Integrating Human- and Computer-Based Approaches to Feature Extraction and Analysis

Citation

Rogowitz, Bernice E. and Alyssa Goodman. 2012. "Integrating Human- and Computer-Based Approaches to Feature Extraction and Analysis." In Proceedings of Human Vision and Electronic Imaging XVII, ed. Bernice E. Rogowitz, Thrasyvoulos N. Pappas, and Huib de Ridder, January 22, 2012, Burlingame, California, (Proc. SPIE): 82910W. Bellingham, WA: SPIE.

Published Version

10.1117/12.915887

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:34778483>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Integrating Human- and Computer-based Approaches for Feature Extraction and Analysis

Bernice E. Rogowitz^a and Alyssa A. Goodman^b

^aVisual Perspectives Research and Consulting (bernice.e.rogowitz@gmail.com)

^bHarvard-Smithsonian Center for Astrophysics (agoodman@cfa.harvard.edu)

ABSTRACT

A major goal of imaging, graphics, and visualization systems is to help doctors, scientists, engineers, and analysts identify patterns and features in complex data. Some feature-extraction systems use algorithms to automatically identify features. Other methods use human observers to identify features visually, or to discover features by interactively manipulating visual representations. Although automatic feature-extraction algorithms are often directed by human observation, and human pattern recognition is often supported by algorithmic tools, very little work has been done to explore how to capitalize on the interaction between human and machine pattern recognition. This paper introduces a preliminary roadmap for guiding research in this space. One key concept is the explicit consideration of the task, which determines which methods and tools will be most effective. The second is the explicit inclusion of a “human-in-the-loop,” who interacts with the data, the algorithms, and representations, to identify meaningful features. The third is the inclusion of a process for creating a mathematical representation of the features that have been “carved out” by the human analyst, for use in comparison, database query and analysis.

Keywords: Visualization, Interactive Data Analysis, Perception, Cognition, Rule-Based Systems, Visual Representation, Analytics, Data Mining, Human-Computer Interaction, Imaging, Visual Algorithms

1. INTRODUCTION

Human vision and cognition have evolved over two million years, providing us with exquisite pattern recognition, feature extraction, and decision-making capabilities. These capabilities have been integrated implicitly and explicitly into imaging, visualization, and visual analysis systems. For example, human observers are routinely called upon to judge the quality of imaging and display systems, image compression algorithms, and visual representations of data. Increasingly, visualization and imaging papers include user studies, in which the effects of different algorithms or design choices are evaluated by observers in a controlled setting. Human capabilities have also served as models for algorithmic operations, influencing machine-learning, computer vision and artificial intelligence methods for extracting features from complex data. Some examples include edge extraction operators in image processing, modeled after the properties of ganglion cell receptive fields [1], algorithms that automatically highlight ‘regions of interest’ in a scene based on the visual response to low-level features in the visual scene [2], image segmentation algorithms that use properties of human color and texture perception [3], and colormap algorithms that integrate guidance about human magnitude perception [4]. In interactive visualization, the analyst’s pattern recognition and decision-making skills drive the analysis process, supported by tools for selecting, filtering and manipulating the representation [5].

Although there are many examples of systems and algorithms that use human observers or characteristics of human perception and cognition to augment computer-based algorithms, and

many examples of imaging systems that use human observers, these are mostly one-off, or “point” solutions, typically focused on a particular data type, for a particular task, for a particular technology. The goal of this paper is to begin building a more fundamental approach for imaging, visualization, and graphics, which integrates human *and* algorithmic feature recognition capabilities. The key components in this system are: (1) the human observer “in the loop,” (2) mathematical algorithms, (3) algorithms that mimic human capabilities, (4) algorithmic and user interface tools that enhance human pattern recognition and problem solving, and (5) methods for capturing and characterizing the patterns and features extracted by the human observer.

