Editor:
Theresa-Marie Rhyne

Can Computers Master the Art of Communication?

A Focus on Visual Analytics

Klaus Mueller, Supriya Garg, Julia EunJu Nam, Tamara Berg, and Kevin T. McDonnell
Stony Brook University

Visual analytics seeks to conduct a discourse with the user through images, to stimulate curiosity and a penchant to decipher the unknown. Figure 1 depicts our view of the visual analytics process. The computer supports the user in this interactive analytical reasoning, constructing a formal model of the given data, with the end product being formatted knowledge constituting insight.

Yet, validation and refinement of this computational model of insight can occur only in the human domain expert's mind, bringing to bear possibly unformatted knowledge as well as intuition and creative thought. So, it's left to this human user to guide the computer in the formalization (learning) of more sophisticated models that capture what the human desires and what the computer currently believes about the data domain, perhaps with an associated confidence level. In visual analytics, the computer uses images and text (and possibly sound and haptics) to exchange information with the user about its view of the domain model.

Obviously, the better a communicator the computer is, the more assistance it will elicit from the user to help it refine the model. This in turn leads to this article's topic—the need for the computer to master the art of interpersonal communication—that is, communication between it and the human analyst.

The Elements of Interpersonal Communication

Obviously, communication is present in many domains, not just in human behavior. Communication protocols are part of many human-made systems, such as computing and telecommunication, and they follow similar definitions. We focus here on human behavior because we aim for the computer to collaborate with the human user.

The *interpersonal-communication protocol*[1] (see Figure 2) always includes

- *a sender*—someone who sends a message verbally or nonverbally to someone else,
- *a receiver*—someone who receives the message,
- *a message*—information in some shape or form,
- *noise*—anything interfering with the exact replication of the transmitted information,
- *feedback*—verbal and nonverbal feedback elicited from the sender or receiver,
- *replication*—one person understanding what's in another person's mind, and
- *understanding*—the receiver's approximation of what the message means to the sender.

The interpersonal-communication framework has three components. *Direct channels* encompass information that the sender directly controls; they're easily recognized by the receiver. *Indirect channels* aren't always under the sender's direct control and are usually recognized subconsciously by the receiver. The *context* is the conditions surrounding the communication from which the receiver can derive the message's meaning.

Communicators use *intonation* or *pitch* to emphasize words and passages. *Brevity* or *economy of words* leads to clear, effective presentations, whereas an *aesthetic choice of words* (good storytelling) can generate more interest, attention, and even fascination. Finally, *personalization* of word choice can target a specific receiver, just as the word choice can indicate a sender's *identity*.

Clearly, some people are more eloquent in these matters than others; the same is true for human-

 Published by the IEEE Computer Society

computer communication interfaces. Next, we map the interpersonal-communication structures and strategies we've discussed into forms realizable by a computer, focusing on visualization and visual interaction.

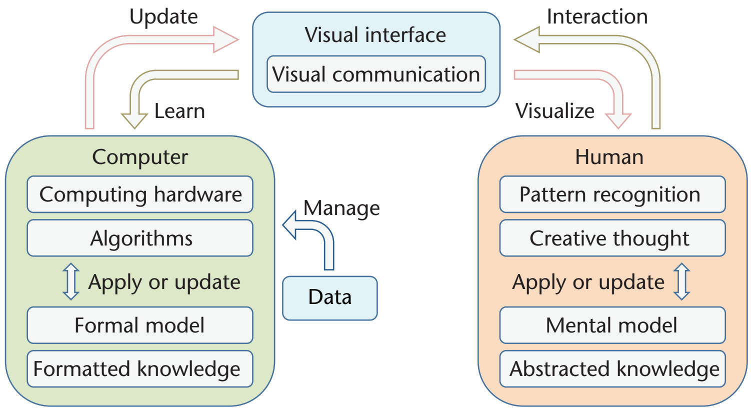

Figure 1. The visual analytics process. The user seeks to create a formal model that captures the analytics problem's underlying mechanisms. In an iterative learning process, the user continuously teaches the computer using his or her informal domain knowledge and expertise until the evolved model sufficiently explains the data.

