

Wakame: Sense Making of Multi-Dimensional Spatial-Temporal Data

Clifton Forlines, Kent Wittenburg

TR2010-031 June 2010

Abstract

As our ability to measure the world around us improves, we are quickly generating massive quantities of high-dimensional, spatial-temporal data. In this paper, we concern ourselves with datasets in which the spatial characteristics are relatively static but many dimensions prevail and data is sampled over different time periods. Example applications include building energy management of HVAC unit diagnostics. We present methods employed in our Wakame visualization system to support such tasks as discovering anomalies and comparing performance across multiple time series. Novel methods include animated transitions that relate data in spatially located 3D views with conventional 2D graphs. Additionally, several components of our prototype employ analytics to guide the user to "interesting" portions of the dataset.

ACM AVI 2010

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Wakame: Sense Making of Multi-Dimensional Spatial-Temporal Data

Clifton Forlines, Kent Wittenburg

Mitsubishi Electric Research Labs
201 Broadway, 8th Floor
Cambridge, MA 02139 USA

forlines@alumni.cmu.edu, wittenburg@merl.com

ABSTRACT

As our ability to measure the world around us improves, we are quickly generating massive quantities of high-dimensional, spatial-temporal data. In this paper, we concern ourselves with datasets in which the spatial characteristics are relatively static but many dimensions prevail and data is sampled over different time periods. Example applications include building energy management and HVAC unit diagnostics. We present methods employed in our Wakame visualization system to support such tasks as discovering anomalies and comparing performance across multiple time series. Novel methods include animated transitions that relate data in spatially located 3D views with conventional 2D graphs. Additionally, several components of our prototype employ analytics to guide the user to “interesting” portions of the dataset.

Categories and Subject Descriptors

H5.2. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Design, Human Factors

Keywords

infovis, spatial-temporal data, multi-dimensional data, radar graph, visual analytics.

1. INTRODUCTION

The rapidly decreasing cost of sensor networks has led to massive quantities of high-dimensional, spatial-temporal data. The quantity of available data pouring in is outpacing our ability to meaningfully display and explore this information. Without corresponding improvements in our ability to visualize this information, the value of this improved sensing of the world around us will be minimized. The need to display, explore, and make sense of massive data collections is driving new research into interaction and visualization technologies. While machine learning and summarization techniques complement (and may one day replace) human operators as they interpret these data collections, for the time being people still perform the sense-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AVI '10, May 25-29, 2010, Rome, Italy

Copyright © 2010 ACM 978-1-4503-0076-6/10/05... \$10.00

making portion of the program. As long as this is the case, researchers will have ample challenges to address in the field of high-dimensional, spatial-temporal data visualization.

This paper presents our initial exploration of this design space. We call our basic visualization objects *Wakame* (pronounced wā-’kā-me, a Japanese word for seaweed, and illustrated in Figures 2-4). Initially developed to display a building’s environmental sensor information, Wakame have the potential to provide benefits in other visualization activities in which multi-dimensional data has spatial and temporal components. Our proposal is to extrude 2D shapes representing multidimensional sensor readings into 3D in order to represent time and to locate these objects in a representation of the space in which the readings occurred. Our methods have been prototyped with datasets containing dozens of sensors, each with as many as 40 dimensions and over 50k samples. Our prototype also provides the user with controls for selecting meaningful time ranges and subsets of the available data dimensions. It also performs analytics on the dataset (such as outlier detection and clustering) with the goal of supporting the user’s investigations. Finally, through dimension reduction and animation, Wakame are related to familiar 2D presentations. Through allowing the user to investigate the dataset with both 2D and 3D presentations, the benefits of each approach can be exploited.

The remainder of this paper is organized as follows. First, we present related work and motivate the need for both 2D and 3D representations. Secondly, we present methods for linking our visualizations to locations in space, as well as relating these 3D visualizations to familiar 2D ones through the use of animation and dimensionality reduction. We then present a prototype system that uses our technique to visualize a building’s environmental sensor information. This prototype employs analytics to support common tasks and guide the user to “interesting” portions of the dataset. We conclude with an example datasets and a description of future work in this line of research.

2. RELATED WORK

2.1 Visualizing Multi-Dimensional, Time-Based Data

Multi-dimensional visualization is a mature field with a long history of research [17]. Parallel Coordinates, Scatterplot matrixes, Nightingale Rose Petals, and Chernoff Faces are among the many techniques to represent high-dimensional data in 2D. Of particular interest to this paper are radar graphs (also known as

star charts, spider graphs, Kiviat diagrams, and circular parallel coordinates). First used by Georg von Mayr [29], radar graphs plot multiple values along radial axes (Figure 1, *top*).

With respect to time-based, multi-dimensional data, several authors have proposed either animating these presentations or extruding these 2D techniques into a third dimension to represent time. Yuhua [9] proposed a 3D Kiviat diagram in which the axes themselves are rolled up and down in 3D, the result of which allows the viewer to better perceive time-varying data. Fanea et al. [6] combine multiple 2D parallel coordinate visualizations into a single visualization, aligning multiple star charts along a time-axis and connecting the vertices with polylines. Most similar to our approach is Hackstadt and Malony's Kiviat Tube [10] and Tominski et al.'s 3D Kiviat Tube [26], both of which display multiple Kiviat charts along a time axis, and connect a single surface along this profile. The result of this extrusion is a tube-like solid, the shape of which illustrates the changes in multiple values over time.

