The Elements of Visualization: A Framework for the Perceptual Optimization of Multi-valued Multi-layered 2D Scientific Visualization Methods

by

Daniel Acevedo Feliz

B. S. Civil Engineering, University of A Coruña, Spain, 1997

M. Sc. Computer Science, Brown University, Providence RI, 2001

A dissertation submitted in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy
in the Department of Computer Science at Brown University

Providence, Rhode Island

October 2006

Abstract

This dissertation describes a framework for the modeling of the effectiveness of visualization methods based upon the quantification of perceptual interactions among visual dimensions and the subjective critiquing of visualization displays by visual design experts. Our aim is to facilitate the visualization process by evaluating the perceptual capabilities of visualization methods from both a psychophysical and visual design perspective. Combining lessons from both disciplines we begin to bridge the gap between the specific nature of psychophysical experiments and the practical needs of scientists exploring real datasets.

Our contributions include new experimental and computational techniques to evaluate how different visualization methods perform when displaying multi-valued scientific datasets in 2D. The work spans several experiments aimed at quantifying the expressiveness of some of the *visual dimensions* that can be used to generate icon-based visualization methods. In particular, we concentrate on five of those dimensions: size, spacing, orientation, brightness and color saturation. Our experiments measure the expressive capabilities of these dimensions based on a set of design factors that we use to characterize our visualization methods. We introduce a formal definition and nomenclature for the space of visualization methods these dimensions form, and for the feature space that the design factor measurements create from our experiments. With these measurements in place, we develop a computational optimization strategy that, given a set of constraints on the design factors for a visualization problem, is able to search that space for effective solutions.

In our experiments, we first explore how individual visual dimensions perform when used to represent scientific data. Then, we study interactions among pairs of visual dimensions by analyzing how each dimension is affected by the others. Given that it is not practical to fully explore all possible combinations among the five dimensions we chose, we must leverage the information gained on these initial experiments. We hypothesize how higher order combinations of visual dimensions would perform by extrapolating the changes in individual performance from single visual dimension methods to methods utilizing pairs of visual dimensions. Finally we evaluate these hypotheses by sampling some of those combinations and having expert visual designers critique the resulting displays.

To evaluate how well our data are able to predict the effectiveness of higher order combinations of visual dimensions, we perform an experiment with real datasets from our driving application domain: remote sensing using synthetic aperture radar polarimetry. This field requires visualizations that allow scientists to explore correlations among multiple variables, creating a perfect test case to evaluate the success of our optimization strategy.

This dissertation by Daniel Acevedo Feliz is accepted in its present form by the Department of Computer Science as satisfying the dissertation requirement for the degree of Doctor of Philosophy.

Date __________ ____________________

David H. Laidlaw, Director

Recommended to the Graduate Council

Date __________ ____________________

John F. Hughes, Reader

Date __________ ____________________

Leslie Welch (Department of Psychology), Reader

Approved by the Graduate Council

Date __________ ____________________

Sheila Bonde
Dean of the Graduate School

Contents

List of Tables

List of Figures

Chapter 1

Introduction and Contributions

"Doctrines and theories are best for weaker moments. In moments of strength, problems are solved intuitively, as if of themselves"
– Johannes Itten in *The Elements of Color*

Defining and exploring the space of possible visualization methods for a given scientific problem has challenged computer scientists, statisticians, geographers, and cognitive scientists for many years; it is still an open challenge. The goal of such models is to describe a searchable space where scientists can find visualization methods that optimally convey the information they require. Our approach to achieving this is to optimize the design of the visualization methods by studying how their visual components affect each other to facilitate, or complicate, data perception and comprehension. The value of identifying the basic dimensions that form visualizations, and their interactions, is that we thus develop a framework to organize knowledge visualization design and predict behavior of data displays [Cleveland and McGill, 1984]. We want to create a way to get scientists, and visualization users in general, closer to an effective visual representation of their data.

The extensive capabilities of current experimental techniques, including computational simulations or sensors sampling some natural phenomena, generate huge amounts of multi-valued multi-dimensional data that need to be represented visually to be explored. The process of effectively representing these scientific data involves several disciplines that must be well understood to create useful displays that allow further scientific analysis: data mining, statistical analysis, visual design, perceptual psychology, computer interface design, and human-computer interaction are some of those disciplines. Rarely does a single person have enough expertise in all of these fields to tackle a visualization problem alone, requiring collaborative efforts among a group of experts.

In our experience, users of visualization software feel overwhelmed by the multiple visualization options available to them. Experience, or a visual designer by their side, are usually the key to a successful visual representation of the data. Our model will not eliminate those, but will get a user with little experience and without access to visual design expertise closer to an effective solution. This will greatly improve the efficiency of the visualization process and let the user concentrate on exploring and analyzing the data at hand.

The main goal of this dissertation is to find a way of choosing a visualization method that will be appropriate to effectively visualize a multi-valued dataset given a set of design goals.

1.1 Scope of this Thesis

We have organized this thesis around the concept of four basic **elements of visualization**:

- The **visualization problem** presented by the end-user,
- The **visual dimensions** used to create the solution,
- The **design factors** that characterize the capabilities of those dimensions, and
- The **effectiveness evaluation** of the solution.

As Watson outlines in [Watson, 2006], to automate the design and creation of visualizations researchers must identify the particular problem and its constraints (our first element), find and capture the heuristics that describe a good solution (our third element) and build a tool that finds one or more of those good solutions in the problem space. To Watson's we add the actual definition of the pieces used to build the visualizations, our visual dimensions, and the measures that characterize the representational capabilities of those dimensions, our design factors.

Once these elements are in place, the process of exploration of the data can really begin by allowing the user to interact with the visualization. By making these four elements explicit, the user is capable of modifying the data mapping parameters and knowingly explore other visualization possibilities according to new questions and hypotheses the analysis of the data might rise.

Visualization Problem

The basic scientific visualization process involves symbolization, the translation of verbal and numerical information into graphic form [McCleary Jr., 1983], and comprehension, the analysis and understanding of the data presented. Our research is oriented towards developing exploratory data visualization methods, with the goal of visually presenting raw data in a way that prompts visual thinking and knowledge construction [MacEachren and Kraak, 1997]. That is our *visualization problem*. Understanding and insight are the main goals of scientific data visualization methods, but methods to represent known phenomena (e.g. turbulence in air flow or stress points in a structure) or geared towards performing specific tasks (e.g. finding extrema or identifying a type of turbulent flow) are qualitatively different from visualization methods designed for exploration of the data. Scientists usually utilize the latter during the early stages of their research, when they require visuals that provide a broad understanding of the data being presented. They begin posing hypotheses and asking questions about the data, which lead them towards task-oriented visualization methods for further analysis. Exploratory visualization methods allow them also, in a first approximation, to qualitatively assess the validity of their experimental and data gathering methods. At this stage, visualization is merely a tool to help scientists think about their problem [Hibbard, 2004].

To that end, we have developed a methodology to create, explore, and evaluate a model for a space of visualization methods where we can search for optimal solutions to a given visualization problem. We focus on visualization methods for multi-valued scalar scientific datasets in 2D. These datasets are widely used in disciplines such as meteorology, geology, cartography, physics, and engineering. Even when scientists are studying three (or higher) dimensional phenomena, they often rely on 2D slices, such as cutting planes or isosurfaces, to explore and study the datasets..

Visual Dimensions

Common practice in scientific visualization is the mapping of scalar quantities to the visual qualities of surfaces containing the data, with color being the predominant example. Other visual qualities that can be used to represent a scalar field on a 2D surface belong not to the surface itself, but to glyphs or icons that can be placed on the surface. Color is again the initial choice for most applications, but size, distribution, and orientation of these icons can also be used to visually represent a scalar field.

The class of visualization methods we are concentrating on includes multi-layered Poisson-disk distributed icons where icon size, spacing, orientation, brightness, and color saturation can be set to a constant or coupled to data values from a scalar field in 2D. Even though we are using this methodology to study a very limited space, this framework could be extended to include more complex visual dimensions and even three-dimensional visualization methods

We call the visual characteristics of the icons our *visual dimensions*, since they are the basic components of our visualization methods. Icons also have the advantage that they can be layered, increasing the number of variables being simultaneously shown.

Design Factors

This dissertation focuses on the creation and evaluation of visualization methods according to a set of *design factors*. These factors relate to the relative importance of the different variables in the dataset, their relationships, and the quality of visualization needed for each one of them.

During our experiments we will measure and model the performance of different methods with respect to our set of design factors. These are important since our visualization problem is the exploration of the data, with no predetermined task in mind. As it will be explained in Section 2.3, our factors need to, therefore, serve as a characterization of how our visual dimensions participate in the process of exploration. There are many possible ways to do this characterization, but we decided on a manageable set of factors so we could perform our experiments in a reasonable amount of time. Yet the results must provide some indication of the expressive power of our visual dimensions and the visualization methods they form.

Effectiveness Evaluation

The *effectiveness evaluation* indicates how well a given method fulfills a set of requirements given in terms of our design factors. These requirements, or design goals, will be used as constraints on an optimization process to find solutions that most effectively fit those goals.

In summary, the effectiveness of a method is measured by its adaptation to the requirements given while, in our experiments, we are measuring and recording the set of values those requirements can take.

Finally, the computational framework we have created to facilitate this process allows for interaction and manipulation of the visualization methods. We do not consider this a separate element of visualization since it is a combination of the two: an iterative process involving the choice of visual dimensions and the evaluation of that choice. Although it is a key part of the data exploration process, the many interface and human-computer interaction aspects of it are beyond the scope of this project.

1.2 Experimental Methodology

Evaluating the effectiveness of visualization methods is difficult because tests to evaluate them meaningfully are hard to design and execute [Kosara et al., 2003].

Our research involves experiments where subjects perform subjective perceptual tasks from which we obtain numerical measures of interactions among visual dimensions. These studies are inspired by psychophysical experiments but geared towards our goal of developing visualization methods for effective data exploration. We also perform subjective studies where expert visual design educators critique visualization methods that use those same dimensions. Our model brings together both experimental approaches by using lessons from the latter to inform the design of the former set of studies.

We must note that we do not aim to find a single optimal solution that will exactly match the visualization problem's description. Perceptual psychophysicists and cognitive scientists have been studying human perceptual capabilities for decades, and there are still many unresolved problems. Even those problems that have been explored find often conflicting experimental results that make it difficult to elaborate a complete and solid theory on how humans perceive visual dimensions. It would not be realistic for this dissertation to try to exactly quantify all possible interactions among visual dimensions. Even limiting ourselves to only five dimensions, there are many elements that affect the reading of visualization displays, such as interaction techniques or display form factors, that we cannot possibly begin to explore if we hope to succeed in our initial goal.

The use of visual design and even artistic expertise to develop visualization methods is widely acknowledged in our discipline. The novelty of this dissertation comes from our goal of quantifying that expert knowledge in a way that we can combine it with perceptual experiments to build our model. We have created a framework for evaluating visualization methods through feedback from expert visual designers and art educators. Our framework mimics the art education process, in which art educators impart artistic and visual design knowledge to their students through critiques of the students' work.

Our characterization of 2D visualization methods acknowledges that the input we get from the designers is directly targeted at the needs of scientists, and is not about artistic qualities, visual appeal, or aesthetics. Our subjects, illustration educators, are experts at evaluating visuals for targeted communication goals; while their results are often appealing and aesthetic, they first have to satisfy those communication goals which, in this case, means presenting scientific data for effective exploration. For our experiments we are utilizing subjects with extensive experience in teaching and critiquing art and design works. They are used to concentrating on the problem at hand, abstracting from aesthetic considerations when they have to focus on what the final goal of the work is.

One of the main advantages of introducing this type of subjective experiment is the fact that expert designers can, not only characterize methods according to our design factors, but they can also tell us *why* a method does or does not work, and in most cases how to fix it by moving along the visual dimensions used in it.

1.3 Contributions and Proposed Work

The expected contributions of this dissertation are:

- The quantitative evaluation of perceptual interactions among visual dimensions for scientific visualization methods.
- The development of a methodology to evaluate visualization design factors from both a perceptual and a visual design perspective.
- The combination of subjective expert visual designer evaluations into the quantitative modeling of a space of visualization methods.
- The development of an effectiveness measure for visualization methods based upon constraints on a set of design goals.
- The development of an optimization strategy based on our experimental data that is capable of incorporating evaluations of complex visualization into our model.
- The application of this new effectiveness modeling and optimization strategy on a scientific application domain.

We have advanced towards accomplishing these contributions by already completing the following elements of the research, which are summarized here and will be further explained in Chapter 4:

- The development of an interactive software environment for creating the visualization displays needed for this research. It provides users with a text-based interface to create and manipulate multi-valued multi-layered visualization methods. We developed it in collaboration with Fritz Drury, professor at the Illustration Department at the Rhode island School of Design (RISD), who advised us on what visual dimensions to implement first and how to organize their interplay in the software. The basic software design is an extensible framework for visualization methods in 3D, and includes other non-icon-based visual elements such as color planes and streamlines. It also includes support for vector and tensor-based datasets.

- The design and implementation of a study comparing 2D vector visualization methods using visual designers as subjects. This study served as a proof-of-concept to evaluate the hypothesis that experiments of this kind, based on subjective designer critiques, can serve as a more efficient means of evaluating visualization methods. Results indicated that designers rated the visualization methods in a pattern similar to the results of the scientists from a previous quantitative study [Laidlaw et al., 2005]. We also found that designer critiques generally took less time and that designers were able to provide methods for improving the visualizations. This result provides key support for using subjective expert design knowledge as the basis for our visualization effectiveness characterization. *This was published as a Sketch in ACM SIGGRAPH'03 [Jackson et al., 2003] and will be submitted, in a more extended form, to IEEE Transactions on Visualization and Computer Graphics Journal.*

- The design and implementation of a study comparing 2D scalar visualization methods using expert visual design educators as subjects. Based on the experience of the previous study, and after the development of our visualization software, we conducted an initial study to evaluate the effectiveness of 2D visualization methods in terms of a set of design factors, which were subjectively rated by expert visual design educators. We successfully characterized a total of 33 visualization methods using 11 different visual dimensions and 6 different design factors for representing single-variable continuous scalar datasets. This study raised the question of using expert designers, specifically educators, versus non-expert designers as in the previous experiment. The level of understanding of the tasks to be performed and the profusion of comments about why and how to improve some methods increased dramatically in this second experiment. Although not empirically evaluated yet, educators seem to be better subjects for evaluating visualization methods than non-experts. We have not yet utilized expert non-educators to complete our sample. *This was published as a Poster in IEEE Visualization'05 [Acevedo et al., 2005] and received the Best Poster*

Award at the conference.

- The formalization of an appropriate mathematical notation for our model of visualization effectiveness. This is a key step in organizing our future experiments and understanding where our research fits in with previous work.

