Practical Scientific Visualization Examples

Russell M. Taylor II
Department of Computer Science
University of North Carolina, Chapel Hill
taylorr@cs.unc.edu http://www.cs.unc.edu/~taylorr/

Abstract

Scientific visualization is not yet a discipline founded on well-understood principles. In some cases we have rules of thumb, and there are studies that probe the capabilities and limitations of specific techniques. For the most part, however, it is a collection of *ad hoc* techniques and lovely examples. This article collects examples where visualization was found to be useful for particular insights or where it enabled new and fruitful types of experiments.

Introduction

Many examples listed here are drawn from Keller and Keller's book Visual Cues, a valuable collection of visualization techniques along with descriptions of which techniques apply each visualization technique. (Keller and Keller 1993) Many of the examples are from recent Case Studies published in the proceedings of the IEEE Visualization conference (a wealth of examples of actual visualization systems in use). The rest of the examples come from colleagues or the author's own research.

This work does not attempt to cover the range of available techniques. Rather, it seeks to collect anecdotal and verified examples of when particular benefits have come from particular visualization techniques. As possible, related examples are grouped into categories. After each category, there is a **Net** section that summarizes the type of result the technique is best at producing.

Examples are restricted to visualizations for which there is an inherent spatial distribution to the data. This includes medical visualization but leaves out the broad fields of algorithm visualization and in-formation visualization that visualize abstract (often many-dimensional) spaces.
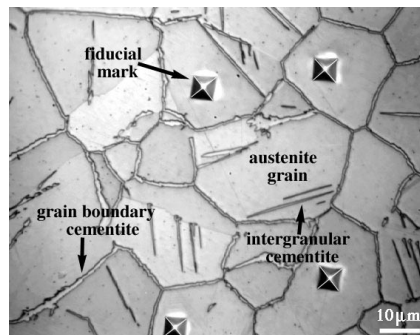
Many excellent systems are unlisted because no published account was found of particular insights gained using them, even when the visualization techniques used are obviously powerful. Space constraints further limit the number of results presented.
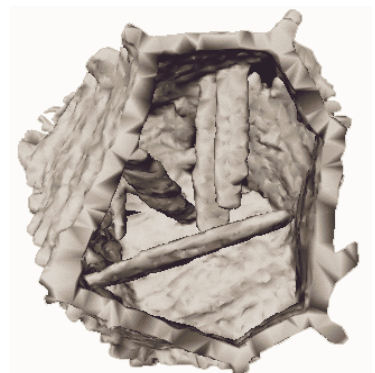
Visual Display

Viewing spatial data as spatial data

Lanzagorta and others at NRL looked at the internal microstructure of steel by polishing down one layer at a time and scanning each with an SEM. (Lanzagorta,

Kral et al. 1998) They packed the slices back into 3D and visualized them using an immersive environment (in the GROTTO). The goal was to understand the morphology and distribution of grains

and precipitates within the steel. The 3D visualization revealed new features: cementite precipitates that had appeared in 2D to be isolated were in fact in contact with austenite grain boundaries, and there was often a hole in the cementite grain

boundary film near the base of cementite precipitates that had gone unnoticed in 2D.

Ross and others at NASA Ames reconstructed an image of the 3D structure of mammalian gravity-sensing organs from multiple electron micrographs. The 3D image revealed for the first time that the elements are morphologically organized for weighted, parallel-distributed processing of information. (Keller and Keller 1993), p 155.

Volume rendering enabled the proper planning of partial, living-donor lung transplants. In this procedure, one of the lobes of the donor's lung is transplanted to the recipient. In order to avoid damaging the neighboring lobe in the donor, surgeons need to know just where to cut the bronchial tubes and blood vessels. When this planning was done based on CT slices, there was often damage to the neighboring lobe. Volume rendering allowed the surgeries to be planned so that this could usually be avoided. (Aylward 1999)

Rupert and others at the Lawrence Livermore national Laboratory used color coding to study the accuracy of a material mix dynamic simulation; displaying the simulation results in the area of interest. Comparisons with expected behavior and experimental results indicated areas where bugs in the program caused unexpected behavior. (Keller and Keller 1993), p. 44. Another example, where Crowley and others found programming errors by visual presentation of the field of study, is found on page 71.