2.2 Spatial-Temporal Visualization

Spatial-temporal visualization in the field of InfoVis has largely concerned itself with explicitly identified events, relationships, and/or movements in space and time. Graph-based visualization methods have been proposed for intelligence analysis [16], communication networks [2], as well as social networks [7] [11]. The visualization of human, animal, or vehicle movement in space and time is also important for building security systems [14], environmental studies [14] and traffic management [1]. However, our primary concern in this line of research is with multi-dimensional sensor streams that are relatively static in the spatial domain. Rather than events or relationships forming the data primitives for visualization, it is the stream of individual sensor readings that serve as the starting point. The discovery of relationships across the streams of data in space and time is the task of the human user of the visualization tool. Examples of application domains that share this characterization include indoor and outdoor environmental monitoring as well as equipment condition monitoring in which space may be a logical map (such as a schematic diagram) rather than a physical one (building floor plan or geographical map). The generalization of geographically-based space-time visualization to diagrams has been previously proposed by [16], but the focus there is still on relations and complex events across entities in space and time rather than on relatively stationary sensor data streams. Tominski et al. [27] describe the use of two different 3D icons that depict multi-dimensional data streams from fixed locations. As with our system, these 3D icons are positioned on a 2D map so that multiple streams can be compared and related to their geographic location.

2.3 Animation in Information Visualization

Previous research suggests that animation provides a better understanding of the relationship among visual depictions of a dataset than do instant transitions [28]. An instant transition forces the viewer to mentally reconstruct the effects of the command, and animation shifts this cognitive task to a perceptual one, thus freeing cognitive resources for the primary task [23]. There are

several overviews and guidelines governing the use of animation in InfoVis (e.g. [12]), and numerous examples of both its positive and negative effects on the understanding of a dataset [28].

2.4 2D vs. 3D Representations

The human-perceptual system is remarkable in its ability to correctly recognize and interpret 3D shapes in the world around us. A good overview of current and past research in this area is provided by [22]. The phenomenon known as "shape constancy" [18] ensures that we are able to recognize the 3D shape of an object (or part) from multiple points of view and that this 3D shape remains constant regardless of the specific 2D shape projected on the retina.

Furthermore, there is evidence from the field of perceptual psychology that repeated exposure to a 3D shape quickly ingrains it in memory [13]. The converse is also true; unfamiliar shapes "pop out" and demand attention. The recognition of these stored shapes occurs "unconsciously" with minimal cognitive effort [3], which is a very desirable quality when one is designing a system for making sense of complex multi-dimensional data. It has been argued that 3D shape is capable of encoding more information than other visual properties like color, weight, and scale [22].

While there are many benefits to 3D presentations of data, there are well-known drawbacks as well. First and foremost is the problem of occlusion. When data is encoded in a 3D scene, foreground objects necessarily block one's view of background objects. Transparency can help, but it increases visual clutter. Furthermore, 3D presentations of data most often have a lower data/pixel ratio than 2D presentations, and thus more severely limit the amount of data visible to the user. It is unclear if 3D presentations allow users to accurately read and compare multiple values when compared to 2D. Finally, interaction and navigation in 3D can be more cumbersome and cognitively demanding, especially when performed on a 2D desktop display.

Given the tradeoffs between 2D and 3D, it seems appropriate to employ both approaches. It is likely that a system that allows the user to easily switch between them while maintaining a link between these different views would be beneficial to gaining insight from the dataset.

3. Wakame

Figure 1 illustrates the basic technique for generating Wakame geometry. A traditional multi-dimensional radar chart (Figure 1, *top*) is drawn on a ground plane in 3D (Figure 1, *middle*). As with other extrusion techniques, time is mapped to the y-axis, and sequential measurements become radar-graph profiles at different heights off of the ground plane (Figure 1, *bottom left*). These profiles are used to create a hollow "tubes" [10][26], the shape of which illustrates the changes in the data over time.

While Figure 1 renders each Wakame in a single color, our prototype system assigns a different color to each of the vertexes that correspond to the different dimensions (Figure 2). With this color, both the overall shape and the changes in shape within each dimension are more easily understood.