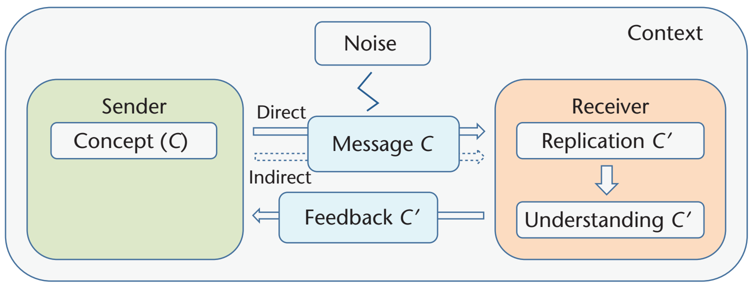

Figure 2. The interpersonal-communication protocol. A sender would like a receiver to comprehend message *C*, conveyed either straightforwardly or via indirect or subconscious mechanisms. However, noise in the communication channel or the receiver's failure to fully comprehend the message's intended meaning can undermine the sender's objective. An iterative clarification process eventually leads to a mutual understanding of the message.

A Visual Human-Computer Communication Protocol

In keeping with this article's general theme, we use the term *human-computer communication (HCC) protocol*. Such a protocol requires certain HCC skills. For example, both participants need to

- master the available devices or resources, such as a monitor (for the computer) and a mouse, keyboard, and multitouch input device (for the human), and
- correctly interpret messages.

As Figure 2 indicates, in HCC both the computer and human can act as sender and receiver. For the computer, the equivalent of a verbal message is a visual artifact (utterance) of some sort (a graph, scatterplot, graphics rendering, picture, and so on). The equivalent of a human verbal message is some sort of interaction, be it a swipe on a multitouch display, a menu item selection, or parsable text.

In the following, we focus on messages sent in the context of some problem-solving task. The ultimate goal is for the computer to acquire a formatted model of the user's unformatted knowledge, while being guided and inspired by the analytical problem at hand. We further assume that the human user is the ultimate judge of the correctness of the messages the computer sends. The computer only *supports* reasoning and diagnosis, and there's always room for improvement (and training).

So, in this communication loop,

- the human sends feedback messages to trigger model refinements (the computer's replication of the human's intuition), and
- the computer responds by sending feedback messages of the model's visual encoding (the computer's understanding of the human's intuition).

On the other hand, the human's replication and understanding of the feedback message are determined by his or her expert knowledge and the visual encoding's quality, which is a direct function of the computer's visual communication skills.

Of course, the human's intuition regarding the model might be partially incorrect or incomplete. So, the computer's understanding expressed by the uttered visual will be necessarily subject to critique. This is what constitutes the iterative nature of the human visual discourse promoted by visual analytics. Here, longer iteration cycles or poorer iteration results might be due to the computer-based reasoning engine's limited skill, the human's poor input devices or gesturing skills, or simply the human's lack of intuition and domain knowledge.

Finally, noise in the visual encodings might stem from the data itself and might prohibit the evolution of a consistent model. Computers deal poorly with noise, whereas humans excel. This is an additional significant strength of visual analytics—by incorporating the human to detect and neutralize noise, we can drive model development forward.

Trust

As with any form of communication—human-human or human-computer—trust is paramount. A visual analytics system's user must feel confident that the software provides trustworthy, unbiased information, guidance, and advice.