Klimenko at Freiburg University and others built a system for the visualization of the sheets swept out by the strings in subatomic string theory. (Klimenko, Nikitin et al. 1999) They adjusted parameters in their visualization system to explore possible configurations, looking for singularities. This allowed them to formulate new hypotheses, which they were then able to prove analytically.

Students of Pizer and Coggins at UNC noted striking similarities between the output of a CORE-based medial axis algorithm (Fritsch, Eberly et al. 1994) and that of code that performs statistical classifications in feature space. "The observation showed us that the outputs of two

computations arising from disparate methodologies with different histories were essentially the same." This result has been replicated, with the result that we know that the medial axes can be found by a specific set of operations in feature spaces. (Coggins 1999)

Net: Whereas projection and measurement techniques can provide better quantitative results to particular known questions, viewing data in its natural spatial extent can provide insight and understanding.

Viewing *transformed* spatial data as spatial data

The UNC *Rspace* program aids in the planning of x-ray crystallography data collection. (Brooks, Thorvaldsdottir et al. 1988) For this task, a detector that collects diffracted electrons must be positioned for multiple collections to obtain the correct amount of overlap between the sampled areas. The geometry of detector attachment limits the possible positions and orientations. When planning data collection, working around an *Ewald sphere* in reciprocal (diffraction) space allows collection over a fraction of the area; the rest follows by symmetry. This task is extremely difficult when done in any straightforward way in 3D Euclidean space. Rspace allows the user to do the planning directly in reciprocal space, making the task easier to do without gaps and with enough overlap. This application allowed faster and better planning of collection strategies, and therefore has been widely used. (Brooks 1999)

Parker and Samtaney at Princeton and Rutgers visualized the results from a simulation of plasma turbulence inside a fusion generator. (Parker and Samtaney 1994) They report that visualization played a critical role in going from the "raw" nonlinear solution of these complex equations into a simplified theoretical model explaining the essential underlying physics. "Through the 3D visualization we observed that the radial structure becomes more elongated during nonlinear saturation, then quickly rips apart." To concentrate the simulation resolution in areas that require it, the calculations are performed on grids that follow magnetic field lines. They report that these "*field-*

line-following coordinates are non-orthogonal and get quite twisted … hence 3D visualization becomes essential in interpreting the results in real space."

Seeger at UNC was developing a suite of haptic visualization techniques (adhesion, friction, bumps, vibration, stiffness) for data display. In order to understand what the forces were doing, he displayed them graphically in a surface coordinate system. This showed the response of the system to forces normal to and tangent to the local surface, and decoupled the surface geometry from the haptic computation. This visualization showed him when the visualization code was misbehaving. (Seeger 1999)

Net: Mapping data from a simulation or experiment into the coordinate system that is most natural for viewing it can make it easier to interpret.

Interaction with the model: natural view changes

Zuiderveld and others at University Hospital in Utrecht studied the effectiveness of a stereo volume-visualization system in clinical use over a period of about three months. (Zuiderveld, Ooijen et al. 1996) They discuss three examples of clinical results found using the system. In the first, an unusual piece of splenic artery was understood in the interactive stereo view, where conventional views had not provided insight. In the second, a small aneurysm was located in the interactive 3D view that had not been seen in 2D and 3D views on film. In the third, stereo viewing provided needed insight into the complex anatomy surrounding an aneurysm of the aortic arch.

Andrzey Wojtczak, a biochemist visiting the UNC Molecular Graphics resource, worked for over two hours in a head-mounted display system with a wide-area tracker looking at his molecule "for the first time", though had seen it many times before on a 2D screen. Due to his improved understanding compared to previous views, he discovered that he had fit part of the molecule incorrectly. (Brooks 1999)

Williams at UCLA had examined one data set using pseudo-color images and cross sections for months. Using the UNC nanoManipulator system (a microscopy visualization system), he discovered an important feature of the data (graphite sheets coming out of the surface) within seconds of viewing a real-time fly-over of the data set. These sheets were visible in only a few frames of the fly-over, when the lighting and eye point were in just the right position to highlight them. (Taylor II, Robinett et al. 1993)

