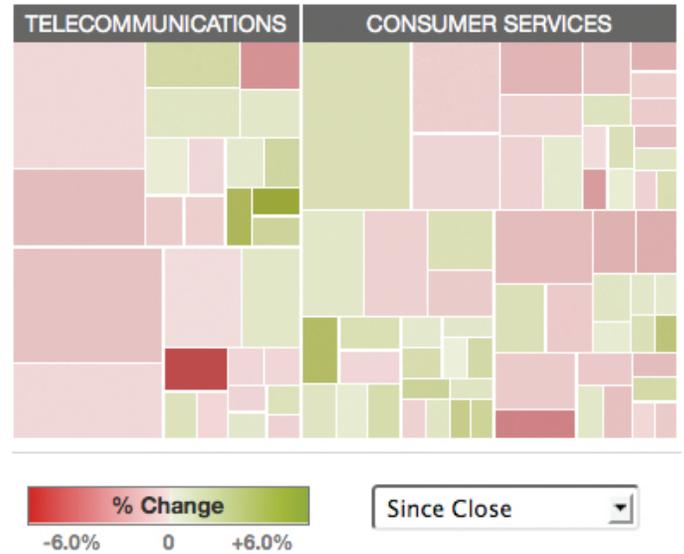


Figure 9: Portion of the Map of the Market on a relatively quiet day (image from June 20, 2012).

“Map of the Market,” presented at Smart Money [45], represents stock market activity. Leaf tree cells represent stocks. The stocks included are fixed, barring navigation. Cell size represents current price as percentage of the total of all included stocks. Coloring represents percentage change. The tree-map layout [8] rearranges the cells as current values change (i.e., it mutates their location and size). The cell coloring range is adjusted to accommodate the largest percentage change present. The positional changing and mutable scales indicate that this visualization belongs in the $\text{MUTATE} \times \text{MUTABLE}$ SCALE category. Visualizations in this category belong to the Immediate identity group. Current state information can be acquired at a glance by examining the current color scale and overall hue/saturation of a region. Comparison across time requires referral to the range of the color scale and potentially shifting correlating positions. More stable comparisons could be made by changing to a fixed layout from the dynamic tree-map or with an EXTREME BINS color scale. However, the chosen visualization is effective for monitoring market status.

7 ALTERNATIVE TAXONOMIES

The taxonomy presented in Section 3 is not the first taxonomy of information visualization, or even the first to examine the effect of dynamics. This section compares our taxonomy to existing ones.

Shneiderman’s Type by Task Taxonomy decomposes static visualizations according to seven data types and seven tasks [37]. Although the tasks focus on static data types, several introduce dynamics through interaction. Our taxonomy looks specifically at dynamic data sources. In many ways these two issues are orthogonal, thus TTT can be used in conjunction with the taxonomy presented in this paper.

More recently, Heer and Shneiderman presented an interactive dynamics for visual analysis taxonomy (IDV) [19]. This taxonomy focuses on the tools of visualizations creation and the interactivity found in those tools. Like TTT, IDV derives most of its dynamics from interaction. Our taxonomy provides details for the “Visualize” category in IDV, where specific encoding is selected. IDV does recognize that visualizations change over time (especially in the “Record” and “Guide” categories), but does little to discuss the implications of those changes on comparison and evaluation.

Heer and Robertson present a taxonomy of animation techniques [18]. Since data dynamics often appear as animation, their taxonomy provides useful context. Specifically, the spatial variation described in our taxonomy provides detail to their “Substrate Transformation” category. Similarly, the retinal variation of our taxonomy provides four subcategories for their “Visualization Change.” Their discussion of congruence [18,42] provides useful information about maintaining comparison capabilities in the face of changes to a visualization.

Keim, et al. indicate temporal data analysis as a significant area of analytics research and a need for “...identification of patterns (...), trends and correlations of the data elements over time...” [22]. They further identify presentation selection as a significant challenge for all visual analytics processes. Our taxonomy presents a means to select visual representations based on the types of comparisons desired (Section 5), and to categorize a visualization to understand its potential weaknesses (Section 3).

Chi presented a taxonomy for describing visualization processes based on the data-state model of visualization [13]. Chi’s taxonomy describes the design space for data transformation by enumerating the different transformation stages. This taxonomy is applied by enumerating the processing performed at each stage of the information visualization pipeline (Figure 1) [12]. Chi’s taxonomy is largely agnostic to data dynamics and does not attempt to provide guidance in design. Our taxonomy treats dynamics specifically. In addition to being descriptive, the search procedure from Section 5 aids directly in design by mapping low-level data characteristics to high-level properties of the visualization.