2. A FRAMEWORK FOR INTEGRATING HUMAN AND COMPUTER FEATURE EXTRACTION AND ANALYSIS

Figure 1 introduces the framework we have been developing for understanding how human and computer capabilities interact across a wide range of visualization, imaging and graphics applications. In this section, we introduce the components of the framework. In section 2, we develop the framework and exercise it with examples from imaging, visualization and graphics.

Figure 1. A draft framework for integrating human and algorithmic pattern recognition. In this framework, characteristics of human perception and cognition influence algorithmic operations on the data ($\mathcal{B}$) and operations for mapping data onto visual representations ($\mathcal{D}$). The human observer ($\mathcal{F}$) is integrated into the system, not only as a passive observer, but as an active force in selecting regions of interest, selecting data mappings, and carving out features. The loop is closed when features that are identified visually ($\mathcal{H}$) are characterized mathematically ($\mathcal{I}$), so they can be used for comparison and database search.

In this framework, the human observer is explicitly “in the loop,” making perceptual, cognitive and aesthetic judgments. The human is also included implicitly, through data operations and transformations based on perceptual and cognitive processing (indicated as “perceptually-based” in the diagram). Mathematical methods that do not rely on human capabilities work in concert with

perceptually-based methods to identify features in the data. And, interaction tools allow the analyst to discover and carve out features in the data. An important emerging opportunity is the coupling of algorithmic methods with human interaction methods to create mathematical representations of features that are identified visually.

The components of the framework

The data and the task ($\mathcal{A}$) are chosen by the human analyst. Algorithms ($\mathcal{B}$) operate on the data, depending upon the task. These algorithms may be purely mathematical, or may implement characteristics of human perception and cognition. The output from this stage is a mathematical representation of the data, including “features” that capture relationships or key characteristics in the data ($\mathcal{C}$). Aspects of the data can be mapped onto visual dimensions, in order to create a *visual* representation, such as an image or a graph. These mapping operations ($\mathcal{D}$) may also be purely mathematical, or may be influenced by characteristics of human perception and cognition. The resulting visual representation ($\mathcal{E}$) is presented to the human observer ($\mathcal{F}$), who makes judgments about the representation. The human may simply observe the output, or may use interaction tools ($\mathcal{G}$) to manipulate the representation in order to discover features. Features that are visually identified by the human observer ($\mathcal{H}$) can be used directly, or, can be characterized mathematically ($\mathcal{I}$), so they can be used for comparison or database query.

2.1 The Data and the Task ($\mathcal{A}$)

In imaging, visualization and graphics, data may be 1-D, 2-D or 3-D arrays, or multidimensional data tables; they may be time-varying, have a graph or network structure, or be symbolic, volumetric, text, field or geographic. The goal of this work is to create a framework that is suited to all types of data.

How these data are operated upon depends not only on the data type, but also on the task, or the goal of the analysis process. This might be to compute the average pixel value in an image, or may involve characterizing the magnitude of a variable, identifying clusters, finding correlations, or identifying interaction effects. For an image, the task may be to identify visual coding artifacts or extract regions of interest; for interactive data analysis, the goal may be to find outliers, uncover structures, compare functional groupings, discover relationships, or track changes over time. The data and the task are represented as overlapping in Figure 1 because the data constrain which tasks can be addressed, and the task constrains the choice of data.

2.2 Algorithmic Operations for Finding and Characterizing Features ($\mathcal{B}$ and $\mathcal{C}$)

Both the data and the task are critical inputs to the feature extraction process. In many applications, a mathematical algorithm is used to analyze the data, such as K-means analysis to extract clusters, segmentation algorithms, neural networks or principal components analysis.

Algorithms are also used to build high-level features, based on many lower-level operators, which can be used to compare data sets or to create a query which can be used to find other instances in a database. For example, a doctor may want to evaluate the growth of a tumor over time, or a scientist may want to search a database for a particular type of vortex in a turbulent, time-varying field flow. Important contributions to this problem come from computer vision, machine vision, image processing and topology, and may often involve a recursive learning process.