Using the UNC *Sculpt* steerable molecular simulation system, Duke biochemist Jane Richardson examined the *Felix* custom-designed protein to understand what could make it stable. The key lies in the cross-bonding between four alpha helices (the bonding of the two pairs making up each subunit was clear; understanding the bonding between the two subunits was not). She couldn't get it to work in the configuration that she initially tried, so she used the online optimization to keep the subunits together while she flipped one whole subunit over to the other side. No automatic minimizer or simulated annealing method would have tried this extreme configuration change, and she discovered a better (lower-energy) solution in the new configuration. The stability is due to a handful of critical disulfide bonds between the subunits. (Brooks 1999)

While building Sitterson Hall, the UNC Computer Science department members were arguing with the architects over whether a gathering space at the receptionist window in the front lobby was wide enough to allow someone to pass a group of people who were gathered there. To better understand the space, a virtual walkthrough was designed based on the blueprints. The flat-shaded walkthrough, running at 1 frame every 4 seconds, convinced the department members that the area had to be widened. This has proven to be a wise decision. (Brooks 1999)

The Scientific Computing and Imaging group at the University of Utah has developed a scientific problem-solving environment, *SCIRun*, that can display (among other things) volume renderings of medical data sets. (Parker, Weinstein et al. 1997) A nearby company named Voxel has created a machine that makes 3D holograms from medical images. The holograms can be overlaid

on the patient's actual anatomy by viewing the patient through the holographic plate. Natalie, the daughter of Voxel's owner, had a tumor right above her brainstem, a type that is nearly always fatal due to the difficulty of removal. After the surgeons had prepared their surgical plan using traditional slice-based visualization, they were asked to revisit the plan using both holographic and immersive stereo volume-rendered display of the data. After an hour of review, the surgeons changed their surgical plan. The operation was a success. (Johnson 1999)

Net: You can think of each frame in a real-time display of the surface as a new filter applied to the data set, with the user in control of the filter parameters (viewpoint, lighting direction) through natural motions of head and hand. People are adept at understanding the structure of 3D surfaces this way, having learned this skill over a lifetime. Intuitive exploration can produce insight.

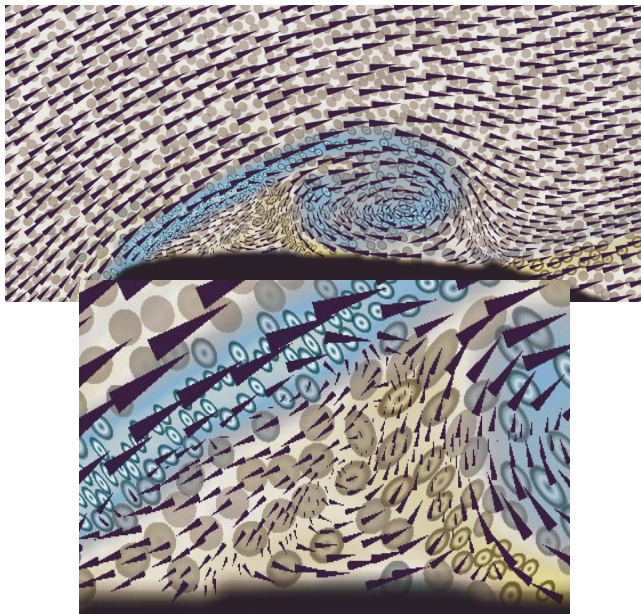**Combining multiple (measured or calculated) data sets**

Figure XXX: Visualization of multiple parameters of measured flow past an airfoil. Velocity, speed, vorticity, rate of strain and divergence shear are all shown. Divergence (shown by ellipse area) should be everywhere the same in an incompressible fluid; variations in ellipse radii are in the lower (zoomed) image indicate out-of-plane flow.

Kirby, Marmanis and Laidlaw at Brown University show a layered visualization, where color lies under ellipses, which lie under arrows. (R. M. Kirby and Laidlaw 1999) They have used this system to look at rate of strain (second-order tensor), turbulent charge (vector) and turbulent charge (scalar) looking for correlations between these and velocity and vorticity. A change of area in the ellipses within the visualization shown in figure XXX revealed that there was out-of-plane flow. They also report that the visualization posed new questions.

