



Beyond Control Panels *Direct Manipulation for Visual Analytics*

Alex Endert

Pacific Northwest National Laboratory

Lauren Bradel and Chris North

Virginia Tech

To tackle the onset of big data, visual analytics (VA) seeks to marry the human intuition of visualization with mathematical models' analytical horsepower. A critical question is, how will humans interact with and steer these complex mathematical models? Initially, users applied direct manipulation to such models in the same way they applied it to simpler visualizations in the premodel era—by using control panels to directly manipulate model parameters. However, opportunities are arising for direct manipulation of the model outputs, where the users' thought processes take place, rather than the inputs. Here we present this new agenda for direct manipulation for VA.

Direct Manipulation for Information Visualization

Direct manipulation specifies three principles for interaction design for information visualization:¹

- continuous representation of the object of interest,
- physical actions or labeled button presses instead of complex syntax, and
- rapid incremental reversible operations whose impact on the object of interest is immediately visible.

Typically, these principles are applied through a control panel, containing visual widgets such as sliders, buttons, or query fields, coupled to the parameters of a visual representation in the main view. For example, in Spotfire, analysts can choose attributes

to map to available visual encodings (node color, size, shape, and so on); select variables for the *x*-, *y*-, and *z*-axes; and adjust sliders to filter by ranges on specific data dimensions (see Figure 1). We contend that for VA, with the introduction of complex mathematical models behind the visualizations, direct-manipulation interaction has the opportunity to evolve beyond the use of control panels.

Spatializations for Sensemaking

Spatializations create a visual representation of information in which data items' relative proximity approximately depicts their similarity. (That is, the “near \approx similar” metaphor holds true.) For example, in Figure 2, clusters of documents represent themes or topics of interest. Such spatializations can be generated manually or computationally.

Manual Generation

Analysts can leverage manually generated spatial layouts to aid their analyses. For example, by organizing spatial layouts, they can externalize their insights about a dataset on the basis of the information's positions.² They frequently organize such layouts according to complex schemas using mixed metaphors, often organized topically according to the semantics relevant to their current analysis needs.

Analysts use tools that support manually constructing spatializations to visually synthesize hypotheses.³ That is, they directly manipulate spatial structures (often mixing clusters, timelines, connections, geography, order of discovery, process waypoints, and so on) that help reveal their sense-making process. Such informal relationships in