Perceptually-based algorithms

Perceptually-based algorithms enhance mathematical methods by implementing algorithms that mimic characteristics of human pattern recognition and feature-extraction. For example, the center-surround structure of retinal ganglion cells, which have a central excitatory region surrounded by an inhibitory surrounding region, has been used in image processing to model edge-enhancement operators [1]. Convolving an image with a difference-of-Gaussian operator, a convenient model for retinal ganglion cells, produces an image with enhanced edges and reduced detail elsewhere, which both increases the information about important contours in the image and reduces the bandwidth required to transmit it. Perhaps the most important visual contribution to image processing algorithms is the human contrast sensitivity function, whose inverted U-shape depicts our lower sensitivity to very low and very high spatial frequencies [6]. This simple function has been used in image coding, digital halftoning, and image quality metrics, providing a guide for gaining efficiency by devoting the greatest bandwidth to regions of greatest spatial-frequency sensitivity, by hiding sampling noise in regions of low contrast sensitivity, and by providing a weighting for spatial information in image quality metrics. Taking this approach a step further, this inverted-U function can be modeled as the envelope of spatial-frequency-selective channels, with differential sensitivity to different, overlapping ranges of spatial information. This model has had practical applications in a wide range of imaging applications, including the JPEG standard for image coding. Psychophysical measurements of spatial-frequency and orientation-selectivity in human vision have played a major role in modern image quality models, such as the Cortex Transform [7], and have been used to evaluate the quality of compressed images, to tune compression algorithms, and to evaluate and hide the effects of noise.

2.3 Mapping Data onto Visual Representations ($\mathcal{D}$)

Data, algorithms, and the results of applying algorithms to data can all be represented visually, so that they can be viewed, analyzed, manipulated and explored by a human observer. This could be a simple table of numbers, an equation, an image, a graphic, or a sophisticated visualization. Visual representations are created by mapping characteristics of the data onto visual parameters. For example, the data value in each location of a digital X-ray can be rendered as a luminance value, magnitude can be represented by mapping values of a colorscale onto a surface or streamline, or multiple characteristics of a complex phenomenon can be mapped onto multiple parameters of a glyph. Computer graphics algorithms can be used to map data structures onto particular geometries (rendering) and graph layout algorithms are commonly used to map a network of interactions onto a graph. For each representation, decisions are made, either explicitly or by default, regarding how the values of the variables are represented, whether by position, relationship, color, length, depth, brightness, movement, *etc.*

Perceptually-based algorithms for mapping data onto visual representations

Perceptually-based mapping algorithms capture some aspect of human perception or cognition. One especially important contribution has been in the design of color maps. Data values are mapped onto values of a color scale, and care is taken for equal steps in the data to correspond to equal steps along the color scale. Less attention has been paid to ensuring that equal steps along the color scale correspond to equal perceptual steps. Based on psychophysical guidance from S.S. Stevens [8], Rogowitz, *et. al.* [4] hypothesized that colormaps with a monotonic luminance component would be most effective in representing magnitude information in a visualization setting. They used color measurement and calibration algorithms to create color maps that had different luminance, hue and saturation profiles, constructed in different standard color spaces, and found that colormaps with a monotonic luminance component produced the most monotonic increment threshold curves. Color maps with a monotonic luminance profile also produced the highest judgments of image quality, when mapped onto images of a face [9]. The “rainbow”

colormap, which is still the default in many visualization systems, produced non-uniform increment threshold curves, where equal steps did not appear to have equal steps. This work formed the basis for one of the perceptual rules in the PRAVDA system [10], which offered users colormap selections based on the type of data and the analysis task. Recently, Borkin, *et. al.* [11] have conducted a study with domain experts making judgments about possible blockages in heart arteries. They found that when the doctors used representations constructed with the rainbow color map, they missed a far greater percentage of pathological instances, but were very accurate with a colormap with a monotonic luminance component. Thus, the perceptual literature influenced the design of these perceptual colormaps, experiments were conducted to measure their effectiveness in a controlled study, the results were used to guide the selection of colormaps in a visualization system, and studies with domain experts demonstrated the real-world impact of making perceptually-inappropriate choices.