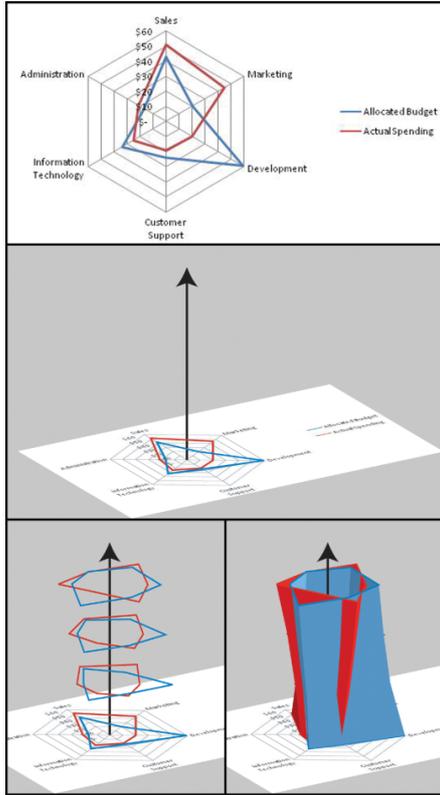


Figure 1. A traditional 2D radar chart (*top*) is drawn on a plane in 3D (*middle*). Time is mapped to the y-axis. Sampled values form the profile at different heights (*bottom left*), and a 3D shape is formed (*bottom right*).

3.1 3D Shape Perception

With our technique, correlations between the changes in values among multiple dimensions create recognizable 3D shapes. For example, a bulge in a Wakame that extends around its circumference indicates a similar change in values across many dimensions, just as a “cinch” in a Wakame indicates the lowering of several values (Figure 2, *right*). Similarly, a “wave” indicates a rising of the values of several dimensions that coincides with a lessening in others (Figure 2, *left*). Our design leverages the idea that events that cause predictable changes in many values simultaneously will result in shapes that become familiar to the practiced viewer [13]. To best take advantage of this property, careful attention should be paid to the ordering of dimensions around the center axis. Toward this end, we look to previous work on dimensional reordering in the field of multi-dimensional visualization (e.g. [21],[29]).

3.2 Viewing and Comparing Multiple Time Ranges and Multiple Datasets

Figure 1 illustrates how multiple overlapping Wakame can be rendered simultaneously in the same visualization. In this case, the red and blue Wakame overlap, and their intersection informs the viewer about the relationship in values between these two time series. The red and blue Wakame could come from two datasets

recorded from different sources, or could even be derived from separate time ranges in the same dataset.

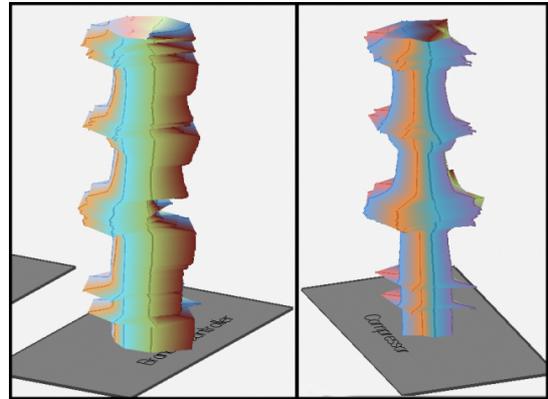


Figure 2. A Wakame's shape implies meaning. (*left*) A “wave” in the Wakame indicates that certain values are rising while others are lowering. (*right*) “Bulges” and “cinches” show correlation in value changes among dimensions.

In the previously cited examples, comparisons are always made among Wakame that contain identical data dimensions. While this is appropriate for a homogeneous sensor network, what about heterogeneous collections of data? Figure 3 shows two Wakame representing recorded data from two components in a system. Because these components perform different functions, they sense and record data in different dimensions. As such, a direct comparison is impossible; however, by rendering them side-by-side, the cross correlation in shape between these heterogeneous data dimensions can be revealed.

3.3 Relating Wakame to Spatial Locations

Figure 4 shows a screenshot of our prototype system. The main window of the application shows the floor plan of a workspace, with 11 individual offices and three HVAC (heating, ventilation, air-conditioning) zones. At each meaningful location on this ground plane, our system renders a Wakame representing the multi-dimensional time-series data associated with that point. In this example, we have simulated sensor data recorded in each room of the building, although any multi-dimensional time-series dataset with spatial locations can be viewed in this way. While

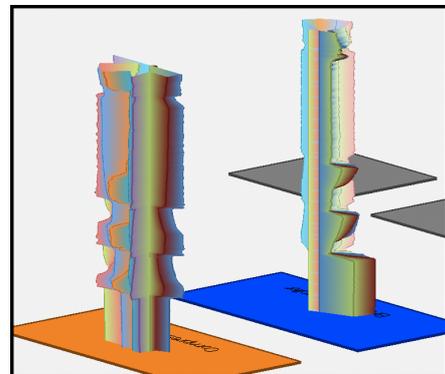


Figure 3. Shape correlation across different Wakame with different dimensions of data.