Nuuja and Lorig at the Pittsburgh Supercomputing Center visualized the results of a time-dependent simulation of ozone over Los Angeles. The simulation was tied to ozone measurements made throughout the day and night. At night, the measurements were dropping, so it appeared that the ozone concentration was dropping at night and increasing again during the daytime. The visualization revealed that at night, ozone merely floats higher in the atmosphere rather than disappearing. (Keller and Keller 1993) p. 88.

The first protein whose structure was solved from crystallography data without first building a physical model was Bovine Cu,Zn-superoxide dismutase in 1974 by Dave and Jane Richardson using the UNC GRIP system (a density fitting system). (Brooks 1999) Now, this is how protein fitting is routinely done. People were using non-computer-based scientific visualization before then; they used regularly spaced sheets of Plexiglas in a 3D cube, with the electron density surfaces drawn on the sheets. A half-silvered mirror superimposed a brass model of the protein (whose linkages are known but whose exact shape in 3-space is not) on the volume. This apparatus was known as a Richards' box, after its inventor Fred Richards. The advantages to the electronic fitting system are that the density data can be recomputed and redrawn automatically when the user has improved the fit and that the coordinates for the atoms are available upon completion (rather than having to be teased out with a plumb bob and meter stick). The 3D embedding of the model within the measured density is what allows the fitting to work at all.

Kim and Bejczy at JPL overlaid a locally-computed simulation of the motion of a remote-controlled robot arm, where the visual from the remote channel had a 2-3 second delay in response to user action. The simulated arm position (which was updated immediately) was displayed in wire-frame on top of the delayed remote video image; this helped the user deal with the extreme latency in the remote feed by seeing a prediction of what the motion would be. (Keller and Keller 1993), p. 171.

Smith and Littlefield at Lowell combined the information in two MRI images to form a new image consisting of *glyphs* (objects whose shape and color is determined by multiple data sets). The shape of these glyphs formed a discernable texture pattern in the combination image that indicated a hot spot (metabolic center) in a tumor. This spot was not visible in either of the original MRI images. (Keller and Keller 1993), p. 177.

Pandit and others at the University of Utah simulated and visualized the energy absorption within a human head (from MRI scans) in response to a cell-phone placed near the ear. (Pandit, McDermott et al. 1996) Seeing the relative location of the phone, antenna, head and energy distribution revealed that the major contribution to power deposition was from the feed-point of the antenna. It also showed that peak absorption rate was where the ear is in contact with the head (rather than the outside of the ear).

Montgomery and others at Stanford used a virtual-environment system to aid in the planning of reconstructive surgery for a 17-year old boy who had a tumor on the left side of his face. (Montgomery, Stephanides et al. 1998) Their system enabled the surgeons to compare the damaged side with mirror-image geometry of the intact side to determine the best fit and to produce a template for reconstruction. They also displayed a section of his hipbone for comparing curvatures with the template and to produce a printed paper template to aid with harvesting of the bone. This improved planning enabled the surgery to be performed in much less time than earlier surgeries,

and allowed the complete reconstruction to be performed in one operation. (Reconstructions often require several operations to make adjustments.)

Multimodal registration of MRI data with MR angiogram data allows the simultaneous viewing of the extent of a tumor within the brain and the patient's arteries and blood vessels. This allows surgeons to determine which arteries should be blocked to cause tumor death and which should be left open to avoid stroke. (Bullitt, Liu et al. 1999)

During surgery, visual overlay of the surgical instrument within a preoperative CT/MRI scan allows the surgeon to understand what is going on inside the real skull. Of course, the more modification (drain fluid or cut), the less registered things are because the brain is distorted compared to the preoperative image. One visualization showed remaining tumor in the map image that had been missed by visual inspection during surgery due to brain swelling. Another case involved a child who had a trapped fluid pocket on the side of his head, hidden behind a jungle of blood vessels. The surgeon used a frameless system to navigate her way through the blood vessels. (Bullit 1999)

Net: The simultaneous, registered display of multiple data sets can provide improved understanding of the relationship between them. This has been especially effective in medical visualization, for the comparison of radiographic data and actual anatomy.

