Two-Handed Volumetric Document Corpus Management

David S. Ebert, Amen Zwa, and Ethan L. Miller
University of Maryland Baltimore County

Christopher D. Shaw
University of Regina

D. Aaron Roberts
NASA

To find a document in the sea of information, you must embark on a search process, usually computer aided. In the traditional *information retrieval* model, the final goal is to identify and collect a small number of documents to read in detail. In this case, a single query yielding a scalar indication of relevance usually suffices.

In contrast, *document corpus management* seeks to understand what is happening in the collection of documents as a whole (that is, to find relationships among documents). You may indeed read or skim individual documents, but only to better understand the rest of the document set. Document corpus management seeks to identify trends, discover common links, and find clusters of similar documents. The results of many single queries must be combined in various ways so that you can discover trends.

We describe a new system called the Stereoscopic Field Analyzer (SFA) that aids in document corpus management by employing 3D volumetric visualization techniques in a minimally immersive real-time interaction style.

This 3D volumetric interactive information visualization system for document corpora combines two-handed interaction and stereoscopic viewing with glyph-based rendering of the corpora contents.

System overview

Our system for minimally immersive document corpus management consists of two tightly coupled components: the minimally immersive visualization engine, SFA, and the document management and information retrieval engine, Telltale. Telltale[1]—a dynamic hypertext environment for text corpora—provides text indexing, management, and retrieval based on *n*-grams—*n* character sequences of text. Telltale provides document similarity measures (*n*-gram-based *m*-dimensional vector inner products) visualized by SFA for analyzing patterns and trends within the corpus. SFA allows the interactive visualization of these trends and similarities through glyph-based volume visualization. Two-handed interaction within SFA lets you interactively select documents for viewing and further analysis within the Telltale system.

In SFA, each document is represented by an icon or *glyph* located and shaded according to document attributes provided by Telltale, such as similarity measures. You can visualize six or more dimensions (attributes) of the information space. The three-space location of each glyph permits the comparative display of three attributes. Glyph size and shape easily encodes two or more additional information attributes. The glyph's color allows the independent display of one dimension, usually one not needing high accuracy of perception. Additionally, opacity serves as a filtering attribute. Careful choice of glyph mappings helps you identify trends and perceive important information from large information sets.

SFA interface

With SFA, the user sits in front of a graphics console that has a screen, keyboard, mouse, stereo glasses, and a pair of six-degree-of-freedom (6DOF) magnetic trackers (*bats*) that provide a two-handed minimally immersive interface, as shown in Figure 1a. Each bat has three buttons glued onto the surface, as shown in Figure 1b. Using the bat lets the user directly manipulate the 3D objects of interest without intermediate steps.

We chose the minimally immersive style because it does not isolate the user from accessing information on the desktop or from traditional I/O devices like the screen, keyboard, and mouse—especially important in text applications. At the same time it provides most of the advantages of immersive techniques.

The user interacts with 3D elements of the system by manipulating the two bats and pressing bat buttons to invoke operations. Each bat has a distinct role, with the dominant hand responsible for picking and manipulation, and the less-dominant hand responsible for context setting of various kinds. Our system is ambidextrous, but 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. Using both hands improves spatial perception through the user's proprioception.[2]

The left hand manipulates the position and orientation of the entire scene and sets the drawing context from

a hierarchical 3D tracker-based sundial menu.[3] The circular sundial puts each menu choice in its own pie-shaped sector. Left-bat orientation controls the shadow stick selector to pick the desired item.

The right hand can select a 3D volume subset by sweeping out a 3D translucent box. To sweep, the user presses a right button to place one corner of the volume to be displayed, drags the right bat to the opposite corner, then releases the button. This automatically culls the glyphs inside or outside of the box (user selectable). As the selection box is swept out, the front and back walls are drawn as translucent rectangles so that the user can see the glyphs being selected.

The right hand can also select a single glyph by pointing a laser-like probe into the volume. The glyph closest to the probe has its value printed and passed through a socket connection to the Telltale system, which displays the corresponding document. 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 volume's boundary walls. The probe, represented by a narrow cylindrical shaft, attaches to the right cursor. The user controls the position and orientation of the probe with the right bat.