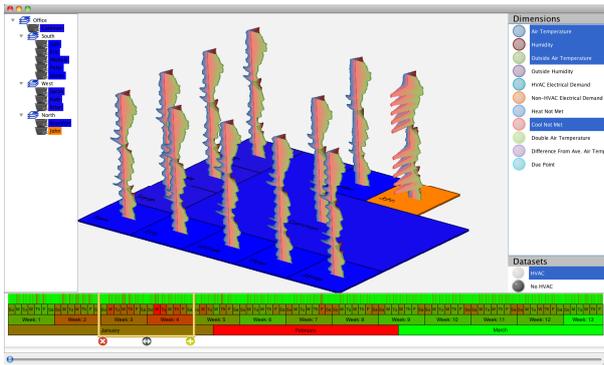


Figure 4. Prototype system overview. From each sensor location in this floor plan, a Wakame grows vertically over time. The profile at every height is derived from the multi-dimensional data collected at that moment. Controls around the border allow for the specification of time ranges as well as the constraining of displayed dimensions and Wakame.

both 3D and 2D presentations of temporal data can reveal temporal patterns, trends that occur in the temporal *and* spatial dimensions are much more difficult to detect with a collection of 2D charts and graphs. For example, a failing HVAC component might cause environmental changes that cascade across offices over time. This said, our 3D presentation makes accurate comparison among Wakame more difficult than comparing axis aligned 2D charts – a drawback that is addressed in the next section.

4. Collapsing for Traditional Views

Wakame are designed to present high-dimensional, time-series data in a manner that helps the viewer “make sense” of complex datasets. However, given the tradeoffs between 2D and 3D presentations that were discussed previously, there are times when a user is better served by a 2D representation of the data. While a 2D line-graph provides an excellent and easily accessible view of a handful of dimensions over time, one loses the spatial relationship among the sensor locations when viewing multiple 2D graphs. Previous systems have proposed simultaneously displaying multiple views of a dataset [20] along with explicit visual linking between corresponding features in the multiple views [4], arguing that multiple views might help a user's understanding of the data. Following this approach, one might choose to present Wakame visualizations alongside traditional 2D views of the data. While promising, multiple views consume more pixels and screen space than their single-view counterpart.

Our compromise is to relate traditional 2D views to the Wakames' locations in the space through the use of camera-view and geometry animation. Our prototype system includes two view transformations--one for collapsing the spatial dimensions of the Wakame into a traditional line chart, and one for collapsing the time-axis of the Wakame into a traditional radar graph. These transitions are designed to help the viewer relate these 2D presentations of the data to their spatial locations.

Figure 5 illustrates the transition between several Wakame and their 2D line-chart counterparts. In our system, this transition is bi-directional--as such, Figure 5 can be read left-to-right or right-to-left. Three characteristics of the visualization animate when this transition occurs. First, the camera rolls to the left so that the

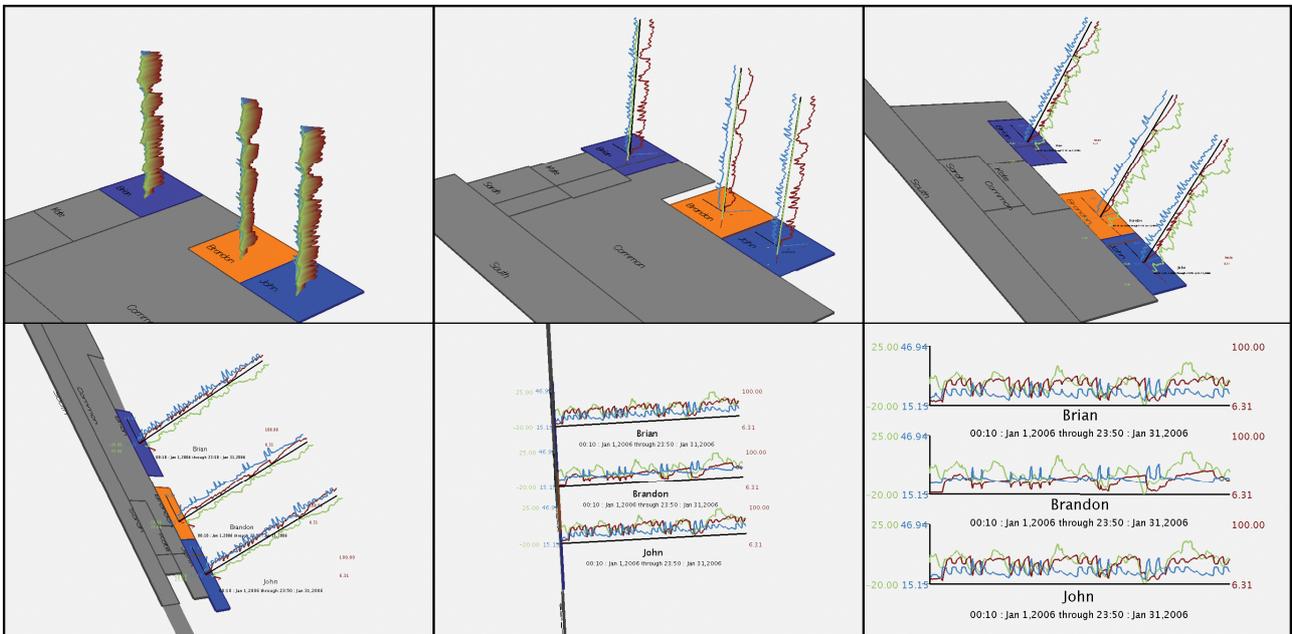


Figure 5. Collapsing spatial dimensions. When viewing the data in 2D is more advantageous, our system animates the appearance and position of the Wakame and camera so that traditional 2D charts are displayed. This transition is bi-directional, so the cells in this figure can be read left-to-right or right-to-left.