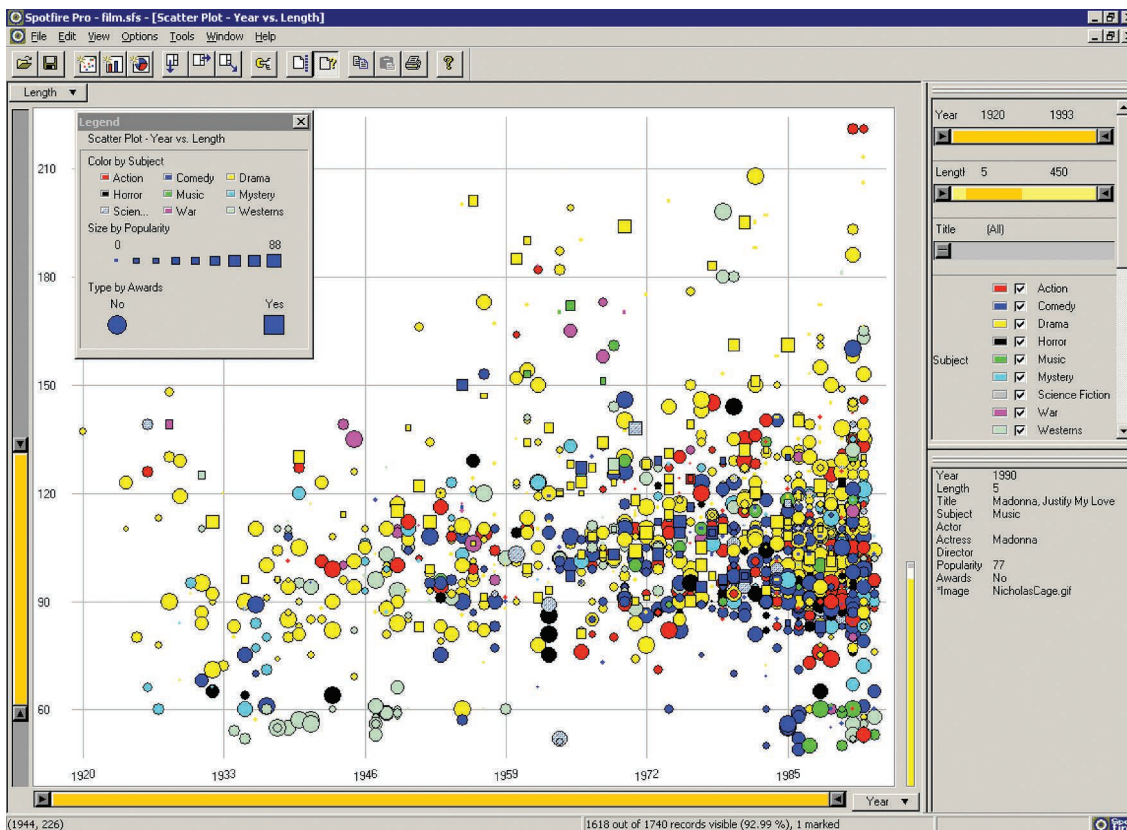


Figure 1. Typical use of direct manipulation. The Spotfire scatterplot view can represent several dimensions of the data through spatial position and visual encodings; users manipulate it through buttons and sliders on control panels.




Figure 2. The In-Spire Galaxy View represents documents as dots. Each cluster of dots represents a group of similar documents.

the spatial layout are beneficial because they don't require analysts to overformalize relationships too early in the process. This process of gradually increasing relationships' formality is called *incremental formalism*.⁴

Computational Generation

Computationally generated spatializations are driven by the recent emphasis on big data and involve complex mathematical models. These models, combined with user intuition and visualizations,

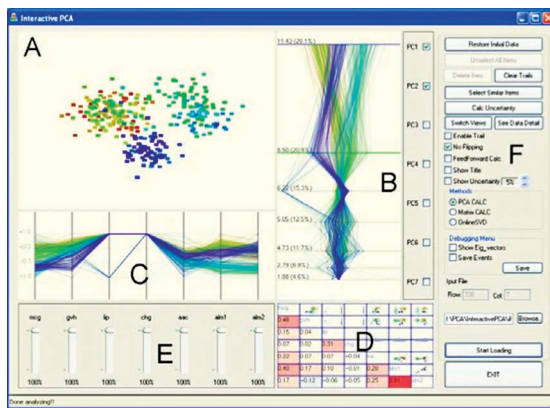


Figure 3. iPCA (Interactive Principle Component Analysis) provides a dimension reduction algorithm that users manipulate through buttons and sliders in a control panel.⁸ (Source: Remco Chang; used with permission.)

form the basis for VA, in which analysts operate dynamic tools that facilitate analysis and sense-making of large, complex datasets.⁵ Models leveraged in VA tools include, but aren't limited to, those for entity extraction, topic modeling, link analysis, dimensionality reduction, clustering, and labeling.

These models employ various distance metrics to measure the similarity between data objects. Analysts can use these metrics to spatialize data. For example, unstructured text can be represented as a “bag of words”—high-dimensional data in which each dimension is a unique keyword or phrase in the text. For example, in In-Spire's Galaxy View layout, nearby points represent similar documents (see Figure 2).⁶ This helps analysts recognize relationships between documents and between clusters of documents.

Designing User Interaction for Spatializations

For computationally generated spatializations, the question arises of how to design user interaction. The complex statistical models that compute similarity using a combination of algorithms have numerous parameters to tune on the basis of the analysis's context. For example, for visual text analysis, users must directly adjust keyword weights (measures of importance for each keyword and how much it influences the overall layout), add or remove documents and keywords, or provide more information on how to parse the documents for keyword entities upon import.