- The design and implementation of an evaluation of a parameterized set of 2D icon-based visualization methods where we quantified how perceptual interactions among visual dimensions (size, spacing, and icon brightness) affect effective data exploration. In the previous experiment, the difficulty and number of the tasks required, the high variance of the responses obtained, and the small subset of visual dimension combinations tested made our results difficult to generalize. This current experiment improved the tasks by making them more accessible to non-experts, lowering the variance between subjects. Of course, this moves away from the critique-inspired methodology towards more quantitative perceptual tasks but, as mentioned before, keeping in mind the application of the results and the type of visualization display we will create. This experiment presents the basic methodology for modeling perceptual interactions among visual dimensions. *This work has been accepted for publication at IEEE Transactions on Visualization and Computer Graphics [Acevedo and Laidlaw, 2006] and will be presented at the IEEE Visualization'06 conference in October.*

After these initial elements have been completed, we propose to conclude this dissertation by completing the following research components, which will be explained in more detail in Chapter 5:

- Augment the last perception-based study to include the other two visual dimensions we are working with in our model (orientation and color saturation). We will have a full set of base line data to characterize perceptual interactions.

- Perform a new experiment to measure the relative saliency among our visual dimensions. We will be able to compare these results with some of the values obtained during the expert visual design educators study. Saliency is one of the design factors that are key in the data exploration process.

- With those two experiments completed we will be able to hypothesize the expected values for the different design factors for visualization methods that display two continuous scalar data variables.

- Evaluate those hypotheses sampling the space of visualization methods. These samples will be sparse, since the number of combinations will be very high and we cannot possibly test all options. We will utilize expert visual design educators for the evaluations and to inform the choice of method samples to be tested. The samples will include controls for combinations with expected results, as well as new more unpredictable combinations.

- Develop an optimization strategy to search for methods based on constraints. Given the results of the previous evaluations, and the reasoning used by expert visual design educators to choose the samples and predict evaluation outcomes, we will develop our optimization scheme. The order of the search, whether we start by constraining the number of layers or the number of active visual dimensions (those mapped to data variables), will be determined by learning from the process experts followed to take their decisions.

- Evaluate our optimization scheme using real data from two scientific domains: Synthetic Aperture Radar (SAR) polarimetry and Computational Fluid Dynamics (CFD). We will have both expert scientists and visual designers evaluate the results of the optimization process based on the constraints posed by scientists.

1.4 Organization of this Document

Along with a more in-depth description of our four elements of visualization, Chapter 2 will introduce some definitions and notation that will be used throughout this dissertation. This chapter clarifies the scope of the thesis and puts each of our elements of visualization in context: the type of visualization problems we are dealing with, the type of visual dimensions we consider, and our definition of visualization effectiveness. Although it introduces very basic concepts, the discussion in this chapter helps the understanding of our research plan and the decisions made along the way.

After this introductory chapter, Chapter 3 will explore the extensive literature related to this project and how we shaped our investigations based on previous work.

Chapter 4 will describe the main results obtained so far and how our experiments have evolved. Although the visualization problem remains consistent throughout all this process, our concept and use of the other two elements of visualization changed based on lessons learned in these initial experiments.

Chapter 5 will outline the design for our next experiments and the description of our optimization process, which we will use to search our space of visualization methods. We will also describe, in detail, the application areas of this work and how our model and optimization will improve the science in those fields. This will be followed by a final discussion of our results, the expected impact of this dissertation, what the future lines of research would be following this dissertation, and some conclusions.

Chapter 2

The Elements of Visualization

Every visualization process starts with a question. A question about some characteristic of the dataset a scientist has just compiled and needs to study [Springmeyer et al., 1992]. After performing some filtering of the data, and maybe some statistical analysis or data mining, it becomes clear that looking at a set of numbers does not help the scientist understand what the data contains. A visual representation is needed. The first element of visualization is realized: the visualization problem.

The first step is figuring out how to translate the numbers into visual entities so it is easy to explore the relationships among all the variables in the dataset. The components used to create that representation are what we call visual dimensions. Color, shape, size and movement are examples of some of those dimensions, and they form what we consider the second element of visualization.

There are many ways of combining those dimensions to show the data, and scientists need some guidance in deciding which mapping, from numerical data to visual dimensions, is appropriate for his or her visualization problem. This guidance comes from our third element of visualization: the design factors that characterize the capabilities each of the dimensions have for representing data.

Finally, evaluating the effectiveness of different methods, our fourth element of visualization, addresses the process of deciding which method fulfills the goals the scientists has for the visual display. It accomplishes this by providing quantitative measures that can be used to evaluate the quality of a given visualization method.

Let us review our four elements of visualization one by one.

2.1 The Visualization Problem

The visualization problem has two distinct components: the goal for the creation of a visual representation of a dataset and the type of dataset being visualized. We will discuss these in the following subsections.

2.1.1 The Goal of the Visualization

What is the goal of the visualization display? This is the first question we must answer when we want to transform the set of numbers that form our dataset into a visual representation. Maybe the goal is to check the dataset for problems or obvious errors. Maybe we want to highlight some parts of it, like extrema or areas below a certain threshold. Maybe we are searching for a specific pattern that indicates some interesting phenomena is happening. All these examples would require, at first glance, a dedicated visualization design that would translate the numerical data into a visual representation that fulfills the requirements. We can classify these visualization problems into two main categories: explanatory problems and exploratory problems.

In cases when the goal is to show specific characteristics of the data and we know how to find them, or when we want to show the results of an experiment that revealed some unexpected patterns, our visualization problem is explanatory. There are certain things in the data that require the viewers attention, and the visualization method used should lead the viewer to them.

On the other hand, there are occasions in which the end-user wants to just see whether the data coming out of the experiment *looks OK* or there are some errors in it. In this case the approach is one of exploration. The end-user wants all the data presented in front of her in an unbiased way. There are no preconceptions about more or less interesting areas that should be highlighted or blurred. These exploratory visualization problems are the ones that we are going to address in this dissertation.

In some sense, the lack of a clear task to be performed by our visualization users makes our job more difficult. The fact that they just want to visually absorb everything the data has to offer without creating biases is a big challenge. For example, local maxima of a dataset can be marked using visually salient icons. Their numerical values can even be displayed beside the icons. There can certainly be a discussion about the design of such icons, the placement of the numerical value, etc., but the visualization problem is clear.

The goal of exploratory visualization is to gain insight into how the data is spatially organized. In the multi-valued case, exploration seeks an understanding of the relationships

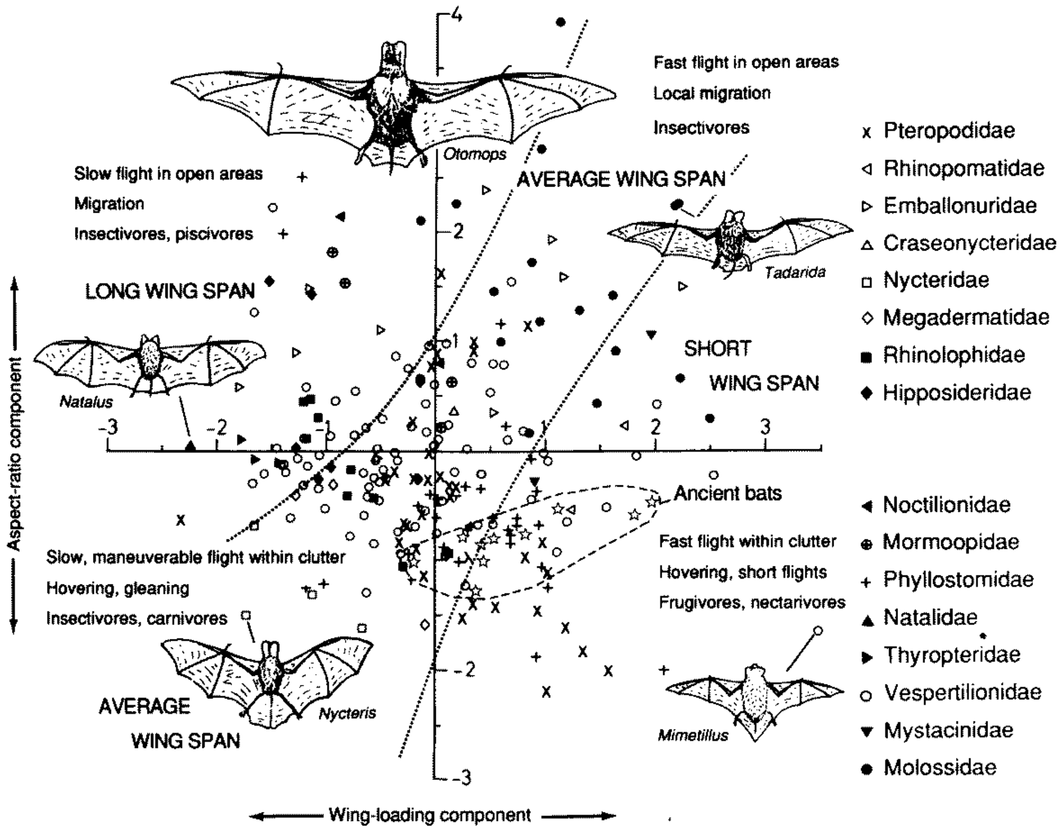

Figure 2.1: An information visualization display showing and example of non-spatial data from [Norberg and Rayner, 1987] presented in a cartesian grid designed ad-hoc for this visualization. It displays different anatomical and flight-related variables for many species of bats to try to discover a correlation between size of the bats, their flight speed and their behavior. Shape and location are used as visual dimensions in this case.

among data variables. Once these are presented, the visualization user will begin asking more explanatory questions, derived from the insight gained and requiring, in general, a different type of visual display that helps support his or her arguments.

2.1.2 The Type of Dataset Being Visualized

Once we have the main goal the visual display must fulfill, we must take into account the scientific problem we are trying to address. In other words: what type of dataset are we dealing with?

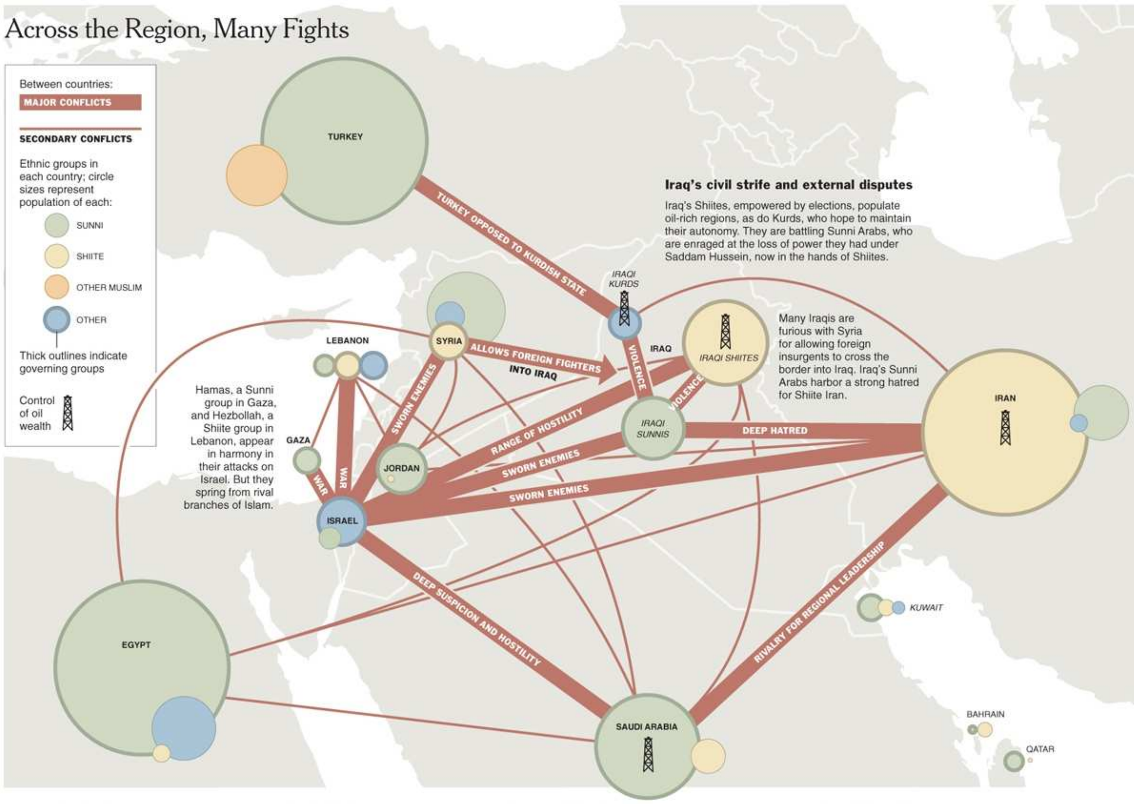

Figure 2.2: An excellent information visualization display from the New York Times that shows several types of data using size, color, and outlines on top of a map, which provides the spatial information component (Copyright 2006 The New York Times Company)

Information Visualization vs. Scientific Visualization

There is much debate in the community about the distinction between these two types of visualization [Rhyne, 2003]. The annual IEEE conference on visualization is divided in two to distinguish between research in one area or the other.

Information visualization deals with datasets that do not have an inherent spatial component or that, having one, represent abstract non-physical data. On the first case, a visual representation of those datasets must be made in an abstract space delimited by some of the variables present in the data (see Figure 2.1). The second case is more debated since it has clear spatial reference, such as the geographical area indicated by the map in Figure 2.2, but includes non-physical information (in the same figure, the types of conflicts are indicated by lines.)

When the dataset contains information that has a clear spatial component and involves physical phenomena, we have a scientific visualization display (see Figure 2.3). Even using

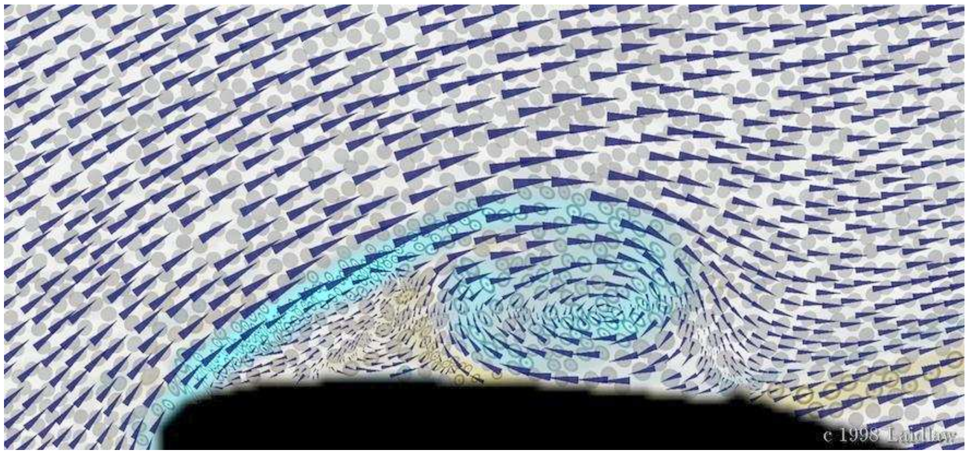

Figure 2.3: Visualization of experimental 2D flow past an airfoil [Kirby et al., 1999]. Six different variables of the flow are visible at every point in this image. It shows relationships among the values that can verify known properties of this particular flow or suggest new relationships between derived quantities.