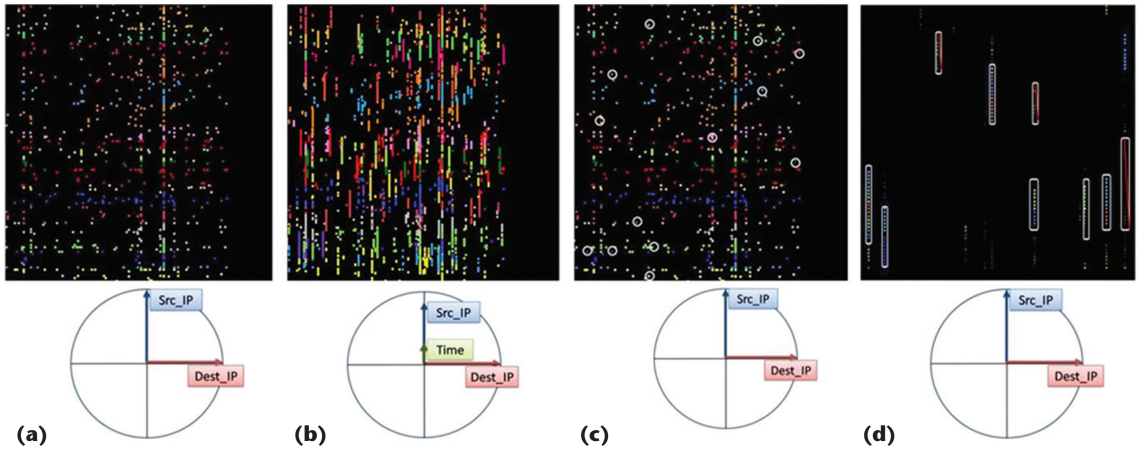

Figure 3. A case study of the iterative cooperative model-learning loop for a network dataset. (a) The initial scatterplot of source IP addresses (dest_IP) versus destination IP addresses (src_IP). All packets projecting onto similar (*x*, *y*) coordinates are the same color. (b) Tilting the projection plane to reveal time relationships. Formerly superimposed packets now appear as shorter or longer streaks along the *y*-axis. (c) Tilting back and marking packets with longer time frames. (d) Visualizing newly classified packets given the learned rule.

For many years, researchers have been considering uncertainty in visualization and to what extent a given visualization accurately reflects reality. Uncertainty can arise from various stages of the visualization pipeline:[2] from data acquisition (for example, calibration, observation, and numerical simulation), to data transformation (for example, interpolation, quantization, and aggregation), to the visualization itself (for example, algorithms, color maps, and rasterization).

Although there's widespread agreement that quantification and visualization of information uncertainty are important, we still lack a standard, general framework for dealing with inaccurate, incomplete, inadequate, or suspect data. This is partly because "uncertainty" is largely a domain-specific term meaning different things to different people.[3] Moreover, in some disciplines, it's possible and desirable to gauge error within a certain tolerance, whereas in other fields it's sufficient to make qualitative statements about the confidence you can place in data accuracy (for example, high, medium, or low). Finally, debate is still ongoing about whether uncertainty is itself a form of data and should be visualized using traditional means or whether we should treat it as metadata for augmenting data visualization. Decisions about these and other aspects of data uncertainty necessarily impact HCC's effectiveness.

Elements of Effective Visual Computer Communication

We now present techniques to enhance the computer's visual communication skills and therefore HCC. To make this discussion more concrete, we provide examples from our own research—further evidence is available throughout the visualization literature.

Communication in the Iterative Cooperative Model-Learning Loop

We start with a practical incarnation of the visual analytics process depicted in Figure 1. Our example, described in more detail elsewhere,[4] is set in network traffic analysis. The (very large) dataset consists of a one-hour snapshot of Internet packets with attribute information describing the source and destination IP addresses and ports, as well as the time stamps, packet IDs, and protocol numbers. Assume that we want to learn the concept of *webpage load* from this data, using Prolog logic programming to formalize this model (which will later be used for automated diagnosis of network traffic).

Identify interesting patterns. The first message interchanges focus on exploring the data. The analyst considers the source and destination IP addresses the most critical attributes. So, the initial scatterplot view is a projection of the source versus destination IP addresses (see Figure 3a). The analyst then hypothesizes that time might be an interesting secondary variable. So, by manipulating our touch-pad multidimensional-navigation interface,[4] the analyst gradually changes the projection plane's tilt (along the *y*-axis) to distinguish the superimposed packets by their time stamp (see Figure 3b).

The motion parallax reveals that some source and destination IP packet interchanges have different time spans. Knowing from experience that webpage loads usually require only short time spans, the analyst marks some of these packet sequences in the original source-versus-destination plot (where they project into single points—see Figure 3c). The analyst then sends these examples to the inductive logic programming (ILP) system, which formulates a rule describing the set of examples, using all their attributes (not just IP addresses and time stamps). The resulting rule (in Prolog syntax) is the computer's current understanding of webpage load:

```
webpage_load(X) :-
same_src_ips(X),
   same_dest_ips(X),
   same_src_port(X,80),
   timeframe_upper(X,10).
```

Assess the constructed model. Here, the analyst verifies this concept understanding. The computer demonstrates to what extent it has replicated the concept by retrieving from the data file a random subset of tuples that fit this rule and showing them in the dynamic scatterplot display (see Figure 3d). The analyst examines these patterns using the touch-pad multidimensional-navigation interface. In this case, the analyst notices that some tuple sequences have a very small number of exchanges, which are probably not due to a webpage load. The analyst concludes that the model isn't fully developed and that the system requires further training, as we describe next.