One such spatialization for streaming text data is Streamit, in which users explore a dataset by directly manipulating keyword weights.⁷ Similarly, iPCA (see Figure 3) is an interactive visualization tool that uses principal component analysis (PCA) to reduce high-dimensional data to a 2D plot.⁸ Users employ sliders and other visual controls to

directly adjust numerous model parameters, such as individual eigenvalues, eigenvectors, and other PCA components. In this way, they can observe how the visualization changes. This lets them gain insight into a dataset, assuming they know enough about the underlying PCA model to understand the implications of changing model parameters.

The straightforward application of direct manipulation suggests creating graphical controls for each parameter. This use of control panels might have been appropriate for early information visualizations in which the controls mapped naturally to dimension filters and plot axes. However, it might be problematic for more complex models used in VA applications.

This approach has three fundamental usability problems. First, many analysts aren't experts in complex mathematical models and thus don't understand the meaning of the parameters for the interactive controls. Second, analysts think about and understand the documents at the semantic level, yet the interactive controls for the models operate at the lower syntactic level of the model parameters. This creates a mismatch. Third, when analysts haven't yet gained a good understanding of the documents and their insights are still informal, they don't yet have a basis for expressing their inputs into the formal model parameters. These problems arise because the focus of direct manipulation in the computationally generated spatializations (the model parameters) differs significantly from that in manually generated spatializations (the documents).

Suppose an analyst recognizes a small set of documents in a spatialization that she believes are related to a semantic topic X of her interest, but the current layout doesn't reflect her hypothesized similarity. She directly increases the weight of term X in the control panel (for example, by directly manipulating the layout parameters). However, this has no effect because X doesn't appear in the documents.

Alternatively, she could move the documents together herself (for example, by directly manipulating the layout output). She could then receive feedback from the models concerning other interesting keywords that do relate to those documents. Also, the layout could be automatically updated to include other relevant documents. This would enable her to gain insight that helps to better formalize her understanding of X.

This approach presents an opportunity to evolve the design of user interaction beyond control panels to achieve direct-manipulation VA. The need exists to cooperatively integrate computationally

generated spatializations with those manually generated. This would shift the focus of interaction from control panels for model input parameters to direct manipulation of the model outputs as represented by the spatialization itself.

Direct Manipulation of Spatializations

A trend is emerging in how VA systems that use complex statistical models handle interaction. This trend stems from letting users directly manipulate the data in a spatialization to guide and improve the layout according to their interests or interpretations. For example, to indicate that two data points in a spatialization differ more than is computationally indicated, users can move them apart directly in the view. So, the model learns about the dissimilarity and updates the spatialization to reflect the desired structure of the data. Thus, users can employ familiar direct-manipulation interactions within familiar spatialization metaphors, enabling them to interact with complex, unfamiliar mathematical models.

Within the spatial metaphor, we see three levels of interactivity that motivate this emerging concept of direct-manipulation VA. These levels are based on the extent to which machine learning steers the model.

The first level is direct manipulation of spatial constraints. These interactions let users place (and move) spatial constraints directly in the spatialization. For example, the Dust & Magnet tool lets users place a series of “magnets” representing specific data dimensions or keywords in the spatialization.⁹ Data objects rich in those dimensions are more attracted to the magnets. Such direct manipulation enables users to guide the spatialization layout by placing additional query-like attractors in the space.

The second level is direct manipulation of parameter weighting. Such data-centric interactions leverage metric-learning techniques to adjust the weighting schema of the dimensions or features used in distance metric calculations.¹⁰ Specifically, updates to the weighting scheme reflect the features emphasized by the user’s interaction (the weight of relevant features of interest increases, and the weight of other features decreases). The weights are adjusted incrementally on the basis of heuristics associated with each type of interaction. For example, ForceSpire tightly couples several interactions related to text analytics, such as repositioning documents, highlighting text, annotating, and searching, to the underlying dimension reduction model.¹⁰ For instance, highlighting a phrase in a document that contains a set of keywords