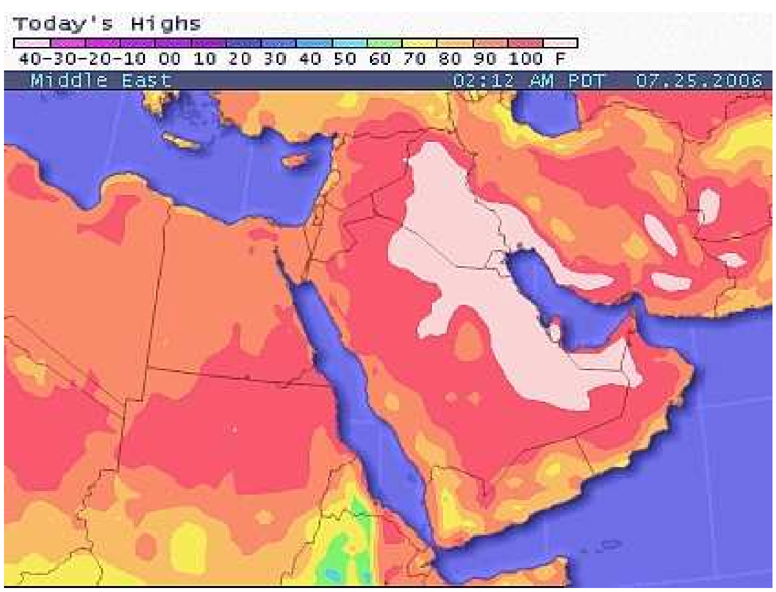

Figure 2.4: Example of a very simple scientific visualization display using the same spatial reference as Figure 2.2.

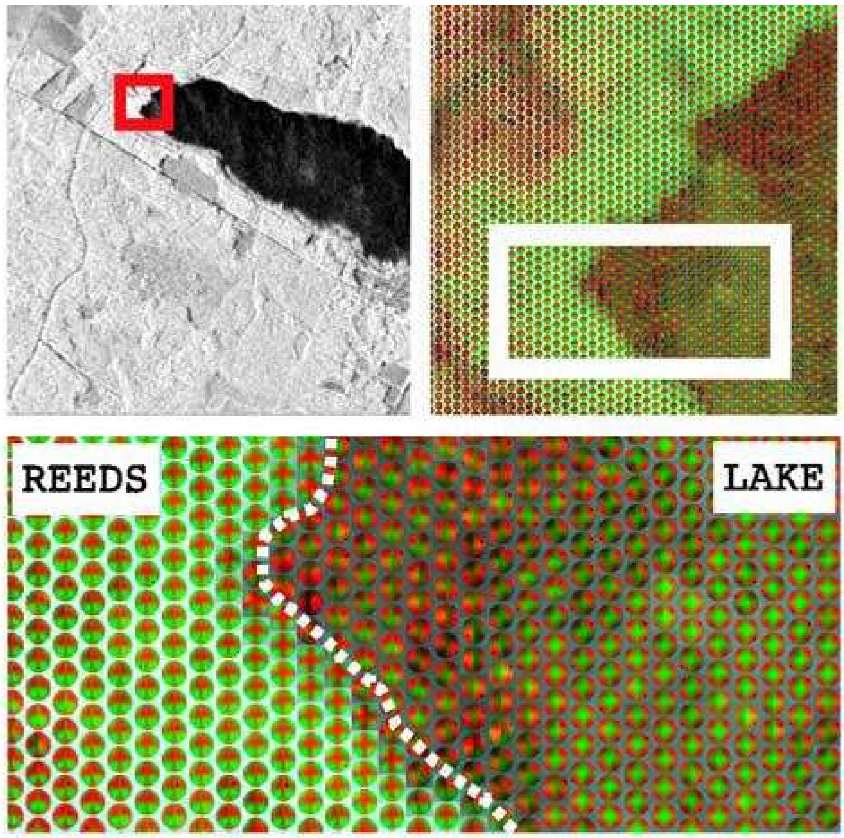

Figure 2.5: Example of continuous data being visualized using discrete glyphs. In this case SAR polarimetric response patterns indicate different types of surface cover. A coherent change in response pattern between the lake surface and reed beds can be detected [Woodhouse et al., 2002].

the same spatial context as Figure 2.2, the origin of the data and their characteristics provide the visual display different qualities (see Figure 2.4)

The distinction is by no means clear cut. The information in the examples of Figures 2.1 and 2.2 is not un-scientific, but the data are qualitatively different from the examples of Figures 2.3 and 2.4. This dissertation is aimed at developing better, perceptually effective scientific visualization methods.

Data Continuity

Another key characteristic of the data is whether it is continuous or discrete. Continuous data can be queried and visually represented at every point in the region of space where it resides. Discrete data, on the other hand, corresponds to measurements at specific spatial locations. The data to be visualized might come from the interpolation of values gathered

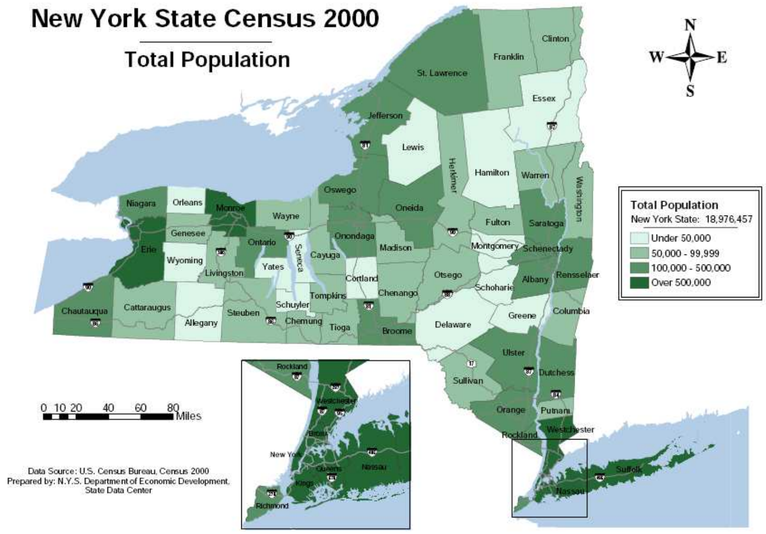

Figure 2.6: Example of discrete non-interpolable data. This is also an example of information visualization data which clearly has a spatial component.

in those discrete points; the temperature readings that originated Figure 2.4 were obviously discrete, but the dataset being visualized is the interpolated one, making it continuous.

There are also continuous data being visualized discretely (see Figure 2.5) but the intention of the visual display is to have the user do the interpolation.

Data continuity can sometimes serves as a distinction between information and scientific visualization displays. Since physical quantities such as temperature or pressure can only be measured in discrete locations but exist in every point in space, we consider them continuous and scientific visualization methods are responsible for visualizing them. Data variables measured discretely but that cannot be interpolated, like number of people living in an area (see Figure 2.6), are the responsibility of information visualization methods.

For this thesis we will consider only interpolated continuous datasets. Since our visualization methods are icon-based, we will be discretizing the interpolated datasets in order to facilitate layering of multiple variables in the same display. Figure 2.5 accomplishes this by creating a complex glyph that incorporates multiple variables. As we will explain in the Section 2.4, some of our design factors deal with the loss of spatial frequency due to this discretization. Perceptually measuring this quality for our visualization methods will allow us to optimize our visualization results depending upon the importance of the different requirements: e.g. is exploring multiple variables more important than higher information frequency?

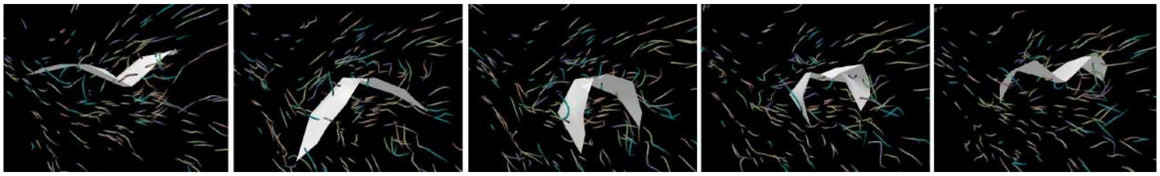

Figure 2.7: Particle flurries in a Virtual Reality visualization of air flow around a bat in flight [Sobel et al., 2004]. Over a short time, the particle animation gives a synoptic visualization of the main features of the three dimensional vector field.

Data Characteristics

There are two more data characteristics that will define the visualization problem. The first one differentiates between qualitative and quantitative data. The former include elements of different classes that may or may not be ordered and that have no inherent numerical relationship among them. Quantitative data, on the other hand, maintains a mathematical relationship among all data elements.

There are three types of quantitative data depending on the number of values each data element has. In a scalar dataset only one value exists at each location in space where data is present. A vector dataset contains elements with as many values as the dimensionality of the space they are in, i.e. in 2D, vector elements have two components; three in 3D, and so on. A tensor dataset contains elements with more values per spatial location, e.g. a second-order tensor has nine values per element.

The last characteristic of the data that defines the visualization problem is the space in which the data live. Whether it is 1D, 2D, 3D, or more, the visualization display must be adapted to the requirements of the data space. Time should also be included as dimension here. Dynamic datasets require a very different treatment than static ones since correlations across time-steps and maintaining temporal coherence in the visualization method become keys to avoid distracting the user with artifacts not really belonging to the data.

Summary

The previously described characteristics of the visualization problem shape the visualization methods that can be applied to it. The discussion up to this point is intended to briefly introduce the vast amount of visualization challenges that datasets in all those different spaces pose for the creation of effective visualization methods.

For this dissertation we chose to constraint our model to continuous scalar datasets in 2D. This might seem, on first glance, a simplification of the problem, given that there

Figure 2.8: Three different visualizations of the diffusion tensor magnetic resonance imaging data of a brain. This tensor field data can be explored using an immersive VR environment such as a CAVE (left) [Zhang et al., 2003], a fishtank-VR setup (center) [Demiralp et al., 2006], or a physical model created through color rapid prototyping (right) [Acevedo et al., 2004].

Figure 2.9: Virtual Reality is again used here to visualize archaeological excavation data. Colors indicate the type of artifact while size and quantity indicate other abstract variables present in the dataset. From left to right, a view of the excavation site, an overview of the full dataset, and two different moments of a user interacting with the system. In this case, the need to perform spatial correlations were the key to using a VR environment [Acevedo et al., 2001; Vote et al., 2002].

are many researchers that have already tackled much more complex types of datasets and developed successful visualization methodologies. Some examples of complex datasets successfully being visualized are shown in Figures 2.7, 2.8, and 2.9, where different visual techniques, display form factors, and interactions, combine to form the various visualization methods.

Our choice is based on the fact that the basic components of all those visualizations visually interact in ways that are still not clearly understood. Almost all great visualization methods must be iterated upon until a solution is reached that effectively shows the information required and minimizes perceptual issues. We want to analyze those issues from the ground up. By constraining ourselves to more manageable datasets, we can be more thorough in the analysis of perceptual interactions, eliminating from the experiments a multitude of dependent variables. This is not to say that those variables, such as the type of display used, the interaction techniques, etc., are not important, but they should be added to the experiment once other basic visual characteristics are better understood.

2.2 The Visual Dimensions

Each of the types of data detailed in the previous section will potentially require a different visualization method. For the purposes of this dissertation we can define a visualization method as follows:

> A **visualization method** is an abstract function that transforms a scientific dataset into a visual representation to facilitate data exploration.

The choice of method must be made based upon their strengths and weaknesses, keeping in mind the underlying pledge that every visualization method must display the data truthfully and avoid misleading the viewer. Fred Brooks summarized this point beautifully in a talk given during IEEE SIGGRAPH 2003:

> *"Visualize to inform, not to impress. If you really inform, you will impress."*

Going from numbers to pictures is usually the first step in the exploration of any scientific dataset, and choosing the right tool for the job can be the difference between success or failure of a scientific enquiry. As an example of this, Tufte shows how the space shuttle Challenger's catastrophic launch in January of 1986 could have been avoided if only the available data had been presented correctly [Tufte, 1997]. The visualization problem, which includes the goal of the visualization and the type of data, will partly determine this choice (the next element of visualization, the effectiveness evaluation, also plays a role in the choice of method). Nevertheless, the available choices are many and their differences in terms of visualization effectiveness are not well understood, especially when used in combination.

The visual dimensions are the toolset we use to create the visual representation of the data. There are non-visual dimensions we could be using to represent data, such as sound or haptic interfaces, but this dissertation is aimed at a small subset of just visual dimensions.

We will define these dimensions as elements that, more or less independently, can be used to create a visual representation of a scientific dataset. This is to say, and for the purposes of this dissertation, that continuous scalar data variables must be mappable to them.

In the previous section we have established that the focus of this dissertation will be on multi-valued continuous scalar datasets in 2D. Given this, there is still a large number of dimensions we could choose from. Some of the most salient ones are shown in Figure 2.10.

Note that even when only a few of these dimensions get mapped to data variables in any particular visualization method, all of them will be present in the final visualization display. For example, it is obvious that even when size is not used to represent any data variable in the dataset, all icons must have a certain size, either constant or randomized, across the display. Our approach to reflecting this into our model will be further explained in Chapter 4.

For practical reasons we chose to limit our experiments to five visual dimensions: icon brightness, icon size, icon spacing, icon orientation and icon color saturation. This decision was made to decrease issues due to subject fatigue during experiments and to provide a relevant sampling of the full space of visualization methods. The reasons to choose these five particular dimensions are diverse. Size and density are elements that received a highly varied set of reactions during our initial experiments [Acevedo et al., 2005]. Also, very few studies have been published exploring these two elements together [Wolfe, 1998]. Icon brightness was chosen because it is an element that has been studied in depth, allowing us to compare our results with previous experiments. Icon orientation has had a lot of attention in the perceptual psychology literature related to texture discrimination but its use for scientific visualization is limited to a few studies. It is, from our experience during the first experiments with visual designers, a difficult element to perceive as a scalar, but it is a very preattentive dimension. Finally color saturation provides our link with the use of color for visualization. Including hue would bring in a lot of different issues related to the use of color spaces that would complicate the main focus of this dissertation. Saturation, being a very much neglected visual dimension for scientific visualization, provides a convenient middle ground to introduce color in our experiments and explore its capabilities for visualization use.