The Role of the Task in Mapping Data onto Visual Representations

Which representation is best depends on the data type and task that the visualization or image is intended to serve [12]. This point is illustrated in Figure 2, which shows four different visual representations of the same three-dimensional table of numbers. This is a Monte Carlo simulation for a risk evaluation model; for each scenario, the spot price for Mexican Peso/US Dollar (column 1) and the Volatility for that trade (column 2) are input to the model, and a Profit and Loss prediction (column 3) is computed. In *a*, the three columns are plotted as a 3-D scatterplot, with spot price and volatility as X and Y, and the value of the trade on Z. *b* shows a heat map, with the value of the trade for each spot price (X) and volatility (Y) depicted as a color along a rainbow colormap. In *c*, the value of the trade is shown both in color and in height, and in *d*, a surface has been computed to show the value of the trade for every X and Y, with shadows cast.

Figure 2. Four different representations of the same data. The “best” representation for extracting meaningful features depends on the task. These representations were designed to help the analyst (A) identify outliers, (B) find hot spots, (C) explore the shape of the profit surface, and (D) express the shape of the data analytically.

The “most effective” visualization depends on the task. If the goal is to identify outliers, then *A* is the most effective. If the goal is to provide a quick qualitative picture of the places where the trade is most lucrative, the heat map (*B*) is the best choice. If it is important to understand the risk profile, and where the values are stable and where they change abruptly, then *C* is the best choice. And, if the goal is to capture this surface analytically, or to see how the individual values are distributed along that straddle, or whether the straddle shape is in price or volatility (see the shadows) then *d* is the best choice.

2.4 The Human Observer (*F*)

Whether it be a squiggle in the margin of a topology book or a dynamic simulation playing out on a wall of displays, the visual representation is designed for human consumption. The image at the far right of Figure 1 represents the “human in the loop,” who brings a full complement of capabilities to the table, including spatial, temporal and color vision, shape and feature perception, attention, pattern recognition, decision making, problem solving, and visual reasoning.

Depending on the application, the human may simply be an observer judging the output of an imaging or visualization system. This might be an engineer providing a subjective evaluation of an image compression algorithm, a stock analyst looking at a time series of mutual fund on a smart phone, or a doctor examining an x-ray for a fracture. User observations can also be measured more formally in user studies, where a controlled set of visual stimuli are evaluated using a controlled set of measurement conditions. User studies can, for example, help evaluate the visual impact of parameters in a particular imaging system or visualization technology. Experimental studies can also be designed to generate results that generalize beyond the particular system under investigation. In particular, psychophysical, perceptual and cognitive studies measure visual responses to controlled visual representations in order to test hypotheses about visual processing.

2.5 Interactive Exploration (*G*)

Increasingly, imaging and visualization systems provide capabilities for manipulating the visual representation. Financial applications routinely allow the user to select different time periods, compare different time series, and access different drill-downs. Computer aided design systems provide 3-dimensional image viewers that allow users to rotate graphical objects and select different illuminations. Interactive visualization applications provide interaction tools that help the analyst to search for features. Using these tools, the analyst can dynamically adjust parameters, explore relationships, transform variables and create new probes and descriptors on the fly. This is an iterative design process, where each manipulation of the representation can stimulate new hypotheses about the structure in the data, which can be followed up with additional operations. The data do not usually contain pre-computed features; it is the analyst’s job to explore, define, and characterize the underlying patterns and features.