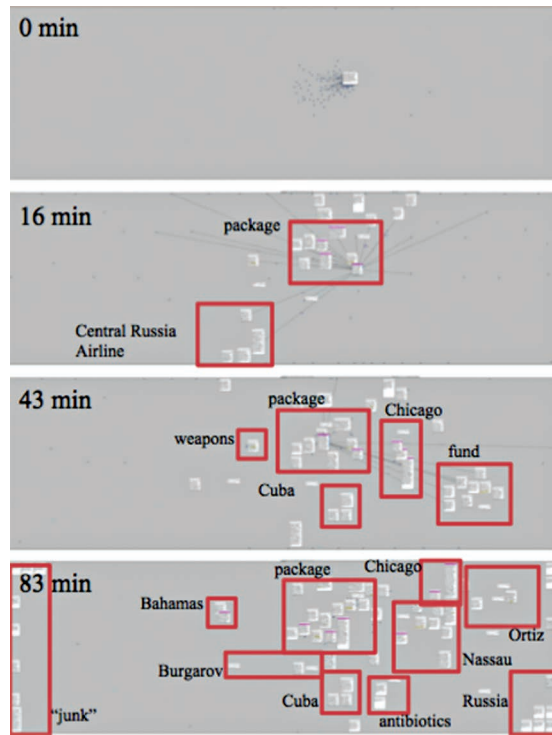


Figure 4. A ForceSpire spatialization’s progression. As the user gains insight, ForceSpire’s model learns to emphasize relevant features.

increases those keywords’ weight in the distance metric.

The third level is direct manipulation for model steering. These interactions leverage machine learning to calculate the amount of change to each feature in the weighting schema. Basically, the VA application receives an updated spatial layout from the user and, given that layout, inverts the model to determine the updated model parameters. This might require an optimization search process to find the best overall fit. Then, the application can apply the updated parameters in the forward application of the model to show how the updated fit changes the layout. For example, observation-level interaction¹¹ and Dis-Function¹² let users move groups of data points in a multidimensional-scaling layout closer together or farther apart to guide machine learning and explore alternative structures in the data.

In summary, all these interactions let users interact directly with the information in context. Over continuous use, the spatialization updates to reflect the incremental insights the user generated (see Figure 4). This creates a symbiotic relationship between the user’s sensemaking process and the system’s machine learning.

Opportunities and Challenges

The following areas provide opportunities and pose challenges for research on direct-manipulation VA.

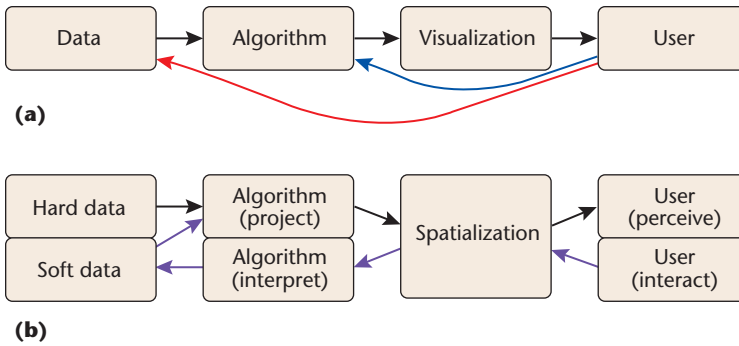


Figure 5. Changing the visualization pipeline to support direct-manipulation visual analytics. (a) In the traditional pipeline, users interact directly with the algorithm (the blue arrow) or data (the red arrow). (b) In the new, bidirectional pipeline, users interact directly with the spatial metaphor; interaction must be interpreted through the model (the purple arrows).

Model Steering

The steering of mathematical models has become a popular way to adapt those models' visual output to the user's domain, task, and workflow. Users can augment the statistical determination of important features and characteristics in a dataset. Because the resulting visualizations include the user's domain expertise, they become more applicable to the domain.

Figure 5 highlights the changes to the visualization pipeline necessary to support such direct manipulation. In the traditional pipeline (see Figure 5a), control panels directly adjust model input parameters. In the new pipeline (see Figure 5b), direct manipulation of the spatialization requires inverting the model to interpret the action's intent, as we mentioned before. The pipeline maps the interaction backward by interpreting the actions and adjusting the parameter data—for example, learning dimension weights. There are many possible approaches to this interpretation step. Additionally, using multiple models would further complicate the pipeline, necessitating a many-to-many mapping of interactions to models.

