Visualization Viewpoints

Editors: Theresa-Marie Rhyne and Lloyd Trinish

Are We There Yet? Exploring with Dynamic Visualization

Penny Rheingans

University of Maryland, Baltimore County

Imagine yourself driving a classic convertible along a meandering road through the summer countryside. You revel in the play of light on the foliage, on the shape and texture of the landscape, and on the feel of the breeze through your hair. Each bend of the road reveals an intriguing and unsuspected surprise about your surroundings. You feel the beauty and peace of the scenery seep into your consciousness as you begin to grasp the essential nature of this corner of the earth. Suddenly, a small voice pipes up from the backseat, Are we there yet?

A majestic, or even ordinary, landscape can't really be captured in a handful of picture postcards. The images, no matter how lovely, are no substitute for an actual journey through that territory, showing us only glimpses of what we can discover there. Even a guided tour from the comfortable seat of a tour bus limits our exploration of the landscape to safe and frequently visited vistas. More adventurous travelers who leave the beaten path may find unique treasures, or they may find roads blocked by flocks of sheep. Either way, they come away with an experience richer and deeper than the perfect scenic snapshot.

Visualization should be like that drive through the countryside, only without the whining from the back seat.

Getting there is half the fun

Interactive graphics practitioners have long understood that viewing a virtual object by controlling the viewpoint dynamically is more illuminating than viewing a still image or even a precomputed animation.[1] Dynamic manipulation engages a viewer's kinesthetic sense in addition to his visual sense, adding an immediacy to the exploration experience.

Finding the right way to represent data has been an active topic of much thought and discussion since the beginnings of visualization. The topic continues to generate interest, as recent Visualization Viewpoints columns in this magazine show. In the January/February 2001 issue, Mike Bailey discusses the power and promise of interactive direct volume rendering, describing the ability to sculpt curves that describe transfer functions.[2] The article, "The Transfer Function Bake-Off," in the May/June 2001 issue compares promising approaches to transfer function design for direct volume rendering.[3] Similarly, visualization tools, both research and commercial, have long provided facilities for changing elements of visualization mapping. In each of these cases, however, interaction serves primarily as a means to the end of finding a good representation.

But what if interaction with the method of visualization were a goal itself rather than just a means to a good isolevel or color mapping? Just as directly controlling a virtual scene's view is valuable in its own right, dynamic manipulation of visualization mapping parameters can spark insight that viewing a single representation or animation does not. A viewer who can directly control an isosurface's threshold, a color scale's components, or a height mapping's scaling factor can interact more easily and broadly with the data. Figure 1 shows four snapshots of the iron protein with a dynamically controlled isolevel. Direct control of the mapping parameters highlights interesting regions where structure appears as isolevel changes (for example, near the "shoulders" of the molecule). Although we can see these regions in an animation (or even in the sequence of stills in Figure 1), we can explore the nature of the regions more fully by controlling our view.

When discussing the power of interaction, I distinguish between dynamic and interactive control. With *interactive* parameter control, the displayed image only updates periodically, such as when the user releases a button or selects a menu. With *dynamic* manipulation, a displayed image changes as the viewer moves a continuous input device, such as a slider, joystick, mouse, or tracker. The researcher not only sees the initial and final representations, but also the representations in between. Dynamic manipulation creates an illusion of directly manipulating the object under study, rather than invoking invisible entities to alter the object. This process of interacting with the data by moving the control devices and seeing the representation change in response, as much as viewing the individual representations contribute to the researcher's understanding of the data.

In visualization, the distinction between an interactive display and a dynamic display can be subtle. For example, if manipulating a virtual dial controls the isolevel value in a volume representation, the display would be interactive if the display updated when the user released the mouse button, but would be dynamic if the display updated as the user turned the virtual dial. While many display techniques exist that employ dynamic elements, especially those offering dynamic viewpoint control, only a few involve dynamic manipulation of the mapping from data to geometry, color, or other visual attributes.

Dynamic visualization can lead to better understanding for several reasons. Most obviously, multiple representations are better than a single representation. For a data set, a certain representation may show one kind of relationship between data elements while another representation better shows a different relationship. Second, dynamic representations more clearly present information about variable spatial derivatives and relative contributions as well as raw variable values. For instance, as a researcher manipulates a color mapping, colors move across the image surface in a continuous manner, which shows the local change rate of variable values. Third, dynamic control of the mapping builds an intuitive link between the control motions that a user performs and the visual results of those control motions.