Critique the constructed model. The analyst marks the erroneously classified packet time sequences as negative examples and returns them to the ILP system for rule refinement. This leads the computer to infer the (final) Prolog rule:

```
webpage_load(X) :-
same_src_ips(X),
   same_dest_ips(X),
   same_src_port(X,80),
   timeframe_upper(X,10),
   length(X,L),
   greaterthan(L,8).
```

Communication via Direct and Indirect Channels

The canonical visual direct-channel communication example is the pictogram, an icon designed to unambiguously convey meaning, often by a pictorial resemblance to a physical object. Pictograms are simple, easy to recognize, intuitive, and don't have to be learned. Many signs in public life are pictograms, such as traffic signs or the trashcan icon on computer desktops. Some patterns in information visualization can also be iconic, such as a straight up- or down-sloped line representing growth or decline.

Any visual sign more complex than a pictogram sends parts of its messages across indirect channels. For example, although a picture of a person can never fully describe the real person, it can show aspects or properties reminding us of that person or what he or she represents (which is also likely subject to personal opinion). For example, a picture of Andy Warhol might signify just him or the entire pop art culture. In the former case, the picture is *iconic*; in the latter, it's *indexical* (see the writings of Charles S. Peirce [1839–1914] on semiotics).

So, a picture serves essentially as an index (or pointer) into the viewer's conscious (and subconscious) memory,[5] triggering further private analytical processes in the viewer's mind. Thus, we might say that the picture transmits information and that the viewer adds further personal information. For example, browsing vacation pictures, an old high school yearbook, or even a picture of some person, place, or object can bring back many memories from the past (and might lead to new, spontaneously generated thoughts).

The same goes for visual phenomena in scientific and information visualization, such as patterns in flow fields or scatterplots. The message they invoke flows across indirect channels, indexing the user's past experiences and stimulating creative interpretation and thought. This then hopefully leads to novel associations and insight in the process, given the user's current frame of mind. So, these indirect channels form the essential communication pathways for visual analytics. However, they require continuous feedback to confirm, refine, and adjust the interpretations; this is the foundation of the analytical human-computer discourse.

Brevity by Mental Indexing

Given that visual signs trigger information retrieval from the brain, we can consider them a form of information compression and therefore use them to efficiently enrich an information display. This has been amply exploited in glyphs often employed in scientific visualization. Also, graphs in information visualization frequently use generic icons for quick recognition, but these have limited capacity to communicate specific information. This inflexibility is a significant shortcoming of using icons to index into the viewer's mental database. In essence, a picture is just an array of pixels—it's the brain that interprets, understands, and indexes into memory to recover the vast amount of information encoded by this array (which is a measurement of the world).

We recently described a framework involving natural-language processing and Web-scale image databases to help users identify metaphors suitable to visually encode abstract semantic concepts.[6] The visual representations resulting from these

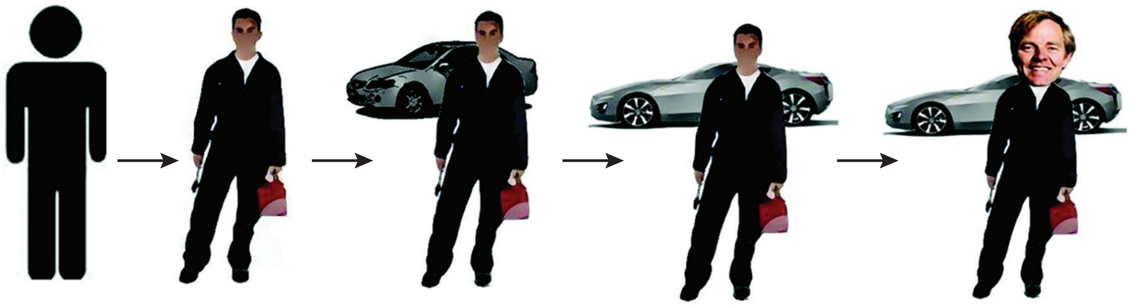

Figure 4. Semantic zooms. These are images with increasing semantic detail, here for the concept "Bill, the sports car mechanic." By embedding these visual encodings into a suitable background scene, users can provide and refine contextual indexing.