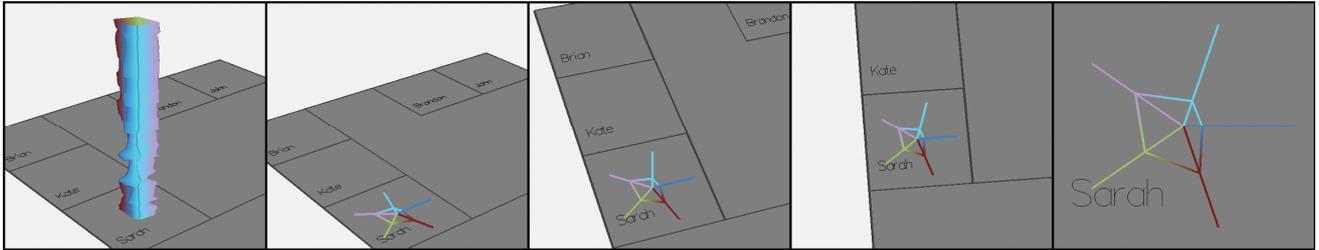


Figure 6. When the user selects a single moment in time, the Wakame collapse along the y-axis and the camera animates to view them from directly above. This animation runs in both directions as single moments or time ranges are selected.

time-axis now runs horizontally across the screen (its most common orientation). Second, the points on the Wakame's radar graphs orient so that they align to a single axis that runs up-and-down in the new view plane of the camera. Finally, the positions of the Wakame on the ground plane animate so that they equally spread out in the camera's new view plane. The result of these three animations is a collection of 2D line-charts stacked in the view plane.

Heer et al. [12] recommend that complex animations be split into steps, an approach that they argue will help the viewer understand the transition. As such, we implemented a variation of our system that executes the camera, dimensional, and positional animation in series.

Figure 6 illustrates the transition between a Wakame and its 2D radar graph counterpart. In this case, the radar graph represents the value of the Wakame at a single instant in time. When this transition takes place, the camera animates so that it faces perpendicularly to the ground plane, and centers itself on the selected Wakame. Meanwhile, the Wakame collapses along the y-axis (time) dimension. As with the previous animation, this transition can occur in either direction as the user issues commands to display either time ranges or single instances in time. When multiple Wakame are viewed in this way, their spatial relationship is maintained. In a variation of our prototype, this 2D radar graph displays the mean and variance of a selected time range.

5. Prototype System Overview

In order to explore and develop our Wakame visualization approach and to develop the analytic features of the user interface, we build a prototype system capable of displaying large amounts of multi-dimensional time-based data. In the figures and examples used in this section, we simulated a large amount of sensor information for a fictional office space under varying environmental conditions.

Figure 4 shows a screenshot of this prototype. Occupying the center and largest region of the screen is the 3D view. In this case, the floor-plan of an office building is shown, with Wakame growing vertically from their corresponding sensor locations in the workspace. On the left, a tree-view shows a hierarchical view of the sensors and their corresponding zones. On the right, our system provides two lists for selecting among multiple dimensions and among entire datasets (in this example, two datasets are loaded representing two separate simulations of the space).

Finally, along the bottom of the application, a timeline shows the available data and allows the user to select one or more time regions for display. These components each respond to input on one another, and their appearance is designed to guide the user toward interesting and meaningful portions of the dataset.

5.1 3D Scene

The 3D Scene contains a plane in which the spatial components of the dataset are grounded. In Figure 4, this plane is the office floor-plan of the building in which the visualized sensor data was collected, and we have experimented with cartographic (Figure 9) and schematic planes as well. The importance of this plane is that it links time-series' locations with their visual representations.

Interaction with the 3D Scene takes place with the system mouse or other pointing device. Clicking directly on regions of the ground plane controls selection, and dragging the background rotates the camera using the ArcBall camera-orientation technique [24].

5.2 Dimension and Dataset Lists

Figure 4 shows the Dimensions List (*top*) and Datasets List (*bottom*) on the right side of the application. These lists are used to hide and show individual dimensions and datasets respectively. As individual dimensions are hidden and shown, the points of the Wakames' radar graphs animate to equally space themselves out around the y-axis. When entire datasets are hidden and shown, they fade in and out.

5.2.1 Derived Dimensions

In addition to the recorded or simulated data dimensions presented in the Dimension List, users are able to define their own derivative dimensions using the Groovy scripting language [7]. These scripts allow users to write simple functions that take the available data as parameters and return derived values allowing the sophisticated user to author their own derived dimensions to aid in their sense-making and exploration of the dataset.