Other Visual Techniques

Highlighting critical values

Mapping a critical value to a color that contrasts with the display of other values can reveal small errors and unexpected structure. Watterberg thus discovered an unexpected Moire-like pattern in a visualization of a 2D field; the pattern indicated that regions which were supposed to be identically zero in fact had areas with values very close to zero. (Keller and Keller 1993), p. 107.

Viewing a time series as a *transformed* time series

Besenbacher and colleagues at the University of Aarhus studied the dynamic properties of surface changes by scanning with a Scanning Tunneling Microscope (STM) at high speed during the changes and then replaying the information as an adjustable-time movie. (Besenbacher, Jensen et al. 1991) The movies in this case are computed off-line, and replayed at the original or other rates. They report that "Such information, which cannot be obtained by any other means, is very decisive for a full understanding of both the growth mode of reconstructed phases and the resulting static structure." They provide several examples of particular insights that were enabled by watching the time series as a movie: They were able to determine that the formation of O-Cu chains proceeds by adding rows to existing areas, using Cu atoms from step edges combining with O atoms diffusing on the surface; they discovered that the outermost row in an island of chains is not stable; they also found that in Cu(110)-c(6x2)O growth, the c(6x2) phase grows isotropically.

Making the Invisible Visible

David Banks and colleagues at Mississippi State developed a system that visualizes the results of applying various lenses, slits and apertures to an optical system. (Banks, Brown et al. 1998) An instructor who used the system reported that it was superior for instruction to hands-on optical benches because it made the invisible (light traveling through the system) visible and it abstracted out the behavior of the system. (McNeil 1999)

Haptic Display

The haptic display section is grouped by project, then by technique. A more in-depth view of the first two projects (and other UNC projects not listed here) can be found in (Brooks, Ouh-Young et al. 1990).

GROPE-1: An early haptic feedback application at UNC allowed the user to feel the effects of a 2D force field on a simulated probe, and was used to teach students in an introductory Physics course. (Batter and F.P. Brooks) These experi-

ments used a 2D sliding-carriage device with servomotors for force presentation. Experimental results showed that haptic feedback improved the understanding of field characteristics vs. a visual-only implementation (for students who were interested in the material). Students reported that using the haptic display dispelled previous misconceptions. They had thought that the field of a (cylindrical) diode would be greater near the plate than near the cathode, and they thought the gravitation vector in a 3-body field would always be directed at one of the bodies.
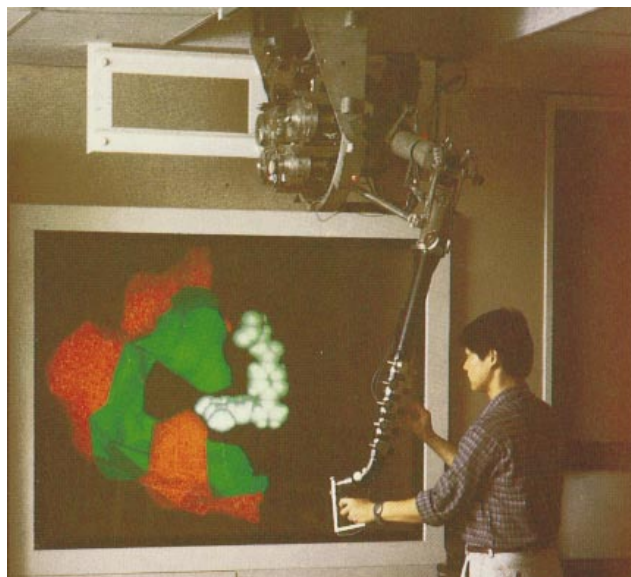

Figure XXX: The *Docker* application simulates the forces between a drug and its receptor site in a protein as the user guides it to the minimum-energy configuration.