Several taxonomies of time-varying data have been presented [1, 2, 14, 47]. However, most of these taxonomies assume that the data is historical, and thus static from the standpoint of analysis and visualization. In the terminology of [47], there is no variation in *user time* driven by changes in the data. These static, time-oriented taxonomies focus on how analysis and representation can be used to effectively convey historical dynamics. Our taxonomy presents a framework for discussing how changes to the visualization itself affect interpretation over time. These other taxonomies do provide insight into mechanisms for effective interpretation. For example, the different options for presenting event-based analysis results presented in [2] can be categorized according to our taxonomy. The process described in Section 5 can assist in matching user tasks and data characteristics to a visual representation.

Aigner, et al. present a taxonomy for time oriented data that specifically includes dynamic representations [1]. However, they treat all dynamic representations as equal and do not “...investigate the subtle details of the variety of visual approaches available.” Our taxonomy addresses this omission directly for dynamic representations in the same way that earlier static taxonomies do for static representations [26, 49]. In their discussion of dynamics, Aigner, et al. (supported by others [19, 38]) indicate that long-term comparisons are often untenable due to the limits of visual memory, noting that “...the animation takes too long for users to remember its course.” This issue is a principal disadvantage of visualizations in the IMMEDIATE category of our taxonomy, but visualizations in the PRESERVING and TRANSITIONAL categories (when treated appropriately) can present the cumulative history, eliminating much of this cognitive loading. Access to images from prior time-steps may also reduce the cognitive load of comparisons across large time ranges by reducing memory effects [18, 42].

In summary, the taxonomy presented in this paper is unique in that it provides an analysis of dynamic data representations. It provides a description of eight different types of variation (four spatial and four retinal) and three higher-level categories. These categories have been useful in exploring the properties of existing visualizations. Our taxonomy does not provide guidance about working with different types of time (cyclic vs. linear, etc.) and does not pro-

vide a detailed analysis of the comparison process (only comparison aids). It also does not provide guidance to the specific analysis to perform, only broadly restricting it based on representational properties. Where many prior taxonomies investigate the entire analytical process from a high-level perspective, our taxonomy focuses on representation and interpretation.

8 EXTENSIONS AND CONCLUSIONS

The taxonomy presented in this paper focused on streaming data. However, the distinct characteristics of other types of dynamics are not addressed. For example, the distinction between user-requested and data-derived changes to the visualization is qualitatively significant but not considered. Though zoom/pan navigation can be thought of as a mutable position technique, zoom and pan do not *modify* the reference system; instead, they modify the viewport onto that reference system. This distinction is significant to interpretation, but not considered in the taxonomy. Similarly, brushing techniques may mutate retinal properties, but the changes are more akin to a transient overlay than source data or encoding changes. Preliminary investigation indicates that our taxonomy can be used to describe some of the dynamics derived from interactivity. However, it is also clear that this taxonomy is incomplete with respect to this expanded task. Similarly shortcomings exist for dynamics introduced at other stages of the InfoVis pipeline.

The suggestions for visualizations presented in Section 5 are not validated by user studies. They are derived from cognitive principles and existing studies of specific visualizations. Each category essentially represents a hypothesis that can be tested by user studies comparing visualizations from each category on relevant tasks.

The taxonomy presented may be extended by investigating different time scales, differences between retinal variables, mixtures of retinal treatments, and additional spatial categories (i.e., “create without mutate”, “delete but not create”, “modifiable spatial scale”, “abstract spatial scale”). Using visualizations from different categories concurrently (i.e., coordinated multiple views), may present interesting combination effects.

In preparing this taxonomy of dynamic visualizations, exemplars were not uniformly distributed. For example, visualization in the CREATE & DELETE row were more plentiful than other rows. Columns IMMUTABLE and MUTABLE SCALE were more prevalent than other columns. Why some visualization categories are more plentiful than others is not known, and deserves further investigation. Possible sources for this skew to/from certain techniques lies in: (1) our survey methods; (2) that the technical barriers of moving between some techniques is lower than for others; or (3) the types of questions being answered with dynamic data visualizations tend to favor certain characteristics.

Dynamic data visualization is a field of growing importance. Understanding the characteristics that make creating and interpreting a dynamic visualization is an important task. Examining analysis from the standpoint of identity and mutability leads to a deeper understanding of existing visualizations. Furthermore, these aspects indicate a number of research areas that are not currently well understood but that show promise for fruitful future investigation.