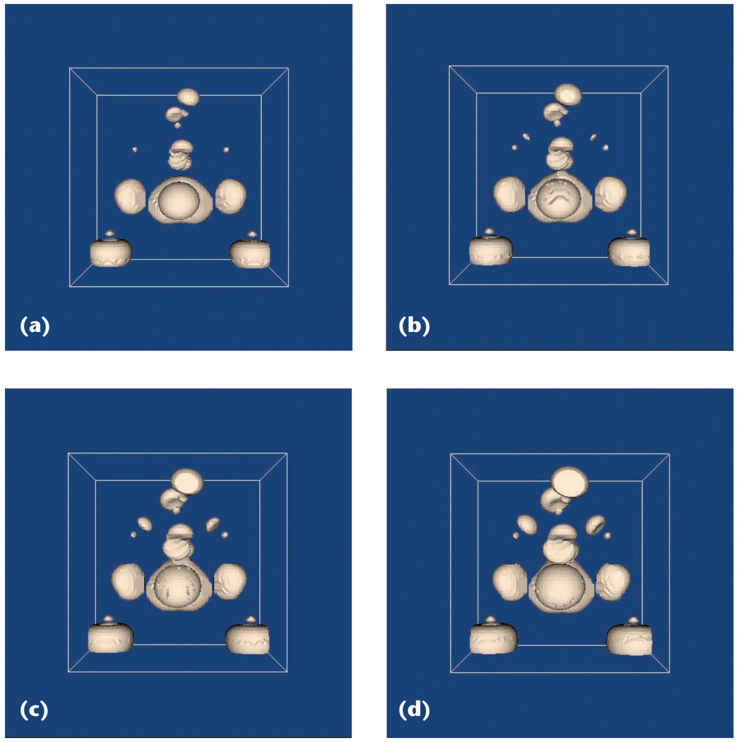

1 Iron protein with a decreasing isolevel. New features become visible during interaction.

Figure 2 shows a sequence of color mappings formed by manipulating the saturation component of all colors in the scale. The data predict income level from other characteristics, displayed on a Kohonen map. The mappings indicate education level by brightness and gender by hue. The black contour line at the right side of the images in Figure 2 separate those likely to make a high income (inside the contour) from those likely to make a