This area involves two main challenges. First, how do you invert models and map interactions to the parameter-learning process? Second, how do you incorporate multiple models into the visualization pipeline?

Feature Selection

A common stage of spatialization is feature selection. Features can be selected algorithmically from most forms of data, such as extracting keywords from text, extracting visual and audio signatures from images and sound, and so on. The purpose is to represent otherwise unstructured data as high-dimensional. For example, a VA application could use a number of natural-language-processing

models to select keywords or key phrases from unstructured text. These models determine keywords that are statistically more expressive than others, for that dataset. A frequent additional step selects features to optimize the signal-to-noise ratio.

This area involves two challenges. First, how do you incorporate users' domain expertise, which includes features that might not be in the dataset? Second, how do you interactively combine features from different data types (for example, text, audio, and video)?

Feature Extraction

Another common stage of spatialization is feature extraction. A high-dimensional representation must be reduced to a low-dimensional spatialization. This process typically applies a weighting schema to the set of selected dimensions to emphasize each dimension differently when projecting it onto the 2D layout. Because the low-dimensional representations are inherently ambiguous representations of high-dimensional data, interactions in these low-dimensional spaces can also be ambiguous. Multiple inferences might be possible, requiring assumptions or more user input.

The challenge here is, how do you accurately interpret the interaction in the spatialization and apply the high-dimensional representation or weighting scheme to it?

Mixed Metaphors

As we mentioned before, users employ different contexts and metaphors to refer to information in different regions of spatializations.^{2,13} Common metaphors include topical clusters, timelines, geo-spatial layouts, social networks, and process history. Users frequently mix these metaphors in the same workspace as either separate areas or nested schemas. These metaphors might be well defined or ambiguous and might evolve.⁴

This mixed-metaphor use of spatializations poses challenges to layout and clustering models that are generally designed to compute one type of layout across the entire visualization. So, you might need to combine multiple types of models in complex ways. For example, you could combine iCluster, which enables direct manipulation of a cluster membership model,¹⁴ with ForceSpire to enable dynamic layouts of clusters in space, in much the same way analysts currently do manually. The space's continuity and flexibility could represent probabilistic membership.

This area involves two challenges. First, how do you detect, interpret, compute, and visualize mixed models that represent mixed metaphors? Second,

Table 1. Using multiscale models to address big-data challenges for direct-manipulation visual analytics (VA).

	Level of scale		
	Display	Database	Cloud
Usage description	The system lays out the data according to the user's spatial-organization feedback.	The system aggregates clusters of data in the layout according to the user's grouping feedback.	The system uses the layout to query very large data and retrieve additional relevant data.
Data scale of manipulation (no. of data items)	<1 million	<1 billion	<1 trillion
Algorithms	Dimensionality reduction	Clustering Classification Topic modeling	Information retrieval Sampling Streaming
Visualization	Spatial layout Visual proximity = similarity	Clusters Hierarchy Containment Visual aggregate = similarity	Salience Depth Visual salience = similarity
Interaction	Moving items	Grouping items Piling	Deleting items Searching
Interactive feedback for machine learning	Similarities Dimension weights Object weights	Cluster counts and contents Centroid landmarks Labels	Object relevance Keyword dimensions and weights

how do you learn which model best captures the user's interaction, on the basis of the layout?

Multiscale Models

To support big data, VA can leverage multiple models that deal with information at multiple scales (see Table 1). For small amounts of data, you could display all the data points on the screen by using dimensionality reduction (DR) models to organize space. At larger scales, cluster models can aggregate data into fewer groups that could then be fed to DR models. At even larger scales, information retrieval (IR) algorithms become essential for streaming or sampling data to dynamically display only relevant data.

You can apply a consistent direct-manipulation approach across all levels of scale by implementing a system of mutual learning across models. For example, the IR model can query for data relevance based on the dimension weights that the DR model learned. Likewise, the IR model can learn from user actions such as placing uninteresting data in the trash.