Results

Using Telltale output, we compared our visualization techniques to 2D techniques used by the Information Visualization and Exploration Environment (IVEE) system.[4] The document corpus consisted of 1,833 articles from the *Wall Street Journal* from September 18, 1989 to October 13, 1989. We generated document similarities to the following "thematic" articles: Manuel Noriega (a coup attempted to overthrow Noriega in Panama during this period), Federal Reserve Bank, foreign exchange rate, and gold prices.

Figures 2a and 2b show the results from each system. In Figure 2a, similarity to gold, foreign exchange rate, and Federal Reserve Bank were mapped to the *X* axis, *Y* axis, and color, respectively. Initial analysis of the figure does not reveal the relationship between the themes of gold price and foreign exchange rate within articles in this corpus. The figure does show a relationship between the themes of foreign exchange rate and Federal Reserve Bank.

1 (a) Using SFA's two-handed stereo interface. (b) A 3D tracker with buttons attached.
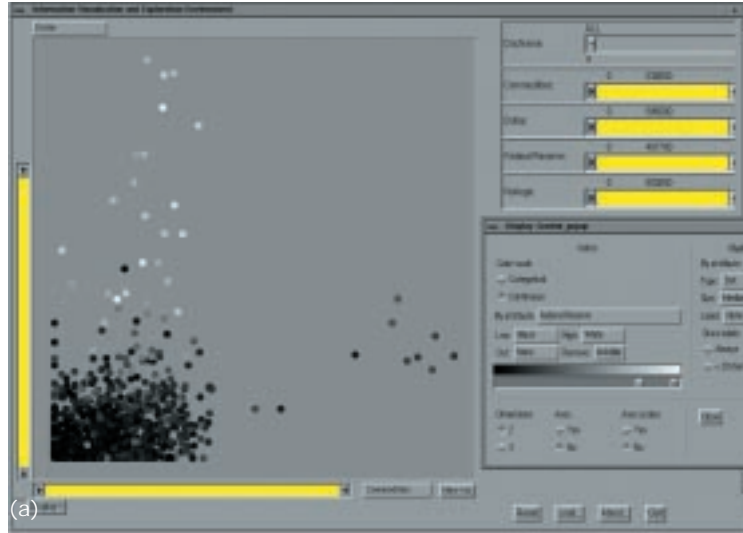
2 (a) 2D visualization of 1,833 documents' relationship to commodity prices, foreign exchange, and the federal reserve using IVEE. (b) 3D visualization using SFA. The probe is selecting one article. See http://computer.org/pubs/cg&a/extras/g4ebe.htm for a video demonstration of this data set.
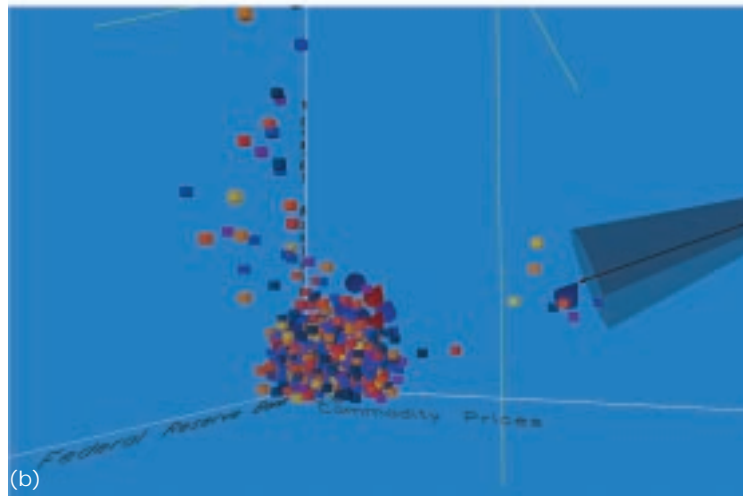
For the visualization shown in Figure 2b, similarity to gold, foreign exchange rate, and Federal Reserve Bank were mapped to the *X*, *Y*, and *Z* axis location of each document glyph, respectively. Glyph color represents article date, ranging from blue (oldest) to red (most current). Glyph shape represents similarity to Noriega, with cubes depicting little similarity and cones greatest similarity. Figure 2b clearly shows the 40 documents most relevant to these themes out of this large corpus. This figure also shows two binary relationships among several articles: the relationship between the themes of

foreign exchange rate and Federal Reserve Bank (left data branch) and the relationship between the themes of gold price and foreign exchange rate (right data branch).

