Int J Digit Libr (1999) 2: 144–156

INTERNATIONAL JOURNAL ON
Digital Libraries
© Springer-Verlag 1999

Interactive volumetric information visualization for document corpus management

Christopher D. Shaw[1], James M. Kukla[2], Ian Soboroff[2], David S. Ebert[2], Charles K. Nicholas[2], Amen Zwa[2], Ethan L. Miller[2], D. Aaron Roberts[3]

[1] Department of Computer Science, University of Regina, Regina, Saskatchewan, S4S 0A2, Canada;
E-mail: cdshaw@cs.uregina.ca
[2] Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, MD, 21250, USA;
E-mail: {jkukla1,ian,ebert,nicholas,zwa,elm}@cs.umbc.edu
[3] NASA Goddard Space Flight Center, Mailstop 692.0, Greenbelt, MD, 20771, USA;
E-mail: roberts@vayu.gsfc.nasa.gov

Received: 15 December 1997/Revised: June 1999

Abstract. This paper describes aminimally immersive three-dimensional volumetric interactive information visualization system for management and analysis of document corpora. The system, SFA, uses glyph-based volume rendering, enabling more complex data relationships and information attributes to be visualized than traditional 2D and surface-based visualization systems. Two-handed interaction using three-space magnetic trackers and stereoscopic viewing are combined to produce aminimally immersive interactive system that enhances the user's three-dimensional perception of the information space. This new system capitalizes on the human visual system's pre-attentive learning capabilities to quickly analyze the displayed information. SFA is integrated with adocument management and information retrieval engine named Telltale. Together, these systems integrate visualization and document analysis technologies to solve the problem of analyzing large document corpora. We describe the usefulness of this system for the analysis and visualization of document similarity within acorpus of textual documents, and present an example exploring authorship of ancient Biblical texts.

Key words: Volumetric information visualization – Minimally-immersive techniques – Two-handed interaction – Glyph rendering

1 Introduction

To find a document in the sea of information, one must embark on a search process, usually aided by computer. If one's goal is to search through a large corpus of documents for the 10 articles most relevant to a single well-defined topic, then a single query yielding a scalar indication of relevance is usually sufficient. In this traditional *information retrieval* model, the final goal is to identify and collect a small number of documents to read in detail. The remainder of the document corpus is viewed as partially or wholly irrelevant to the task, like so much ore from which gold must be extracted.

By contrast, *information analysis* seeks to understand what is happening in the collection of documents as a whole. Individual documents may indeed be read or skimmed, but only as a means to gaining a greater understanding of what is happening in the rest of the document set. In this situation, the results of many single queries must be combined in various ways so that trends may be discovered. *Information analysis* is a general term for the analysis of many documents in a collection, and is commonly used in the intelligence and defense communities. Information analysis seeks to identify trends, discover common linkages, and find clusters of similar documents. People interested in discovering emerging trends in science, medicine, technology or business may also engage in information analysis. Our term *document corpus management* refers to the use of automatic text analysis and visualization tools in the information analysis task. Document corpus management therefore shares some of the goals of Data Mining, or Knowledge Discovery from Data [16], which is the nontrivial extraction of implicit, previously unknown, and potentially useful information from data.

This paper describes a new system that aids in the document corpus management task by employing various types of document analysis. The analytical outputs are visualized using 3D volumetric visualization techniques in a minimally-immersive real-time interaction style. An intuitive two-handed interaction metaphor [17, 31] is created using three-space magnetic trackers. This two-handed interaction, in combination with stereo viewing, provides a minimally-immersive system that increases

the user's three-dimensional perception of the information space. Glyph-based volume visualization [15] and hardware acceleration techniques are used to provide interactive rendering speeds for multivariate time-varying information spaces.

2 Prior art

This paper describes a document corpus management system that is a new integration of a highly-interactive 3D visualization system named SFA with a document management and information retrieval engine named Telltale. Together, these systems integrate visualization and document analysis technologies to solve the problem of analyzing large document corpora. We now review some of the prior art for these components.

Most systems for information visualization to date have been 2D or surface-based (e.g., [2, 10, 19, 27, 37, 38]). Many of these systems provide useful visualization tools and are good for showing two-dimensional relationships. Some can show the relationship among three parameters of the information space. For example, the SPIRE system [37] shows the similarity among topics in an information space projected onto two dimensions. The frequency of documents within a topic is displayed graphically by the height of mountains in their *Themescapes* component, creating a landscape of document topics. However, to understand complex multi-dimensional relationships, users must mentally fuse several two-dimensional images.

These 2D and surface-based systems, however, fail to capitalize on the human perception system's ability to understand full three-dimensional volumetric space [23, 34, 35]. With volume visualization and careful mapping of information to visualization attributes, more dimensions of the information space can be perceived than with two-dimensional surfaces. To visualize the volumetric data, our system uses glyph rendering, displaying a small shape (glyph) at each location within the three-dimensional volume. Glyph, or iconic, visualization is an attempt to comprehensibly encode more information in the visualization [26, 28] and is an extension to the use of glyphs and icons in fields like cartography and statistics. The use of 3D glyph visualization allows the user to visually compare the results of 6 or more queries as they apply to the entire document set. For each document, the results of three of these queries determine the location of a glyph, while the remaining queries are assigned to glyph size, color, opacity, and shape. This system is, therefore, capable of visualizing complex multi-dimensional relationships vital to information analysis and document corpus management.

In *The Elements of Graphing Data* [11], Cleveland cites experimental evidence that shows that the most accurate method to visually decode a quantitative variable is to display position along a scale. This is followed in decreasing order of accuracy by length, slope angle, area, volume, and color. For distinguishing a categorical variable, shape and color can be quite effective. A 3D glyph display provides three scales along which a glyph can be measured, while a 2D display provides only two. Terrain height does contain 3D information, but only a single surface is visible. Translucent surfaces may allow scenes of higher depth complexity, but the third or fourth surface encountered along the line of sight is usually quite indistinct, and, therefore, hard to measure and compare. A 3D glyph display has the unique capability of allowing the most effective methods of visual perception, because a glyph's 3D location can give the user a reading on 3 separate scales. Because the space surrounding a glyph is unoccluded, glyph displays allow the user to see the other glyphs at greater depths, allowing potentially high depth complexity.

Several researchers [4, 8, 17] have examined the use of virtual reality environments for visualization; however, little work exists on applying these techniques to information visualization. Our two-handed minimally-immersive environment extends the interface ideas of Feiner and Beshers [7] to two-handed interaction and volume rendering.

2.1 Text indexing

In our document corpus management systems, we have combined our information visualization techniques with two emerging text indexing techniques – Latent Semantic Indexing and n-gram analysis. These techniques seek to overcome the limitations of keyword-based text analysis systems by applying statistical approaches to text processing.

In the standard approach to text indexing, documents are keyed by literal words in the document. A user's query must contain words that exactly match words in the document in order for the document to be returned to the user. This method will of course fail to retrieve documents that do not contain any of the user's suggested words. Moreover, many documents will be retrieved that use the query words in a semantically different way than what the user had in mind. The standard retrieval models implicitly treat words as if they were entirely independent, each with a single, unique meaning. We briefly review two methods that we employ that attempt to improve retrieval performance by overcoming this assumption of a one-to-one correspondence between words and meanings.

2.2 Latent semantic indexing

Deerwester et al. [13] discuss a method for automatic topic-based indexing of documents called *Latent Semantic Indexing* (LSI), which takes advantage of semantic relationships between terms and documents. These relationships are derived mathematically from implicit linkings of related terms in documents. Typically, an author

will use a particular set of words when discussing the same idea, and will use this set in a small number of contexts. As the topic changes, the word set changes. Latent semantics is simply the idea that these clusters of word usage imply a meaning that each word in isolation might not.