5.3 Timeline

The bottom of Figure 4 shows the timeline widget for selecting time ranges. A full description of the features of this widget is outside of the scope of this paper; however, two properties of the timeline are important for understanding its use in the prototype system. Firstly, the timeline contains a hierarchical collection of selectable time ranges, which allow the user to quickly select individual or multiple months, week, days, and so on. Secondly, the color of these time ranges indicates their difference from

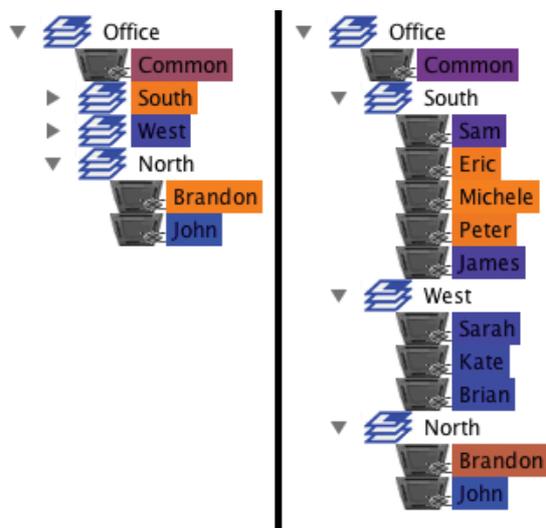


Figure 7. The Wakame tree shows the hierarchical relationship among Wakame. In this example, Wakame representing sensors placed in individual offices are organized into a collection of zones. When collapsed, these zones display the numerical averages of their sub-Wakame.

similar time ranges in the dataset, a feature that is explored in more depth in the following section.

6. Analytics and the GUI

When given a large multi-dimensional dataset to work with, people turn to analytics to help sort, summarize, and filter this data so that it is more digestible to the human mind. Our prototype uses analytics to alter the appearance of the GUI components with the aim of guiding the user to interesting portions of the dataset in which meaningful discoveries can be made.

Our approach is simple: any change in selection of the time range, dimensions, datasets, or Wakame triggers an update to the appearance of the timeline, Wakame tree, camera position, and 3D scene. In this section, we describe this use of analytics and the guidance it provides.

6.1 Coloring the Timeline

The use of analytics to color our timeline follows Lesh et al. [19] who used a similar approach to highlight cells in a user-defined spreadsheet. The basic approach is to allow the computer to quickly perform an exhaustive investigation of the space, and to use the results of this calculation to direct the user's attention to certain areas of the dataset. The user then, in turn, selects from among these "interesting" areas and the cycle begins again.

As shown in Figure 4, the cells in our timeline (representing hours, days, weeks, and so on) are each drawn using a color between red and green. These colors reflect the cell's "difference" from neighboring cells for the currently selected dimensions and Wakame. "Difference" might mean different things for different tasks, and thus should be highly tailored to specific applications and their tasks. In our prototype, we measure difference from the mean, but any number of more sophisticated measures are possible and are supported through our scripting system.

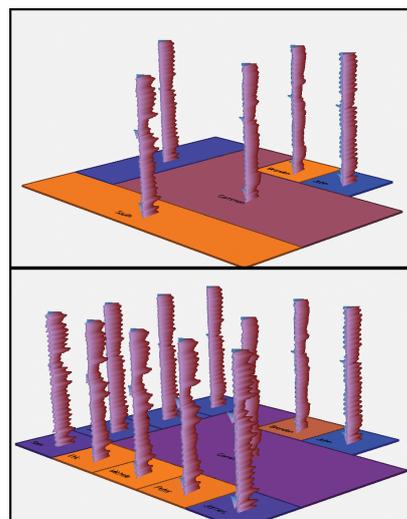


Figure 8. The top and bottom portions of this figure correspond to the left and right portions of Figure 7. By collapsing items in the tree, the user can reduce the visual complexity of the 3D scene. Analytics inform the presentation of the 2D zones in these figures. In this example, the Wakame are clustered into groups as represented by their color.

The result of these analytic methods is that the user is drawn towards meaningful time periods in the dataset. In addition to investigating these time periods themselves, larger cycles in the dataset become visible when the timeline is viewed as a whole. For example, vacations and monthly meetings (during which individuals' offices are unoccupied, but common rooms are crowded) become visible in the timeline, as do seasonal variations such as HVAC upticks during the hot summer months.

6.2 Wakame Tree

The Wakame Tree (Figure 7) is used both for displaying the hierarchical relationship among Wakame and for managing selection. There is a one-to-one mapping between nodes in this tree and Wakame in the 3D scene, with only the selected nodes being visible in the 3D window.

There are two types of nodes--leaves and branches. Leaves in the tree represent locations in the workspace for which there is actual recorded or simulated time-series data to be displayed. In the example shown in Figure 7, these leaves represent environmental sensors placed in the individual offices of this building, and are named after the office residents. Branches in the tree represent conceptual groupings of the leaves. In the tree shown in Figure 7, there are three building zones (west, north, and south) that each contains a collection of offices. These three zones are in turn placed within the root "office" node, which also contains a leaf corresponding to the large common area in the middle of the floor plan. Figure 8 (*top* and *bottom*) show the appearance of the 3D Scene for the two states of the Wakame Tree shown in Figure 7 (*left* and *right*).