ACKNOWLEDGEMENTS

This work is partially by a grant from the Lilly Endowment and NSF Grant No. 1111888.

REFERENCES

- [1] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski. Visualizing time-oriented data—a systematic view. *Comput. Graph.*, 31(3):401–409, June 2007.
- [2] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski. Visual methods for analyzing time-oriented data. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):47–60, Jan. 2008.

- [3] C. R. Aragon and M. A. Hearst. Improving aviation safety with information visualization: a flight simulation study. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '05, pages 441–450, New York, NY, USA, 2005. ACM.
- [4] D. Archambault, H. Purchase, and B. Pinaud. Animation, small multiples, and the effect of mental map preservation in dynamic graphs. *IEEE Transactions on Visualization and Computer Graphics*, 17(4):539–552, Apr. 2011.
- [5] D. Archambault, H. C. Purchase, and B. Pinaud. Difference map readability for dynamic graphs. In *Proceedings of the 18th international conference on Graph drawing*, GD'10, pages 50–61, Berlin, Heidelberg, 2011. Springer-Verlag.
- [6] R. Azuma, H. Neely, M. Daily, and R. Geiss. Visualization tools for free flight air-traffic management. *IEEE Comput. Graph. Appl.*, 20(5):32–36, Sept. 2000.
- [7] P. Baudisch, D. Tan, M. Collomb, D. Robbins, K. Hinckley, M. Agrawala, S. Zhao, and G. Ramos. Phosphor: explaining transitions in the user interface using afterglow effects. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, UIST '06, pages 169–178, New York, NY, USA, 2006. ACM.
- [8] B. B. Bederson, B. Shneiderman, and M. Wattenberg. Ordered and quantum treemaps: Making effective use of 2d space to display hierarchies. *ACM Trans. Graph.*, 21(4):833–854, 2002.
- [9] J. Bertin. *Semiology of Graphics*. Reprinted by University of Wisconsin Press, 1967.
- [10] K. Börner, C. Chen, and K. Boyack. Visualizing knowledge domains. In B. Cronin, editor, *Annual Review of Information Science & Technology*, volume 37, pages 179–255, Medford, NJ, 2003. American Society for Information Science and Technology.
- [11] C. Brewer. *Designing Better Maps: A Guide for Gis Users*. Environmental Systems Research, 2004.
- [12] S. K. Card, J. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufman, 1999.
- [13] E. H. Chi. A taxonomy of visualization techniques using the data state reference model. In *INFOVIS '00: Proceedings of the IEEE Symposium on Information Visualization 2000*, page 69, Washington, DC, USA, 2000. IEEE Computer Society.
- [14] C. Daassi, L. Nigay, and M.-C. Fauvet. A taxonomy of temporal data visualization techniques. *Revue Information Interaction Intelligence*, 5(2):41–63, 2006. Paru en 2006 pour un Volume en 2005.
- [15] Digg Labs and Stamen Designs. Digg: Arc. www.visualcomplexity.com/vc/project.cfm?id=585, June 2012.
- [16] Dynamic Diagrams. Digital Orrery. www.dynamicdiagrams.com/work/orrery/, June 2012.
- [17] Google. Gapminder: World. www.gapminder.org/world/, 2008.
- [18] J. Heer and G. Robertson. Animated transitions in statistical data graphics. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)*, 13:1240–1247, 2007.
- [19] J. Heer and B. Shneiderman. Interactive dynamics for visual analysis. *Queue*, 10(2):30:30–30:55, Feb. 2012.
- [20] A. L. Hill, D. G. Rand, M. A. Nowak, and N. A. Christakis. Emotions as infectious diseases in a large social network: the SISa model. *Proceedings of the Royal Society (B)*, July 2010.
- [21] S. Kamvar. *We Feel Fine: An Almanac of Human Emotion*. Scribner, 2009.
- [22] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. Visual analytics: Definition, process, and challenges. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, editors, *Information Visualization*, pages 154–175. Springer-Verlag, Berlin, Heidelberg, 2008.
- [23] A. Koblin. Visualizing amsterdam sms messages. sand-box.aaronkoblin.com/projects/amsterdam/index.html, June 2012.