LSI examines the similarity of word usage in documents, and creates a reduced-dimension feature-space representation in which words that occur in similar contexts are near each other in the feature space. LSI first generates a representation that captures the similarity of usage of terms, and then retrieves documents based on this representation. LSI uses singular value decomposition (SVD) to generate the feature space. No external dictionaries or knowledge bases are used to explicitly determine *word associations* because the word associations are derived from a numerical analysis of many texts. The learned associations are derived automatically by the SVD machinery. However, a set of keywords or terms is required.

Specifically, a term-by-document ($t \times d$) matrix M is first constructed, where $M(i, j) =$ the number of occurrences of term t_i in document d_j.

This matrix is decomposed using singular value decomposition

$$M = T\Sigma D' . \tag{1}$$

The columns of T and D are orthonormal, and are called the *left* and *right singular vectors*. Σ is a diagonal matrix containing the *singular values* σ, by convention all positive and ordered by size. If M is $t \times d$ and of rank r, T is a $t \times r$ matrix, D is $d \times r$, and Σ is $r \times r$. The SVD "projects" both terms and documents into an $r-$dimensional space – the feature space of the document set. One can choose several interesting dimensions and view documents and terms graphically in the feature space [5].

2.3 N-gram processing

N-grams are overlapping n-character sequences of text in a document. In a typical document processing system using n-grams, a document is processed by sliding an n-character window across its text [12, 25]. During this process, all alphabetic characters are converted to lower case, non-alphabetic characters are converted to spaces, and multiple spaces are collapsed into a single space. For example, the first 5-grams in this sentence would consist of "for e", "or ex", "r exa", and so on. This process produces a list of overlapping n-grams that are used to generate a vector $d_i = (d_{i,1}, d_{i,2}, \cdots)$. Each element of the vector is calculated by $d_{i,k} = c_{i,k}/m_i$, where $c_{i,k}$ is the count of n-gram k in document i and $m_i = \sum_k c_{i,k}$ is the total number of n-grams in document i. While there are 27^n possible unique English n-grams, experimentation has shown that relatively few of them occur in any corpus [12]. For example, a 40 MB collection of articles from the *Wall Street Journal* has about 270 000 unique 5-grams (out of a possible total of 7.5×10^{18}), and this number increases very slowly as the corpus increases in size.

N-gram processing tools are somewhat language independent because they do not need stop word lists or other language-dependent features. The use of n-grams rather than words provides an additional benefit: indexing is robust in the face of degradation such as that caused by scanning in text using optical character recognition, speech recognition or handwriting recognition.

Telltale [25] is a dynamic hypertext environment that provides text indexing via a hypertext-style user interface for text corpora. This indexing is done using statistical techniques on n-grams to create links associating documents that are "similar".

Telltale maintains an "average" document vector a, called the *centroid*, whose elements a_k are calculated by averaging $d_{i,k}$ over all documents i in the corpus. The similarity is then calculated by taking the cosine of the representation vectors after subtracting the centroid:

$$similarity(d_i, d_j) = \frac{\sum_{k=1}^{t}((d_{i,k} - a_k) \cdot (d_{j,k} - a_k))}{\sqrt{\sum_{k=1}^{t}(d_{i,k} - a_k)^2}\sqrt{\sum_{k=1}^{t}(d_{j,k} - ak)^2}} . \tag{2}$$

The numerator in (2) is the dot product of the vectors $d_i - a$ and $d_j - a$, representing documents i and j respectively. The denominator in (2), the product of the sum of squares of each term in the respective vectors, is used to normalize the result.

A user wishing to find documents using Telltale starts a search by entering a section of a document that the user considers "good". Telltale then converts the query to an n-gram vector in the same way as it would for a document, and scores this vector against all documents in the corpus displaying the similarities in ranked order. The user provides a sample of the "content" of interest and the hypertext engine supplies a list of relevant documents by using the similarity computation.

The similarity computation works because documents about the same subject tend to use the same vocabulary and thus are likely to contain many of the same n-grams. Subtracting the centroid from each document's vector before taking the cosine allows the comparison to occur relative to the "average" document in the corpus. This subtraction is similar to removing stop words, but is more flexible because the user does not need to specify stop words.

While checking for similarity against a sample document yields better statistical performance, Telltale can also search using a simple query string. In this case, Telltale checks just the query string's n-grams for presence in the document. It scores similarity using a modified version of (2) that counts only the n-grams q_k that are

present in the query string q. $q_k = 1$ if the n-gram k is present in the query string, $q_k = 0$ otherwise.

$$simQuery(q, d_i) = \frac{\sum_{k=1}^{t}(d_{i,k} \cdot q_k)}{\sum_{k=1}^{t}(q_k)} \,. \tag{3}$$

Rather than the unidimensional list present in the original Telltale, we are experimenting with methods that allow a user to explore a text corpus using two, three or more documents as the "query". We have also extended its facilities by employing LSI to generate a feature space using the n-gram set in place of terms. Because n-grams span adjacent words in a document, we have an additional layer of granularity in term relationships by capturing local linkages of words into phrases. This provides a mechanism for LSI to discern writing style between documents.

3 Volumetric visualization with SFA

Our system for minimally-immersive document corpus management is composed of two tightly-coupled components: the minimally-immersive visualization engine, SFA, and the document management and information retrieval engine, Telltale. Telltale provides document similarity measures that are visualized by the SFA system for analyzing patterns and trends within the corpus. During visualization within SFA, the user can view the selected documents within the Telltale viewer and use these for further similarity comparisons and trend analysis.

In SFA, each document is represented by an icon or *glyph* that is located and shaded according to document attributes, such as similarity measures provided by Telltale. Six or more dimensions (attributes) of the information space can easily be visualized. Careful choice of glyph mappings is useful to identify trends and display important information from large information sets. It has many of the advantages of direct volume rendering, such as the ability to represent information on 3 orthogonal scales, while avoiding the limitations of iso-surface rendering.

3.1 Mapping information to glyphs

One advantage of glyph rendering is the ability to map characteristics of the information space to attributes of each glyph. The first mapping is to determine the glyph's location. The three-space location of each glyph allows the comparative display of three attributes of the information space. The size and shape of the glyphs can easily encode two or more information attributes. The color of the glyph allows the independent display of one dimension, usually one in which high accuracy of perception is not needed. Opacity is another glyph attribute that can encode information attributes. Care must be taken in mapping an attribute to the opacity and size of a glyph. Small values of this attribute can make the glyph completely transparent, preventing the comparison of the other attributes for this particular glyph. However, opacity mappings can be very useful to eliminate uninteresting information. For instance, if the user is only interested in documents in a corpus which talk about U.S. foreign policy, using the similarity to this theme as the opacity of the glyph will aid in discerning the relevant documents. The mapping of information attributes to the glyphs' three-dimensional coordinates can be used in a similar fashion, as can be seen in Fig. 1. This figure is a visualization of the relationship of 1833 documents to three themes. A quick examination of the image allows you to identify the 10–20 documents out of 1833 that are of interest.