Interactive Exploration to Discover Semantic Features in Data

Figure 3 shows an interactive visualization technique that takes advantage of the finding that color acts as a natural semantic marker for human observers. In this technique, first introduced by Tukey [5], the analyst uses color to highlight or “brush” important regions in one visualization and the system automatically paints that color onto corresponding regions in other linked visualizations. Figure 3 demonstrates this technique, showing two visualizations of data from a finite element model of the heart [13]. In the left-hand figure, the analyst noticed that the calcium in the subspace variable (CaSS) was bimodal, and used two different colors to distinguish between the modes. The right-hand figure shows the spatial location in the 3-D model of the semantically-tagged elements.

The analyst was surprised to discover that the red and yellow voxels defined a spatially coherent region, and that within this region, the voxels from the two ranges of CaSS were interdigitated.

From a cognitive perspective, the color brushed onto data points gives the analyst a way of marking regions that are semantically important, organizing those points into a category. The analyst can mark a region in one view, examine the implications in other views, and then interactively change that region, mark additional semantic regions (in different colors) in any of the linked representations, and iteratively identify meaningful multivariate patterns.

Figure 3. Color is a natural semantic marker for human observers and can be used to highlight corresponding regions across multiple visualizations. In this example, the analyst controlled the painting of colors interactively, trying out different hypotheses visually. Here, the analyst used two colors to mark the two modes in a histogram (left panel) and discovered that the cells that had low and high values of CaSS occupied the same region of the heart and were interleaved (right panel).

Interactive Analysis Supported by Algorithmic Approaches

The human analyst and mathematical algorithms can interact to identify features in complex data. In the top-left panel of Figure 4 (*a*), the analyst, in [13], has used color to semantically delineate two populations in the scatterplot, one with high spatial current (red) and another with high voltage (green). Different geometrical mapping algorithms (*b*, *c*, and *d*) were developed in order to allow the analyst to explore the behavior of the two groups. The point-to-point mapping (*b*) allowed the analyst to examine the 3-D structure of the red population of elements. The interpolation of the individual voxels onto a surface (*c*) allowed the analyst to see that the red region encircles a high voltage region (green). The interpolation onto tensors derived from diffusion imaging (*d*) gave the analyst direct insight into how these parameters played out onto the structure of the heart fibers.

In this example, the analyst used color to mark semantic regions, and was able to view the projection of those regions in other linked representations. Although in principle, the analyst can apply color in any of the representations, applying and removing color in a 3-D object is currently quite awkward. One idea for carving out structures from 3-D models is to use haptic interfaces, building on our exquisitely fine visually-guided motor skills [14, 15]. The future, then, could include tools that allowed the analyst to sculpt the data representations in order to reveal their hidden structures.

Figure 4. Brushing across multiple geometric representations. The human analyst has used color to select regions of interest in a scatterplot (a). Corresponding regions are automatically highlighted in three different 3-D representations of the heart. These different algorithms allowed the analyst to explore how the different regions in parameter space (a) played out in the volume (b), on the surface (c), and along the fibers (d) of the heart.

2.6 Capturing Visually-Extracted Features ($\mathcal{H}$, $\mathcal{I}$)

Once the analyst, alone, or aided by algorithms, has identified a feature of interest, it could be captured as an image for publication and communication. However, it would be extremely valuable if that feature or structure could be expressed mathematically. The feedback loop at the bottom of Figure 1 expresses this goal, which is to complete the data representation and analysis cycle with an analytical or mathematical representation of the features discovered visually. In some cases, solving this problem is easy. For example, in Figure 4, the analyst can express the complex 3-D geometry in panel B as a simple Boolean query in two variables—the structure is simply all those voxels which have a spatial current above a certain value and voltage below a certain value. However, since the two other representations (C and D) involve interpolation in geometric space, no simple query captures the mathematical structure of these features. And, if the analyst had modified these 3-D shapes, based on domain knowledge, or had been free to carve out features by transforming variables or by selecting and de-selecting regions in multiple linked representations, it may be necessary to model the emergent 3-D structure explicitly, or to develop a model based on parametric as well as geometric variables. If the visually-discovered feature could be expressed mathematically, it could be used to measure changes over time, or could be used as a query to find related features in a database.

