Exploratory Visualization of Surgical Training Databases for Improving Skill Acquisition

David Schroeder ▪ *University of Minnesota*

Timothy Kowalewski and Lee White ▪ *University of Washington*

John Carlis, Erlan Santos, and Robert Sweet ▪ *University of Minnesota*

Thomas S. Lendvay ▪ *University of Washington*

Troy Reihsen and Daniel F. Keefe ▪ *University of Minnesota*

A new system has visualized force, position, rotation, and synchronized video data from 300 bimanual laparoscopic-surgery tasks performed by more than 50 surgeons. Insights and feedback from an interdisciplinary iterative design process and use case studies support such visualization's utility.

Minimally invasive surgical techniques, such as laparoscopic surgery, typically offer important patient benefits. Unfortunately, the psychomotor and cognitive skills needed to become proficient at laparoscopic surgery remain extremely challenging to learn and teach. In addition, objective means of measuring surgical performance have been limited to summary, one-dimensional data such as task time or total errors, making it difficult to evaluate surgeons' skills and provide feedback to them.

The advent of new technologies such as robotic surgical instruments, computer simulators, and computer vision algorithms has created an opportunity to record multidimensional data including tool paths, forces, deformations, and physiological information during surgery and training (see the "Related Work in Surgical Visualization" sidebar). These new data sources have great potential to lead to objective evaluation metrics. However, incorporating these data into a new paradigm of data-intensive surgery raises many important research challenges, including how to store, query, analyze, interpret, and display large-scale time-varying multidimensional surgical datasets (see the "Obtaining Surgical Data" sidebar). In particular, little research has focused on visualizing data collected during surgical procedures (including training) to identify potential areas for error reduction and potentially accelerate learning.

Here, we explore visualizing trends in large databases of multidimensional time series data from surgeons performing a basic laparoscopic-surgery training task. Previous researchers have applied machine-learning techniques to these quantitative performance data—for example, classifiers can help distinguish novice and expert surgeons.[1] We believe visualization could augment computational tools such as these to elucidate specific factors or trends that could inform the design of improved training tools. Furthermore, visualization could provide a particularly effective method for conveying data trends to surgeons.

From the visualization research community's perspective, as supported by a panel discussion at VisWeek 2010, these goals relate to the grand challenge of supporting analysis at scale. Regarding

 0272-1716/12/$31.00 © 2012 IEEE

Related Work in Surgical Visualization

Fueled by advances in computer graphics, VR, robotics, haptics, visualization, and related hardware technologies, surgical simulation tools have experienced tremendous growth and reached new levels of realism.[1] To complement simulators, visualization research has contributed many tools to help us understand imaging and related data. Real-time clinical surgical visualization tools have also been a focus of research[2] and, more recently, commercial consoles such as Intuitive Surgical's da Vinci surgical robot. (For more on da Vinci, see the "Obtaining Surgical Data" sidebar.) Although many simulator tools can report summary statistics at the end of a trial, we know of no previous exploratory visualization systems designed to examine aggregate multidimensional high-resolution surgical performance data.

Regarding scaling visualizations of time-varying data to work across a large set of motions rather than just a single motion, our research is similar to recent multiview visualizations applied to studying animal biomechanics.[3] Surgical applications raise new questions in terms of incorporating video into multidimensional visualizations, developing abstract 3D representations to encode bimanual tool trace data, and designing visualizations that can convey appropriate feedback to surgeons.

References

1. N. Zhang et al., "Volumetric Modeling in Laser BPH Therapy Simulation," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 6, 2010, pp. 1405–1412.
2. H. Fuchs et al., "Augmented Reality Visualization for Laparoscopic Surgery," *Proc. 1st Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, Springer, 1998, pp. 934–943.
3. D. Keefe et al., "Interactive Coordinated Multiple-View Visualization of Biomechanical Motion Data," *IEEE Trans. Visualization and Computer Graphics*, vol. 15, no. 6, 2009, pp. 1383–1390.

surgical training applications, specific visualization-at-scale challenges arise in the areas of organizing and querying databases of movies, forces, positions, high-level descriptions of tasks, and other data and in recognizing that the results generated during data exploration are useful data themselves. Another core challenge is supporting visualization of coordinated complex bimanual activities based on multivariate, time-varying data. Scaling visualization applications to support aggregate analysis across hundreds or thousands of instances of complex time series data such as these remains a major research challenge.

Working toward these grand challenges, this article contributes what is, to our knowledge, the first interactive exploratory visualization of a database of surgical data on the scale of the Surgical Genome Project (see the "Obtaining Surgical Data" sidebar). Our visualization system employs three novel visualization techniques iteratively designed for analyzing surgical data:

- a smart brushing (selection) interface,
- multiview video overlay, and
- an animated 3D tool trace visualization.

To store surgical training data, the system uses an extensible formal data model, and it links exploratory visualizations to a robust database infrastructure.

The Task and Data

To illustrate how our system works, we visualized surgical data consisting of 300 peg-transfer tasks performed by more than 50 surgeons. The peg-transfer task is one of five standard tasks in the ubiquitous Fundamentals of Laparoscopic Surgery training module.[2] In this task, the surgeon tries to pick up and move a series of six blocks from the pegs on one side of a board to the pegs on the other side and then move them back. The blocks must be picked up by one hand's laparoscopic tool and then transferred in the air to the other hand's tool. Each tool has a grasper resembling a small alligator clip. For the task to be considered successful, the surgeon must complete it without errors (such as dropping a block) under a given time threshold.

Besides recording a video for each attempt, we captured the

- Cartesian tool position,
- tool rotation around the shaft axis,
- tool grasp force,
- grasper angle, and
- time.