This area involves two challenges. First, how do you coordinate direct manipulation to steer models across multiple levels of scales for big data? Second, how do you enhance algorithm performance to support real-time direct manipulation of big data?

Implicit and Explicit User Interaction

With direct-manipulation VA, the system must infer user intentions from user interactions. How-

ever, one action could have multiple possible intentions. For example, dragging a document out of a cluster might indicate that it didn't belong in that cluster, that the user is establishing a new cluster with new nearby documents, or nothing at all. More implicit or explicit user input might be needed to accurately represent the user's actual reasoning process. The amount of approximate or specific input needed might vary.

These options imply the possibility of many parameters for the interaction. Too much explicit input might pull the analyst out of his or her cognitive zone. Analysts should be able to focus on the task, not the tool, using interaction to support their reasoning process.

This area involves two challenges. First, how can the user interface balance explicit and implicit user interaction for model feedback? Second, how can users easily undo or revise direct-manipulation interactions?

Multiparameter Interaction

Novel input modalities might offer more powerful ways for users to express their complex intentions. For example, multitouch interfaces can provide richer interaction for individuals and groups by providing more simultaneous input points with which to express parameters. For instance, in the machine-learning step, a user could move a data point with one hand while specifying target data points with the other hand to indicate which similarity relationships he or she intends. The added

Table 2. The principles of direct manipulation for information visualization are recast for VA.

Direct manipulation for information visualization ¹	Direct-manipulation VA
Continuous visual representations of objects and actions	Spatializations provide a common ground between models and cognition. Users are shielded from the complexity of underlying models and parameters.
Physical actions or button presses instead of complex syntax	Interactions occur in the visual representation. Interactions are tightly coupled between the spatialization and the underlying models.
Rapid, incremental, and reversible actions with immediately visible effects	Models incrementally learn from interactions throughout the analytic process. Visual feedback of the updated model is displayed in the visual metaphor.

bandwidth of these multitouch interactions can more accurately define such manipulations' analytical reasoning.

Large, high-resolution displays can provide more area with which to construct spatial relationships. They give the analyst real, meaningful space as a communication medium and as common ground between the human and model. For example, distances between documents can imply a similarity measure, whereas the absolute location of information can serve as a landmark for themes and concepts. Direct-manipulation VA might be the killer app for these novel hardware technologies.

This area involves two challenges. First, how much user input is needed to convey intention to the models? Second, how can the system provide real-time visual feedback regarding the interpreted actions?


Bias

Model steering potentially introduces user biases into visualizations. Researchers have attempted to address this challenge. For example, capturing interaction data over time¹⁰ can reveal new keywords added to the model. The distribution of weight between these user-derived keywords and those extracted from the data might indicate how much the user's domain expertise influences the spatialization.

Furthermore, the temporal history of keyword weighting can indicate trends in the analysis. Converging trends in the weighting of entities might indicate confirmation bias, whereas diverging weights might represent an analysis involving multiple hypotheses. In particular, it might be possible to quantify specific biases such as confirmation bias¹⁵ and alert users to them in real time. Biases are also opportunities to steer algorithms toward a user's expression of interest, but downsides such as overfitting and missing other interesting insights could occur. Such data could also be used to compare multiple analysts' processes or support collaborative methods.

The challenge here is, how do you illuminate the potential bias associated with introducing the user's domain expertise into the model?

Direct manipulation is familiar to information visualization designers, given graphical controls over direct visual mappings (for example, x- and y- axes on scatterplots, dynamic queries of value thresholds, and so on). However, as visualizations employ increasingly complex mathematical models, interaction designers face the challenge of maintaining the intrinsic principles that make direct manipulation successful, while adapting it to control complex model parameters that might not clearly map to the visual representation. As we showed, for VA, the goal of providing direct manipulation isn't fully realized through control panels for model parameters.

Direct manipulation of the visual representation itself (see Table 2) will enable users to test hypotheses, discover relationships, and input their domain expertise into the calculations used to produce the view. Tools should strive to strike a balance between fully automated and fully manual solutions. In other words, a balance must exist between cognition and computation in VA. By leveraging the information-rich medium of a spatial layout as the primary communication method between the user and system, researchers will be able to realize direct-manipulation VA. We hope that the research opportunities and challenges we presented will help establish a firm science of interaction in VA. 