In addition, a set of features identified by human observers could be used as input to a machine learning algorithm, which could create a more general classifier. Systems like this exist today, but the human is used as a dynamic classifier, telling the algorithm whether, for example, a tumor does

or does not exist in a medical image. If the analyst could communicate to the machine-learning algorithm by carving out multiple instances of the tumor, using data from multiple sources, this could potentially lead to much more nuanced and sophisticated descriptors. Moreover, these archetypes could be used to train future generations of pathologists, and to compare mathematically-described features to real data in a diagnostic setting.

In addition, it would be interesting to ask what mathematical representations are best for characterizing different types of features, what representations would make the semantic meaning of these features more salient to the human analyst, and how to train analysts to interact effectively with visual and algorithmic tools.

3. RESEARCH DIRECTIONS

In this paper, we have looked at how human- and computer-based approaches participate in the process of identifying and characterizing features in data, and have developed an initial framework for representing this complex interplay. The work in this paper is based on consideration of a wide range of data types and imaging applications, across many domains. Wherever we looked, we found fertile ground for explicitly coupling the natural capabilities of human perception with the analytical tools of modern computer –based systems.

Our next step is to exercise this framework with scenarios from different domains to see whether the components and the flow accommodate all the ways in which algorithms and humans interact to extract meaning and structure from imaging, visualization, graphics, and visual analytics data.

Our longer-term goal is to help drive a cross-disciplinary dialogue on the interaction of human and algorithmic capabilities, much in the spirit of the white paper produced by Thomas and Cook [16]. We hope this framework will serve as a springboard for developing a research agenda in this area. Some initial ideas include:

Using human pattern recognition and algorithmic tools synergistically. What computer-vision, image processing, and computer vision enhancements could be made to imaging systems to augment human pattern recognition? How can algorithmic tools be used more seamlessly in the exploration environment? How does this depend on the analysis task?

Human perception and cognition. How do humans perceive features? What is the relationship between the underlying structure in the data, the parameters of the visual representation, and what is perceived? What visual methods could be encapsulated in algorithms?

Human interaction with data and models. How can we improve how observers explore, interact and navigate through visual representations of data and models? There are enormous individual differences in the ability to see patterns and make inferences from visual data, and it has been argued that many of these differences have to do with innate analytical abilities. What advances are needed to train analysts to see patterns in data and to think visually?

Carving out features in high dimensional data. What methods will help observers more effectively extract meaningful patterns? What new techniques will be required to discover features across multiple data sets or representations? What novel visualization and user interface (*e.g.*, haptic) methods will be required?

Capturing the visual mathematically. Once an analyst has carved out a feature, how can it be expressed mathematically so that features can be compared, or similar phenomena can be searched out from a large database?

4. CONCLUSION

Human observers are exquisitely adept at seeing and interpreting visual patterns, which explains the important role human pattern recognition has played in imaging, visualization and visual analysis. Computer vision, machine learning and statistical algorithms have also played a critical role in identifying meaningful features in visual data. Although isolated point solutions have been proposed that combine human and computer capabilities, there is no general, cross-disciplinary science uniting these approaches. Our goal is to create a framework that explicitly captures the interplay between human- and machine-based approaches, the influence of human perception and cognition on computer-based algorithms, the ways algorithmic tools can support human judgment.