The EDGE (Electronic Data Generation and Evaluation) platform captured a total of 428,924 frames of multivariate data at 30 Hz; watching all the video collected would take almost four hours. (For more on EDGE, see the "Obtaining Surgical Data" sidebar.) The database infrastructure we developed derived additional quantities, such as velocity or block transfer events, and passed them on to the visualization system.

Our Visualization System

Inspired by Ben Shneiderman's information-seeking mantra ("overview first, zoom and filter, then details on demand"),[3] Figure 1 illustrates how several complementary interactive data views support exploratory visualization. Traditional 2D plots support graphing any dimension of the data against any other dimension and support analysis and filtering. Video augmented with data overlays provide

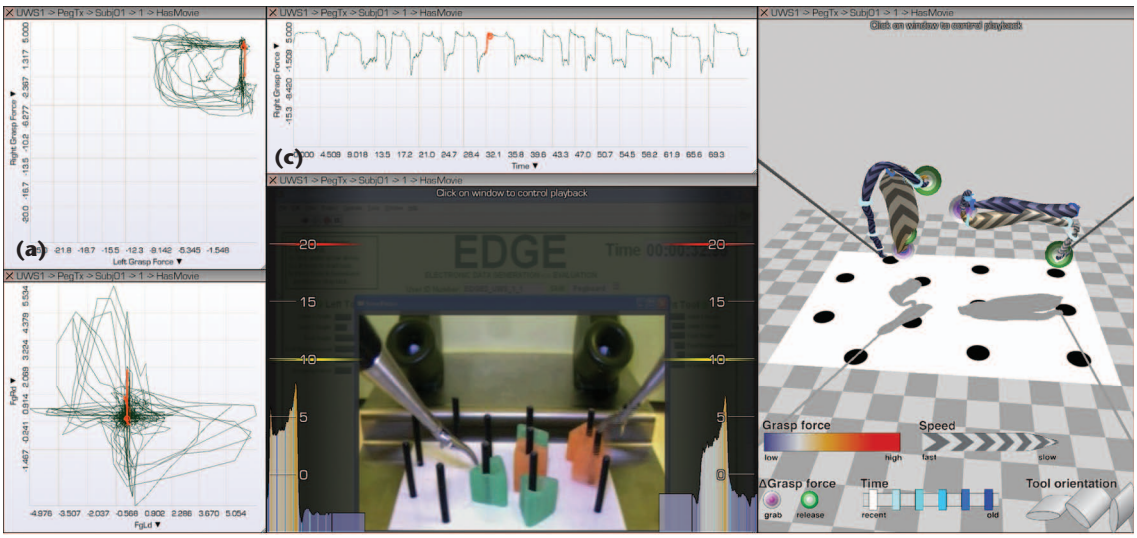

Figure 1. Visualizations to explore multidimensional data collected during laparoscopic surgical training. (a) The left and right grasp force. (b) The time derivative of these quantities. (c) The right grasp force over time. (d) Video recorded during the trial. (e) Detailed 3D tool trace information. All views are linked; multiple windows can refer to the same set or subset of data, and interactions in one view affect the others.

contextual information required to interpret trends, and a 3D visualization supports analysis of multivariate tool path data in a spatial context.

All views are linked; multiple windows can refer to the same set or subset of data, and interactions in one view affect the others. Brushing techniques let users select a subset of data displayed in the current window and then filter the data on the basis of this selection, activating new visualization windows to display just the relevant subset. Users can save these selections, allowing repeatable data analysis.

Smart Brushing

The basic interaction for controlling selection and data filtering is click-and-drag. The user moves the mouse to create a selection rectangle, which selects any data points falling within it. After the user makes a selection, he or she can right-click to activate a contextual menu, which contains options for creating new visualization windows tied specifically to the selected data points.

We developed several extensions to this basic brushing interface to support visualization of surgical data. For example, Figure 2a shows a plot of the left and right grip force. The user has selected an interesting subset of the data that usually corresponds to the subtask of transferring the block from one hand to the other. In a traditional visualization, to understand what's really happening at these points in time, the user would typically activate a video window, which would then play the sequences of highlighted frames. The problem is that the highlighted frames represent very small time windows—just the split

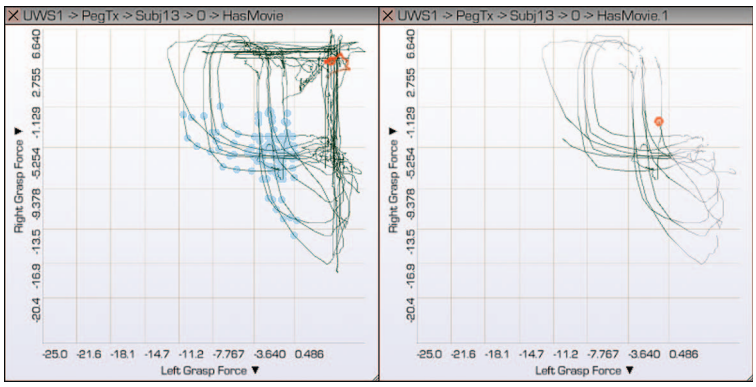

Figure 2. A brushing interface intelligently extends mouse selections forward and backward in time to capture appropriate context. (a) The user highlighted the blue points in this plot. (b) The curves show the selection as interpreted by the interface.

second when the block transfers from one hand to the other. Watching just these frames would provide little context to interpret the data.

Our solution (see Figure 2) automatically extends the selection forward and backward in time to provide more context to interpret the data. Actually, the selection shown in Figure 2b would be impossible to make without this smart interface. The user wouldn't be able to select the appropriate subset without also selecting additional, irrelevant frames.