Ming Ouh-young at UNC designed and built the Docker haptic feedback system to simulate the interaction of a drug molecule with its receptor site in a protein. Docker computes the force and torque between the drug and protein due to electrostatic charges and inter-atomic collisions, presenting the force and torque vectors both visually and using haptic feedback. Experiment showed that chemists could perform the rigid-body positioning task required to determine the lowest-energy configuration of the drug up to twice as quickly with haptic feedback turned on compared to using the visual-only representations. (Ouh-young, Beard et al. 1989; Ouh-young 1990) The most valuable result from using this system is probably the radically improved situation aware-

ness that serious users report. Chemists reported that when they felt the forces, they better grasped the details of the receptor site and of why a particular drug fit well or poorly.
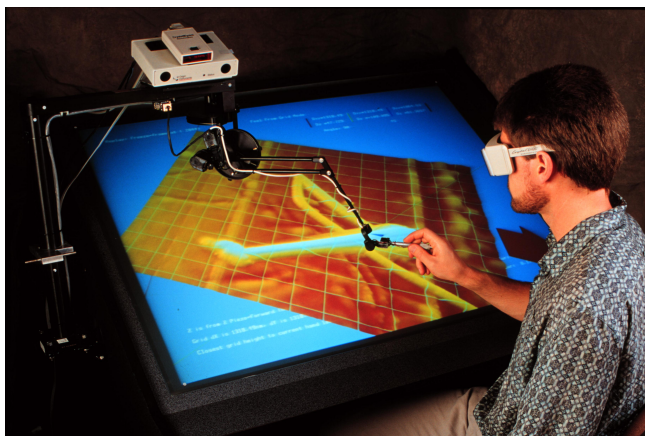

Figure XXX: The *nanoManipulator* application adds a virtual-reality interface to a scanned-probe microscope. Haptic feedback allows the user to feel and modify the microscopic surface.

The nanoManipulator: The nanoManipulator (nM) provides an intuitive interface to scanned-probe microscopes, allowing scientists from a variety of disciplines to examine and manipulate nanometer-scale structures. (Taylor II, Robinett et al. 1993) The nM displays a 3D rendering of the data as it arrives in real time. (Figure XXX) Using haptic display and control, the scientist can feel the surface representation to enhance understanding of surface properties and can modify the surface directly. Studies have shown that the nM greatly increases productivity. (Finch, Chi et al. 1995)

The haptic feedback component of our system has always been exciting to the scientists on the team; they love being able to feel the surface they are investigating. However, it is during modification that haptic feedback has proved itself most useful, allowing finer control and enabling whole new types of experiments. Haptic feedback has proved essential to finding the right spot to start a modification, finding the path along which to modify, and enabled greater precision than permitted by the standard scan-modify-scan experiment cycle. (Taylor II, Chen et al. 1997)

Finding the right spot

Haptic feedback allows the user to locate objects and features on the surface by feel when they move the tip around near the starting point for a modification. Surface features marking a desired region can be located without relying only on visual feedback from the previous scan (which is often misaligned due to drift and other nonlinearities in the positioners). This allowed a collaborator to position the tip directly over an adenovirus particle, then increase the force to cause the particle to dimple directly in the center. It also allowed the tip to be placed between two touching carbon filaments in order to tease them apart. (Taylor II, Chen et al. 1997)

Finding the right path

The scanned image shows only the surface as it was before a modification began. There is only one tip on an SPM: it can either be scanning the surface or modifying it, but not both at the same time. Haptic feedback during modification allows one to guide changes along a desired path.
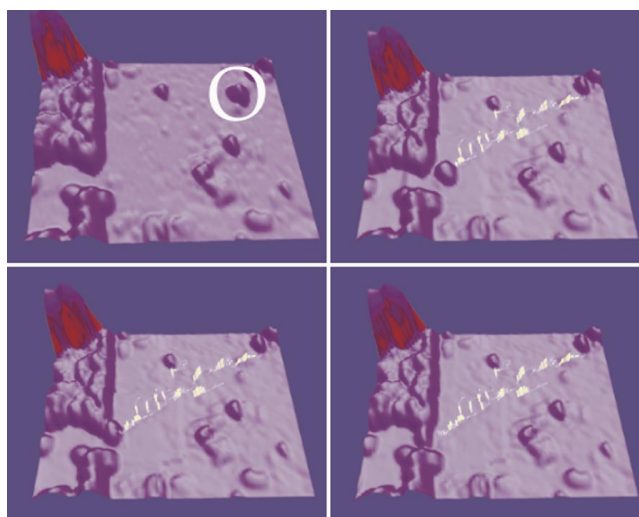

Figure XXX: In this sequence of images from left to right, a 15nm gold ball (circled) is moved into a test rig. Haptic feedback is used to feel as the ball is pushed and to determine when it has slipped off the tip.