Apart from these five dimensions mentioned, another one that we will consider will be the number of layers a visualization method utilizes. For example, if two data variables have to be represented in the same display and we want to use size and color saturation, these can be accomplished on one or two layers, as Figure 2.11 shows. We are including the number of layers as a visual dimension in our model. Furthermore, to differentiate among layers, shape will also be used when the measured saliency differences among the visual dimensions is not sufficient given the visualization requirements. This point will be further explained in the next section, and Figure 2.11 shows an example of its use.

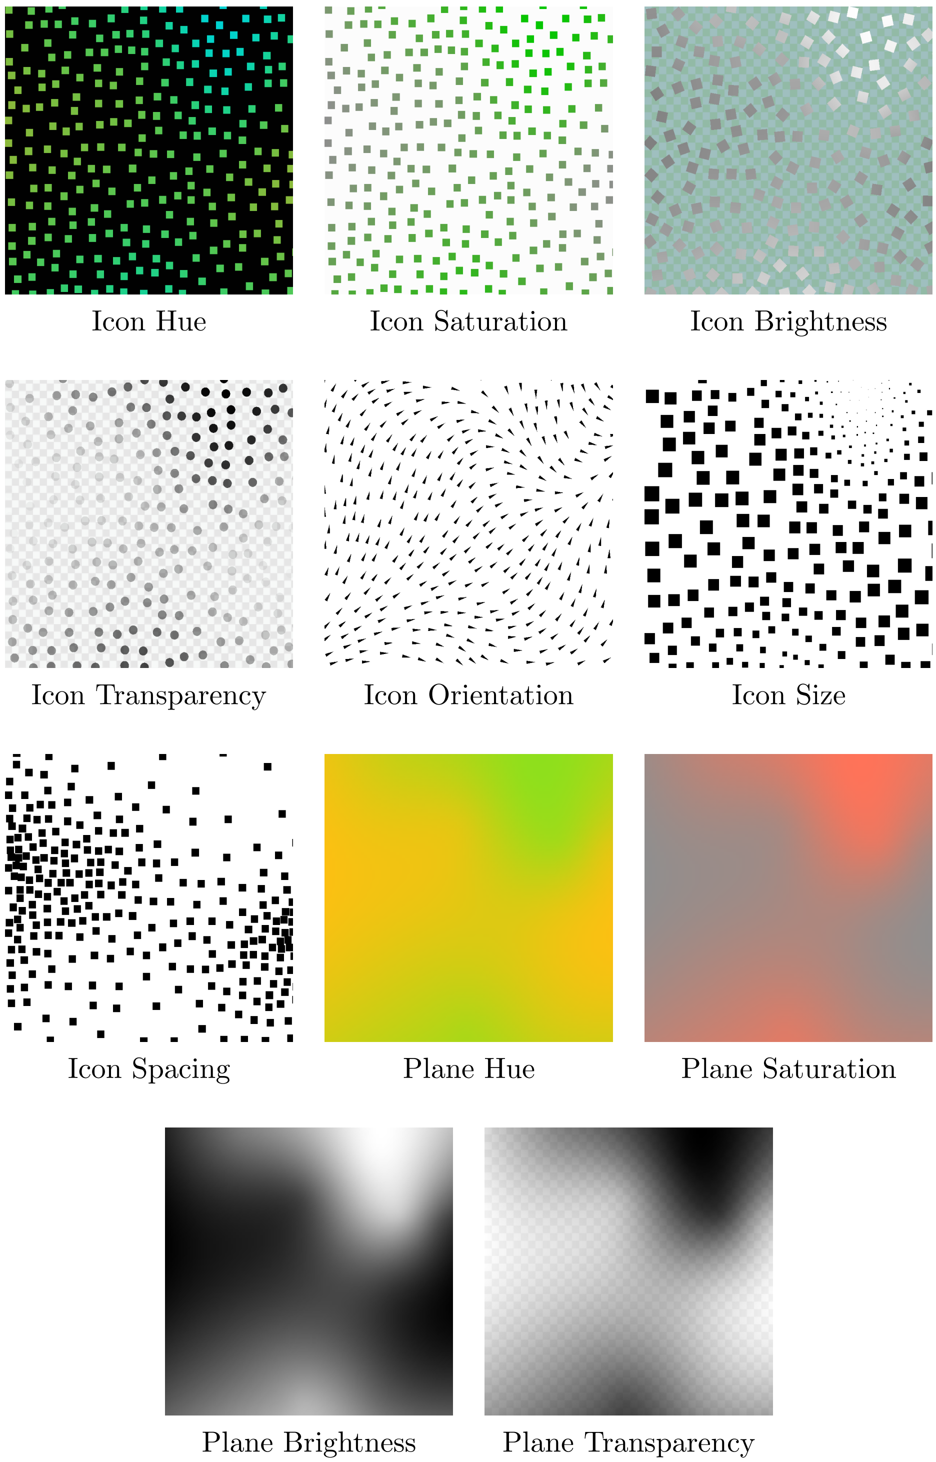

Figure 2.10: Visual dimensions. Eleven different visualization methods that represent the same continuous scalar dataset. We are quantifying the perceptual interactions among size, spacing, orientation, icon brightness and icon saturation when representing scalar datasets in 2D. Our methodology could be extended to explore the rest of these visual dimensions.

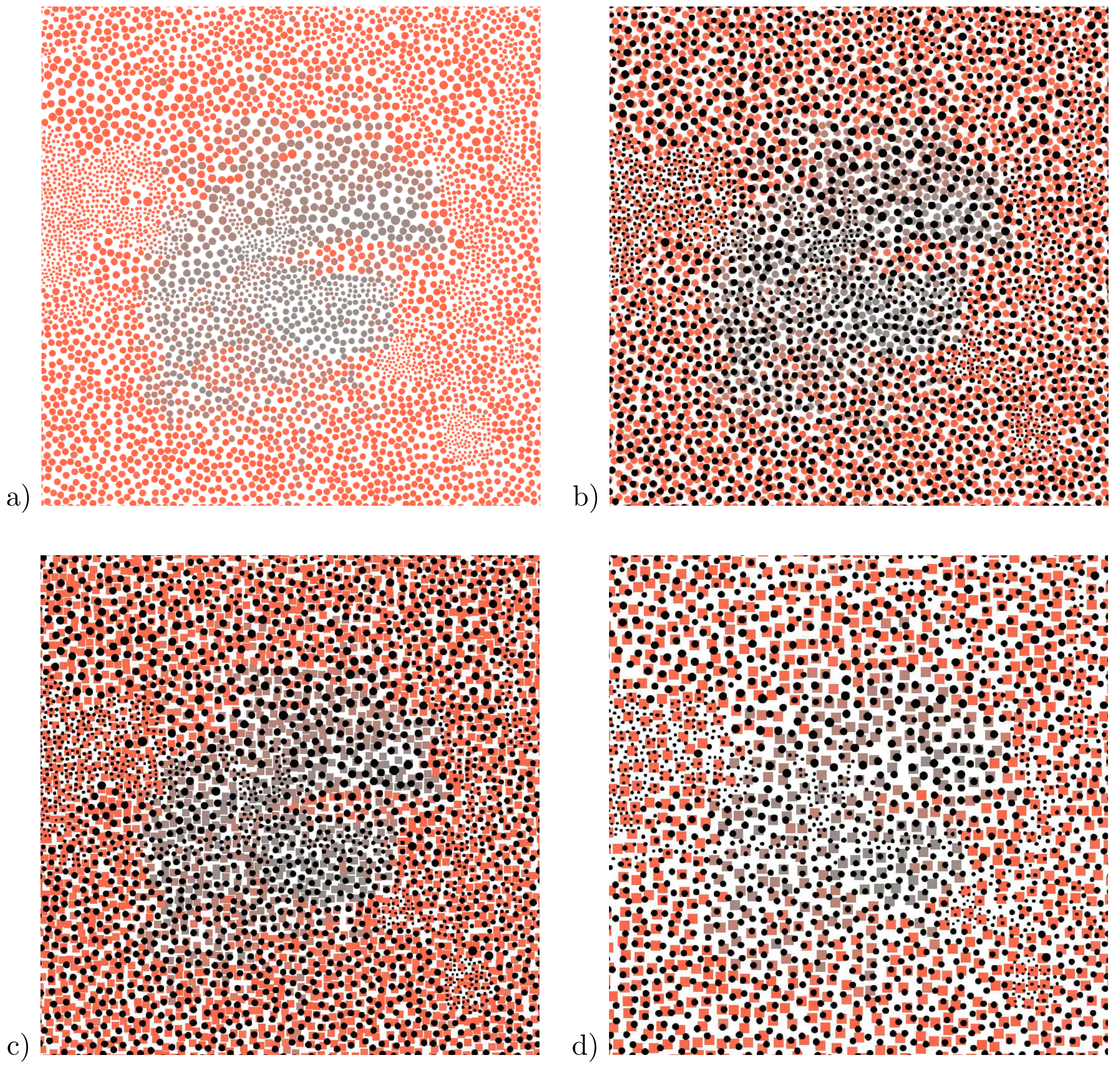

Figure 2.11: Example of a two-variable visualization using (a) a single layer, (b) two layers with the same icon shapes, (c) two layers with different icons shapes, and (d) slightly changing the spacing parameters from (c). Our experiments will quantify which of these representations works best depending upon what the visualization requirements are.

To summarize, our analysis of interactions among visual dimensions represents the study of how independent these dimensions are of each other. Since these dimensions represent the axes of a space of visualization methods, finding and quantifying those interactions would lead to an understanding of how orthogonal those axes are. An ideal space would have orthogonal axes that could be used independently of each other when creating visualization methods.

2.2.1 Definitions and Nomenclature

This nomenclature is meant to provide a framework for a full exploration of visualization space, even though this dissertation is constrained to exploring a few dimensions.

A visualization method takes a scientific dataset and produces a visualization. We define a space, $\mathbb{V}$, of scientific visualization methods. In general, our space includes layered iconic and RGBα representations of 2D multi-valued data. The visual dimensions that are present in each layer are:

- Icon color hue (v_0)
- Icon color saturation (v_1)
- Icon color lightness (v_2)
- Icon transparency (v_3)
- Icon orientation (v_4)
- Icon size (v_5)
- Icon spacing (v_6)
- Background color hue (v_7)
- Background color saturation (v_8)
- Background color lightness (v_9)
- Background transparency (v_{10})
- Icon shape (v_{11})

Note that icon shape is included here for the purposes of differentiating icon layers, as shown in Figure 2.11.

A visualization method, $v \in \mathbb{V}$ maps data values to visual dimensions. We can combine multiple layers, which we will denominate l_k. Each one of these layers will contain all 12 visual dimensions defined above. The subscript $_k$ of the layer indicates its order in the final visualization, from back to front:

$$v = \{l_0, l_1, l_2, ...\} \in \mathbb{V}$$

where,

$$l_k = \{(m_0, m_1, ..., m_{11}), (r_0, r_1, ..., r_{11})\}$$

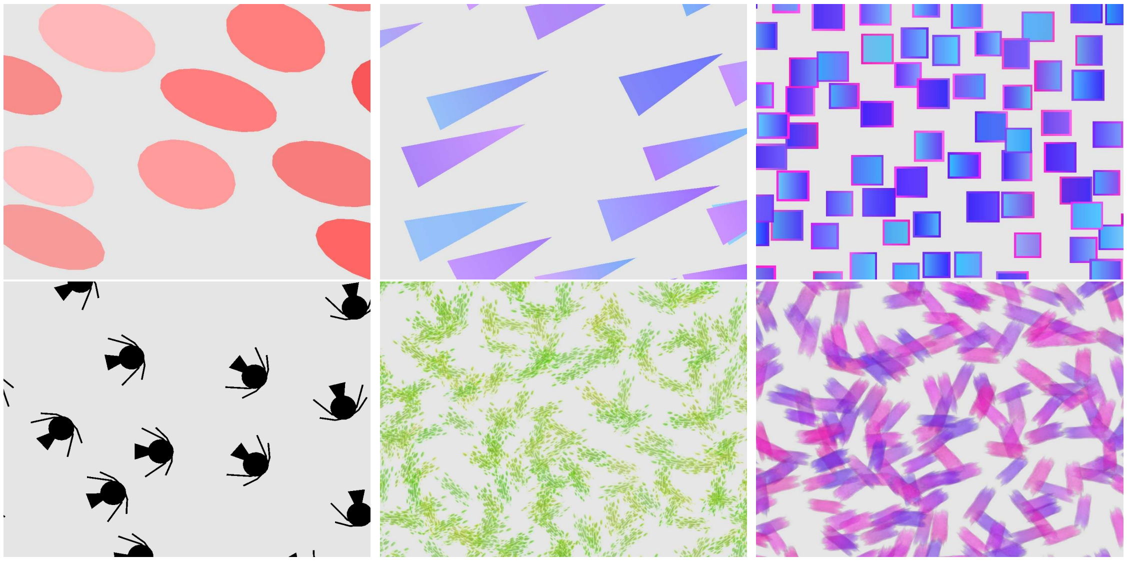

Figure 2.12: Different discrete icons. The top three show each basic primitive (ellipse, triangle and rectangle) provided in the language. The bottom left image shows a composite icon created using a combination of primitives. The last two images show how texture mapping the icons can generate arbitrary shapes.

Each component of l_k refers to one of the 12 visual dimensions v_i:

$$m_i = \begin{cases} 0 & v_i \text{ is not mapped} \\ d_i & v_i \text{ is mapped} \end{cases}$$

$$r_i = \begin{cases} c_i \in \mathbb{R} \in [0,1] & m_i = 0 \\ (b_i, e_i) \in \mathbb{R}^2 \in ([0,1],[0,1]) & m_i \neq 0 \end{cases}$$

where d_i is the index of the data variable mapped into v_i.

2.2.2 EVOLVIS: A Visualization Language

To create the actual visual displays for this dissertation we developed a language to describe scientific visualizations of multi-valued, two-dimensional datasets. Our goal was to create a language with which a user could quickly create complex and precise data-driven visualizations as well as facilitate their modification during the iterative design process. We called this language and the rendering system to display it EVOLVIS.

EVOLVIS is a general tool that can combine three types of basic visual elements – discrete icons, color planes, and streamlines– into layers. A text file fully describes the resulting method and controls the layering of the different elements, their appearance, and their spacing, including the mapping assignments of any of these visual dimensions to data variables. In addition the language supports extensions to accommodate scalar, vector, and tensor data in 2D and 3D.