Figure 5. Detail abstraction using multiscale edge detection. (a) The original purse image. (b) A generic purse image with weak edges removed. (c) Strong and weak multiscale edges.

encodings can serve as stand-alone expressive icons or clip art in visualization and visual analytics applications. Furthermore, just like verbal communication, visual communication provides for a semantic level of detail. Our interactive system lets users create visual semantic zooms, using lexical databases such as WordNet and Lexical FreeNet to shape a text query in terms of the search concept's scope and detail and then retrieve the corresponding image. Figure 4 shows such an example for the "Bill, the sports car mechanic" concept. By embedding these visual encodings into a suitable background scene (for example, a semantic object), users can provide and refine contextual indexing.

Brevity by Abstraction

The essence of a given concept is some instance of it with irrelevant details abstracted away. You can achieve this by generalizing its semantics or trimming away specifications at some level of detail. Just as in verbal communication, this can be achieved in visual communication, and just as savvy speakers effectively use illustrative language, visualization can use effective graphics techniques.

Here, we go far beyond standard image-processing techniques such as linear filtering, instead using statistical analysis appropriately. Besides reducing the visual to the detail essential for the given task, the ensuing abstraction can visually indicate that the missing detail is unknown, thus sparking curiosity and the desire to obtain it. For instance, our sports-car-mechanic example could be the visual equivalent of verbally stating "some mechanic," which might trigger an investigation that yields more and more detail until eventually Bill is identified.

More specifically, for visual signs, we seek a picture of a given concept that unifies all the concept's known facts while abstracting away the unknown facts. In images, a fact is expressed as a visual feature or collection of features. We can construct an average image for a category by looking for common features across a set of queries. We might know only that our suspect is a male blue-collar worker. So, we can formulate a corresponding search query of this term and use our lexical resources to retrieve images of instances of the term—mechanics, plumbers, electricians, and so on. We can then remove the difference in detail across images by finding the level of image-based abstractions at which all images appear similar. (The visual analytics process would then further investigate this subject, refining the corresponding visuals until it identifies a particular individual.)

We implemented this method using multiscale edge detection to compute shape context histograms at multiple scales (see Figure 5). We then packed these into high-dimensional feature vectors and clustered them for all images in the set. At some level of scale, the different expressions for a concept are the most similar; this forms the scale on which we base the abstraction. The cluster exemplar (an image closest to the cluster center) then represents this concept's average instance at that scale, and we create the final icon by abstracting away all details at higher scales. For this, we use Poisson blending guided by the noise-free edges to create an abstracted illustration that contains only features of known facts.

For information visualization, the abstraction must be derived directly from the data (and background information if available) and can be obtained using statistical analysis. Previously, we described an illustrative information visualization framework for high-dimensional data that used parallel coordinates.[7] In this framework, we employed *k*-means clustering, mean, standard deviation, and correlation for analysis, and we used the results to determine visual salience. As illustrative style elements, we employed halos, shadows, color, shading, and filled contours. Figure 6 shows an example. Here, one data layer might be the one of interest, whereas the others serve as context.

As in verbal communication, abstraction in the visual domain can lead to salient facts being abstracted away accidentally (or purposely, which often occurs in advertising or magazines). Because data analysis drives the abstraction in visual analytics, more sophistication can alleviate these problems.

Intonation and Highlighting

The visual equivalent to intonation is highlighting. Proper coloring of the feature of intended emphasis can focus attention, called *pop-out*. Pop-out exploits the visual system's low-level mechanisms, which let humans detect visual properties rapidly (although not necessarily consciously). Pop-out is strongly related to a color patch's vividness, its size, and the degree to which it differs from the vividness of other colors in the scene (and local brightness contrast).

Figure 7 illustrates results from a color design framework we recently devised.[8] It shows 2D colorization of an image of biological cells, generated through transmission electron tomography. This colorization also demonstrates how the original image gray-level intensities modulate the lightness of the assigned object colors (here, the cells). In this figure, the labels denote the object classes, colored according to their importance.