Figure XXX shows haptic feedback being used to maneuver a gold colloid particle across a mica surface into a gap that has been etched into a gold wire. The colloid is fragile enough that it would be destroyed by getting the tip on top of it with modification force or by many pushes. This pre-
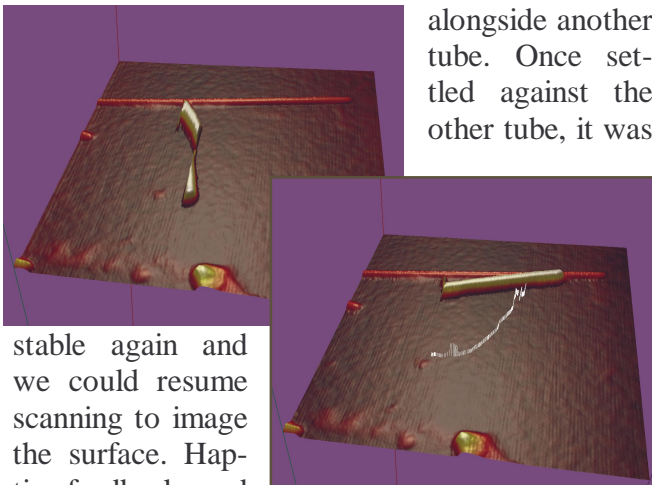
vents attempts to move it by repeated programmed "kicks". Haptic feedback allowed the user to tune the modification parameters so that the tip barely rode up the side of the ball while pushing it. This allowed the guidance of the ball during pushing so only about a dozen pushes were required. (Taylor II, Chen et al. 1997)

Haptic feedback was also used to form a thin ring in a gold film. A circle was scraped to form the inside of the ring, leaving two "snow plow" ridges to either side. By feeling when the tip bumped up against the outside of the outer ridge, another slightly larger circle was formed. This formed a thin gold ring on the surface. (Taylor II, Chen et al. 1997)

A light touch: observation modifies the system

When deposited on the surface, carbon nanotubes are held in place by residue from the solution in which they are dispersed. On some surfaces, the tubes slide freely once detached from the residue until they contact another patch of residue or another tube. Even the light touch of scanning causes them to move. Guiding the tip by hand and switching between imaging and modification force, we have been able to move and re-orient one carbon tube across a surface and into position alongside another tube. Once settled against the other tube, it was stable again and we could resume scanning to image the surface. Haptic feedback and slow, deliberate hand motion allowed us to find the tube at intermediate points when we could not scan. The fact that the surface cannot be imaged at intermediate stages prevents this type of experiment from being performed using the standard scan-modify-scan cycle. (Taylor II, Chen et al. 1997)

Net: We have found that haptic feedback gives improved situational awareness in the applications we have built. It can give students a better understanding of force fields and build a scientist's intuition about a problem. Force feedback is especially good for manipulation tasks, where the user is moving a real or simulated object as part of the task.

Acknowledgments

References

Aylward, S. (1999). Personal Communication.

Banks, D. C., J. M. Brown, et al. (1998). "Interactive 3D Visualization of Optical Phenomena." IEEE Computer Graphics & Applications **18**(4): pp. 66-69.

Batter, J. J. and J. F.P. Brooks (1971). GROPE-I: A computer display to the sense of feel. Information Processing, Proc. IFIP Congress 71. pp. 759-763.

Besenbacher, F., F. Jensen, et al. (1991). "Visualization of the Dynamics in Surface Reconstructions." *Journal of Vacuum Science Technology* **B 9**(2): pp. 874-877.

Brooks, F. P., Jr. (1999). Fourteen Years of Interactive Walkthroughs. ACM SIGGRAPH '99 Course Notes, course #20, Los Angeles, CA, ACM SIGGRAPH. pp. H1-H22.