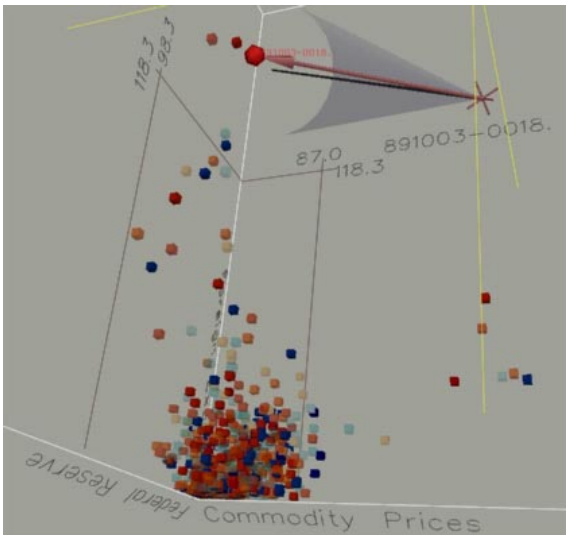

Fig. 1. Three-dimensional visualization of 1833 documents' relationship to commodity prices, foreign exchange, and the federal reserve. The probe is selecting one article

3.2 3D Rendering and perception

Upon initialization, SFA presents the user with a 3D rendering of the volume of glyphs, and a few 2D control panels (Fig. 1). The 3D volume is enclosed by a wireframe box delineating the coordinate system, with labels indicating the textual attributes that are assigned to the X, Y and Z axes. There is also a legend panel within the 3D scene that lists the complete mapping of text attributes to glyph attributes. The use of vector (Hershey) fonts for this helps maintain a strong 3D impression (Fig. 2). SFA uses 3D perspective projection, animation, plus stereo viewing through the use of Liquid Crystal Shutter Glasses. Most users find this combination to be extremely helpful in generating a 3D impression.

The volume of glyphs can be manipulated directly by user interaction. As described in Sect. 4.1, the user can rotate, translate, and scale the 3D glyph volume using either the mouse or one of two 3D trackers. When the user rotates the volume of glyphs, it rotates rigidly as if it were

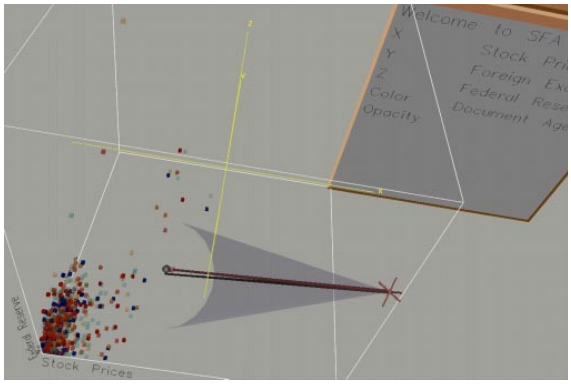

Fig. 2. Three-dimensional visualization of 1833 documents' relationship to stock prices, foreign exchange,and the federal reserve. The legend panel is also visible

a solid piece of glass with the glyphs rigidly suspended inside it. The combination of perspective projection and animation provides visual motion parallax, which arises when objects moving at the same speed parallel to the screen appear to move faster the closer they are to the viewer. The human visual system correctly perceives the different depths [23].

Similarly, the kinetic depth effect is a property of the human visual system which maintains the perceptual rigidity of objects. Even minimal visual stimuli such as a collection of dots will be perceptually maintained as being points on a rigid 3D object undergoing a rigid transformation [34]. Under normal viewing conditions, the visual system is quite resistive to distortions in the image, and in many cases where the 3D shape undergoes extensive non-rigid transformation, subjects perceive a shape-preserving 3D motion. Since the glyphs are fixed within the enclosing box, the kinetic depth effect is strongly present when the user moves the 3D volume, and the user is able to correctly understand the 3D positions of the glyphs.

SFA interactively animates the 3D volume of glyphs in a perspective view, which directly provides motion parallax and allows the kinetic depth effect to occur. Together with stereo viewing, users get a strong sense of the 3D space in SFA. When the 3D volume is stationary, the user has a weaker 3D stimulus because motion parallax is not present, but the perspective projection of the 3D volume maintains the scene's proper 3D relationships. In an experimental evaluation of a similar 3D system for viewing 3D connected graphs, Ware and Franck showed that "motion cues combined with stereo viewing can substantially increase the size of the graph that can be perceived" [35].

3.3 Animation control

To maintain a fluid interaction when the user is moving the data set, the scene must be animated in real time. SFA sets a minimum update rate of 10 frames per second, which requires certain optimizations to be performed when thousands of documents are to be displayed. To achieve this update rate, SFA only displays partially or fully opaque glyphs, caching these glyphs for quick access. SFA also draws the glyphs in order from largest to smallest, stopping when the 0.1 second time limit is reached. This allows the most visible portion of the volume to be drawn first within the interactive time constraint. This allows good interaction on machines with at least the graphics capability of an SGI O2. Depending on the geometric complexity of the glyph, between 2000 and 8000 glyphs can be animated at 10 frames per second on an O2. When the 3D volume is not being moved, there is no need to draw the glyphs at 10 frames per second, so the entire set of visible glyphs is drawn. This allows the user to rotate the largest glyphs in the scene at interactive rates, then freeze the volume and watch the entire set of glyphs appear.

4 SFA interface

SFA provides both a traditional two-dimensional mouse-keyboard interface and a two-handed minimally immersive interface. For the traditional two-dimensional interface, several control panels allow the user to control the display, including such attributes as glyph type, the current data file, and the current color and transparency map.

4.1 Two-handed minimal immersion

In addition to the 2D interface, SFA has a two-handed minimally immersive interface which uses a pair of 3D magnetic trackers. The user sits in front of a graphics console that has a screen, keyboard, mouse, and the two 3D sensors. Figure 3 shows a user using the system. Each 3D sensor, or *Bat*, has three buttons glued onto the surface, as shown in Fig. 4. The user interacts with the system by manipulating the two trackers, and pressing the tracker buttons to invoke operations. Each 3D sensor has

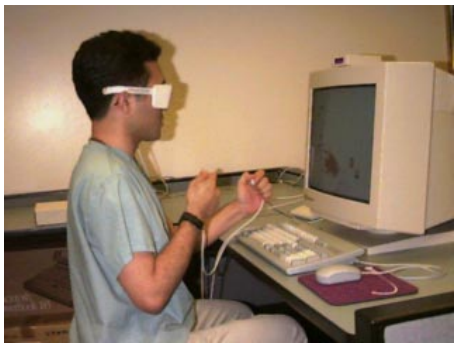

Fig. 3. A user using the two-handed stereo interface to SFA

Fig. 4. Polhemus sensor with attached buttons

a distinct role, with the dominant hand being responsible for picking and fine manipulation, and the less-dominant hand being responsible for moving the 3D volume around and setting spatial context. For the sake of rhetorical convenience, we will refer to the dominant hand as the right hand and the less-dominant hand as the left, but the system is ambidextrous because the Polhemus trackers are symmetric and can be handled with equal ease by either hand.

One of the main justifications for this approach is that complex 3D objects can be more easily manipulated by 3D tracking devices than with a mouse [36], because the user does not have to mentally break down the 3D task into a sequence of 2D operations. The use of a 3D device allows the user to directly manipulate the objects of interest without intermediate steps. Two 3D devices gives the user access to double the spatial bandwidth, since both hands can be employed in parallel to quickly achieve the desired operation.

We chose the minimally-immersive style because it does not isolate the user from traditional I/O devices like the screen, keyboard, and mouse. In our document corpus management application, text input is quite important, so ready access to the keyboard is vital. An alternative style is the Virtual Reality, or *fully immersive* style. The VR style has the drawback that the user must wear a head-mounted display, which does not allow the user to see the surrounding physical workspace of desk, screen, keyboard, and mouse. We find that the minimally-immersive style offers the advantages of true 3D input devices without throwing away the usefulness of the keyboard and mouse.