Another effective way to alert a receiver to an important fact in verbal communication is exaggeration. In visual communication, an artful embodiment of this concept is caricatures. Peter Rautek and his colleagues recently described a framework that uses caricatures to focus on shape deviations that would be difficult to convey with color.[9]

Visual Aesthetics

Speakers and authors who deliver their message in a vibrant, imaginative, and appealing manner will keep audience attention levels high and will likely generate better audience response and recall. Likewise, good visual aesthetics and well-designed interaction frameworks make visual exploration an enjoyable task while reducing stress, which can be a significant factor in today's fast-paced data analysis scenarios. Stress caused by unpleasant design makes people less able to cope with difficulties, less flexible, and less creative. So, aesthetic design plays an important role in the problem-solving task that's at the core of visual analytics systems. For more detail on academic research and

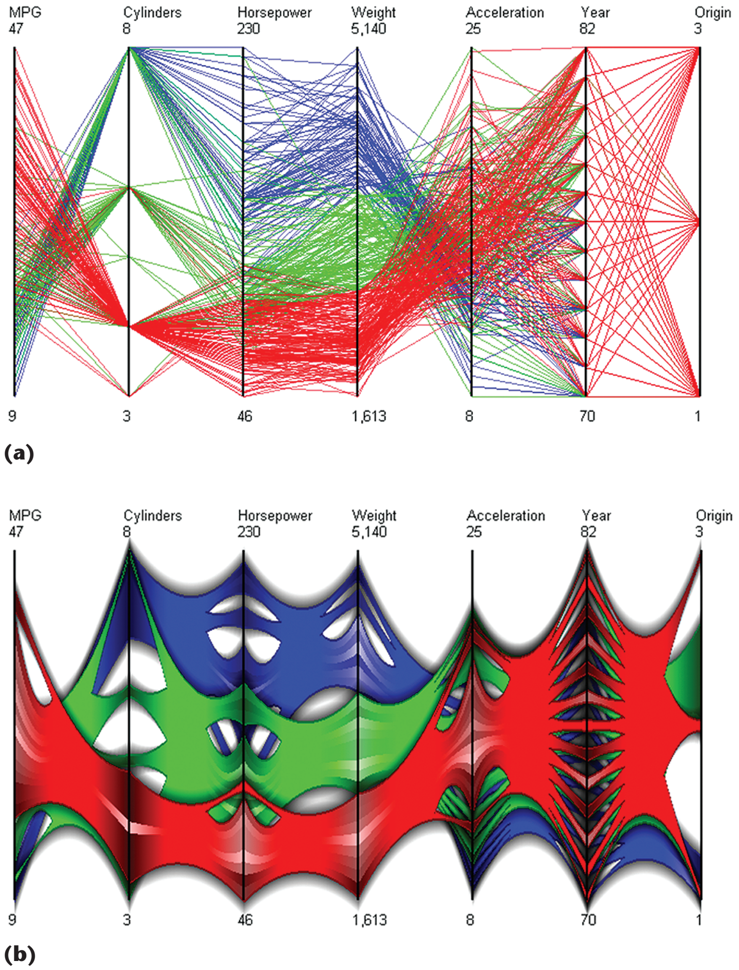

Figure 6. Abstraction for information visualization. (a) Traditional parallel coordinates. (b) Illustrative parallel coordinates with abstraction derived via statistical data analysis.

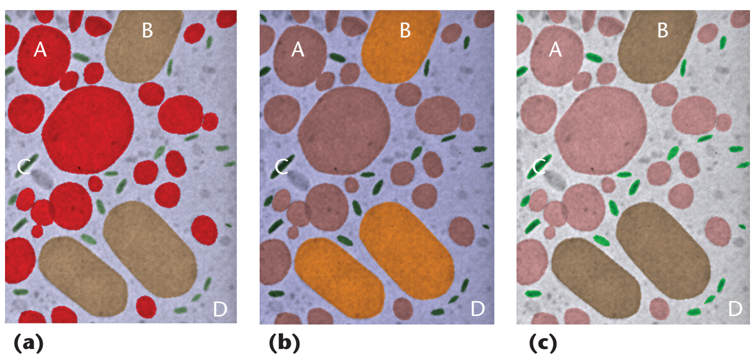

Figure 7. Colorization of transmission-electron-microscopy data to indicate relative importance, where (a) object class A is the most important, (b) class B is the most important, and (c) class C (the long, elliptical cells) is the most important.