We experimented with several techniques to determine how far in time to extend the user's selection. The simplest is to expand by a fixed time such as 0.5 second. More data-driven approaches are also possible. In Figure 2, for each sequence of

Obtaining Surgical Data

Modern surgery has been derided as more art than science, perhaps because surgeons lack quantitative tools during operations and training. Pervasive instrumentation, simulators, and widespread robotics are fundamentally altering surgery, enabling an unprecedented degree of precision and quantitative rigor that promises to improve patient outcomes and reduce critical costs. We call this budding field *computational surgery*.

The Current State

A growing trend in the operating room (OR) is platforms for collecting rich, quantitative surgical data. An early example from the University of Washington's BioRobotics Lab (BRL) was BlueDragon (see Figure A1). It recorded data from mechanically instrumented surgical tools while surgeons practiced common laparoscopic procedures on anesthetized pigs. BlueDragon recorded video and 12 degrees of freedom for position, orientation, forces, and torques, plus grasp force and tissue contact for each hand at 30 Hz. A desktop version of BlueDragon, Simulab Corporation's award-winning EDGE (Electronic Data Generation and Evaluation; see Figure A2), has recorded more than 400 surgical training tasks.

The most prevalent data source is Intuitive Surgical's da Vinci surgical robot (see Figure A4). With more than 1,500 units deployed in hospitals worldwide and rapid expansion into additional surgical specialties, it's poised to replace several manual laparoscopic procedures. This establishes computational surgery on a firm quantitative foundation in the OR, generating more than 80 dimensions of sensor data along with high-definition stereoscopic video. To ease data access, BRL developed SurgTrak, a da Vinci add-on, to record 20 dimensions of position and orientation data and synchronized video.[1] SurgTrak has recorded more than 1,000 iterations of robotic surgical tasks in dry-lab settings. An appealing alternative is the Raven surgical robot (see Figure A3), an open, academic research platform that yields all its internal signals[2] (http://brl.ee.washington.edu/laboratory/node/26).

The Surgical Genome Project

To advance, computational surgery must collect ethically sound quantitative data from practicing surgeons; however, this process can be forbiddingly arduous and expensive. To this end, the Surgical Genome Project (SGP) is an open, Web-based repository of relevant, de-identified surgical data for use by academic researchers and clinicians (http://surgenome.ee.washington.edu). The SGP intends to spur progress by allowing easy participation and a common framework for comparing analytical results. Yet, the same richness and breadth that make this data so beneficial to surgery is also a hurdle to its use. As this new branch of science emerges, scalable exploratory visualization tools could uniquely contribute to and catalyze its success.

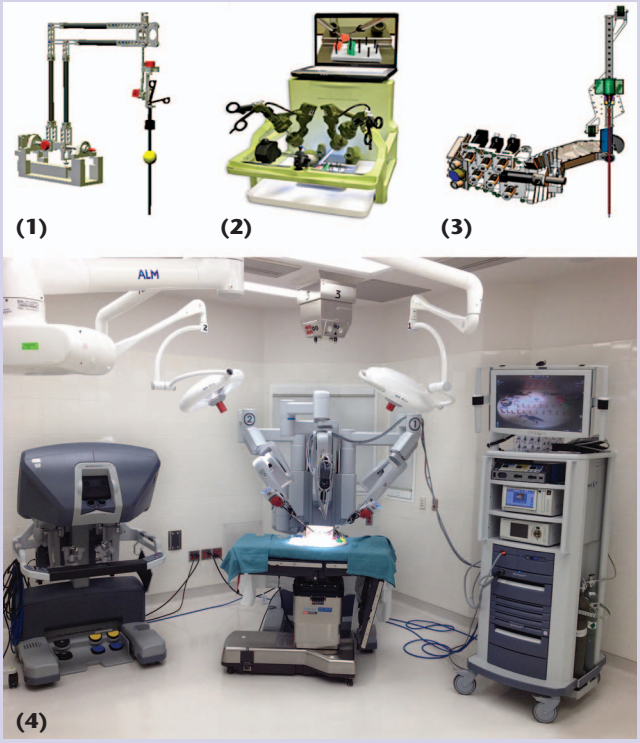

Figure A. Surgical data capture platforms. (1) BlueDragon. (2) EDGE (Electronic Data Generation and Evaluation). (3) The Raven surgical robot. (4) The da Vinci surgical robot.

References

1. L.W. White et al., "SurgTrak: Affordable Motion Tracking and Video Capture for the da Vinci Surgical Robot," *Proc. SAGES Ann. Meeting*, vol. 1, Soc. Am. Gastrointestinal and Endoscopic Surgeons, 2011, p. 204.
2. C. Reiley et al., "A Review of Methods for Objective Surgical Skill Evaluation," *Surgical Endoscopy*, vol. 25, no. 2, 2011, pp. 356–366.

frames selected, the algorithm traverses the frames before and after to identify data discontinuities (for example, a jump from high to low force). It sets the selection endpoints to correspond to these instances.

Video Overlays

Abstract data plots alone typically don't provide enough context to interpret data. Videos, on the other hand, are particularly valuable for providing contextual information but don't provide concrete quantitative data. One of our design goals was to tightly integrate these two complementary data displays.

With this motivation, we overlaid animated

Surgical Training and Formative Feedback

Laparoscopic surgery is one of the most common minimally invasive surgical techniques. Current laparoscopic procedures require surgical skills far beyond the basic set needed for the peg transfer task. (For an explanation of this task, see the section "The Task and Data" in the main article.) Achieving the correct use and control of force is one of the most challenging skills for surgeons to learn and instructors to teach. Although metrics exist for describing tissue damage after it has occurred (such as the area of tissue tear), it's difficult for a mentoring surgeon to observe inappropriate use of force until it's too late.[1,2]

In contrast, an ideal training scenario would include formative feedback—that is, useful feedback delivered in real time during training and perhaps even during clinical use. To enable a new approach to data-driven formative feedback for surgical training, we need new tools to

- analyze surgical data in real time,
- relate real-time performance data to metrics derived from massive datasets, and
- provide appropriate feedback to training surgeons via visual, auditory, or other channels.