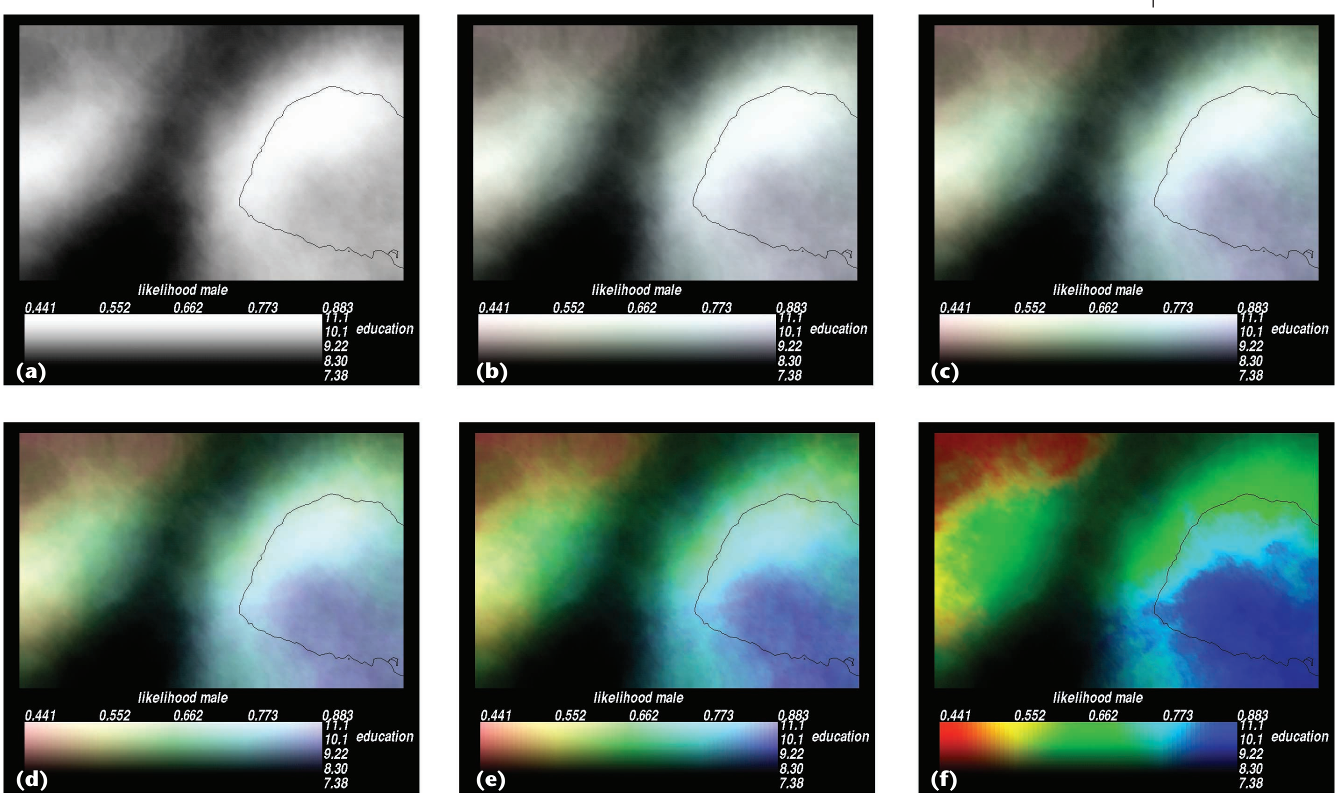

2 Six subtly different income data visualizations formed by manipulating the saturation component in the color scale.

low income (outside the contour). In Figure 2a, with zero saturation, only the education level is visible, giving no clear explanation why people in the cluster of high education to the right make high incomes while those in the cluster to the left don't. As increasing saturation makes hues more apparent, the two regions become differentiated by gender. Income data visualizations formed by manipulating the saturation component in the color scale regions become differentiated by gender. Specifically, highly educated men (in bright blue and cyan) are likely to make good incomes; highly educated women (in bright yellow and red) might not. Although both variables are visible in Figure 2f, the process of seeing the education distribution—both with and without the gender information, and particularly manipulating them in smoothly varying combinations—sparks questions and insights that come more slowly from static images.

A trip report

To explore dynamic visualization's power, I constructed a tool (see the "Calico" sidebar) for creating and manipulating bivariate color mappings using several different color models.

Using Calico, I conducted two experimental studies of the effects of control over the color mapping on accuracy, confidence, and preference.[4,5] The first study investigated the effects of control (animated, interactive, and dynamic) and update rate on understanding quantitative bivariate socioeconomic data. Figure A in the "Calico" sidebar displays three samples and shows income and education levels.

I asked subjects to answer questions about the value of one of two displayed variables in a particular US county. I found that increasing control increased subject preference, accuracy, and confidence while changing the update rate had little effect. In fact, several subjects expressed frustration with animated representations that changed in ways they couldn't control. The experiment showed that control increased accuracy (an average of 39 percent) and confidence in tasks

Calico: A Dynamic Colormap Manipulation Tool

A color space is represented by colored samples scattered through the portion of the space that's displayable on the current display. These samples are regularly spaced in the color space's underlying coordinate system. In Calico, we can create, view, and manipulate color mappings using any one of the RGB, hue, lightness, and saturation (HLS), or CIE LUV color models. When desired, the color space appears in the screen's center, with the color scale represented by a curve or surface within the space. In Figure A, I turned off the color space display and only the color sheet appears. The color scale appears in the screen's lower right. The screen's upper left portion contains the image space representing a 2D data set using the current color scheme. Users can make changes to the mapping in the color space and the image space reflects them in real time.

Users can implement color paths and sheets using splines through a set of control points. They can also generate color sequence control points from parametric expressions that specify the color component values in terms of data variable values and physical input device positions. The parametric expressions can contain symbolic data variables (u, v), dynamic input device variables, arithmetic operators, and built-in function calls.

Expressions containing input device variables are re-evaluated by the interpreter when the user moves corresponding input device. Users can edit the color sheet by affine transformations or by grabbing and dragging control points of the sheet with a joystick. During such free-form deformations, a physical slider controls the degree to which the grabbed point drags along nearby control points. Both the example image and the color sequence's geometry change dynamically as users manipulate the input devices.