Brooks, F. P., Jr. (1999). Personal communication.

Brooks, F. P., Jr., M. Ouh-Young, et al. (1990). Project GROPE - Haptic displays for scientific visualization. Computer Graphics: Proceedings of SIGGRAPH '90, Dallas, Texas. pp. 177-185.

Brooks, F. P., Jr., H. Thorvaldsdottir, et al. (1988). Fourteenth Annual Report Interactive Graphics for Molecular Studies. Chapel Hill, University of North Carolina Department of Computer Science.

Bullit, E. (1999). Personal Communication.

Bullitt, E., A. Liu, et al. (1999). "Registration of 3D Cerebral Vessels with 2D Digital Angiograms: Clinical Evaluation." Academic Radiology **6**: pp. 539-546.

Coggins, J. (1999). Personal Communication.

Finch, M., V. Chi, et al. (1995). Surface Modification Tools in a Virtual Environment Interface to a Scanning Probe Microscope. *Computer Graphics: Proceedings of the ACM Symposium on Interactive 3D Graphics*, Monterey, CA, ACM SIGGRAPH. pp. 13-18.

Fritsch, D. S., D. Eberly, et al. (1994). Simulated cores and their applications in medical imaging. Information Processing in Medical Imaging, Proc. IPMI '95, Kluwer, Dordrecht, the Netherlands. pp. 365-368.

Johnson, C. (1999). Personal Communication.

Keller, P. and M. Keller (1993). Visual Cues: Practical data visualization. Los Alamitos, CA, IEEE Computer Society Press.

Klimenko, S., I. Nikitin, et al. (1999). Visualization in String Theory. Hot Topics Proceedings of IEEE Visualization 1999**:** 29-32.

Lanzagorta, M., M. V. Kral, et al. (1998). Three-Dimensional Visualization of Microstructures. IEEE Visualization, Research Triangle Park, NC. pp. 487-490,574.

McNeil, L. E. (1999). Personal Communication.

Montgomery, K., M. Stephanides, et al. (1998). A Case Study Using the Virtual Environment for Reconstructive Surgery. IEEE Visualization, Research Triangle Park, NC. pp. 431-434.

Ouh-young, M. (1990). Force Display In Molecular Docking. Computer Science. Chapel Hill, University of North Carolina**:** Tech Report #90-004.

Ouh-young, M., D. V. Beard, et al. (1989). Force Display Performs Better Than Visual Display in a Simple 3-D Docking Task. Proceeding of IEEE Robotics and Automation Conference, Scottsdale, AZ. pp. 1462-1466.

Pandit, V., R. McDermott, et al. (1996). Electrical Energy Absorption in the Human Head From a Cellular Telephone. IEEE Visualization, San Fransisco, CA. pp. 371-374.

Parker, S. E. and R. Samtaney (1994). Tokamak Plasma Turbulence Visualization. *Visualization '94*, Washington, D.C., IEEE Computer Society Press. pp. 337-340.

Parker, S. G., D. M. Weinstein, et al. (1997). The {SCIR}un computational steering software system. Modern Software Tools in Scientific Computing. E. Arge, A. M. Bruaset and H. P. Langtangen, Birkhauser Press**:** 1-44.

R. M. Kirby, M. M. and D. H. Laidlaw (1999). Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting. IEEE Visualization 99, San Fransisco, CA, IEEE Press. pp. 333-340.

Seeger, A. (1999). Personal Communication.

Taylor II, R. M., J. Chen, et al. (1997). Pearls Found on the way to the Ideal Interface for Scanned-probe Microscopes. Visualization '97, Phoenix, AZ, IEEE Computer Society Press. pp. 467-470.

Taylor II, R. M., W. Robinett, et al. (1993). The Nanomanipulator: A Virtual-Reality Interface for a Scanning Tunneling Microscope. *SIGGRAPH 93*, Anaheim, California, ACM SIGGRAPH. pp. 127-134.

Zuiderveld, K. J., P. M. A. v. Ooijen, et al. (1996). Clinical Evaluation of Interactive Volume Visualization. IEEE Visualization, San Fransisco, CA. pp. 367-370.