For the purposes of this dissertation, only icon elements will be used, although color planes were also used for some of the initial pilot studies with expert visual designers. Other visual dimensions that are possible with our language include border width and color, texture mapping, and more complex icon shapes that can be controlled by the data. Figure 2.12 shows some examples of the possibilities EVOLVIS provides.

2.3 The Design Factors

So far we have defined two of our four elements of visualization, making clear in each of the previous two sections the scope of this dissertation: we want to create multi-layered exploratory scientific visualization methods for multi-valued continuous scalar datasets in 2D.

To achieve this, the first thing we need to establish is a set of design factors that, without constraining our model to any specific task, represent the exploratory nature of our visualization methods. Once we have these factors, we will be able to quantify how the different visual dimensions express those and how, when combined, perceptual interactions among dimensions affect that expressiveness.

2.3.1 Definitions and Nomenclature

Developing upon the nomanclature from the previous section, $\mathbb{W}(v)$ provides an evaluation of a visualization method, $v \in \mathbb{V}$. It produces a vector of values each of which quantitatively characterizes the visualization method with respect to a specific design factor.

$$\mathbb{W}(v) = (w_0(v), w_1(v), w_2(v), ...)$$

In order to generate an effectiveness model for our visual dimensions we need to specify the type of design factors we will be accepting. We consider three different ones:

- *Data Resolution (w_0):* The number of different values of a data variable we want to be able to differentiate in the visualization.

- *Spatial Feature Resolution (w_1):* The size of the features we want to be able to see in the visualization.

- *Saliency (w_2):* How much we want a data variable to pop-out among the rest of the variables.

We derived these factors from our experience creating scientific visualizations for our collaborators in many disciplines and from our pilot study on designer-critiqued visualization methods [Jackson et al., 2003]. In that study we utilized a superset of factors that, given that experience, we have narrowed down to these three described here.

2.3.2 Data Resolution

This design factor refers to the number of different values that should be visible for a given variable. Although we are dealing with continuous scalar data, scientists often bin their data for easier comprehension.

In our experiments we measure the number of levels a visual dimension has by counting *jnd* (just noticeable difference) units. This means that, although we are not explicitly binning the data for display, we are quantifying how many bins will be perceived. Our result is the limit of the number of bins possible for a certain visual dimension in a given range.

Some dimensions have a limited range, such as brightness or orientation, while others like size or spacing, do not have specific limits. As we will explain for the spatial feature resolution design factor, we limit these dimensions so we can provide higher spatial feature resolution, hence constraining the range and the perceivable bins. This limitiation of the range has consequences for the data resolution itself, since there is a limited number of jnd's perceivable in a given range. Measuring data resolution for different ranges will be useful so we can combine visual dimensions using ranges that conflict the least between them.

We measure this factor as the total number of levels visible.

2.3.3 Spatial Feature Resolution

This design factor is aimed at giving the visualization user some control over how much information is lost when the icon-based representation is used. Size and spacing between icons are the main dimensions that affect this factor, but our visual system processes visual dimensions differently. Some dimensions will be easier to interpolate than others so, for example, larger spacing and smaller icons could provide similar spatial feature resolution results for icon orientation than for icon color saturation using closer together and larger icons.

Important features of a dataset might be lost if the size of those features are beyond the capabilities of a certain visualization method. When icons are used for representing continuous data, it is unavoidable that gaps will be present in the final display. These gaps are what make a multi-layered visualization possible but, at the same time, they create a challenge. There is a trade off between showing a higher spatial resolution (smaller features) and providing a comparative view with other layers in the visualization (larger holes to see through the layers).

We also limit the range of some visual dimensions so we can visualize smaller features. Size and spacing do not have specific limits for their maximum values. Since icons can potentially be made as large as the full display size, this would clearly limit the available spatial feature resolution. We will provide numeric values for our available ranges when we define the different experiments in Chapter 4.

We measure this factor as the size of the smaller feature a method can represent, measured as a percentage of the image width. With this measure we try to abstract from the actual physical size of the stimulae used for our experiments, because we roughly know the distance our subjects stand from the screen when they perform our experimental tasks (either on paper or on the computer screen).

2.3.4 Saliency

With this factor we address the level of importance a data variable must have among the other variables present in the visualization. There might be times when the user wants to highlight a particular variable and keep others as a context. There might also be cases when all variables must be visualized at the same level of importance, leaving the highlighting and backgrounding of some of them for a later stage of exploration.

We measure this factor using direct comparison between methods and asking subjects which method dominates the composition.

2.3.5 Capturing Designer Critiques

As we mentioned before, the advantage of utilizing expert visual design educators as subjects for some of our experiments is that they can provide reasons for the success or failure of a certain visualization method. During our experiments we will try to capture this information numerically by asking them for estimates on how much a design factor would change if a certain change in one of the visual dimensions is performed. We can indicate this as the derivative of our vector of design factors $\frac{d\mathbb{W}(v)}{dv}$:

$$\frac{d\mathbb{W}(v)}{dv} = (\frac{dw_0(v)}{dv}, \frac{dw_1(v)}{dv}, \frac{dw_2(v)}{dv})$$

where,

$$\frac{dw_j(v)}{dv} = (\frac{\partial w_j(v)}{\partial v_0}, \frac{\partial w_j(v)}{\partial v_1}, ..., \frac{\partial w_j(v)}{\partial v_{11}})$$

Some of these derivatives will be obtained from the analysis of videotapes recorded during the experiments and from interactive sessions where experts modify visualization methods according to different requirements. It is usually easier to let subjects explain their reasoning as they perform the experimental tasks than providing them with a form to fill out. This is a more difficult method to obtain quantitative results, but it allows subjects not used to numerically critique designs to feel more confident about their decisions.

Gathering this information, we will not only model how different methods perform for each design factor, but we will be able to better guide our optimization process in search for solutions for a visualization problem.

2.4 The Effectiveness Evaluation

Given a multi-valued dataset and a set of design goals for each individual variable, the end result of the use of our model will be to obtain an effectiveness evaluation over the space of visualization methods.

With the exploratory goal in mind, our measure of effectiveness has to come either from an increase in insight for the user or from an optimal and expressive use of perceptual resources. The former is hard to measure and non-standard to quantify [North, 2006], since it could mean different things in different scientific domains. There have been some attempts at this, most notably [Saraiya et al., 2005], using protocol analysis [Owen et al., 2006] as the experimental methodology.

Since the goal of exploration per se does not contain any specific tasks we can measure performance from, our approach is to present the data in a way that allows for maximal perceptual bandwidth. This is to say, we will measure effectiveness by how well the perceptual capabilities of the different visual dimensions are used. Having measured these capabilities, our design factors, we can search for an optimal solution to a visualization problem posed in terms of those capabilities.

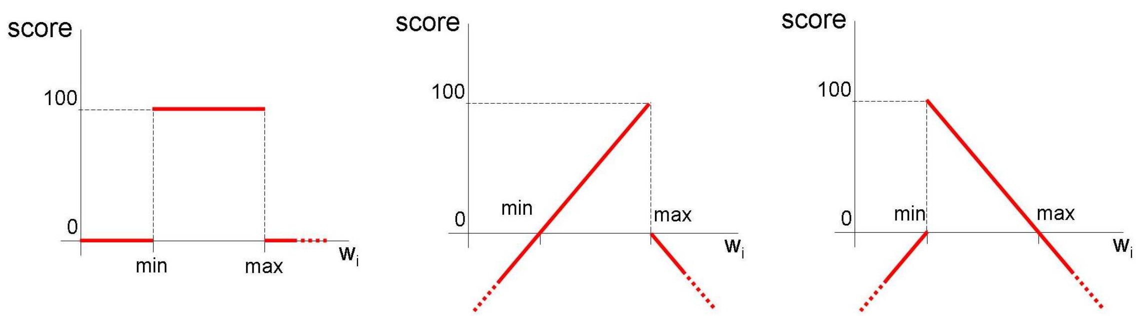

Figure 2.13: Scoring schemes for three different types of requirements. The horizontal axis is always the values for design factor w_i. From left to right: constant score within valid range, increasing score, and decreasing score. Note that for the last two schemes points outside and farther away from the valid range score worse than closer ones.

Along these lines, there is a term in art and visual design called *economy of line*. It means that, when creating a charcoal or pen-and-ink drawing, the least amount of lines should be used to show the pose. The expressive power of a well drawn single line is huge. In our case, we will know we have an effective method when the expressive power of our visual dimensions is used in the right amount to convey the message in the dataset.

The Effectiveness Ratings

The quantitative characterization of each of the previously described design factors allows us to measure the effectiveness of a visualization method based on how well it fulfills a set of given design goals upon those factors.

The effectiveness rating mechanism is based on assigning scores, from 0 to 100, to a visualization method based upon whether it satisfies a certain requirement, e.g. if a scientist needs a method capable of representing spatial features smaller than 2% of the image width. There are three different ways we generate scores, and they are summarized in Figure 2.13. Following this scoring scheme we can select the method that most efficiently fulfills a set of requirements. Importance weights for each design factor must also be given. This is why we measure effectiveness: the highest scoring method will be that which doesn't *waste* visual resources and satisfies the requirements at the same time.

2.5 Summary of this Chapter

In this chapter we have defined our *elements of visualization*: the visualization problem, the visual dimensions, and the effectiveness evaluation. For each of these we have defined the scope of this dissertation. We have also introduced the nomenclature and mathematical notation we will utilize throughout the thesis.

The purpose of this chapter was to characterize the different components of the research and make clear what our assumptions and limitations are. With all these framework in place we will now be able to present our previous results (Chapter 4) keeping in mind the scope and goals we want to achieve and the type of problems we are addressing. Chapter 5 will then introduce the proposed next steps for this dissertation and how they will build upon previous work to achieve our goal: the modeling of the effectiveness of 2D multi-valued multi-layered scientific visualization methods.

The next chapter will put our framework in perspective with respect to the state of the art in visualization research, as well as perceptual and design literature.

Chapter 3

Literature Review

Our work is related to three main research areas: visual design, visual perception and scientific visualization design. We will address them separately in this chapter.

In this thesis we are combining all three of these disciplines to facilitate the synthesis of effective visualization methods. Visual design informs our work through techniques for critiquing and works related to image composition and how visual components work together to convey a message. Visual perception, on the other hand, is our main source of low level characterizations for our visual dimensions. The experimental techniques used by perceptual psychologists help us design our own studies, targeted towards more higher level practical applications. Finally, we are trying to contribute to the advancement of the field of scientific visualization, and there are many researchers whose work inspires and complements our own. We summarize here the main sources of knowledge from all three of these areas of research.

3.1 Visual Design

Our main point of connection with visual design is the quantification of how the different dimensions that form our basis for visual data communication perform and interact together. There are many authors that have approached the problem of classifying those dimensions and providing guidance for their use, but we are providing a bottom-up approach that numerically quantifies individual performance first and moves on to combinations and their interactions.

Visual designers and artists are trained on how to communicate messages visually. In our case the message is a scientific dataset. We have previously researched, and continue to pursue, the idea of using artistic techniques for scientific visualization [Laidlaw, 2001; Vote et al., 2003; Laidlaw et al., 2004; Kirby et al., 2004; Keefe et al., 2005]. Our experience in this area, and our ongoing collaboration with the Rhode Island School of Design, helped us select the set of visual dimensions that form the *means* by which we communicate our message.

In looking at art and design critique, there are many principles of design and composition expert critics look for when critiquing a certain piece. Rhythm, repetition, balance, proportion, scale, variety and unity are some examples [Sayre, 1995]. In our case, some of those come defined by the data itself, hence not subject to evaluation. For example, rhythm and balance in a visualization display are, for the most part, controlled by how the values are distributed across the spatial range of the data. A certain balance could still be judged based on an overall visual balance of the display (whether there are areas that attract attention or not), but it could be argued that finding those areas is precisely the goal of the initial data exploration. Our initial set of design factors does not address this discussion, which we could study further in later stages of the project.

Another aspect of the critique looks at the set of visual elements a piece uses to convey the message intended. There is no agreement among researchers about what is the basic set of visual dimensions that can be used as a basis to create visualization methods. Ontologies about what visual dimensions are most commonly used served as inspiration for us to come up with a testable set. There are many publications used in art and design schools that deal with specific visual dimensions, but Wallschlaeger and Busic-Snyder in [Wallschlaeger and Busic-Snyder, 1992] provide a very comprehensive classification of the different elements involved in the communication process. Their work spans visual principles for architecture, art and design, and demonstrates the commonalities among those disciplines in this context. Our approach is similar in the sense that we are applying these concepts to an area that makes use of them [Swan et al., 1999], but has not had many researchers studying the formalization of their use. Although Wallschlaeger and Busic-Snyder provide a very clear description of each element (color, shape, texture, etc.), they fail to formalize the interaction among them and the issues arising from their simultaneous use, a key component in our research.

In the classification and analysis of visual dimensions for data representation, one of the first and most cited works outside of the academic literature for art and design is Bertin's *Semiology of Graphics* [Bertin, 1983]. Our approach is very similar to his in that we are trying to characterize the capabilities of each of our visual dimensions individually, and then build up a model of how they perform in combination. He acknowledges that any combination of dimensions is possible but he dedicates very few pages to formalizing the use of those combinations. Our studies are designed to gather knowledge and provide a basis for a formal model for the effective combination of visual dimensions. Our work also presents an opportunity to address a main criticism of Bertin's work, that he lacks experimental results for his factual presentation of visual properties, by providing quantifiable evidence of his theories.

Bertin describes in extensive detail the associative and selective qualities of what he calls *retinal variables*: location, size, value, texture, color (hue), orientation and shape. Compared to our visual dimensions, he combined saturation and hue into a single variable (color), and used texture in a similar way that we use icon spacing. In the case of the location visual dimension, the datasets we will be working with have an inherent spatial component which prevents us from using the location of the visual elements as a carrier of any other information. An example of this are the points on a map: they indicate precise geographical locations, so distortions from the real locations would change the map itself and modify the message (cartograms are an exception to this, in which the visible space changes its meaning [Bertin, 1983, p.120]).

A significant difference with Bertin's work is that we are dealing with continuous datasets instead of discrete ones. In his book, the main examples are centered around map displays, although charts and information visualization displays are also discussed. Maps are very close to the kind of display we are going to be analyzing. In our case, datasets are assumed to be continuous across the spatial range of the data variables. Although maps can contain that type of dataset (e.g interpolated temperature or precipitation readings in meteorological maps), there are many discrete variables (e.g. labor statistics or population maps) that are not usually interpolated. Our research will extend Bertin's results to continuous data.

Many researchers have followed and applied Bertin's work, and map making is one area that has used his work and inspiration extensibly. MacEachren presents an excellent summary of previous research in cartographic visualization [MacEachren and Kraak, 1997]. He expanded Bertin's visual variables to include crispness, resolution, transparency and arrangement. He also divided Bertin's color into hue and saturation for a total of 12 visual variables. Although his classification is better supported by experimental references from map makers and perceptual scientists, we miss some discussion about the specific use of each variable, both individually and in combination (combinations of hue, lightness and saturation are briefly presented). He provides clues towards the generation of rules for map-making but does not go as far as presenting such rules.