Finally, this figure shows that only one document discusses the effect of the coup attempt against Noriega on the price of gold (upper right cone in Figure 2b).

These initial results show that volume visualization and interactive stereoscopic manipulation of some data sets can highlight the similarities of topics better than 2D visualization. Volume visualization represents trinary relationships better because spatialization is a more significant visual cue than color. Glyphs are also a valuable tool for visualizing large corpora because of the improved spatialization compared to 2D techniques. We need more rigorous studies to fully characterize the relative advantages of 2D and 3D techniques for information visualization.

Conclusion

These visualizations have shown how real-time, two-handed, interactive volume visualization of document corpora can quickly convey trends and multidimensional relationships among the data. Careful attribute mapping and interactive volume culling quickly eliminate useless data, allowing better analysis of the most interesting information.

Our experience shows that the two-handed interaction metaphor is a natural way to interact with volumetric data, and a naive user can become comfortable and proficient using the system within 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 large time-varying multivariate information spaces. ∎

References

1. C.E. Pearce and C. Nicholas, "Telltale: Experiments in a Dynamic Hypertext Environment for Degraded and Multilingual Data," *J. Am. Soc. of Info. Science (JASIS)*, Vol. 47, Apr. 1996, pp. 263-275.
2. K.P. Hinckley et al., "Attention and Visual Feedback: The Bimanual Frame-of-Reference," *Proc. ACM Siggraph Symp. On Interactive 3D Graphics,* ACM Press, New York, 1996, pp. 121-126.
3. D. Ebert et al., "Two-Handed Interactive Stereoscopic Visualization," *Proc. IEEE Visualization 96*, ACM Press, New York, Oct. 1996, pp. 205-210.
4. C. Ahlberg and E. Wistrand, "IVEE: An Information Visualization and Exploration Environment," *Proc. Information Visualization 95*, Oct. 1995, IEEE Computer Soc. Press, Los Alamitos, Calif., pp. 66-73.

David S. Ebert *is an assistant professor of computer science and electrical engineering at the University of Maryland Baltimore County. His research interests include scientific, information, and medical visualization, realistic rendering and animation of gases, and procedural modeling and animation. He received his BS, MS, and PhD in computer and information science from the Ohio State University in 1986, 1988, and 1991, respectively.*

Chris Shaw *is an assistant professor of computer science at the University of Regina. His research interests focus on virtual reality and two-handed 3D user interfaces for scientific visualization and for the design of free-form surfaces. He is co-architect of the MR Toolkit. He received a PhD in 1997 and an MS in 1988 in computing science from the University of Alberta. He received a BMath in computer science from the University of Waterloo in 1986.*

Ethan Miller *is an assistant professor at the University of Maryland Baltimore County, where he is working on massive and high-speed storage systems, information retrieval, and new directions in operating systems. He received a BS from Brown University in 1987 and an MS and PhD from the University of California, Berkeley in 1990 and 1995, respectively.*

D. Aaron Roberts *is an astrophysicist at NASA Goddard Space Flight Center, specializing in the computer simulation and observational analysis of interplanetary turbulence. He has worked for many years on the design of stereo visualization methods to enhance the scientific return from his physics research. He received his BS and PhD degrees in physics from Massachusetts Institute of Technology.*

Amen Zwa *is a PhD student in computer science at the University of Maryland Baltimore County. He received his BSEE in 1991 from the Catholic University of America. After working in the software industry for four years, he went back to obtain his MS in computer science from UMBC in 1996.*

Readers may contact Ebert at ebert@umbc.edu, Shaw at cdshaw@cs.uregina.ca, and Miller at elm@acm.org. Or, write Ebert at Computer Science and Electrical Engineering Dept., Univ. of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250. Contact Shaw at the Dept. of Computer Science, University of Regina, Regina, Saskatchewan S4S 0A2, Canada.