The simultaneous use of two 3D sensors takes advantage of the user's innate proprioceptive knowledge of where his/her two hands are in space. Guiard [18] shows how the left and right hands quite often act as elements in a kinematic chain. For right-handed people, the left hand acts as the base link of the chain.[1] The right hand's motions are based on this link, and the right hand finds its spatial references in the results of motion of the left hand. The right and left hands are involved in asymmetric temporal-spatial scales of motion (right hand for high frequency, left hand for low frequency). Finally, the left hand usually moves first, setting up the spatial context for the right hand to follow.

[1] For left-handed people, the roles of right and left are reversed.

In most two-handed tasks, there is a natural division of labor [18], which SFA exploits by assigning the (low-frequency) setting of spatial context to the left hand, and the (high-frequency) selection and picking operations to the right. This division of labour is more than just a slogan for the interface style. In addition to providing a criterion for assigning tasks to hands, involving both hands in this way improves spatial perception. This perception is not provided visually, but is, instead, provided by the means of the user's proprioception [20].

4.1.1 Left-hand operations

The left hand has three tasks to perform:

- Manipulate the position and orientation of the entire scene.
- Select the drawing context from a 3D tracker-based hierarchical menu.
- Select the constraint axis for bounding-box adjustment.

Scene orientation is a toggle controlled by left button 3 (the button near the wire). Clicking button 3 attaches the volume to the left cursor, and clicking it again leaves the volume in place. This clutching mechanism allows the user to spend a large amount of time moving the workpiece around without the burden of continuous button pressure. Thus, the left hand acts as a spatial context setter and sensor, allowing the user to get a better sense of the 3D space.

The middle button (button 2) on the left bat pops up a hierarchical sundial menu [31], as shown in Fig. 5, which is controlled by the orientation of the left bat. The menu choices are arrayed on the circular plate of the sundial, each on its own pie-shaped sector. The desired item is picked by pivoting the *shadow stick* about its base so that the stick's endpoint lies visually in front of the sector. The user pivots this shadow stick by manipulating the bat's orientation. This main menu in SFA contains a submenu

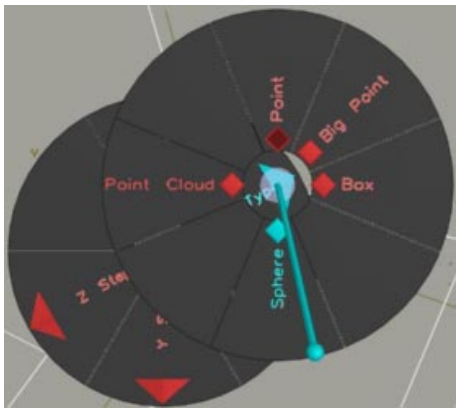

Fig. 5. A 3D hierarchical sundial menu

to select the glyph type, a submenu for each of X, Y, and Z step interval selection, and cursor control menu that scales the scene up or down and reorigins the 3D cursors to the center of the workspace.

The Sundial menu is provided to allow the user to hold the bats for a long period of time. Without the sundial menu, the right hand would have to drop the right bat and use the mouse to pick items from a regular mouse-based pull-down menu, and then pick up the right bat again. Moving the hand from one device to another takes time, so the bat-based sundial menu provides ready access to the wide range of controls that menus traditionally provide without forcing the user to continually switch devices. Assigning a menu to a bat button also allows that button to have access to more than just one fixed command when the user presses the button. Because the user maintains the bats in a steady grip, the user is able to concentrate on the information analysis task at hand, without having to think about switching devices.

4.1.2 Right-hand operations

The right hand has two tasks to perform:

- Select 3D volume bounding boxes.
- Pick a glyph to display in Telltale.

The user may decide to concentrate on a subvolume of the 3D space, and temporarily stop drawing the glyphs outside this subvolume. To select this subvolume, the user presses button 1 (furthest from the wire) on the right bat to place one corner of the subvolume bounding box, drags the right bat to the opposite corner, then releases the button. While the button is pressed, the bounding box is drawn as a semitransparent box, allowing the user to see which glyphs will be enclosed by the bounding box as it is dragged out. The bounding box disappears when the button is released, and SFA draws only the glyphs that were within the bounding box. The user may restate the bounding volume at any time by sweeping out a new bounding box.

4.1.3 Multiple subvolumes

SFA actually makes up to 8 bounding boxes available, and displays the glyphs in the union of these boxes. This is useful for eliminating documents that the user may deem irrelevant to the analysis task. For example, the large cluster of documents near zero at the bottom of Fig. 1 is not relevant to any of the given search terms. Each bounding box can be individually resized, reshaped, or deleted as the user sees fit.

This bounding box control is provided by a menu of thumbnails in a column on the right of the main window. The currently-selected bounding box has its thumbnail highlighted, and the user can restate the currently-selected bounding box by sweeping out a new box in the main window. The user may select another bounding box for editing by entering its thumbnail and pressing right button 1. The user invests some effort in stating a bounding box, so the ability to edit it allows the user to preserve that investment.

4.1.4 Glyph selection

To select a single glyph, the user uses the right bat to orient a probe into the volume, and the glyph closest to the probe has its value printed and passed through a socket connection to the Telltale system. The corresponding document can then be displayed in Telltale. The document ID is also displayed in the 3D scene using screen-aligned text at the glyph, and at the corresponding 2D locations on each of the boundary walls of the volume. Thus, the box that encloses the entire 3D volume provides 3 orthographically projected 2D surfaces which can be used to present information. In addition to the document ID, the appropriate projected X and Y location of the glyph at each boundary wall is displayed, along with an axis-aligned line drawn towards the origin of the coordinate system. This helps to provide detailed spatial measurements as the need arises.

A probe, represented by a narrow cylindrical shaft, is attached to the right cursor, and the user controls the position and orientation of the probe with the right bat. The distance from this probe is computed for each control point using a specialized distance metric called the *probe metric*, and the grid point that generates the smallest value is highlighted.

Around the probe axis is drawn a translucent half-cone which has a radial spread angle of 10 degrees. Only the back face of the cone is drawn, so that it looks like a translucent scoop, which helps to disambiguate the scene. Objects that are behind the scoop are dimmed by the translucency, thereby highlighting the closer objects.

4.2 Discussion

Together, these interaction techniques form an integrated two-handed system that can be readily used to examine and query 3D volumetric glyph displays. The SFA interface builds on an experimental two-handed 3D CAD system named THRED [31].

To evaluate THRED's ease of use, Shaw conducted a user test that had subjects use THRED to perform a 2D editing task in which an irregular 2D grid of rectangles was stretched until it was made into a regular grid of squares [32]. Subjects were also asked to perform this task using a standard mouse-based 2D computer-aided design system that was a simplified clone of the Alias editor. Both the mouse-based Alias clone and THRED had an identical underlying CAD geometry management system, identical display methods, and identical interaction operations. In a between-subjects test, subjects able to

perform this 2D grid-stretching task more quickly with THRED than with the Alias clone. We are currently in the process of conducting a formal study on the usability of SFA.