Cartographic data visualization is an area that uses similar datasets to the ones we utilize. Techniques and classifications of visual elements for map communication form an important basis for our model and our choice of dimensions for the experiments. Along these lines, MacEachren also describes the three main components of map communication as the data, the graphical elements, and the user. He contends that a characterization of all of them must be obtained to create effective maps. We are constraining our research to a very specific type of data and a predefined set of graphical (visual) elements that can potentially be used to represent that type of data. The end-user's characterization is represented by our set of design factors.

Finally, one of the most cited works on effective visual design for communicating scientific data has to be Tufte's series of three books on information visualization [Tufte, 1983, 1990, 1997]. In his books he discusses many examples of cases where bad choices in the design of visualization displays caused problems and misinterpretations. He also presents alternatives for how to fix those issues and introduces alternative designs that, for the most part, are effective in conveying the original ideas. It has been said that with Tufte's work bad practice has been uncovered. Even though this is a very valuable step to take in any scientific field, the recognition of existing flaws, Tufte does not take the next step and tries to formalize his views into a coherent comprehensive model. It is very difficult to connect all the extensive advise given in his work and categorize it in a homogeneous way. It seems clear that was not the intent of his work, and that is where our approach fits in. We are trying to build that model and, for that, we are starting the same way Bertin did in his work: from the ground up. The examples Tufte introduces present multiple combinations of visual dimensions and other factors that are very difficult to isolate. We are starting with the study of visual dimensions in isolation and trying to quantify the expressiveness of each of them as they get combined in increasingly complex situations.

In summary, Tufte does an excellent job at critiquing and analyzing finalized visualization displays, but fails at exploring how much or how little each component of the displays participates in its success or failure.

3.2 Visual Perception

At the other end of the spectrum are perceptual psychologists and psychophysicists. They are interested in studying how our eyes and brains perceive, process, and store visual information. For that, they utilize very basic visual displays that are design to trigger very low level responses on the viewer. This helps them isolate how individual pieces of information are processed and build a model for how we perceive the world around us. In our case we are quantifying how visual dimensions are perceived by visual experts and measuring the effect that their combination has in the perception of scientific datasets.

Ware [Ware, 2004] provides an excellent reference towards the understanding of all perceptual processes involved in information comprehension. Color, texture, form, and motion are the main elements discussed in his work, beginning from the physiological elements involved in perceiving each of those, up to a series of recommendations for their use in displaying abstract information. Ware takes a broad approach at information visualization and, although continuous data is discussed in the book, it is not its main focus. He provides a very good introduction to the theory of integral and separable dimensions for visual attributes, but provides little quantifiable evidence for his classification. Our work will provide such evidence for the displaying of continuous scalar data and for how separable dimensions combine to form complex visualization displays.

Along these lines, we have found little experimental evidence about the perception of combinations of visual dimensions. Callaghan studied how hue and lightness interact in a texture segregation task [Callaghan, 1984]. She also compared, in pairs, hue with form and line orientation [Callaghan, 1989]. Although she reached valid conclusions about which variables dominate and when they interfere, her stimulae were limited to two levels of the visual variable being analyzed (e.g. horizontal and vertical for the oriented lines), while the potentially interfering variable was randomized or kept constant. In general, given that our data is continuous, more than two levels of each visual variable will be displayed. We have not found any studies for the interaction of more than two visual variables. Note that, in our displays, all visual dimensions present in our language must be set. Even when only a single data variable is being mapped to a single visual dimension, there are a whole set of other dimensions that can potentially interfere with it.

Our experiments are very much inspired by Carswell and Wickens's work [Carswell and Wickens, 1990] in which they classify different graphical attributes into integral, separable or configurable dimensions depending upon how each attribute's reading is affected by the others, taken pairwise. They found that visual elements can help each other when displaying the same information (redundancy or performance facilitation), or inhibit each other when only one element is changing (filtering interference). They also describe a third type called condensation in which opposite variation of each variable occurs simultaneously. Their experimental displays are based on single icons, looked at in isolation. We are extending their experiments to more complex displays and, for now, limiting our analysis to filtering interference analysis (see Section 4.3.)

Our goal is to find the visual characteristics of different visual dimensions when displaying quantitative information, where visual saliency is the property that makes one data value different from the next. The measurement of saliency or texture contrast thresholds is common in texture segregation studies [Bergen, 1991]. Those studies utilize stimulae with regions where the particular visual element differs in some amount with respect to the surrounding region [Landy and Bergen, 1991]. This is similar to our research in that we are also measuring *jnd*'s for visual dimensions, but our stimulae include overlapping textures. Our texture segregation is a more complex problem, since even a single layer of icons can contain two or more textures, e.g. one for spacing and one for icon brightness. We are interested in measuring the segregation between these textures, so we cannot directly apply *jnd* values from those texture segregation experiments to our model. Also, many stimulae are required to explore the full range of a visual dimension, and even more to include interference analysis with secondary elements. Our studies are designed to evaluate, with less iterations, a larger portion of the range for each element.

Moving a little closer to experiments directly applicable to scientific visualization synthesis, most of the literature about perceptually effective data representation is based on experience. Authors define sets of guidelines that, in the absence of visual perception theories [Senay and Ignatius, 1994], follow common practice and established knowledge [Eick, 1995]. In general, these approaches rely on a clear definition of the task a visualization must fulfill, making them difficult to apply in our research. Our exploratory visualization methods are geared towards presenting the data as clearly and unbiased as possible for scientists to explore.

In our case, instead of evaluating error or speed for a specific task, we qualify the different methods based on design factors present during exploratory analysis. Also, since our subjects are expert visual designers with years of experience in design critiques, we are able to simultaneously evaluate multiple visual dimensions from our language. They are not new elements for our subjects, so we can exploit their expertise in a more efficient way. This methodology allows us not only to understand how our visual dimensions are interpreted by our expert designers, but also how the individual visual dimensions are combined by an observer into coherent percepts [Landy and Movshon, 1991].

Dastani [Dastani, 2002] takes a different approach. His goal is to match the structure of the datasets, relational databases in this case, with the perceptual structure of the visual dimensions used in the visualization display. Again, this is difficult to apply to our scientific datasets, but it is interesting to note that he includes in his discussion the choice of values for the visual dimensions not mapped to any data variables but still present in the displays. We also keep track of these when designers evaluate our visualization methods, and they comment on them, so we can build a complete model of effectiveness. We, like Dastani, try to avoid methods with unwanted visual implicatures by non-mapped visual dimensions.

Our evaluation approach comes closest to the work of Healey. He has studied extensively the application of preattentive processing to visualization [Healey et al., 1993]. Preattentive processing allows detection of visual elements in a display without focusing attention on them. Initially, he focused on experiments comparing hue and orientation [Healey et al., 1996]. Subjects in his experiments were asked to perform numerical estimation tasks with varying hue and orientation differences, as well as varying display time. Based on this discriminability experiments, he identified guidelines for color selection [Healey, 1996] that we used for our studies.

He also proposed ViA, a visualization system based on perceptual knowledge [Healey et al., 1999]. The goals of this system are very similar to the ones in our research. He builds, by hand, the perceptual knowledge-base used to suggest a visualization method, while we are gathering that knowledge through subjective evaluations. Finally, Healey successfully used perceptually-based visualization displays to visualize datasets with up to 4 data variables [Healey et al., 2004].

Finally, in the case of multi-valued visualizations, our layering of icon-based visualizations takes note from Watanabe's texture laciness studies [Watanabe and Cavanagh, 1996]. Texture laciness defines the phenomenon that occurs when two textured surfaces are overlapped and the top one becomes perceptually transparent so the bottom one can be perceived without interference. He identified icon similarity as the main element affecting laciness. In our case, we want some amount of laciness to be present, itself controlled by the saliency required by the visualization. Spatial feature resolution was not studied by Watanabe as a factor of laciness, but it is something we are including in our studies.

It is important to note that this combination of visual dimensions into perceptually relevant entities has been studied for decades, starting with the Gestalt psychologists and their laws of perception and grouping [Ellis, 1939]. These laws are one of the earliest attempts to qualify how the human visual system recognizes relationships among visual dimensions. We are trying to quantify some of those relationships and apply that knowledge to the effective visualization of scientific data.

3.3 Data Visualization

We titled this section Data Visualization to combine both information and scientific visualization literature. Hanrahan [Hanrahan, 2005] recognizes the artificial and somewhat unclear nature of the separation between information and scientific visualization, but acknowledges that most of the research aimed at the definition and characterization of a space of visualization methods has been done in the information visualization field. Our work is very much inspired and guided by the classification models developed for information visualization.

Many researchers in information visualization have followed and applied the previously mentioned work by Bertin on graphic semiology. Cleveland was one of the first in ordering what he called perceptual tasks (out visual dimensions) based on their accuracy when users read visualization displays [Cleveland and McGill, 1984]. [Mackinlay, 1986] augmented his classification by including expressiveness and effectiveness as the two main measures to evaluate how well a certain dimension performs representing data. Mackinlay went as far as to develop a compositional algebra that would describe how dimensions and tasks were matched to choose a certain method. He also acknowledges the existence of situations were visual dimensions will interfere with each other, throwing off the original classifications, but he does not study those cases. This thesis tries to build a similar classification of visual dimensions and complete the quantification of those perceptual interactions.

In our case the task is exploratory, but many researchers have approached the modeling of the space of visualizations based on a taxonomy of tasks [Casner, 1991; Springmeyer et al., 1992; Shneiderman, 1996]. Their results vary and are appropriate for the tasks represented, but they all fall short in the study of perceptual interactions when multiple variables need to be represented. Furthermore, even in the cases were they approached the issue, the type of relational or nominal data they deal with makes the extrapolation of their findings into our domain really difficult.

Other examples of modeling the space of visualizations are [Robertson, 1991; Miceli, 1992; Laper, 1995; Lange et al., 1995; Card and Mackinlay, 1997; Nowell, 1997; Andrienko and Andrienko, 1999; Chi, 2000; Nagappan, 2001; Salisbury, 2001; Jankun-Kelly, 2003]. Most of these works have the commonality of being rule-based: they rely on building a set of rules that will guide the visualization synthesis process. In our case we do not build rules before hand and rely on experimental evidence to build our knowledge of expressiveness and effectiveness for each of our visual dimensions. Our model will provide a dynamic representation of our space of visualizations so that new knowledge learned during its use will require only a local refitting of the model.

Indeed, Johnson, in his latest list of top scientific visualization problems [Johnson, 2004], recognizes the quantification of the effectiveness of visualization methods as one of the major research areas in this field. He also included perceptual issues, multi-field visualization and theory of visualization, all areas that we are addressing in this thesis. [Tory and Moller, 2004] also talks about some of the new challenges the field of visualization has to tackle and she concentrates on human factors. She includes perceptual measurement of effectiveness and the need for a formal modeling of these to really anchor the field and move forward. We believe we are modestly addressing those issues in this thesis and we will make a valuable contribution to the visualization field.

Our contribution will become a practical application in the area of multi-valued visualizations. [Weigle et al., 2000; Taylor II, 2002; Bokinsky, 2003] are works closely related to ours in the sense that they deal with scalar fields in 2D and try to develop new techniques to display them effectively. They heavily rely on experimental evaluations to validate their techniques, but they do not explore what are the fundamental expressive characteristics of the visual dimensions that form their visualizations. They build up their techniques based on previous work and their own experience, obtaining valid results and developing a layering technique that lets them present effectively multiple variables simultaneously.

Finally, our work has also some similarity to the research on multiple surface visualization [Interrante et al., 1997; House and Ware, 2002; House et al., 2006]. These approaches try to visualize a 3D object by placing glyphs on its surface. Since the point of view of the object is expected to change, these glyphs do not usually change based on the characteristics of the surface. In our case, the overall textures created by mapping a scalar field to visual dimensions could create the illusion of a surface and might actually be representing one, but the icons are used to indicate the very values being explored.

Chapter 4

Experimental Evaluation of Perceptual Interactions

In this chapter we will go, in more detail, through the three experiments we have performed so far in this dissertation. We will present the results obtained and explain how they motivate our research plan.

The first study was a comparison of 2D vector visualization methods using visual designers as subjects [Jackson et al., 2003]. The results from this experiment encouraged us to develop a methodology to evaluate scientific visualization methods using expert visual designers. This led to our second study where expert illustration educators evaluated multiple 2D scalar visualization methods [Acevedo et al., 2005]. In this latter experiment we began including visualization methods showing two scalar fields simultaneously. We were trying to evaluate how the different visual dimensions interacted. The results led to the third experiment, where a design inspired more on psychophysical studies allowed us to quantify directly the perceptual interactions among three visual dimensions [Acevedo and Laidlaw, 2006].

4.1 Experiment 1: Designer Critiques of 2D Vector Visualization Methods

This pilot study examined how graphic design knowledge could provide a fast and robust visualization evaluation methodology, one that assessed scientific visualizations for their scientific value while also improving the design and composition of the visualization methods. Since graphic designers, particularly illustrators, are trained to judge how well visual designs convey specific pieces of information, it was our hypothesis that they could evaluate scientific visualization methods for how well they fulfill design goals based both on the scientific task represented and the actual visual design.

In a previous user study [Laidlaw et al., 2005], several 2D vector field visualizations were quantitatively evaluated, where three tasks were defined to understand the flow in a bounded 2D vector field. The three tasks were counting the number of critical points in the flow, identifying the type of critical point in one of them, and performing an advection task where they had to follow a particle moving through the flow. For each task, performed on the computer, the accuracy and execution time for each observer were measured. The results showed that particular visualization techniques were more suited to certain tasks than others. Details of that study are beyond the scope of this dissertation, but the reader is encouraged to read the above mentioned publication for a detailed explanation of the methods and tasks used.

4.1.1 Our Approach

Utilizing the same flow visualization methods, our study had designers grade scientific visualization methods based on their subjective estimates of user performance for certain tasks, and give verbal feedback (i.e., critique) on the effectiveness of each method for fulfilling the same three tasks. We hypothesized that designers would rank the methods similarly to the objectively measured task performance from the original experiment.

Having the critiques would also enable us to understand why methods work well and to synthesize better visualization techniques. This was the first experiment where we hoped to identify which elements of these methods worked best for the given tasks. We hoped to combine these elements, or dimensions, and create new visualization methods that used those to allow better performance in all tasks.

To have valid comparative data, we ran our six subjects, all graphic and visual designers with various levels of experience, through the original computer-based user study. This also served as a training session so they could familiarize themselves with both the tasks and the visualization methods.