Video overlays, audible feedback modes, and 3D multivariate surgical-tool trace displays can serve as a starting point for developing these tools. For example, we believe incorporating these visualization strategies into interactive stereoscopic training environments will be a particularly valuable direction for future research. This approach has great potential to improve the effectiveness of surgical procedures and patient health while accelerating skill acquisition.

References

1. L.P. Sturm et al., "A Systematic Review of Skills Transfer after Surgical Simulation Training," *Annals of Surgery*, vol. 248, no. 4, 2008, pp. 690–679.
2. L. Sturm et al., *Surgical Simulation for Training: Skills Transfer to the Operating Room*, tech. report 61, Australian Safety & Efficacy Register of New Interventional Procedures—Surgical, Royal Australasian College of Surgeons, 2007.

data plots on top of video replays, visually bringing quantitative data displays as close as possible to the videos that help us interpret them. Figure 3a shows a screenshot of a video of a trial performed by an expert; Figure 3b shows a screenshot for a novice. The bar plots on the left and right depict the instantaneous grip force applied by the left and right hands. To further illustrate patterns in the use of force, the display includes an animated line graph capturing the last several seconds of force measurements for each hand.

The augmented video visualization also includes aural and visual indicators activated when tool forces exceed either a warning or danger threshold. The audio tones are localized to the speaker corresponding to the tool that exceeded the threshold (that is, the left or right speaker). Using the same thresholds, the corresponding side of the video is tinted yellow (warning) or red (danger). Currently, these features are used for post hoc data visualization, but we believe that there's great potential to utilize similar techniques for real-time formative feedback (see the "Surgical Training and the Importance of Formative Feedback" sidebar).

These augmented video visualizations enable two new analysis modes. First, we can observe patterns in the use of force from one hand to the other (that is, bimanual coordination) by looking at the data displayed in a single video. Second, we can compare multiple tasks performed by the same surgeon or the same task performed by multiple surgeons.

Multivariate 3D Visualization

It's challenging to identify correlations among multiple data variables and understand how the variables change relative to the 3D tool position. To facilitate this, we developed a series of exploratory 3D visualization techniques.

Our approach is inspired by the success of the textured and color-coded 3D glyphs that Colin Ware and his colleagues introduced for visualizing the diving and feeding activities of whales as they move through the ocean.[4] Our data also represent motion paths through a 3D space, but three important differences informed our rendering algorithm's design (see Figure 4). First, the surgical tools often return to the same positions as the surgeon works, so path traces frequently overlap. Second, the data can be noisy, and some portions have little motion. Finally, we're interested in analyzing coordinated motions—that is, motions of two tools working in tandem.

The legend in Figure 4 describes our current mapping from data to visualization. A textured ribbon traces the 3D tool path. The ribbon's orientation, texture, and color are all tied to data variables; the space surrounding each ribbon includes additional glyphs. The color mappings match those used in the video overlays, enabling associations between the two types of displays when they're animated synchronously.

Our design began with a rendering algorithm nearly identical to Ware and Wiley's; we then

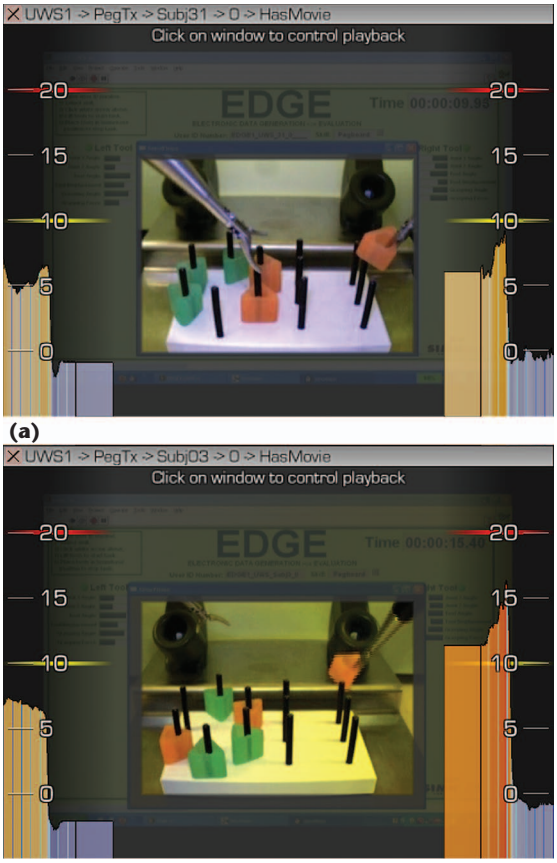

Figure 3. Video overlays show the grip force for each hand for (a) an expert and (b) a novice surgeon. In both cases, the display is animated and synchronized with video playback. The large bars on each side of the window indicate the current forces; the history of change in forces is plotted as an animated graph extending off the screen from each bar. The high force visible in Figure 3b relative to Figure 3a is indicative of a difference observed between the novice and expert.

refined the strategy over several iterations to work more effectively for this surgical application. Initially, the ribbon form displayed in the visualization was a fixed width, which could facilitate depth judgments. However, this design was sensitive to changes in the tool's position and orientation resulting from noisy data or muscular jitter. These portions of the data created discontinuities in the ribbon form that ended up being more visually salient than other, more important regions of the data.