One common concern with this two-handed style is that the user may experience pain and fatigue while using this system. A result published by Shaw [33] compared the Alias clone modeler, THRED, and a single-tracker version of THRED. With the single-tracker editor, the user holds the tracker in the right hand, while the left hand presses keys on the keyboard to enter other commands. The experiment showed that subjects did not experience pain in any of the three interfaces. Subjects reported a statistically significant increase in fatigue in the *left hand* while using the One-Handed editor, probably because of the relative immobility of the left hand over the keyboard. Subjects did not report an increase in fatigue for THRED, the precursor to SFA. In a follow-up study comparing the Alias clone to THRED, subjects using THRED did not experience an increase in fatigue, while those using the Alias clone did experience a statistically significant increase in fatigue in the right shoulder.

The result of the pain and fatigue experiments shows that for short-term use (2 hour sessions), fatigue is more significant in a mouse-based interface than a two-handed interface in the style used by SFA.

5 Shape visualization

Cleveland [11] cites experimental evidence that shows the most accurate method to visually decode a quantitative variable in 2D is to display position along a scale. This is followed in decreasing order of accuracy by interval length, slope angle, area, volume, and color. Bertin offers a similar hierarchy in his treatise on thematic cartography [6]. However, these orderings of visual effectiveness are based on two-dimensional visualization systems. Some recent work has been performed on three-dimensional perception of shape, indicating that shading based on shape is perceptually significant [21].

SFA employs glyph position in 3D, 3D scale (corresponding to Cleveland's length, area and volume) and color, and in the vector-based flow visualization, slope angle. The next opportunity for encoding values is shape. One of the most difficult problems in glyph visualization is the design of meaningful glyphs. Glyph shape variation must be able to convey changes in associated data values in a comprehensible manner [26]. This difficulty is sometimes avoided by adopting a single base shape and scaling it non-uniformly in 3 dimensions. However, the lack of a more general shape interpolation method has precluded the use of shape beyond the signification of categorical values [6]. We have chosen the procedural generation of glyph shapes using superquadrics [3] because superquadrics offer the required shape interpolation mechanism.

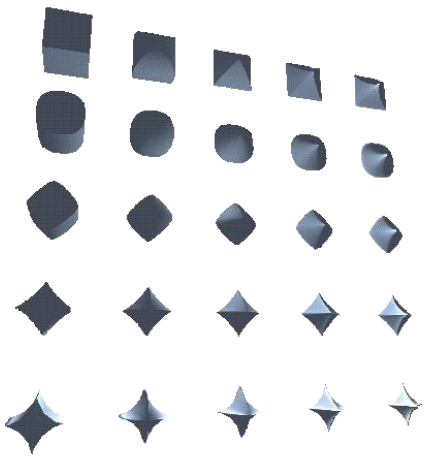

Fig. 6. Example superquadric shapes created by varying each exponent from 0 to 4

5.1 Procedural shape visualization using superquadrics

Because of the need for meaningful glyph design and the complexity of the problem, we opted for a procedural approach, which allows flexibility, data abstraction, and freedom from specification of detailed shapes [14]. Our goal for glyph design was to allow the automatic mapping of data to shape in a comprehensible, easily controllable manner. Superquadrics are a natural choice to satisfy this goal. Superquadrics [3] are extensions of quadric surfaces where the trigonometric terms are each raised to exponents. Superquadrics come in four main families: hyperboloid of one sheet, hyperboloid of two sheets, ellipsoid, and toroid. For our implementation we have chosen superellipses due to their familiarity, but the system can be easily extended to use other types of superquadrics as well as combinations of types. For example, supertoroids could be used for negative values and superellipsoids for positive values.

In the case the of superellipsoids, the trigonometric terms are assigned exponents as follows:

$$\underline{x}(\eta,\omega) = \begin{bmatrix} a_1 & \cos^{\epsilon 1}\eta & \cos^{\epsilon 2}\omega \\ a_2 & \cos^{\epsilon 1}\eta & \cos^{\epsilon 2}\omega \\ & a_3 & \sin^{\epsilon 1}\eta \end{bmatrix}, \quad \begin{matrix} -\pi/2 \le \eta \le \pi/2 \\ -\pi \le \omega < pi \end{matrix}$$

These exponents allow continuous control over the characteristics (in some sense the "roundness" or "pointiness") of the shape in the two major planes which intersect to form the shape, allowing a very simple, intuitive, abstract schema of shape specification. For example, as Barr states, $\epsilon 1 < 1$ and $\epsilon 2 < 1$ produces cuboid shapes, $\epsilon 1 < 1$ and $\epsilon 2 \sim 1$ produces cylindroid shapes, $\epsilon 1 > 2$ or $\epsilon 2 > 2$ produces pinched shapes while $\epsilon 1 = 2$ or $\epsilon 2 = 2$ produces faceted shapes. As can be seen in

Fig. 6, varying the exponents achieves smooth, understandable transitions in shape. Therefore, mapping data values to the exponents provides not only a continuous, automatic control over the shape's overall flavor, but a comprehensible shape mapping as well. By using superquadrics, we can provide the appropriate shape visual cues for discerning data dimensions mapped to glyph shape while not distracting from the cognition of global data patterns.

Glyph shape is a valuable visualization component because of the human visual system's pre-attentive ability to discern shape. Shapes can be distinguished at the pre-attentive stage [24] using curvature information of the 2D silhouette contour and, for 3D objects, curvature information from surface shading [21].

Unlike an arbitrary collection of icons, curvature has a visual order, since a surface of higher curvature looks more jagged than a surface of low curvature. Therefore, generating glyph shapes by maintaining control of their curvature will maintain a visual order. This allows us to generate a range of glyphs which interpolate between extremes of curvature, thereby allowing the user to read scalar values from the glyph's shape. Pre-attentive shape recognition allows quick analysis of shapes and provides useful dimensions for comprehensible visualization.

Our use of glyphs is related to the idea of marks as the most primitive component that can encode useful information [6]. Senay points out that shape, size, texture, orientation, transparency, hue, saturation, brightness, and transparency are retinal properties of marks that can encode information [29, 30]. To produce understandable, intuitive shapes, we are relying on the ability of superquadrics to create graphically distinct, yet related shapes. We are encoding two data dimensions to glyph shape in a manner that allows the easy separation of the shape characteristics. Not much research has been done on 3D glyphs, and superquadric glyphs are our first step in exploring this area.

We have conducted a pilot user study that indicates that the average viewer can easily distinguish at least 20 shapes along the diagonal axis of the superquadric superellipse array in Fig. 6. Subjects were presented with a target shape and a different edit shape, and were asked match the target shape by sliding a slider to adjust the edit shape. Subjects were presented each of 25 different target shapes in random order, and were given a 10 second time limit to complete each edit trial. For each target shape, the standard deviation of all the responses was used to compute the range of parameters for similar-looking shapes. This standard deviation was used to compute the difference in parameter values to the next distinguishable target. That is, the minimum distance between target A and target B should be:

$$targetParameterA + StdDev(A) + StdDev(B) = targetParameterB\,. \quad (4)$$

The result of this computation yielded a total of 20 distinguishable superellipse shapes.

Other glyph shapes have been explored in 2D displays, such as metroglyphs, star glyphs, and stick figure icons. These glyphs are not directly portable to 3D because they are flat. Superquadrics have the advantage of being quickly perceptible, and they allow a quick, intuitive mapping from information values to shape. Stars and stick figures require a certain amount of cognitive effort to decode, and therefore take time to use effectively. For this reason, they are not suitable for quick information analysis.

Since size and spatial location are more significant cues than shape, the importance mapping of data values should be done in a corresponding order. In decreasing order of data importance, data values were mapped to location, size, color, and shape. In our experience, shape is very useful for local area comparisons among glyphs: seeing local patterns, rates of change, outliers, anomalies. We chose to map either one independent variable to both glyph exponents or two related variables to each glyph exponent to ensure the understandability of the shapes.