After completing the first study designers were shown all 6 different methods to visualize 3 different datasets. This was done on paper, as shown on Figure 4.1. Different datasets were shown for each of the three tasks. Subjects had to subjectively rate the methods based on how accurate they believed a real user would be performing the task at hand. They also had to rate the methods for how long they felt a user would take performing the

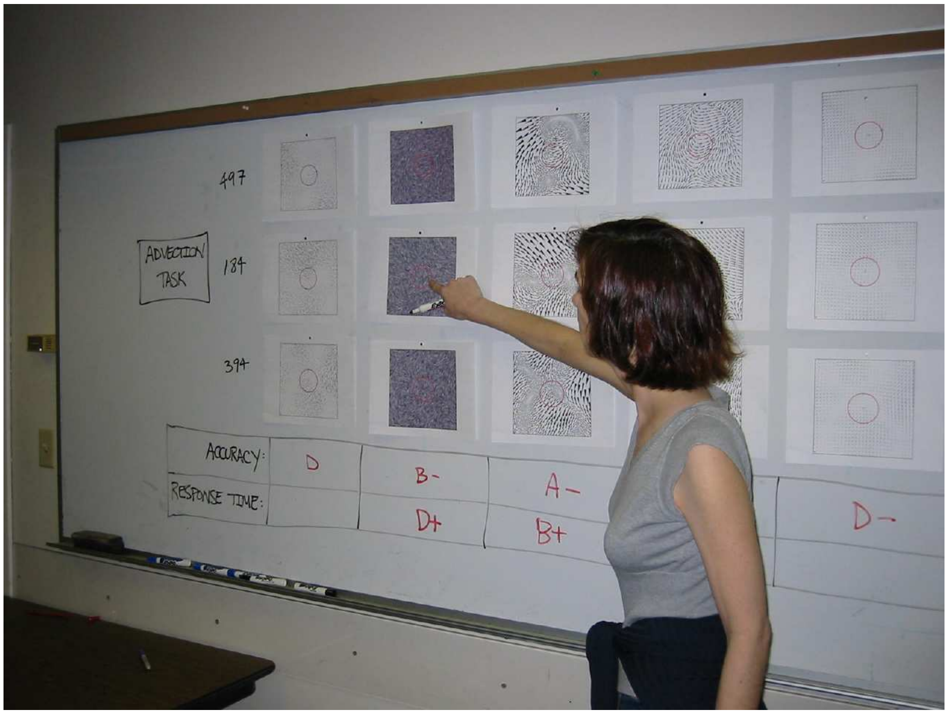

Figure 4.1: During the study, designers rated the different methods subjectively, based on accuracy and time to perform the task. They were also able to comment on how the visual elements used on each method would affect their performance. Subjects were also asked to rank the methods in order so we could correlate their perceived order with the accuracy and time ratings.

task. Showing three datasets per method allowed them to eliminate the possibility of an easy dataset making the method perform better. This was also avoided when they ran the full computer-based study, since they actually saw nearly 300 different datasets in sequence and could gain an overall sense of performance for each method.

4.1.2 Results and Discussion

Overall, designers rated the visualization methods in a pattern similar to what was found from the more objective and quantitative experiment. In the original study, seventeen subjects, twelve students and five flow experts, performed the almost two-hour long computer-based experiment. Our five subjects took much less time and got comparable results.

Not only that, they provided us with many insightful pieces of advice about why methods performed like they did and how to improve them. They even created new methods, like the one in Figure 4.2, that would overcome most of the issues they found and perform well on all the given tasks.

Figure 4.2: After the experiment, designers were asked to create a whole new visualization method design that would outperform the six methods they saw during the experiment. This image shows an example of one of the results. Black flow lines help on the advection task, where white marks indicate direction of the flow. Critical points are clearly marked by big white dots, and the type of critical point is indicated by a series of arrows around their locations.

These results were encouraging but left some open questions. Although the tasks are relevant scientifically, designers seemed to have a very easy time evaluating them. The concepts they used during the critiques were very basic, even though all six methods we used were state of the art for 2D flow visualization. Basically, the crit task did not seem hard enough to tax the designers' brains. We believe we could get much more out of them if we could design an experiment with harder tasks.

Also, in our experiment none of our subjects were active professionally or were teaching design courses. While the initial study found no differences between experts and non-experts performing the quantitative tasks, our subjective tasks may however create some differences between subjects with different levels of expertise.

Finally, another question raised by these results was whether there is a point at all to performing quantitative studies. Our answer is yes, since the ratings provided by designers are more qualitative and it is hard to extract numeric values from them that could be used to design a visualization method.

4.1.3 Lessons Learned

The main take-home message from this pilot study is that designers can effectively evaluate scientific visualizations. They actually provide extra information, like reasons for the good or bad performance of visualization methods, not possible to obtain using subjects from the specific scientific field.

Our conclusion is that, while using designers, combining objective and subjective experiments seems to be the way to go. This would provide us with numeric estimates of performance while, at the same time, provide guidance as to what elements of different visualization methods help or impede the performance in certain tasks. Issues with training and understanding of the scientific goals of the visualization methods will come up using designers, but the payoff seems big enough to overcome those.

Finally, the specificity of the tasks created a situation that did not tax the potential design expertise of our subjects. As it can be observed by the design in Figure 4.2, a task oriented design query yields, naturally, a very explanatory type of visualization method where answers to all three tasks are very explicitly depicted. Our hypothesis is that concentrating on a more holistic task for overall exploration of the datasets will bring out the best from the designers, and the results will be more effective the more expertise subjects have.

A more exploratory task will also provide more broad utilization and explanation of the different visual dimensions used by the methods. Subjects' comments during this study were naturally limited to the capabilities of the different dimensions in relation with the tasks. An exploratory goal will allow designers to think about and compare the full range of visual dimensions available and their different characteristics. To make the experiments possible, a more formal definition of the visual dimensions and their interactions was needed. This was partially realized in our next experiment.

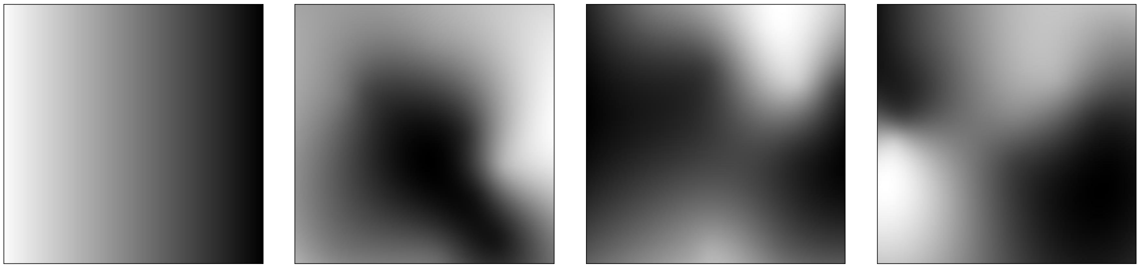

Figure 4.3: During the study, subjects were presented with multiple visualization methods representing these four single-variable datasets. The first one is a linear dataset, while the rest are general, continuous and smooth changing height fields.

4.2 Experiment 2: Evaluation of 2D Scalar Visualization Methods by Illustration Educators

Based on the conclusions from the previous study, we developed the previously introduced methodology to evaluate how individual visual dimensions perform when representing a single scalar dataset in 2D. The formalization of the visual dimensions and the design factors was explained in the previous chapter, but this was the first experiment where we introduced our parameterized space of multi-layered multi-valued scientific visualization methods and the language to define it.

In this particular study we had five subjects, all expert educators from the Illustration Department at RISD, evaluating 33 different visualization methods. The number of methods comes from 11 visual dimensions (see Figure 2.10) and 3 different parameterizations for each one. To follow our nomenclature, for each visual dimension v_i, we generated 3 single layer methods $v \in \mathbb{V}$ in the form:

$$v = \{l_0\} = \{(0, 0, ..., d_i, ..., 0), (c_0, c_1, ..., (b_i, e_i), ..., c_{11})\}$$

where each method has a different mapping $((b_i, e_i))$ of data variable d_i. The datasets were continuous, smoothly changing, scalar fields in 2D. Following what we did for the previous study, we tried to eliminate the dataset from the analysis by showing 4 different single variable datasets, shown in Figure 4.3.

In summary, we created a total of 132 images that we printed and placed on the wall for our subjects to evaluate. The setup is shown in Figure 4.4.

The evaluation was based on 6 different design factors that provide a characterization of the exploratory nature of our visualization methods while, at the same time, allow for numerical analysis of the experimental data. In this sense, the factors we define here provide

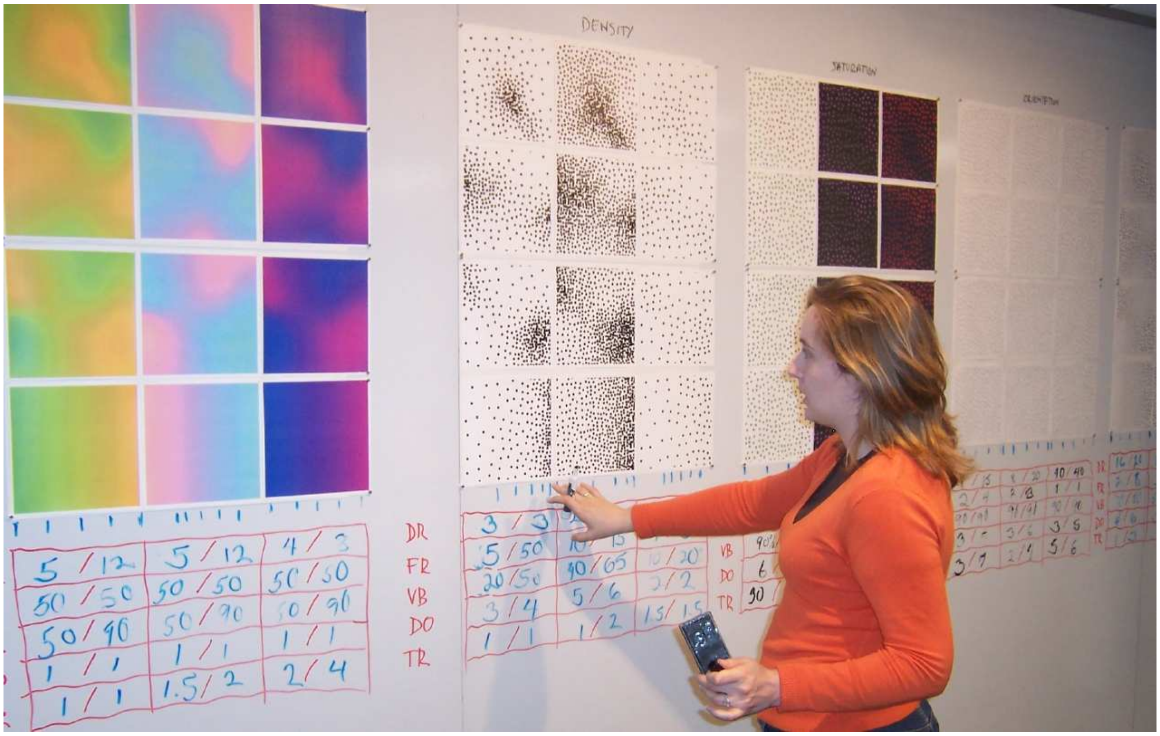

Figure 4.4: A subject in our study critiques visualizations of 2D datasets with a single scalar variable. Illustration educators were shown a total of 132 visualizations corresponding to 3 parameterizations of 11 visual dimensions and 4 different datasets. For each parameterization, they evaluate all 6 of our design factors (bottom of the image).

information about the quality of the data presented and the capability of a visualization method to work in combination with other methods. For this study the factors were:

- *data resolution:* the number of different levels of a data variable that can be distinguished by a viewer;

- *spatial feature resolution:* the minimum spatial feature size that can be reliably represented with a method, expressed as a percentage of the image width;

- *visual linearity:* the perceptual linearity of the mapping from data value to visual dimension; this factor is measured by asking subjects to indicate the locations where they see the values of 0, 0.25, 0.5, 0.75, and 1.0 along the image for a linear dataset visualization;

- *visual bandwidth:* the percentage of a method that can be covered when combined with other methods but still remain readable. This design factor is aimed at estimating how different visual dimensions will perform in multi-layer methods;

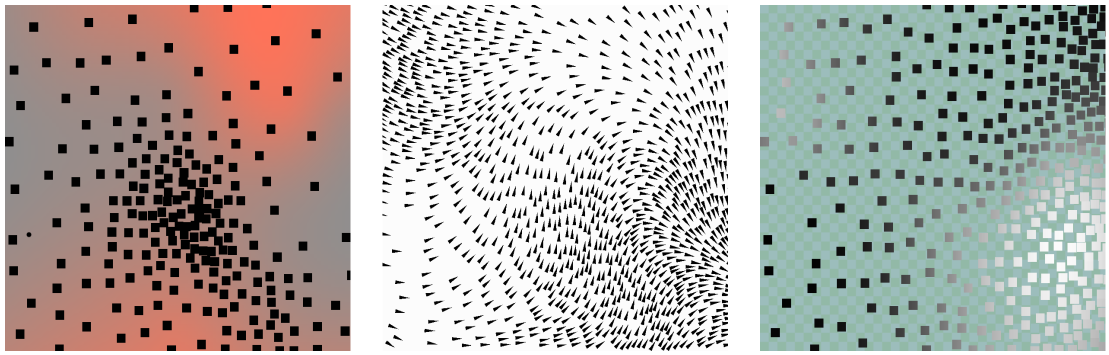

Figure 4.5: For our pilot study, we showed subjects two-variable visualization methods. From left to right: colorplane saturation and icon spacing, icon orientation and spacing, and icon brightness and spacing. Note that, for the purposes of formalization of our model, the combination of an RGBα colorplane and the icons form a single layer.

- *dominance:* the forcefulness or *punchiness* of the data mapping. This indicates how much a method would dominate the composition when combined with other methods, measured as a value from 0 to 10;

- *time to read:* the time it takes an average user to comprehend the data, measured in seconds.

For this experiment we created a novel experimental methodology for capturing quantitative knowledge from visual design experts. This is a clear improvement with respect to the previous study, since we are trying to provide designers a way to convey their critiques through the use of our design factors. During their critiques, our subjects provided three different kinds of information for each design factor: numeric ratings, specific suggestions for directions of improvement, and explanations of their ratings. We videotaped the sessions, which last approximately 3 hours, and we encouraged in-depth explanations of the numerical ratings.

For their training, subjects were introduced to all methods and design factors the day before their critique, when they were given the instructions for the experiment. A webpage [1] was available for them to review all visual dimensions we were interested in, as well as the goal of the experiment and the introduction to the different factors. They were encouraged to familiarize themselves with all parts of the experiment and write down any questions they had. They were instructed not to actually perform the design factor estimations. Before and during the actual crit sessions, subjects were allowed to ask any questions or make any comments about the study.