Our solution first applies a lightweight smoothing filter to the data, then maps the ribbon width to speed (using a nonlinear mapping). This design emphasizes the form of the ribbon in areas of decisive motion while diminishing its visual saliency in areas of little motion. It also helps simplify the geometry in the display so that the forms of overlapping motion traces are easier to discern.

Figure 4 reflects several other important design decisions. Shadows on a textured ground plane convey spatial relationships. Rather than displaying the entire tool motion over the course of a trial, we display only the last five seconds of data. When the display is animated, this produces an effect akin to motion blurring. These animated displays are often useful in establishing a clear connection to the video data. However, the 3D visualizations are often easiest to interpret when the playback is paused and the user navigates around the scene with mouse-based camera controls.

We use two visual strategies to link the two tool path traces, helping us to understand how motion is coordinated between the two hands. First, a series of bands around each ribbon function as tick marks indicating the passage of time. Second, glyphs denote discrete events, such as discontinuities in the derivative of force data, which tend to correspond to grabbing or releasing a block. We found that including these discrete markers on both tool traces dramatically improves the viewer's ability to understand how one tool's motion relates to the other's motion spatially and temporally.

Scalability and the Need for Good Data Models

Another challenge arising in this interdisciplinary work is planning for successful scaling as dataset size and complexity increase.[5] When we're visualizing only a few tasks, understanding what we're seeing and what data it comes from is relatively easy. However, as more data are gathered, derived, and aggregated, datasets become more voluminous and varied. Managing results, insights, and the processes used to get them becomes a struggle; we need a consistent, consensus vocabulary. We took a two-pronged approach to this.

First, following John Carlis and Joseph Maguire's notation and disciplined naming methodology,[6] we jointly built our data model. Figure 5 shows a portion of that model, which represents the raw surgical data as follows: a subject has trials, a trial has a sequence of frames, and a frame has handed instruments, each of which has position and force data.

Because we're interested in analyzing how position, force, and other variables change over time, we included derived data:

- an *earlier-frame–later-frame–distance triple*, in which a distance of 1 between frames means the frames are adjacent and other distances connote a sampling;
- an *abutting pair* of those triples; and

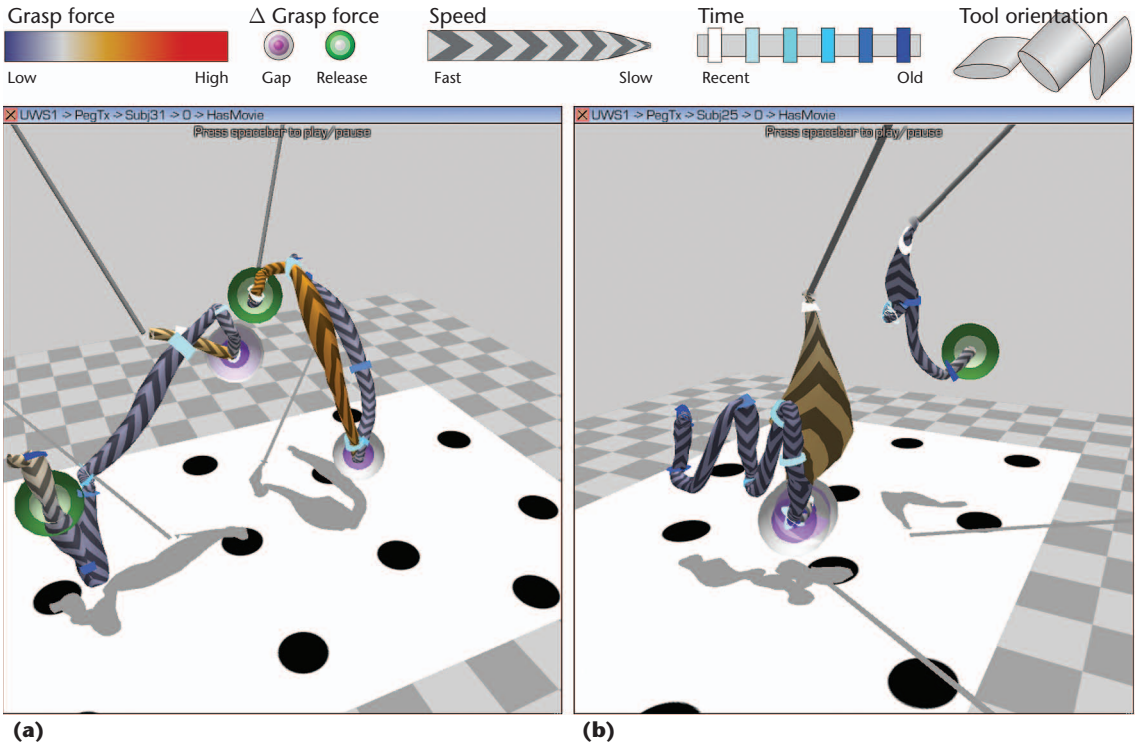

Figure 4. Visualizations of two peg-transfer trials. (a) The expert is moving relatively quickly and decisively. He or she set down a block on the left of the screen, picked one up from the middle, and transferred it between tools in the middle of the screen. (b) The novice is having trouble judging depth. The ground plane shows the location of the block that the novice was trying to pick up; the purple spheres show where he or she eventually picked it up.