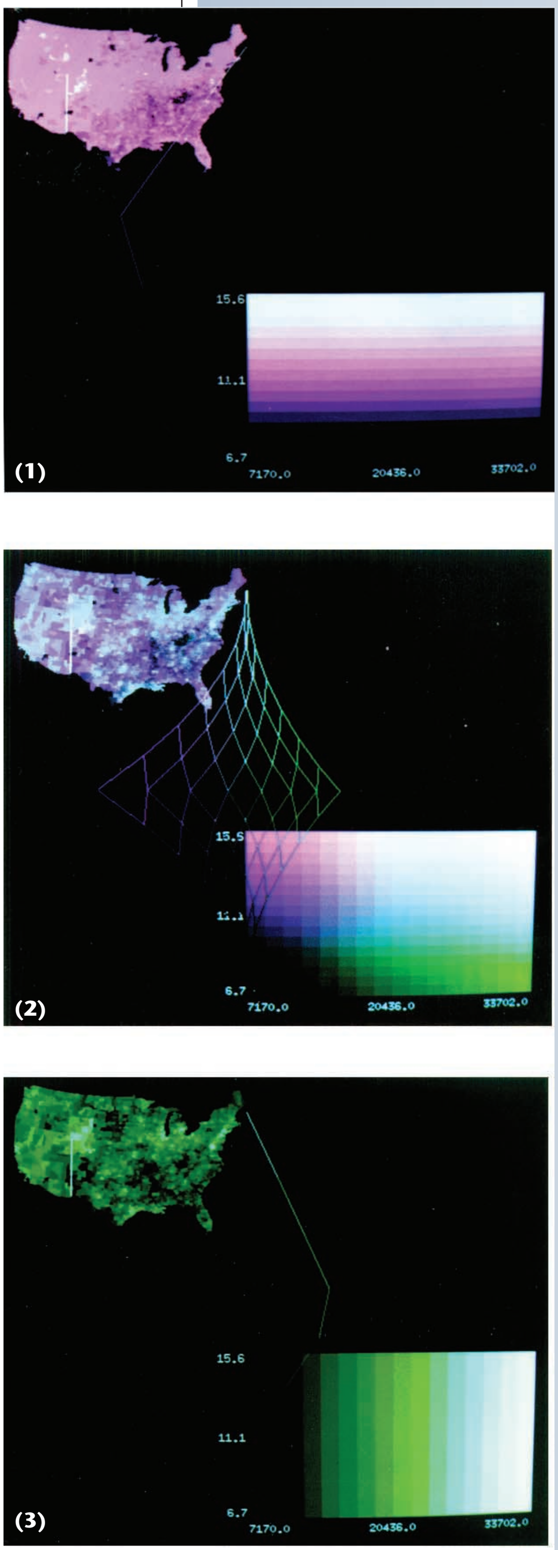

A Three Calico screens showing differing balance between education and income.

requiring judgments about the data value at a specified position (such as quantitative understanding).

The second study looked at judgments about the data's qualitative nature (feature shape, position, and height), comparing a dynamic visualization to a static bivariate display. The results showed that dynamic representations offer significant advantages for shape identification tasks without sacrificing accuracy in comparing heights or positions. Shape identifications were 49 percent more accurate in dynamic displays than in static ones, a statistically significant difference. Confidence in the subjects' answers also increased when they had dynamic control.

These experiments echoed what Held and Hein showed with kittens—that control over what we see has an importance and power of its own (for more information on this, please see the "Further Reading" sidebar). In the first experiment, dynamic representations outperformed animated representations showing all the same images. In the second experiment, the dynamic representation topped the most informative static view. In both cases, meandering along a scenic route that the viewer chose provided the richest experience and greatest understanding.

Further Reading: Some Experimental Findings about Dynamic View Control

Psychological research supports the importance of dynamic control. Held and Hein [1] showed that control over visual experience, rather than just the visual experience itself, is necessary for normal development in cats' visual systems. Kittens that passively received the identical visual stimulation never developed the ability to perform visually guided behaviors. More recent research in interactive graphics also shows the importance of direct dynamic control in shape recognition and structure comprehension.[2,3] Other research reinforces the importance of immediate and natural interaction for both a sense of presence and an ability to find things in a virtual space.[4,5]

Reference

1. R. Held and A. Hein, "Movement-Produced Stimulation in the Development of Visually Guided Behavior," *J. Comparative and Physiological Psychology*, vol. 56, no. 5, 1963, pp. 872-876.
2. W. van Damme and W. van de Grind, "Active Vision: Exploration of Three-Dimensional Structure," *Vision Research 11*, 1994, pp. 1581-1587.
3. C. Ware and G. Frank, "Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions," *ACM Trans. Graphics*, vol. 15, no. 2, Apr. 1996, pp. 121-140.
4. M. Usoh et al., " Walking > Walking-in-Place > Flying, in Virtual Environments," *Proc. Siggraph 99*, ACM Press, New York, 1999, pp. 359-364.
5. R.P.D. Proffitt and G. Williams, "Quantifying Immersion in Virtual Reality," *Proc. Siggraph 97*, ACM Press, New York, 1997, pp. 13-18.