6 Comparison to 2D visualizations

We have applied these visualization techniques and the two-dimensional IVEE system [1] to the output of Telltale. Currently, the user may select document similarities from the Telltale system to be visualized using SFA, can specify documents of interest within SFA, and see the document text within Telltale through the use of socket connections.

The document corpus utilized was 1833 articles from the *Wall Street Journal* from September 18, 1989 to October 13, 1989. Several document similarities to the following example "thematic" articles have been generated:

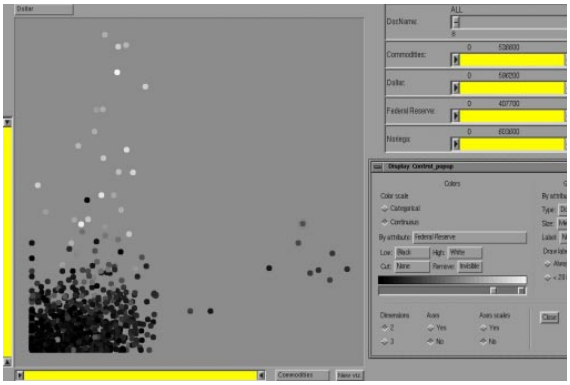

Fig. 7. Two-dimensional visualization of 1833 documents' relationship to commodity prices, foreign exchange, and the federal reserve using the IVEE system

– Manuel Noriega[2]
– Federal Reserve Bank
– Foreign Exchange Rate
– Commodity and Gold Prices
– Stock Market

Figures 7 and 1 show the visualization of similarities to commodities and gold prices, foreign exchange rate, and the Federal Reserve Bank. In Fig. 7, similarity to commodities, foreign exchange rate, and Federal Reserve Bank were respectively mapped to the X axis, Y axis, and to color. Initial analysis of the figure shows that there is not a direct relationship between the themes of commodity prices and foreign exchange rate. More detailed analysis of the figure shows that there is a relationship between the themes of foreign exchange rate and Federal Reserve Bank within this corpus. For the visualization shown in Fig. 1, the similarity to the commodity article was used as the X coordinate of the point, the similarity to the article discussing the foreign exchange rate of the dollar was used as the Y coordinate, and the similarity to the Federal Reserve article was used as the Z coordinate of the glyph. The date of the article was used in determining the color, ranging from blue (oldest) to red (most current). Figure 1 clearly shows the most relevant 25 to 30 documents to these themes out of this corpus of 1833 documents. This figure also shows two similarity patterns among several articles. The pattern in the left of the image shows several documents that discuss both the foreign exchange rate and the Federal Reserve Bank. The pattern to the right of the image clearly shows the relationship of commodity (gold prices) and the foreign exchange rate within several articles.

The volume visualization and interactive stereoscopic manipulation of this data set highlights the similarities of these topics significantly better than a 2D visualization, as can be seen by comparing Figs. 7 and 1. The two relationships shown in Fig. 1 are not immediately clear from examining Fig. 7. In fact, even after detailed examination, the relationship within several articles of commodity (gold prices) and the foreign exchange rate does not become apparent because the similarity of the documents to commodity prices (displayed as color) is not as high as some of the other articles within the corpus.

The visualization shown in Fig. 2 shows the similarity to the stock market, foreign exchange rate, and Federal Reserve Bank, respectively mapped to the X, Y, and Z axes. Again, the date of the article is visualized as the color of the glyph. The transparency of the glyph was also controlled by the article date, with the most recent and oldest articles being opaque and the middle articles semi-transparent. Through interactive visualization of this data, it is clear that there are several articles that lie along the diagonal from the origin to the front, top, right corner of the data volume, indicating a strong relationship of these three themes in several documents in this corpus. In contrast, the visualization in Fig. 8, does not clearly show the ternary relationship among these themes.

6.1 Discussion

One potential disadvantage of a 3D glyph display is that some glyphs may occlude others, thereby obscuring part of the information space. However, as Figs. 7 and 8 show, occlusion is not a problem unique to 3D. These pictures clearly show that in the high-density region near the origin, many document glyphs overlap. The only solution to this problem in 2D (or 3D) is to scale the glyphs down or scale the space up so that the glyphs may be separated in space. In 3D, if there is space between glyphs, then the user can also rotate the volume to find that space. Moreover, in a 2D display, projecting a glyph in parallel onto a 2D plane may indeed *create* overlaps in situations where there had been space between the glyphs in 3D.

Another potential area where 2D views might be more effective is if the multiple dimensions being displayed in a 3D volume are easily separable into two or more 1D or 2D orthogonal information dimensions. In this case, there is likely to be no profit in a 3D display. The information attributes are independent from each other, and the user is unlikely to learn anything new from an integrated display. For example, presumably the birthdate of the author and the similarity of the article to Federal Reserve Bank are not related to each other. A 2D scatterplot display of articles with X = Birthdate and Y = Federal Reserve Bank will not tell the user too much about the articles that could not be learned from 2 separate 1D lists. However, the information analysis task for which SFA is designed is one in which the similarity attributes of each document are unknown a priori. The task is to find "interesting" articles that have more than one significant

[2] A coup attempt in Panama to overthrow Noriega occurred during this time period.

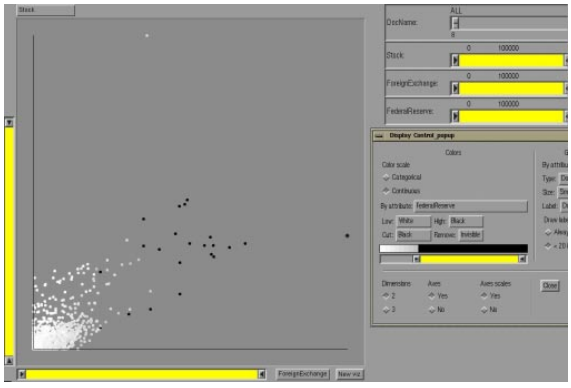

Fig. 8. Two-dimensional visualization of 1833 documents' relationship to stock prices, foreign exchange, and the federal reserve using the IVEE system

similarity rating. For this reason, a 3D display helps to visually group interesting articles together in an intuitive manner.

7 Visualizing authorship with N-grams and LSI

In 1992, Kjell and Frieder [22] analyzed the frequency of 2-grams in The Federalist Papers, in order to determine the authorship of eleven unattributed papers. In their survey, capitalization, punctuation, and spaces are ignored, so that in the phrase "Hello World", *ow* would be a 2-gram.

Kjell and Frieder compared eleven papers of unknown authorship to two prototype documents, constructed by concatenating a number of papers by two different authors. The cosine similarity was taken between an unknown-author document and both prototypes, and authorship was determined by the highest cosine similarity. Using this measure, Kjell and Frieder arrived at conclusions comparable to a classic scholarly study of The Federalist Papers. Additionally, by applying principal components analysis to a set of feature vectors, graphs were produced which demonstrate visually the stylistic relationships between the texts.

Several interesting questions are raised by Kjell and Frieder's work. Are 2-grams large enough to capture writing style? Should inter-word spaces have been included? Is cosine similarity a good measure for determining authorship? Were the feature vectors chosen for visualization good summaries of the documents? We believe that the LSI approach is a better mechanism for deriving feature vectors, which in this case are the actual documents projected into a reduced space. Moreover, the choice of length of n-grams to use is not well understood, and it is probable that preserving word and sentence divisions will improve results.