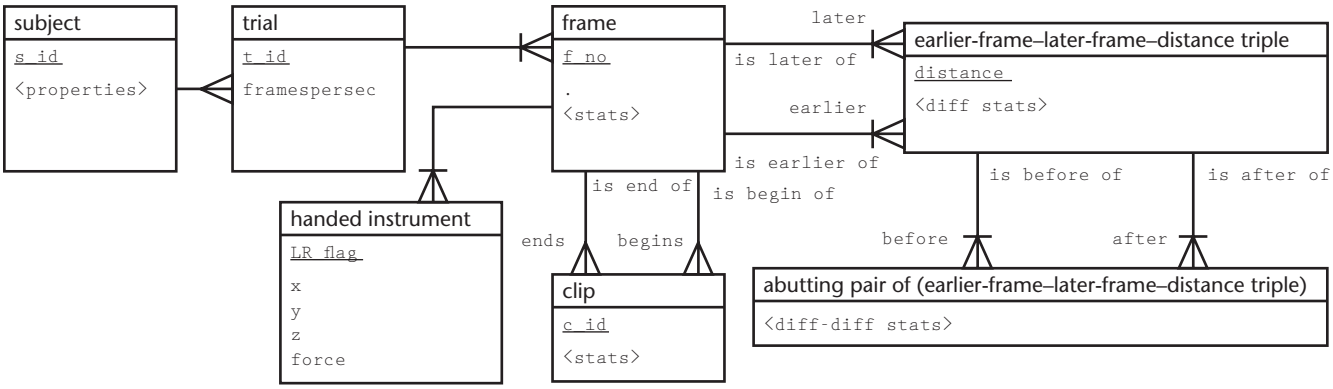

Figure 5. An extensible data model for surgical performance data mapped to a relational schema in a database management system. By clearly separating the storage of derived data from the pristine data while still keeping the association between these data clear, this model allows for easy, rapid, repeatable exploration of the surgical data.

- possibly, depending on the context, a contiguous *clip* bounded by beginning and ending frames.

For the triple, we can calculate the position and force differences and the differences of the differences of one hand. For the abutting pair, we can calculate the relationships between the hands. A user could define a clip directly or by looking at analyses for peaks, troughs, or trends of activity, and then remembering the clips between or around them, perhaps with an enclosing context.

Second, we employ database management system (DBMS) software to manage the data. Where possible, we shift the burden of calculating the derived data and managing the results onto the DBMS. We mapped the data model of Figure 5 to a relational schema, loaded the surgery data, wrote simple DBMS command scripts to calculate the derived data, and visualized the data. We've already seen the benefits of this part of our approach. With the beginnings of a lexicon in hand, we tested ideas on extracted data subsets and readily scaled up to more comprehensive datasets.

This approach also supports an iterative, exploratory process. Early on, we noticed the regularity of the 12 transfers by an expert in contrast to a novice's irregularity. To explore this difference, we computed the frame-to-frame force difference for each trial. The resulting derived data provide an initial means of defining clips via a simple force difference threshold. Our visualizations employ these derived data; for example, in Figure 4, the discrete grabs and releases denoted by the spheres are placed at clip boundaries defined via the DBMS command scripts.

Use Case Studies

To assess the system's utility and potential impact, we conducted data analysis sessions with expert practicing surgeons and the engineers embedded in their teams to develop new surgical training tools. We evaluated two groups, one from each of the two institutions collaborating on our long-term project. Each group included one senior surgeon leader and two engineers. Both senior surgeons had a background in urologic surgery, including minimally invasive techniques. One of the groups included a second surgical fellow with similar areas of expertise.

Our approach included structured, preselected data analysis and interactive-visualization tasks as well as opportunities for free-form exploration, discussion, and other qualitative feedback. Before each session, together with the collaborators, we brainstormed a set of data analysis questions, including several that are challenging or impossible to answer using current tools. Two of the most important questions were these:

- How does the use of force change between more and less experienced surgeons?
- When the user approaches a target, does the velocity indicate any skill-based patterns?

These and other questions provided structure and a relevant "real-world" context for the use-based evaluations.

During each data analysis session, which lasted one to two hours, we began with a 10- to 15-minute explanation of the system. We then presented an example analysis of the use of force, which we introduced by using the brushing interface to select data subsets. We then addressed velocity patterns. Finally, depending on the surgeons' interests, we addressed other questions we had brainstormed previously or analyzed new hypotheses generated during the session. As the surgeons used the system, we gathered qualitative data via a think-aloud protocol.

User Feedback and Assessment of Potential Impact

Feedback on the system was consistent between the two groups of surgeons. The clear high-level response was that the system is a drastic improvement over current practice. One surgeon commented, "This makes objective assessment of surgical skills so much easier than going through the video."

Similarly, both groups emphasized that the ability to quickly filter through data and create (even traditional 2D) plots is a drastic departure from current analyses that use programs such as Matlab or Excel. With the new system, the analysis was immediate. A surgeon remarked, "The beauty is that I know what I want on the left and I [can go directly from these 2D plots to the linked video] on the right." So, we believe one critical advance in this system is the ability to link video data with interactive 2D and 3D plots, which enables a rapid style of data navigation that wasn't previously possible.

Discoveries Made and Trends Confirmed

One interesting discovery enabled by this system is *sympathetic grasp-release*. When a surgeon applies force to one tool (such as in picking up a block), the force applied to the tool controlled by the other hand often decreases slightly, even though at that moment the second tool isn't performing any task (see Figure 6). Neither group of surgeons had heard of such a pattern. However, once we discovered it, within minutes we found more than 10 instances of it in the dataset.

This led to a series of follow-on questions and hypotheses:

- Does this pattern correlate with reported skill level?
- When might this be useful in surgery? (One surgeon suggested that releasing pressure at one location on a tissue while grabbing another location might sometimes decrease the possibility of damage to the tissue.)
- When might this be detrimental? (One surgeon suggested that while cutting with one hand, it might be important to maintain a constant grasp force with the other to hold the tissue taut to make a controlled cut.)
- Does the same pattern appear in data for other tasks and in more realistic surgical situations? (Both groups wondered whether the pattern would apply in suturing and other tasks.)