In this paper, we have crafted a preliminary framework (Figure 1) that integrates the human problem solver, algorithms based on human perception and cognition, and purely mathematical algorithms. In this framework, the data type and the task play a pivotal role, since they drive what types of algorithmic operations and visual representations are appropriate. It also includes a human-in-the-loop component, where the human is an active participant in the interactive exploration of the data. In interactive visual analysis, the human observer manipulates and interacts with representations of the data, using algorithms to create new variables or provide variable transformations to mark semantic regions of interest. The framework also includes a loop for using mathematical tools to capture and characterize visually-identified features. Once a feature or pattern has been identified, computer vision, neural networks, topological operators, or machine learning algorithms can be used to provide a mathematical representation of the feature, which can be used as a query to a larger database.

The goal of this framework is to provide a mechanism for formalizing the role of the human in imaging, and to stimulate research into human pattern recognition, perceptually-based algorithms, and interactive methods that integrate human and algorithmic approaches to extracting and characterizing semantic features in data.

5. REFERENCES

- [1] Young, R.A. "The Gaussian derivative model for spatial vision: I. Retinal mechanisms, spatial vision, Volume 2, Number 4 (1987).
- [2] Itti, L., and Kock, C. "Computational modeling of visual attention," Nature Neuroscience, Volume 2 (2001).
- [3] Chen, J., Pappas, T.N., Mojsilovic, A, and Rogowitz, B., "Adaptive Perceptual Color-Texture Image Segmentation," IEEE Transactions on Image Processing, Volume 14, Issue 10 (2005).
- [4] Rogowitz, B.E., Kalvin, A.D., Pelah, A., Cohen, A. "Which trajectories through which perceptually uniform color spaces produce appropriate color scales for interval data?" Invited paper, Proceedings of the IS&T Color Imaging Conference 7, Scottsdale, AZ (1990).

- [5] Fishkeller, M.A., Friedman, J.H., and Tukey, J.W., "Prin-9, an interactive multidimensional data display and analysis system," in *Dynamic Graphics for Statistics*, Pacific Grove, CA: Wadsworth and Brooks/Cole, 91-109 (1975).
- [6] Barten, P.G.J., *Contrast Sensitivity of the Human Eye and its Effects on Image Quality*. SPIE Optical Engineering Press (1999).
- [7] Watson, A.B. "The cortex transform: Rapid computation of simulated neural images," *Computer Vision, Graphics and Image Processing*, Volume 39, Issue 3 (1987).
- [8] Stevens, S.S. "On the psychophysical law," *Psychological Review* (1957).
- [9] Rogowitz, B.E. and Kalvin, A. "The Which Blair project; A quick visual method for evaluation perceptual color maps," *Proceedings of IEEE Visualization* (2001).
- [10] Bergman, L.D. Rogowitz, B.E. and Treinish, L.A. "A rule-based tool for assisting colormap selection," *Proceedings of the IEEE Conference on Visualization*. 118-125 (1995).
- [11] Borkin, M., Gajos, K., Peters, A., Mitsouras, D., Melchionna, S., Rybicki, F., Feldman, C., and Phister, H., "Evaluation of artery visualizations for heart disease," *IEEE Transactions on Visualization and Computer Graphics*, Volume 17, Issue 12, 2011.
- [12] Wilkenson, L., *The Grammar of Graphics, Second Edition*, New York: Springer-Verlag (2005).
- [13] Gresh, D., Rogowitz, B.E., Winslow, R.L., Scollan, D.F., Yung, C.K., "WEAVE: A System for Visually Linking 3D and Statistical Visualizations, Applied to Cardiac Simulation and Measurement Data," *Proceedings of the IEEE Conference on Visualization*, 489-492 (2000).
- [14] Keefe, D.F. "Free-form VR interactions in scientific visualization, SIGGRAPH Poster (2005).
- [15] Rogowitz, B.E., and Borrel, P, "Virtual Hand: a 3D tactile interface to virtual environments," *Proceedings of the SPIE Conference on Human Vision and Electronic Imaging* (2008).
- [16] Thomas, J.J. and Cook, K, "Illuminating the Path: The Research and development agenda for visual analytics," *IEEE Computer Society* (2005).