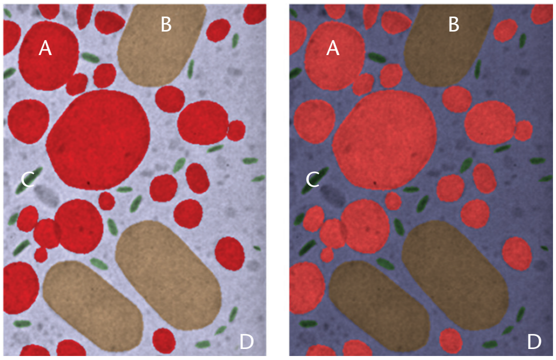

Figure 8. Aesthetics versus highlighting. For both images, feature A (in red) is the most important, and both images were among the top four in terms of correct identification of the most important feature. However, the right image was one of the four rated least aesthetic, so the left image's coloring is preferred.

practical guidelines relevant to visual aesthetics, see the "Related Work" section in the paper "Color Design for Illustrative Visualization."[8]

Keeping up good aesthetics while pursuing (and achieving) a goal is a challenge for every public speaker. It requires a well-balanced interplay of emphasis and aesthetic design. Otherwise, conflicts emerge—speakers don't want to overemphasize the intent and make the overall experience less enjoyable. The previous section discussed color as an important means for emphasis. A basic design rule is that viewers perceive excessive use of vivid colors as unpleasant and overwhelming. So, you should use them between duller background tones.

We recently conducted a user study to gain insight on how to use color for highlighting while preserving an aesthetic design (Figure 8 shows some results). We formulated this (and other) insight and color design principles into a rule-based system for data colorization.[8]

Sender and Identity

Speakers and authors typically have recognizable styles that provide them a sense of identity. For visual analytics, knowing the origin of an original or derived fact or hypothesis can be interesting and informative. Illustrative visualization offers tremendous opportunities to watermark personal identities into a visual rendition of these elements. We could embed this information either as a frame around the corresponding visual encoding or directly into the abstraction style. The latter approach would give each abstraction style a specific personality that can be identified with a specific information source (such as an agency or an analyst).

Human Feedback

So far, we've focused mostly on how computers can communicate visually with human observers. The visual analytics communication loop in Figure 1 also requires human feedback. Besides traditional media such as keyboards, mice, and voice, multitouch (gesture) interfaces have become increasingly popular, thanks mostly to Apple's iPhone and iPad. However, exploitation of gestures as a feedback device in visual analytics is still at an early stage.

We augmented our touch-pad multidimensional-navigation interface[4] with a two-finger multitouch interaction framework that lets users position and orient the two orthogonal projection axes in high-dimensional space simultaneously. This affords much more direct and fluid interaction with the virtual spaceship, thus fostering more direct high-dimensional-space exploration and a better understanding of high-dimensional relationships.

How to Train Your Computer

How can a computer actually become, or evolve into, a better communicator? The best way is through formal user studies involving teaching sessions with the intended user group. A visual reasoning environment, particularly a collaborative one, will potentially incorporate very different people with considerably different qualifications, preferences, and backgrounds. These parameters will determine the visualization's inherent complexity and style. The "Measuring the User" sidebar provides further thoughts and information on this topic.

For example, we devised a preference-measuring methodology based on conjoint analysis, which we demonstrated in a study on measuring the perceived quality of volume rendering.[10] We also used our methodology to devise the color design rules we mentioned earlier.[8] These studies produced many interesting insights. For example, we found that participants disliked visualizations of engine blocks when the engine is oriented as if it's standing on a corner, whereas they didn't object to such a view for medical foot datasets. We concluded that feet are often oriented at arbitrary positions in everyday life, whereas an engine flying through the air appears threatening. Likewise, we found that participants considered black backgrounds more aesthetic but deemed white backgrounds better for studying detail.

Computers, through visualization, have access to a rich repertoire of communication patterns. Although more research is necessary until a computer can truly master the art of communication, we believe that the growing ease of online user studies will greatly facilitate the training required.