This result demonstrates that the system can elicit valuable new insights. It also provides evidence that surgeons can relate patterns elucidated by

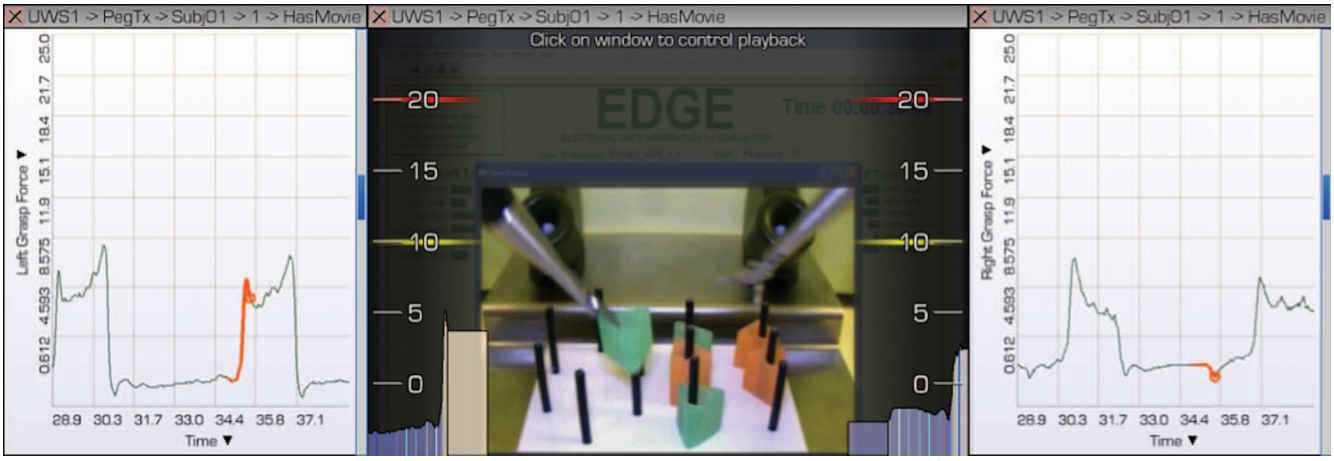

Figure 6. Sympathetic grasp-release. When the left tool grabs the block to pick it up from the peg board, the force on the right tool momentarily decreases. In each 2D plot, the orange circle corresponds to the current video frame, and the highlight shows the past few seconds of data. The large drop in the left plot is the grasp; the small dip in the right plot is the release.

the visualization with their own deep domain knowledge to generate new hypotheses—one of the most important goals of exploratory visualization. In this case, we first noticed sympathetic grasp-release while using 2D force-over-time plots together with a linked augmented video to look at individual transfer events. Through the video, we observed that the force plot for one hand varied when the other hand was picking up the block. Now that we know what to look for, we're confident we can find this pattern through more traditional data analyses. However, without the coordinated linked data displays, we doubt we would have ever discovered it.

We also identified an initial spike in force when a user first grabs a peg, followed by a decrease to a more constant force. We noticed this pattern while developing the video overlays and assumed that it was representative of bad practice. However, during the use case studies, the surgeons in one group provided a different interpretation: this pattern could reflect an experienced surgeon practicing "respect for tissue" (that is, the surgeon quickly reduces the applied force after making a successful grab).

A well-known problem involves moving a surgical tool beyond the intended target; we found that our visualization system can provide an understanding of that situation as well. The depth dimension is typically the most difficult to align, and the 3D visualizations provide surgeons with new virtual viewpoints (for example, a top-down view) to understand 3D tool path data.

Feedback on Specific Visualization Features

The 3D multivariate visualization features are the system's most experimental visual aspect. To make these displays easier to understand, the system includes a simplified 3D tool trace visualization

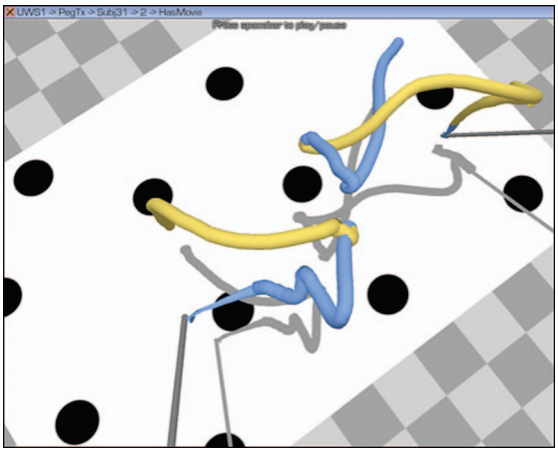

Figure 7. A simplified 3D tool trace visualization. These 3D views provide useful camera angles for understanding 3D tool paths.

(see Figure 7). These 3D views provide useful camera angles for understanding 3D tool paths.

Beyond the simplified views, feedback on the glyph-based multivariate visualization focused on new applications for this type of display. Both groups of surgeons recognized that all the information collected appears in these 3D views and that you could potentially adjust the mappings from data to form according to importance. One surgeon suggested using this display for goal-oriented training, which assigns a readily understandable visual task, such as "reduce the red areas," to a surgeon. To achieve this goal, he or she then would use the multivariate visualization to learn to adjust force, speed, or some other variable.

The smart brushing interface for selecting subsets of data and then calling up new data plots and videos for that subset is one of the system's most important features. All the surgeons said this was an improvement over current methods of examining data. However, surgeons want this

type of selection to work as naturally as possible. One surgeon wanted to be able to say, "I want 12 transfers," or perhaps even, "I want the six left-to-right transfers for this participant." Although selecting transfers such as these using the current interface is possible, there's room to improve the interface to make it more natural for a surgeon's mindset, as opposed to, say, an engineer's. This is an exciting direction for user interface research.