- [24] Lion Entertainment and British Broadcasting Corporation. *Britain from Above*, chapter Taxis in London. Pavilion, August 2008. www.bbc.co.uk/britainfromabove/stories/visualisations/taxis.shtml.
- [25] Y. Loukissas and D. Mindell. A visual display of sociotechnical data. In *Proceedings of the 2012 ACM annual conference extended abstracts on Human Factors in Computing Systems Extended Abstracts*, CHI EA '12, pages 1103–1106, New York, NY, USA, 2012. ACM.
- [26] J. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2):110–141, 1986.
- [27] A. V. Moere. Time-varying data visualization using information flocking flocks. In *Proceedings of the IEEE Symposium on Information Visualization*, INFOVIS '04, pages 97–104, Washington, DC, USA, 2004. IEEE Computer Society.
- [28] National Aeronautics and Space Administration (NASA). Perpetual ocean. svs.gsfc.nasa.gov/vis/a000000/a003800/a003827/, June 2012.
- [29] New York Times Labs. Project Cascade. nyt-labs.com/projects/cascade.html, June 2012.
- [30] D. Offenhuber. Visual anecdote. In *ACM SIGGRAPH 2010 Art Gallery*, SIGGRAPH '10, pages 367–374, New York, NY, USA, 2010. ACM.
- [31] M. Ogawa and K.-L. Ma. code_swarm: A design study in organic software visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1097–1104, Nov. 2009.
- [32] C. Ratti, K. Kloeckl, M. E. Haller, F. Rojas, F. Calabrese, A. Koblin, A. Vaccari, Yahoo! Design Innovations Team, W. J. Mitchell, and S. Sassen. *Design and the Elastic Mind*, chapter The New York Talk Exchange. The Museum of Modern Art, New York, NY, USA, March 2008. senseable.mit.edu/nytel/.
- [33] G. Robertson, R. Fernandez, D. Fisher, B. Lee, and J. Stasko. Effectiveness of animation in trend visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1325–1332, Nov. 2008.
- [34] J. F. Saldarriaga, E. Montgomery, and T. McKeogh. Taxi! juanfrans.com/taxi/, Oct 2011.
- [35] Sexperience and Ipsos MORI. The sexperience 1000. sexperienceuk.channel4.com/the-sexperience-1000, June 2012.
- [36] S. Shekhar, C. Lu, R. Liu, and C. Zhou. Cubeview: A system for traffic data visualization. In *Proceedings of the IEEE 5th International Conference on Intelligent Transportation Systems*, pages 674–678, Washington, DC, USA, 2002. IEEE Computer Society.
- [37] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages*, pages 336–343, Boulder, CO, September 1996. IEEE.
- [38] D. J. Simons and R. A. Rensink. Change blindness: past, present, and future. *Trends in Cognitive Sciences*, 9(1):16–20, 2005.
- [39] Süddeutsche Zeitung Digitale Medien GmbH. Zugmonitor: So pünktlich ist die bahn. zugmonitor.sueddeutsche.de/, June 2012.
- [40] J. Sundararaman and G. Back. HDPV: interactive, faithful, in-vivo runtime state visualization for C/C++ and Java. In *Proceedings of the 4th ACM symposium on Software visualization*, SoftVis '08, pages 47–56, New York, NY, USA, 2008. ACM.
- [41] E. R. Tufte. *The visual display of quantitative information*. Graphics Press, Cheshire, CT, USA, 1986.
- [42] B. Tversky, J. B. Morrison, and M. Betrancourt. Animation: can it facilitate? *Int. J. Hum.-Comput. Stud.*, 57(4):247–262, Oct. 2002.
- [43] J. J. van Wijk. Image based flow visualization. *ACM Trans. Graph.*, 21(3):745–754, July 2002.
- [44] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufman, San Francisco, 2000.
- [45] M. Wattenberg. Visualizing the stock market. In *CHI '99 extended abstracts on Human factors in computing systems*, CHI EA '99, pages 188–189, New York, NY, USA, 1999. ACM.
- [46] M. Wattenberg and F. B. Viegas. Wind map. hint.fm/wind/, June 2012.
- [47] M. Wolter, I. Assenmacher, B. Hentschel, M. Schirski, and T. Kuhlen. A time model for time-varying visualization. *Computer Graphics Forum*, 28(6):1561–1571, September 2009.
- [48] N. Yau. *Visualize This: The FlowingData Guide to Design, Visualization, and Statistics*. Wiley, 2011.
- [49] G. Zelazny. *Say It With Charts*. McGraw-Hill, New York, USA, 3rd edition, 1999.