References

1. K. Hogan, "Communication," 2010; www.kevinhogan.com/communication-influence-persuasion.htm.
2. A. Pang, C. Wittenbrink, and S. Lodha, "Approaches to Uncertainty Visualization," *The Visual Computer*, vol. 13, no. 8, 1997, pp. 370–390.
3. M. Skeels et al., "Revealing Uncertainty for Information Visualization," *Proc. Working Conf. Advanced Visual Interfaces*, ACM Press, 2008, pp. 376–379.
4. S. Garg et al., "Model-Driven Visual Analytics," *Proc. 2008 IEEE Symp. Visual Analytics Science and Technology* (VAST 08), IEEE Press, 2008, pp. 19–26.
5. C. Koch and G. Tononi, "Can Machines Be Conscious?" *IEEE Spectrum*, vol. 45, no. 6, 2008, pp. 54–59.
6. S. Garg, T. Berg, and K. Mueller, "Iconizer: A Framework to Identify and Create Effective Representations for Visual Information Encoding," poster presented at 2009 IEEE Visualization Conf., 2009; www.cs.sunysb.edu/~mueller/papers/iconizerVis09.pdf.
7. K. McDonnell and K. Mueller, "Illustrative Parallel Coordinates," *Computer Graphics Forum*, vol. 27, no. 3, 2008, pp. 1031–1038.
8. L. Wang et al., "Color Design for Illustrative Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 14, no. 6, 2008, pp. 1739–1754.
9. P. Rautek, I. Viola, and E. Gröller, "Caricaturistic Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, 2006, pp. 1085–1092.
10. J. Giesen et al., "Conjoint Analysis to Measure the Perceived Quality in Volume Rendering," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, 2007, pp. 1664–1671.

Klaus Mueller *is an associate professor in Stony Brook University's Computer Science Department and the director of the university's Visual Analytics and Imaging Lab. Contact him at mueller@cs.stonybrook.edu.*

Supriya Garg *is a PhD student in Stony Brook University's Computer Science Department. Contact her at sgarg@cs.sunysb.edu.*

Julia EunJu Nam *is a PhD student in Stony Brook University's Computer Science Department. Contact her at ejnam@cs.sunysb.edu.*

Tamara Berg *is an assistant professor in Stony Brook University's Computer Science Department. Contact her at tlberg@cs.sunysb.edu.*

Kevin T. McDonnell *is a visiting professor in Stony Brook University's Computer Science Department and an associate professor of computer science and the chair of the Department of Mathematics and Computer Science at Dowling College. Contact him at mcdonnek@dowling.edu.*

Measuring the User

A main obstacle to achieving true human-like (artificial) intelligence is that human consciousness depends greatly on highly personal semantic models, knowledge, experiences, skills, preferences, and the like. These things are hard to capture and to encode in machines. We face the same obstacles when encoding data and information into visual representations. These differences are expressed horizontally (same complexity, but different representation) and vertically (reduced complexity). Whereas the former is more a function of personal preferences, possibly motivated by professional or community background, the latter is a function of educational background, classification of the information visualized, and task-mandated (minimal) requirements. So, no single visual reasoning environment will fit all participants, yet it must let all participants communicate with each other and the computing engine.

The key is to develop parameterized models of users and tasks, methodologies to acquire and test those models, procedures to generate the user- and task-suitable visualizations, and appropriate means to translate one representation into another (also called grounding[1,2]). For this, we must capture personal preference vectors (in terms of visualization paradigms) and correlate them with other user information, such as background and education level. We can then use these frameworks to parameterize tasks and knowledge.

To provide these models, we need a rich suite of user studies. Market research has developed statistical frameworks (such as conjoint analysis[3,4]) to efficiently acquire user studies, with a minimal set of users and minimal user involvement. Such models will yield adaptive user interfaces that can eventually predict the best visual representations of the information at hand. Once we've formulated these models, systems can automatically coach analyst users in developing strategies for more complex analyses. Furthermore, we'll be able to generate templates covering the best strategies for the most efficient analysis.

References

1. H. Clark and S. Brennan, "Grounding in Communication," *Perspectives on Socially Shared Cognition*, L.B. Resnick, J. Levine, and S.D. Teasley, eds., Am. Psychological Assoc., 1991, pp. 127–149.
2. S. Brennan et al., "Toward a Multi-analyst, Collaborative Framework for Visual Analytics," *Proc. 2006 IEEE Symp. Visual Analytics Science and Technology* (VAST 06), IEEE Press, 2006, pp. 129–136.
3. J. Giesen et al., "Conjoint Analysis to Measure the Perceived Quality in Volume Rendering," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, 2007, pp. 1664–1671.
4. A. Gustafsson, A. Herrmann, and F. Huber, "Conjoint Analysis as an Instrument of Market Research Practice," *Conjoint Measurement: Methods and Applications*, A. Gustafsson, A. Herrmann, and F. Huber, eds., Springer, 2000, pp. 5–45.

Contact department editor Theresa-Marie Rhyne at theresamarierhyne@gmail.com.