Feedback Related to Iterative System Design

We significantly revised several aspects of the system's design on the basis of collaborative iterative design. During the use case studies, we confirmed the improvement in these features.

We refined the 3D visualization to include a pegboard image, which establishes a spatial context for the data. We simplified shadows to serve as only spatial cues rather than to encode data. We also refined the glyphs representing grab and release events to more concisely convey these events.

The 2D video overlays are the most immediately understandable views in the system. So, we're motivated to make them even more useful and to explore their potential use in follow-on systems designed for real-time formative feedback. For example, one group's surgeons and engineers thought the tinting was undesirable—in operating environments, the surgeon's visual channel is already nearly saturated, so feedback via that modality likely isn't desirable. The groups suggested that audio feedback could be a more promising avenue for future applications to real-time visualization for training.

Our research has focused on improving surgical training and assessment through a more robust, data-intensive approach to analyzing high-resolution surgical performance data. We know from our collaborators that "a good surgeon knows when to go slow and when to go fast." We also know that an expert surgeon can often identify through video whether an expert or novice is performing a surgery. However, two expert surgeons will often approach the same task differently; it's not clear what quantitative, objective, data-driven metrics can best be used to assess skill or to improve surgical training. On the basis of our experiences with our system, we believe visualizations in this style can be catalysts for enabling new discoveries with quantitative surgical data, new forms of communication between surgeons, and ultimately a new paradigm of data-driven surgical training and assessment.

Acknowledgments

The University of Minnesota Digital Technology Center's Digital Technologies Initiative and Washington Technology Center Research Technology Development grant RTD11 UW SB01 supported this research.

References

1. J. Rosen et al., "Generalized Approach for Modeling Minimally Invasive Surgery as a Stochastic Process Using a Discrete Markov Model," *IEEE Trans. Biomedical Eng.*, vol. 53, no. 3, 2006, pp. 399–413.
2. J.H. Peters et al., "Development and Validation of a Comprehensive Program of Education and Assessment of the Basic Fundamentals of Laparoscopic Surgery," *Surgery*, vol. 135, no. 1, 2004, pp. 21–27.
3. B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations," *Proc. IEEE Symp. Visual Languages*, IEEE, 1996, pp. 336–343.
4. C. Ware et al., "Visualizing the Underwater Behavior of Humpback Whales," *IEEE Computer Graphics and Applications*, vol. 26, no. 4, 2006, pp. 14–18.
5. J.V. Carlis, "Pitfalls and Paths to Success in Interdisciplinary Scientific Database Research," *IEEE Eng. Medicine and Biology*, vol. 21, no. 6, 2002, pp. 117–122.
6. J. Carlis and J. Maguire, *Mastering Data Modeling: A User-Driven Approach*, Addison-Wesley, 2000.

David Schroeder *is a doctoral candidate in computer science and engineering at the University of Minnesota. His research interests include how to extend interactive visualization systems to use the data being visualized to inform the interpretation of user input. Schroeder received a BA in computer science from Coe College. Contact him at daschroe@cs.umn.edu.*

Timothy Kowalewski *is an assistant professor of mechanical engineering at the University of Minnesota. His research interests include quantitative surgical-skill evaluation via hidden Markov modeling. He received a PhD from the University of Washington's BioRobotics lab. He originated the Surgical Genome Project. Contact him at tmk@uw.edu.*

Lee White *is a PhD candidate at the University of Washington's BioRobotics Lab. His research interests include skill assessment and training for surgery and medicine and using technology to treat disease and impairment. White received a BSE in biomedical engineering from Tulane University. Contact him at leewhite@uw.edu.*

John Carlis *is a professor of computer science and biomedical informatics and computational biology at the University of Minnesota. His research interests include developing*

better practices in data modeling and querying and developing extensions to database management systems in the context of complex evolving biomedical applications. Carlis received a PhD in business administration from the University of Minnesota. Contact him at carlis@umn.edu.

***Erlan Santos** is a general surgeon working toward developing a simulation center in Lima, Peru. He recently completed an international fellowship at the University of Minnesota. His research interests include surgical simulation, surgical education, and upper gastrointestinal diseases. Santos received an MD from the National University of Trujillo-Peru. Contact him at erlansantos1@hotmail.com.*

***Robert Sweet** practices urology, focusing on robotics and endourology. His research interests include developing a human-tissue property database, advanced real-time and predictive modeling of human-tissue and tool interactions, and systems for disseminating simulation technologies. Sweet received an MD from the University of Minnesota. Contact him at rsweet@umn.edu.*

***Thomas S. Lendvay** is an associate professor of urology at the University of Washington. His research interests include minimally invasive surgical approaches to pediatric patients and surgical-simulation and robotic-technology development. Lendvay received an MD from Temple University. Contact him at thomas.lendvay@seattlechildrens.org.*

***Troy Reihsen** is the director of operations for the University of Minnesota's Simulation PeriOperative Resource for Training and Learning and the Center for Research in Education in Simulation Technologies. His research interests include anaplastology, usability and efficiency metrics, medical education, and tactical medicine. He served more than 22 years in the US Army and recently retired. Contact him at reih0005@umn.edu.*

***Daniel F. Keefe** is a McKnight Land-Grant Assistant Professor in the Department of Computer Science and Engineering at the University of Minnesota, where he directs the Interactive Visualization Lab. He specializes in visualization and interactive computer graphics. Keefe received a PhD in computer science from Brown University. Contact him at keefe@cs.umn.edu.*