SFA can also be used as a tool for understanding and exploring various information analysis algorithms, such as latent semantic indexing. The example presented here uses latent semantic indexing, with n-grams as terms, to visualize authorship among biblical Hebrew texts. The texts consist of the books of Daniel, Song of Songs, and Ecclesiastes. In this set, which has 32 total chapters, Song of Songs and Ecclesiastes are traditionally attributed to King Solomon. Those books differ significantly in content, and both have a different style than that of Daniel. Moreover, half of the chapters of Daniel are primarily written in Aramaic, a language related to but different from biblical Hebrew. Each chapter has between 680 and 4291 characters. The text was acquired from the Snunit Educational Information System[3], part of the Israeli English Teachers Network. Inter-word spaces are preserved as a single space, and punctuation is converted to spaces, except for periods which mark the ends of verses. In the Snunit text, punctuation is present which is not standard across printed texts, so it is ignored.

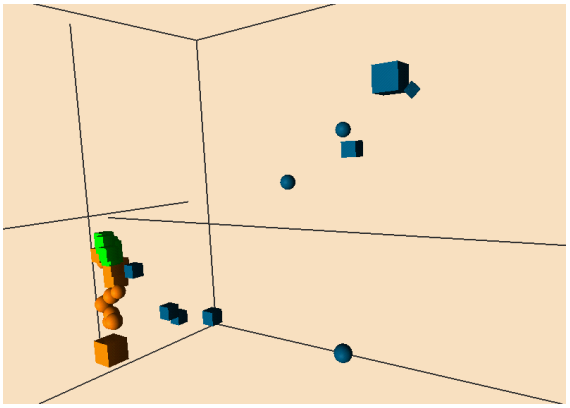

Fig. 9. A visualization of the books of the Bible. Each glyph is a book chapter, with green for Song of Songs, orange for Ecclesiastes, and blue for Daniel. The first 3 dimensions of the LSI space are mapped to X, Y, and Z, respectively. LSI dimension 1 (X, horizontal) divides the chapters according to author. Dimension 2 (Y, vertical), divides Daniel's chapters according to language: the upper right chapters are in Aramaic, while all the others are in Hebrew. Dimension 3 (Z, coming out of page to the left) shows a stylistic difference between one of the chapters of Ecclesiastes and the rest. Dimension 6 maps to size, and dimension 7 to shape

Figure 9 shows a picture of the results of the LSI analysis. Each glyph represents a chapter from one of the three books. Color represents the known authorship of each chapter, with green for Song of Songs, orange for Ecclesiastes, and blue for Daniel. The spatial location of each glyph is determined by the first three columns of the D matrix. Each column of the D matrix is scaled by its corresponding singular value, and plotted in x, y, and z respectively.

The three most significant dimensions of document interrelationships, as derived by the LSI are shown against each other. Note that in the x direction (horizontal in Fig. 9), corresponding to the first dimension, all chapters of Daniel fall to the right of the Solomon texts, and that Daniel falls into two groups. The second dimension, along the y axis (vertical in Fig. 9), cleanly divides the books such that half of the chapters of Daniel are separate from the others. Probing the Daniel glyphs to the upper right shows that these isolated chapters are the second through the seventh of Daniel, which are the ones written in Aramaic.

In this document set, the most prominent differentiating characteristic found by LSI is writing style; the second most prominent is language (based on the size of each D matrix entry). While the texts with the most outstanding values in the first dimension are the Aramaic chapters of Daniel, only one Hebrew chapter of Daniel lies among the Solomon texts. This shows that although language is an influence in the first dimension, it is not the differentiating factor here. The second dimension shows a clear

[3] http://www.snunit.k12.il/

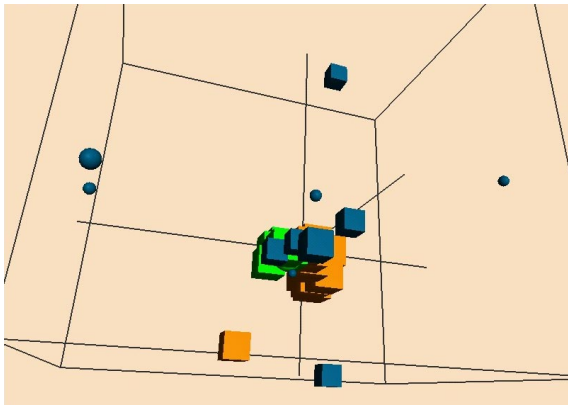

Fig. 10. A second visualization of the books of the Bible, using the same colors as in Fig. 9. Here, the fifth, sixth, and seventh LSI dimensions are used as positions. The first dimension is assigned to size and the second is assigned to shape, with the Aramaic Daniel chapters drawn as spheres. The spatial dispersion of Daniel chapters indicates many stylistic differences among them. The single outlyer Ecclesiastes chapter is again evident

division between those chapters containing Aramaic and those with only Hebrew. The third dimension also shows a differentiating effect for the Aramaic chapters of Daniel, but the size of the effect is somewhat smaller.

In addition, the seventh dimension controls glyph shape, while the sixth dictates their size. Thus, this is image shows six dimensions of the data; five from the LSI and one imposed from external knowledge. We can see that several Solomon documents have a similar glyph shape to some chapters of Daniel, for example.

SFA allows us to dynamically reassign visualization parameters to dimensions of the data. Figure 10 shows a completely different view of the same data, achieved by changing the dimensional assignments. Here, the fifth, sixth, and seventh LSI dimensions are used as positions. The first dimension is assigned to size and the second is assigned to shape, and color again shows the known authorship. This allows us to keep track of the language and author differences while we explore the document relationships in other dimensions.

These figures were taken from data which uses 3-grams. N-grams of length four or five show similar images, with subtle differences. For 4-grams, the first dimension divides more sharply along the Hebrew Aramaic difference. The use of 5-grams shows the Solomon texts grouping still more cohesively, and show the chief distinguishing factor to be language. N-grams of length two, similar to Kjell and Frieder [22] but including inter-word spaces, did not show as clean a differentiation for style as did 3-grams.

7.1 LSI Discussion

Using latent semantic indexing with n-grams as terms can group documents by authorship. This method allows elements of language and style to be indexable attributes of the corpus. Moreover, latent semantic indexing produces several dimensions of possible interest, many of which can be explored simultaneously using SFA's interactive visualization techniques.

One significant complication is that the technique is sensitive to length of the n-grams. Varying the length of n can change the LSI space considerably, because not only do the n-grams themselves carry more (and possibly different) information when they are larger, but there are simply more terms to work with for larger n. The choice of "good" n depends on language and goals. It seems that smaller n are better at capturing style characteristics, because whole words only emerge statistically rather than as literal n-grams. Longer n-grams, especially in Hebrew, can hold one or even two words in some cases.

Understanding of the LSI-generated space is not yet sophisticated to the point that a program could automatically classify documents by writing style. We are able to decide which dimension delineates style because of external knowledge; the LSI creates the space only based on the n-grams from the documents. While algorithms for spatial clustering may be applicable, it is not clear how to algorithmically choose the dimension which highlights writing style. We hope to explore these issues within the scope of visualizing corpus changes over time.

8 Discussion and conclusion

The results displayed in these figures clearly show that volumetric glyph visualization is better for visualizing and analyzing complex relationships among documents. For analysts who must manage large corpora and have many data sets to be analyzed each day, the benefit of our SFA system over traditional 2D visualization systems is that the user can readily understand complex relationships. Three-dimensional spatial analysis and two-handed interaction with stereoscopic viewing allow complex relationships can be analyzed more quickly. Perceptually meaningful mappings of data to glyph color, size, transparency, and shape also aid in the understanding of multi-dimensional data relationships.