References

1. B. Shneiderman and C. Plaisant, *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, 4th ed., Pearson, 2005.
2. C. Andrews, A. Endert, and C. North, "Space to Think: Large, High-Resolution Displays for Sensemaking," *Proc. 2010 ACM Conf. Human Factors in Computing Systems (CHI 10)*, ACM, 2010, pp. 55–64.
3. W. Wright et al., "The Sandbox for Analysis: Concepts and Methods," *Proc. 2006 ACM Conf. Human Factors in Computing Systems (CHI 06)*, ACM, 2006, pp. 801–810.
4. F. Shipman and C. Marshall, "Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in

Interactive Systems," *Computer Supported Cooperative Work*, vol. 8, no. 4, 1999, pp. 333–352.

5. F. Tyndiuk et al., "Cognitive Comparison of 3D Interaction in Front of Large vs. Small Displays," *Proc. 2005 ACM Symp. Virtual Reality Software and Technology (VAST 05)*, ACM, 2005, pp. 117–123.
6. J.A. Wise et al., "Visualizing the Non-visual: Spatial Analysis and Interaction with Information for Text Documents," *Proc. 1995 IEEE Symp. Information Visualization (InfoVis 95)*, IEEE CS, 1999, p. 51.
7. J. Alsakran et al., "Streamit: Dynamic Visualization and Interactive Exploration of Text Streams," *Proc. 2011 IEEE Pacific Visualization Symp. (PacificVis 11)*, IEEE, 2011, pp. 131–138.
8. D.H. Jeong et al., "iPCA: An Interactive System for PCA-Based Visual Analytics," *Computer Graphics Forum*, vol. 28, no. 3, 2009, pp. 767–774.
9. J.S. Yi et al., "Dust & Magnet: Multivariate Information Visualization Using a Magnet Metaphor," *Information Visualization*, vol. 4, no. 4, 2005, pp. 239–256.
10. A. Endert, P. Fiaux, and C. North, "Semantic Interaction for Sensemaking: Inferring Analytical Reasoning for Model Steering," *IEEE Trans. Visualization and Computer Graphics*, vol. 18, no. 12, 2012, pp. 2879–2888.
11. A. Endert et al., "Observation-Level Interaction with Statistical Models for Visual Analytics," *Proc. 2011 IEEE Conf. Visual Analytics Science and Technology*

(VAST 11), IEEE, 2011, pp. 121–130.

12. E.T. Brown et al., "Dis-Function: Learning Distance Functions Interactively," *Proc. 2012 IEEE Conf. Visual Analytics Science and Technology (VAST 11)*, IEEE, 2012, pp. 83–92.
13. A.C. Robinson, "Design for Synthesis in Geo-visualization," PhD thesis, Dept. of Geography, Pennsylvania State Univ., 2008.
14. S. Drucker, D. Fisher, and S. Basu, "Helping Users Sort Faster with Adaptive Machine Learning Recommendations," *Human-Computer Interaction—Interact 2011*, LNCS 6948, Springer, 2011, pp. 187–203.
15. R. Heuer, *Psychology of Intelligence Analysis*, Center for the Study of Intelligence, 1999.

Alex Endert is a visualization scientist at the Pacific Northwest National Laboratory. Contact him at alex.endert@pnnl.gov.

Lauren Bradel is a PhD student in computer science at Virginia Tech. Contact her at lbradel1@vt.edu.

Chris North is an associate professor in Virginia Tech's Department of Computer Science. Contact him at north@vt.edu.

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

Showcase Your Multimedia Content on Computing Now!

IEEE Computer Graphics and Applications seeks computer graphics-related multimedia content (videos, animations, simulations, podcasts, and so on) to feature on its Computing Now page, www.computer.org/portal/web/computingnow/cga.

If you're interested, contact us at cga@computer.org. All content will be reviewed for relevance and quality.

IEEE
Computer Graphics
AND APPLICATIONS