Whenever the user makes a change in the selected dimensions, dataset, time period, or Wakame, the selected Wakame in the tree view and their corresponding zones in the 3D scene are colored to

reflect their similarity to one another. Again “similarity” will have different meanings to users involved in different tasks, and for the purpose of our prototype, we use a simple multi-dimensional clustering to assign Wakame into groups. Figure 7 and 8 show the grouping for two different selections of Wakame. This clustering guides the user to focus their attention on the most different Wakame, or to group similar Wakame for the purpose of simplifying the dataset and the task at hand.

6.3 Intelligent Camera Positioning

When working with an early version of the system, it became obvious that a great deal of effort was needed to correctly position the camera in order to best view the selected Wakame. Toward this end, we implemented an automatic camera positioning system that attempts to position, orient, and zoom the camera to best view the selected Wakame. This approach does an adequate job of “getting a good look at” the geometry one wants to see. When used in combination with the ArcBall rotation around the point of focus, less effort is spent positioning the camera and, hopefully, more cognitive resources are available for interpreting the data.

7. Example Dataset

The majority of figures and examples in this paper come from our exploration of large datasets generated from a simulation of an office building under various environmental and equipment conditions. We used the US Department of Energy's EnergyPlus [5] building simulation program to simulate an eleven-office floor plan, and recorded 12 virtual “sensor” measurements at an interval of 5 minutes over the course of a year.

It was encouraging that some characteristics of the office were immediately apparent. For example, the simulation includes a particularly cold January day, during which the building's HVAC system could not provide a comfortable working environment. Similarly, the cooling system conditioning the corner office that is used for monthly meetings cannot compensate for the additional thermal load added by the meeting's participants. This monthly occurrence resulted in a predictable bump on the office's Wakame the first Monday of each month. While interesting, it is worthwhile to point out that both of these findings could have been made through the use of 2D data representations.

In other cases, the spatial/temporal presentation provided by Wakame visualization lead to more subtle findings and a deeper understanding of the building's HVAC performance. For example, offices on the south-side of the building received more sunlight, which greatly affects their environmental conditions and required more cooling in the summer and less heating in the winter. Similarly, a gradual failure in the cooling system for one of the offices not only led to a rise in the temperature of that office, but also a cascading effect over time to the neighboring offices (Figure 9). In this case, a single failure lead to additional strain on neighboring units, which eventually lead to higher energy consumption, inefficiencies, and a reduction in comfort. It is clear that the spatial facet of the presentation makes the relationship between the performance of neighboring zones more clear than a collection of 2D charts, and we feel that this example lends weight to the argument for hybrid 2D/3D data visualization.

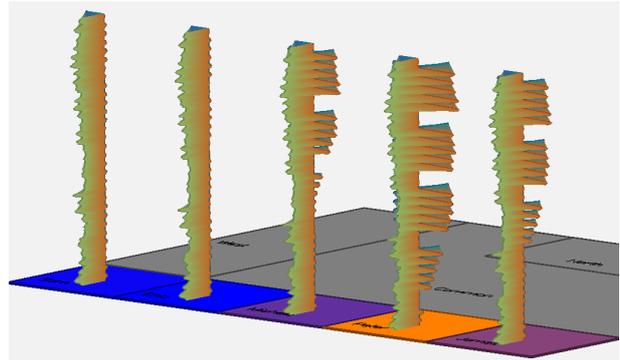


Figure 9. The effects of a failed component cascade to neighboring zones over time.

8. Conclusion

In this paper, we have presented the Wakame visualization technique in the context of previous research into multi-dimensional, spatial-temporal data and have illustrated our approach with a prototype application. We have argued for the use of our approach from the perspective of human perceptual psychology, and have given examples in which Wakame visualization may help users make sense out of large, complex datasets. Some of the tradeoffs between 2D and 3D presentations of data having been discussed, and we presented methods for collapsing our 3D presentations along the temporal and spatial dimensions into familiar 2D views. Finally, our prototype includes several examples of analytics in the UI that were designed to guide the user to interesting parts of the dataset. These first steps seem promising; however, the methods still require experimental validation in order to draw firm conclusions.

9. REFERENCES

- [1] Andrienko, G., Andrienko, N., and Wrobel, S. 2007. Visual analytics tools for analysis of movement data. *SIGKDD Explor. Newsl.* 9, 2 (Dec. 2007), 38-46.
- [2] Becker, R. A., Eick, S. G., and Wilks, A. R. 1995. Visualizing Network Data. *IEEE Trans. on Vis. and Comp. Graphics* 1, 1 (Mar. 1995), 16-28.
- [3] Boring, E. *A History of Exp. Psy.* Appleton, 1929.
- [4] Carpendale, S. 2007. VisLink: Revealing Relationships Amongst Visualizations. *IEEE Trans. on Visualization and Computer Graphics* 13, 6 (Nov. 2007), 1192-1199.
- [5] Energy Plus. US Department of Energy. <http://apps1.eere.energy.gov/buildings/energyplus/>
- [6] Fanea, E., Carpendale, S., and Isenberg, T. 2005. An Interactive 3D Integration of Parallel Coordinates and Star Glyphs. In *Proc. of the 2005 IEEE Symposium on Information Visualization* (October 23 - 25, 2005), 20.
- [7] Gloor, Y., Zhao, P. Tecflow - a temporal communication flow visualizer for social networks analysis. *ACM CSCW Workshop on Social Networks. ACM CSCW Conference, Chicago, Nov. 6. 2004.*
- [8] Groovy. <http://groovy.codehaus.org/>.