These visualizations have shown that real-time two-handed interactive volume visualization of document corpora can quickly convey trends and multi-dimensional relationships among the information data. Careful attribute mapping and interactive volume culling can be used for quick elimination of useless data, allowing better analysis of the most interesting information. The two-handed interaction metaphor is a very natural way to interact with volumetric data, and a naive user can become comfortable and proficient using the system within 5 to 10 minutes. The current system permits interactive picking of document glyphs, allowing the retrieval of the document text in the Telltale system for further exploration. SFA provides great flexibility in information mapping for

improved perception of relationships and also allows for the real-time interaction, navigation, manipulation, and increased understanding of time-varying multivariate information spaces.

References

1. Ahlberg, C., Wistrand, E.: Ivee: An information visualization and exploration environment. In: Proc. Information Visualization '95, pp. 66–73, October 1995
2. Anupam, V., Dar, S., Leibfried, T., Petajan, E.: Research report: DataSpace: 3-D visualizations of large databases. In: Proceedings Information Visualization '95, pp. 82–88, October 1995
3. Barr, A.: Superquadrics and angle-preserving transformations. IEEE Computer Graphics Applications 1(1):11–23, 1981
4. Bergman, L.D., Richardson, J.S., Richardson, D.C., Brooks, Jr., F.P.: VIEW – an exploratory molecular visualization system with user-definable interaction sequences. In: Kajiya, J.T. (ed): Computer Graphics (SIGGRAPH '93 Proceedings), vol. 27, pp. 117–126, August 1993
5. Berry, M.W., Dumais, S.T., O'Brien, G.W.: Using Linear Algebra for Intelligent Information Retrieval. SIAM Review, 37:573–595, 1995
6. Bertin, J., : Semiology of graphics, 1983
7. Beshers, C., Feiner, S.: Autovisual: Rule-based design of interactive multivariate visualizations. IEEE Computer Graphics and Applications 13(4):41–49, 1993
8. Bryson, S., Levit, C.: The Virtual Wind Tunnel. IEEE Computer Graphics and Applications 12(4):25–34, 1992
9. Callahan, J., Hopkins, D., Weiser, M., Shneiderman, B.: An Empirical Comparison of Pie vs. Linear Menus. In: Proceedings of ACM CHI '88 Conference on Human Factors in Computing Systems, pp. 95–100, ACM SIGCHI, Washington, DC, 1988
10. Chuah, M.C., Roth, S.F., Mattis, J., Kolojejchick, J.: SDM: Malleable information graphics. In Proceedings Information Visualization '95, pages 36–42, October 1995
11. Cleveland, W.S.: The Elements of Graphing Data. Wadsworth Advanced Books and Software, Monterey, Ca: 1985
12. Damashek, M.: Gauging similarity with N-Grams: Language-independent categorization of text. Science 267:843–848, 1995
13. Deerwester, S., Dumais, S.T., Furnas, G.W., Landauer, T.K., Harshman, R.: Indexing by Latent Semantic Analysis. J. Am. Soc. Information Sci. (JASIS), 41:391–407, 1990
14. Ebert, D.S.: Advanced geometric modeling. In: Tucker Jr., A. (ed): The Computer Science and Engineering Handbook, chapt. 56, CRC Press, 1997
15. Foley, J.D., McMath, C.F.: Dynamic process visualization. IEEE Computer Graphics Applications 6(3):16–25, 1986
16. Frawley, W.J., Piatetsky-Shapiro, G., Matheus, C.J.: Knowledge discovery in databases: An overview. AI Magazine 13:57–70, 1992
17. Goble, J.C., Hinckley, K., Pausch, R., Snell, J.W., Kassel, N.F.: Two-Handed Spatial Interface Tools for Neurosurgical Planning. Computer 28(7):20–26, 1995
18. Guiard, Y.: Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. J. Motor Behavior 19(4):486–517, 1987
19. Hedley, R.J., Drew, N.S., Wood, A.M., Beale, R.: Case study narcissus: Visualizing information. In: Proceedings Information Visualization '95, pp. 90–96, October 1995
20. Hinckley, K.P., Pausch, R., Proffitt, D.: Attention and visual feedback: The bimanual frame-of-reference. Symposium on Bimanual Specialization, Aug 17, 1996
21. Interrante, V., Rheingans, P., Ferwerda, J., Gossweiler, R., Filsinger, T.: Principles of Visual Perception and its Applications in Computer Graphics. In: SIGGRAPH 97 Course Notes, No. 33. ACM SIGGRAPH, August 1997.
22. Kjell, B., Frieder, O.: "visualization of literary style". In: IEEE International Conference on Systems, Man, and Cybernetics. IEEE, October 1992
23. Kling, J.W., Riggs, L.A., 17 contributors: Woodworth and Schlosberg's Experimental Psychology. 3rd edition edition, New York, Holt, Rinehart and Winston, Inc., 1971
24. Parker, A.J., Christou, C., Cumming, B.G., Johnston, E.B., Hawken, M.J., Zisserman, A.: The Analysis of 3D Shape: Psychophysical Principles and Neural Mechanisms. In: Humphreys, G.W. (ed): Understanding Vision, chapter 8. Blackwell, 1992
25. Pearce, C.E., Nicholas, C.: TELLTALE: Experiments in a Dynamic Hypertext Environment for Degraded and Multilingual Data. J. Am. Soc. Information Sci. (JASIS) 47, 1996
26. Post, F.J., van Walsum, T., Post, F.H., Silver, D.: Iconic techniques for feature visualization. In: Proceedings Visualization '95, pp. 288–295, October 1995
27. Rennison, E.: Galaxy of news: An approach to visualizing and understanding expansive news landscapes. In: UIST '94, ACM Symposium on User Interface Software, pp. 3–12, 1994
28. Ribarsky, W., Ayers, E., Eble, J., Mukherja, S.: Glyphmaker: creating customized visualizations of complex data. IEEE Computer 27(7):57–64, 1994
29. Senay, H., Ignatius, E.: A knowledge-based system for visualization design. IEEE Computer Graphics Applications 14(6):36–47, 1994
30. Senay, H., Ignatius, E.: Rules and principles of scientific data visualization. ACM SIGGRAPH HyperVis Project, `http://ironduke.cs.gsu.edu/classes/hypervis/percept/visrules.htm`, 1996
31. Shaw, C., Green, M.: THRED: A Two-Handed Design System. Multimedia Systems J. 3(6):to appear, 1995
32. Shaw, C.D.: Two-Handed 3D User Interfaces. PhD thesis, University of Alberta, Edmonton, Alberta, Canada, 1997
33. Shaw, C.D.: Pain and Fatigue in Desktop VR: Initial Results. May 1998
34. Wallach, H., O'Connell, D.N.: The Kinetic Depth Effect. J. Experimental Psychology (45):205–217, 1953
35. Ware, C., Franck, G.: Evaluating stereo and motion cues for visualizing information nets in three dimensions. ACM Trans. Graphics 15(2):121–140, 1996, ISSN 0730-0301
36. Ware, C., Jessome, D.R.: Using the bat: A six-dimensional mouse for object placement. IEEE Computer Graphics Applications 8(6):65–70, 1988
37. Wise, J., Thomas, J., Pennock, K., Lantrip, D., Pottier, M., Schur, A., Crow, V.: Visualizing the non-visual: Spatial analysis and interaction with information from text documents. In: Proc. Information Visualization '95, pp. 51–58, 1995
38. Wright, W.: Research report: Information animation applications in the capital markets. In: Proc. Information Visualization '95, pp. 19–26, 1995