Off-road package for visualization

Just as leaving the beaten path while touring can require special equipment, so can venturing off the road through data space. Well-equipped visualization software can enhance dynamic rambles through this landscape and hopefully keep travelers from getting stuck in gigantic information potholes. We can already find some off-road options in existing visualization systems. Others would be easy to add to existing systems, and yet others might require a completely new architecture. There are a few key options:

- *Precision handling*. Wherever possible, continuous mapping parameters should be directly and dynamically controllable. Obvious candidates include all scaling factors; color and opacity transfer function components; probe placement and seeding; and levels for contours, isosurfaces, and thresholding. Conceptually, the space of all possible visualizations of a data set is a high-dimensional space (one dimension per visualization parameter) and each specific visualization has a location in that design space. Moving from point to point in that space should be easy and natural.
- *Cruise control.* Just as directly manipulating individual engine cylinders may not be the best way to control speed, individually manipulating each visualization parameter (for instance, each control point of a transfer function) may not be the most intuitive or efficient way to manipulate visualizations. Ideally, a dynamic visualization would have a set of composite control surfaces corresponding to a car's accelerator, brake, and steering wheel. Designing useful metaphors for a visualization's metacontrols is a wide open research area.
- *Onboard assistance.* As users' freedom to manipulate visualization parameters grows, so does their need for guidance in how to manipulate them. Such expert guidance might encapsulate perceptual considerations or task-based requirements. Ideally, this assistance should make optional suggestions that the user may ignore.
- *Reconfigurable drivetrain.* The classic visualization process is modeled as a sequence of stages, each building on the last. Making a change to an early stage necessitates rerunning the entire pipeline. Flexible dynamic control over visualization parameters requires a software architecture that enables modification of visual attributes assigned in an earlier mapping stage, ideally without requiring the re-execution of the entire pipeline. Such reconfiguration might require deferring the commitment of mapping choices into geometry or maintaining hooks into the geometry to allow efficient manipulation.
- *A good map*. Unfortunately, not every visualization parameter choice will be useful, and backtracking from blind alleys can be an arduous task. A user interface like the one Ma[6] suggests includes directions to a few popular landmarks (good standard visualizations) and visually documents territory already explored. This reduces time spent trying to recapture something interesting that you've already passed.

Bon voyage

Data is a landscape we can explore. The postcards of inspired visualizations, the guideposts of well-chosen techniques and color scales, and the guided tours of prearranged animations are valuable starting places for a journey of discovery, but they shouldn't be the end.

During the discovery process, it's important that we get off the bus and move beyond the scenic overlooks to walk along a few unexplored contours, float down some uncharted flows, and hunt for new species under a cool, midnight color map. Even though each step along the way might not be independently enlightening, walking paths of our own choosing through the data landscape is a valuable experience. As visualization toolsmiths, we need to equip our software with the gear we need to travel off the road. As data explorers, it's time to stop and smell the voxels. ■

Acknowledgments

The ideas in this article have grown out of many conversations with insightful thinkers over a period of many years. The earliest and most crucial of these conversations was with Frederick P. Brooks, Jr. I'd like to thank him for continuing to ask questions that get people thinking. The National Science Foundation (ACIR 9996043 and 0081581), National Institute of Health (5-R24-RR-02170), and Office of Naval Research (N00014-86-K-0680) financially supported this work.

References

1. F.P. Brooks, Jr., "The Computer 'Scientist' as Toolsmith—Studies in Interactive Computer Graphics," *Proc. Info. Processing*, North Holland Publishing, Amsterdam, 1977, pp. 625-634.
2. M. Bailey, "Interacting with Direct Volume Rendering," *IEEE Computer Graphics and Applications*, vol. 21, no. 1, Jan./Feb. 2001, pp. 10-12.
3. H. Pfister et al., "The Transfer Function Bake-Off," *IEEE Computer Graphics and Applications*, vol. 21, no. 3, May/June 2001, pp. 16-22.
4. P. Rheingans, "Color, Change, and Control for Quantitative Data Display," *Proc. IEEE Visualization*, IEEE CS Press, Los Alamitos, Calif., 1992, pp. 252-259.
5. P. Rheingans, "Dynamic Color Mapping of Bivariate Qualitative Data," *Proc. IEEE Visualization*, IEEE CS Press, Los Alamitos, Calif., 1997, pp. 159-166.
6. K.-L. Ma, "Visualizing Visualizations," *IEEE Computer Graphics and Applications*, vol. 20, no. 5, Sept./Oct. 2000, pp. 16-19.

Readers may contact Penny Rheingans at the Dept. of Computer Science and Electrical Engineering, Univ. of Maryland, Baltimore County, 1000 Hilltop Cir., Baltimore, MD 21250; email rheingans@cs.umbc.edu.

Contact department editors Theresa-Marie Rhyne and Lloyd Treinish by email at rhyne@siggraph.org and lloydt@us.ibm.com.