- [9] Yuhua, G. Implementation of 3d kiviati diagrams. Thesis Report. Vaxjo University. 2008.
- [10] Steven T. Hackstadt, S., Malony, A. Visualizing Parallel Programs and Performance. *IEEE Computer Graphics and Applications*, vol. 15, no. 4, pp. 12-14, July, 1995.
- [11] Heer, J. and Boyd, D. 2005. Vizster: Visualizing Online Social Networks. In *Proc. of the 2005 IEEE Symposium on information Visualization* (October 23 - 25, 2005). INFOVIS. IEEE Computer Society, Washington, DC, 5.
- [12] Heer, J. and Robertson, G. 2007. Animated Transitions in Statistical Data Graphics. *IEEE Trans. on Vis. and Comp. Graphics* 13, 6 (Nov. 2007), 1240-1247.
- [13] Helmholtz, H. *Treatise on Pysiological Optics*. (Translated from German, J.P.C. Southall.) Bristol: Thoemmes, 1867/2000.
- [14] Hewagamage, K. P. and Hirakawa, M. (2000) Pattern browser: Spiral-based interactive visualization using timelines on a geographical map. *IEICE Transactions on Information and Systems*, pp. 1679-1686.
- [15] Ivanov, Y., Wren, C., Sorokin, A., Kaur, I. Visualizing the History of Living Spaces. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (Nov. 2007), 1153-1160.
- [16] Kapler, T. and Wright, W. 2004. GeoTime Information Visualization. In *Proceedings of the IEEE Symposium on information Vis.* (October 10 - 12, 2004), 25-32.
- [17] Keim, D. A. 2002. Information Visualization and Visual Data Mining. *IEEE Transactions on Visualization and Computer Graphics* 8, 1 (Jan. 2002), 1-8.
- [18] Koffka, K. *Principles of Gestalt Psychology*. Harcourt, Brace, 1935.
- [19] Lesh, N. and Mitzenmacher, M. 2004. Interactive data summarization: an example application. In *Proc. of the Working Conference on Advanced Visual interfaces* (Gallipoli, Italy, May 25 - 28, 2004), 183-187.
- [20] North, C., Shneiderman, B. Snap-together visualization: a user interface for coordinating visualizations via relational schemata. In *AVI '00: Proceedings of the working conference on Advanced visual interfaces*, pages 128-135, New York, NY, USA, 2000. ACM.
- [21] Peng, W., Ward, M. O., and Rundensteiner, E. A. 2004. Clutter Reduction in Multi-Dimensional Data Visualization Using Dimension Reordering. In *Proceedings of the IEEE Symposium on information Visualization* (October 10 - 12, 2004). INFOVIS. IEEE Computer Society, Washington, DC, 89-96.
- [22] Pizlo, Z. 3D Shape: Its Unique Place in Visual Perception. The MIT Press, 2008.
- [23] Robertson, G., Card, S. K., and Mackinlay, J. D. 1989. The cognitive coprocessor architecture for interactive user interfaces. In *Proc. of the 2nd Annual ACM SIGGRAPH Symp. on User interface Software and Tech.* (Williamsburg, Virginia, United States, November 13 - 15, 1989). UIST '89. ACM, New York, NY, 10-18.
- [24] Shoemake, K. 1992. ARCBALL: a user interface for specifying three-dimensional orientation using a mouse. In *Proc. of the Conf. on Graphics interface '92* (Vancouver, British Columbia, Canada), 151-156.
- [25] Kapler, T. Eccles, R. Harper, R. Wright, W. Configurable spaces: Temporal analysis in diagrammatic contexts. In *Proc. of the IEEE Symposium on Visual Analytics Science and Technology*, 2008.
- [26] Tominski, C., Abello, J., and Schumann, H. 2005. Interactive poster: 3d axes-based visualizations for time series data. In *Poster Proceedings Compendium, IEEE Symposium on Information Visualization*, pages 49-50.
- [27] Tominski, C., Schulze-Wollgast, P., and Schumann, H. 2005. 3D Information Visualization for Time Dependent Data on Maps. In *Proc. of the Ninth international Conf. on information Visualization* (July 06 - 08, 2005). IV. IEEE Computer Society, Washington, DC, 175-181.
- [28] Tversky, B., Morrison, J., Betrancourt, M. Animation: can it facilitate? *Int. Journal of Human-Computer Studies*, 57(4):247-262, 2002.
- [29] Mayr, G. (1877). Die Gesetzmäßigkeit im Gesellschaftsleben. Oldenbourg. p.78.
- [30] Yang, J., Peng, W., Ward, O., Rundensteiner, E. 2003. Interactive hierarchical dimension ordering, spacing and filtering for exploration of high dimensional datasets. In *Proc. IEEE Symposium on Info. Vis.*, 105-112.