[1] http://www.cs.brown.edu/people/daf/study

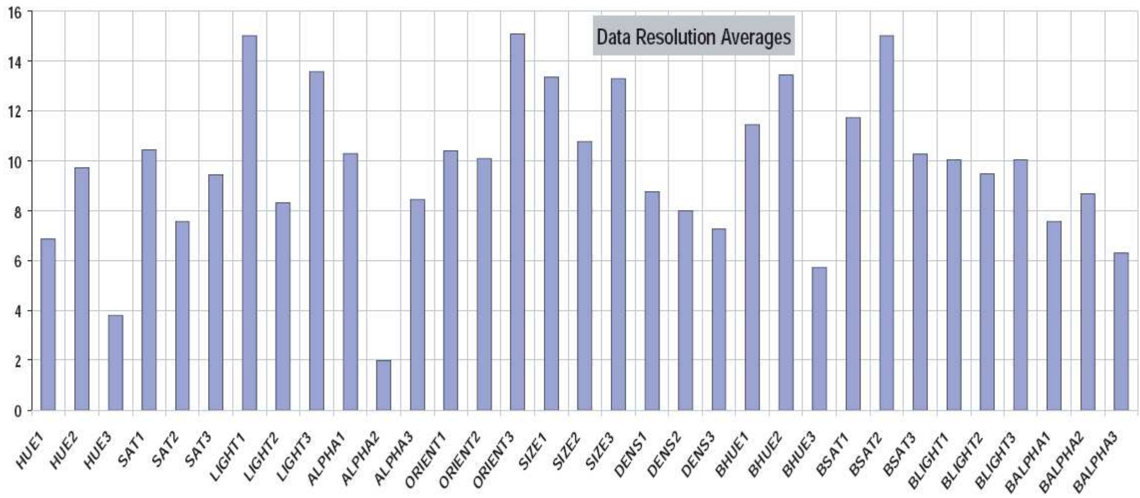

Figure 4.6: Results for the Data Resolution deign factor. All five subjects characterized all 33 methods. The standard deviation for these is around 5 levels.

Apart from these single-variable visualization methods, we also began our exploration of how combinations of visual dimensions work together. In this case, our subjects critiqued and rated combinations of visual dimensions to capture the ability of cue combinations to represent complex relationships within multi-valued data sets. The images on the Figure 4.5 show two examples of visualizations of two-variable datasets.

4.2.1 Results and Discussion

The results of this study allowed us to characterize the effectiveness of individual visual dimensions for visualizing single-variable scalar datasets for exploratory visualization. This was something we wanted to capture during our first experiment but failed to do so due to the tasks chosen and the non-uniformity, in terms of the visual dimensions utilized, of the six flow visualization methods used. Separating each participating dimension and exploring them in comparison with all others allowed experts to concentrate on their individual strengths and weaknesses.

Figure 4.6 shows the results for the data resolution factor where the average number of data levels visible for all 33 methods are shown (11 dimensions x 3 parameterizations per element.) We have analogous characterizations for all design factors.

One advantage of using design experts to do the evaluations is that they can pinpoint reasons why a method does not work. They commented about the neutrality of the shapes used, the effects of negative space, the distinctive features a method like orientation creates that makes it readable, or their technique of squinting their eyes to perceive the overall composition without focusing on single icons.

The main problem with our results was the high variance we obtained among the five experts. Our hypothesis is that this could be due to their slightly different interpretations of some of the factors. Also, although our use of expert educators, as opposed to students, was based on the expectation that they would be able to abstract themselves better from their personal taste, this is very difficult to achieve.

Given this high variance, trying to make inferences from the two-variable results and how they compare to the single-variable ones would not be useful. Designers commented on the intrinsic value of evaluating those combinations, but trying to hypothesize a formal model of how the different dimensions interact would not be possible with the current experimental results.

Effectiveness Measure

With the characterization of all six design factors at hand, we devised five different scenarios, shown on Table 4.1, setting goals for each design factor. These scenarios represent requirements expressed by potential scientists exploring the data. We measured the effectiveness of each method based on how it fulfills the desired requirements. To do this we followed the schemes introduced in Section 2.4. Figure 4.7 shows the results for one of these scenarios. Note that only three factors are constrained. For the other three factors, linearity can be fixed by adjusting the data-to-visuals mapping, while visual bandwidth and dominance are not required for single-variable examples.

Sample Requirement Scenarios			
	DR	FR	TR
Scenario 1	max	max	min
Scenario 2	DR<3	max	min
Scenario 3	4<DR<6	2<FR<5	min
Scenario 4	-	-	min
Scenario 5	4<DR<6	5<FR<10	-

Table 4.1: This table shows five possible sets of requirements placed on Data Resolution (DR), Spatial Feature Resolution (FR), and Time to Read (TR).

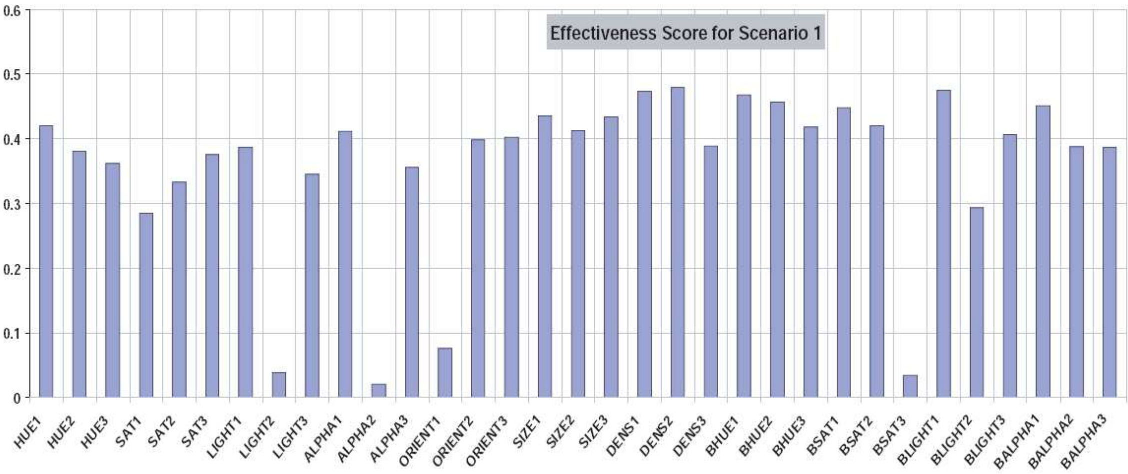

Figure 4.7: Given the requirements of Scenario 1 from Table 4.1, we can measure how well each of the 33 methods are able to fulfill those requirements. A score of 1.0 would be a perfect score.

4.2.2 Lessons Learned

In general, all our subjects felt this is an important and very interesting line of research. None of them were used to making numeric judgments about tasks they perform from experience. They understood our goal of trying to extract that expert knowledge but we felt that, in this case, we over-taxed their brains by making them concentrate on numbers. They enjoyed the freedom of asking questions and explaining their decisions, but the ultimate need for a numerical estimation created problems.

The study setup was also well received. They are used to comparative critiquing in class, and being able to do that helped self-balance the evaluations within each subject's results. This still did not solved the high variance issue, but it made subjects much more confident of their own evaluations.

Finally the length of the study was generally excessive for all subjects. Some of them took as long as 6 hours to complete all evaluations. Even when we moved the training session to the day before the critiques and provided them with an online resource for reference, forcing them to do a continuous critique session probably impacted the results due to fatigue: subjects had to provide a total of 165 numerical estimates plus the linearity ratings.

4.3 Experiment 3: Quantification of Perceptual Interactions

To solve some of the issues introduced by the previous studies we turned to a more psychophysically oriented design for our experiment. Relying on simpler, more perceptual tasks would make it easier to get reliable quantitative data, but it will lead us away from the benefits of using subjective critiques from expert visual designers.

Our goal with this study is to analyze the quality of the data we can obtain and how we can build upon it to accomplish our overall goal of modeling the effectiveness of visualization methods. To that end, we performed an evaluation of a parameterized set of 2D icon-based visualization methods where we quantified how perceptual interactions among icon size, spacing and brightness affect effective data exploration.

During the experiment, subjects quantified three different design factors for each method: the spatial resolution it could represent, the number of data values it could display at each point, and the degree to which it is visually linear. These form a subset of the design factors we evaluated before, but represent the basic factors for single-variable visualization methods.

We devised three different tasks that subjects would have to perform in order to provide us with their evaluations. In the previous experiments we asked subjects to judge how easy or hard it would be for a real user of a visualization to perform a certain task. This is meaningful from the point of view of an illustration expert, but we want to measure the actual perceptual capabilities of our visualization methods. To accomplish this we must test the subject's perceptual system and extract our characterization based on their performance on those tests. These indirect perceptual tasks should make the experiment easier on the subjects but still powerful and generalizable from our perspective.

4.3.1 Methodology

Table 4.2 shows the values we chose for each of the three visual dimensions involved. Icon size and spacing are both measured in pixels. Size indicates the diameter of the circular icons, while spacing indicates the distance between two icons. As mentioned before, we utilize a Poisson-disk distribution to randomly place icons across the image. We experimentally chose the upper limits for size and spacing so we could explore methods with reasonable spatial feature capabilities.

With these parameters we defined six possible value ranges, pairs (b_i, e_i), for each visual dimension: $(0.00, 0.33)$, $(0.00, 0.66)$, $(0.00, 1.00)$, $(0.33, 0.66)$, $(0.33, 1.00)$, and $(0.66, 1.00)$. For icon brightness methods we combined these six ranges with all possible combinations

		values			
		0.00	0.33	0.66	1.00
V_0	brightness	0.00	0.33	0.66	1.00
V_1	size (pixels)	2	5	7	10
V_2	spacing (pixels)	0	3	6	10

Table 4.2: Values used for each of our visual dimensions.

for the other two dimensions, creating a total of 72 visualization methods that we will evaluate. For icon size and spacing methods we kept icon brightness constant at 1.00, so 24 combinations ($6x4$) can be defined for each of those two dimensions. Note that even constraining our experiment to a small number of elements, and only four possible values per element, the number of combinations is quite large: 120 different visualization methods.

Data Resolution Identification Task

For this task we are asking subjects to evaluate how many different levels of the data variable a method is able to represent. Figure 4.8 explains how we created the stimulae for this task. The question subjects must answer is: in what region of the image do you see the sine-wave pattern? Since they are told the pattern will be more pronounced at the top left corner of the image, they just need to place 3 marks to approximately bound the region where they perceive the pattern.

To obtain actual data resolution values we developed the following process. The marks placed by a subject delimit a region on the image where the pattern is visible (see Figure 4.8 b) . The right and bottom boundaries indicate lines where the difference between the extremes values of the sine-wave are last perceived by the subject, i.e. the *just noticeable differences* (*jnd*) boundary. The basic idea to obtain data resolution values is to follow these boundary lines, starting from the top mark, jumping from one level to the next a distance equal to the amplitude at each point. With this process we will also obtain actual values, in the range $(0.00, 1.00)$, for each level identified. Details on how amplitude and wavelength were set can be found in [Acevedo and Laidlaw, 2006].

Spatial Feature Resolution Identification Task

For this task we are asking subjects to evaluate the size of the smallest spatial feature a method can represent. Our approach for this task was to indirectly ask the question by exploring the limits of each subject's visual perception. In this case our datasets are vertical

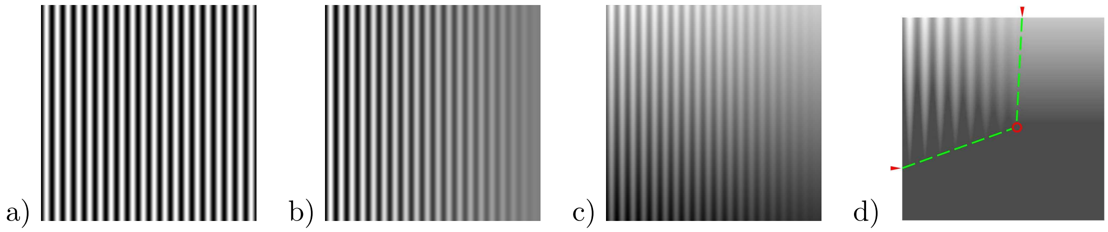

Figure 4.8: Process for creating the stimulae for the data resolution identification task. (a) Shows a vertical sine-wave dataset. (b) Shows the same dataset with amplitude values a linearly decreasing from left to right. (c) Shows the final appearance of the datasets used for this task, where we also linearly move the zero value of the sine-wave from $a/2$ at the top of the image to $1 - a/2$ at the bottom. (d) Shows how subjects would mark the area where they perceive the sine-wave pattern.

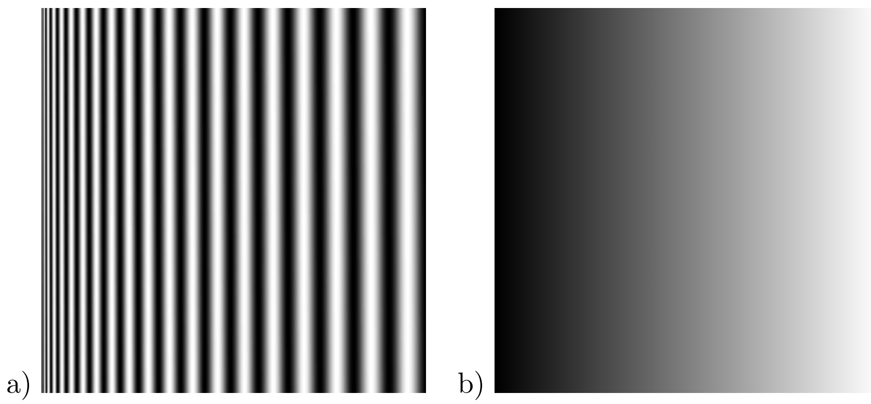

Figure 4.9: (a) Shows an example stimuli for the spatial feature resolution identification task dataset, with wavelength λ linearly decreasing from right to left. (b) Shows the stimuli for the visual linearity perception task.

sine-wave patterns that maintain constant amplitude a but linearly change their wavelength λ from left to right across the image. Figure 4.9 (a) shows an example of this dataset using brightness values from 0.0 to 1.0 ($a = 1.0$).

By asking subjects to place a mark when they stop perceiving the sine-wave pattern, we are obtaining our minimum feature size measurement. $\lambda/2$ at that point will be the minimum spatial feature a method can represent. For this task we use all 120 visualization methods mentioned before. The amplitude for each display is indicated by the range (b_i, e_i). Figure 4.10 (d), (e) and (f) show examples of images used for this task for each of the three visual dimensions.

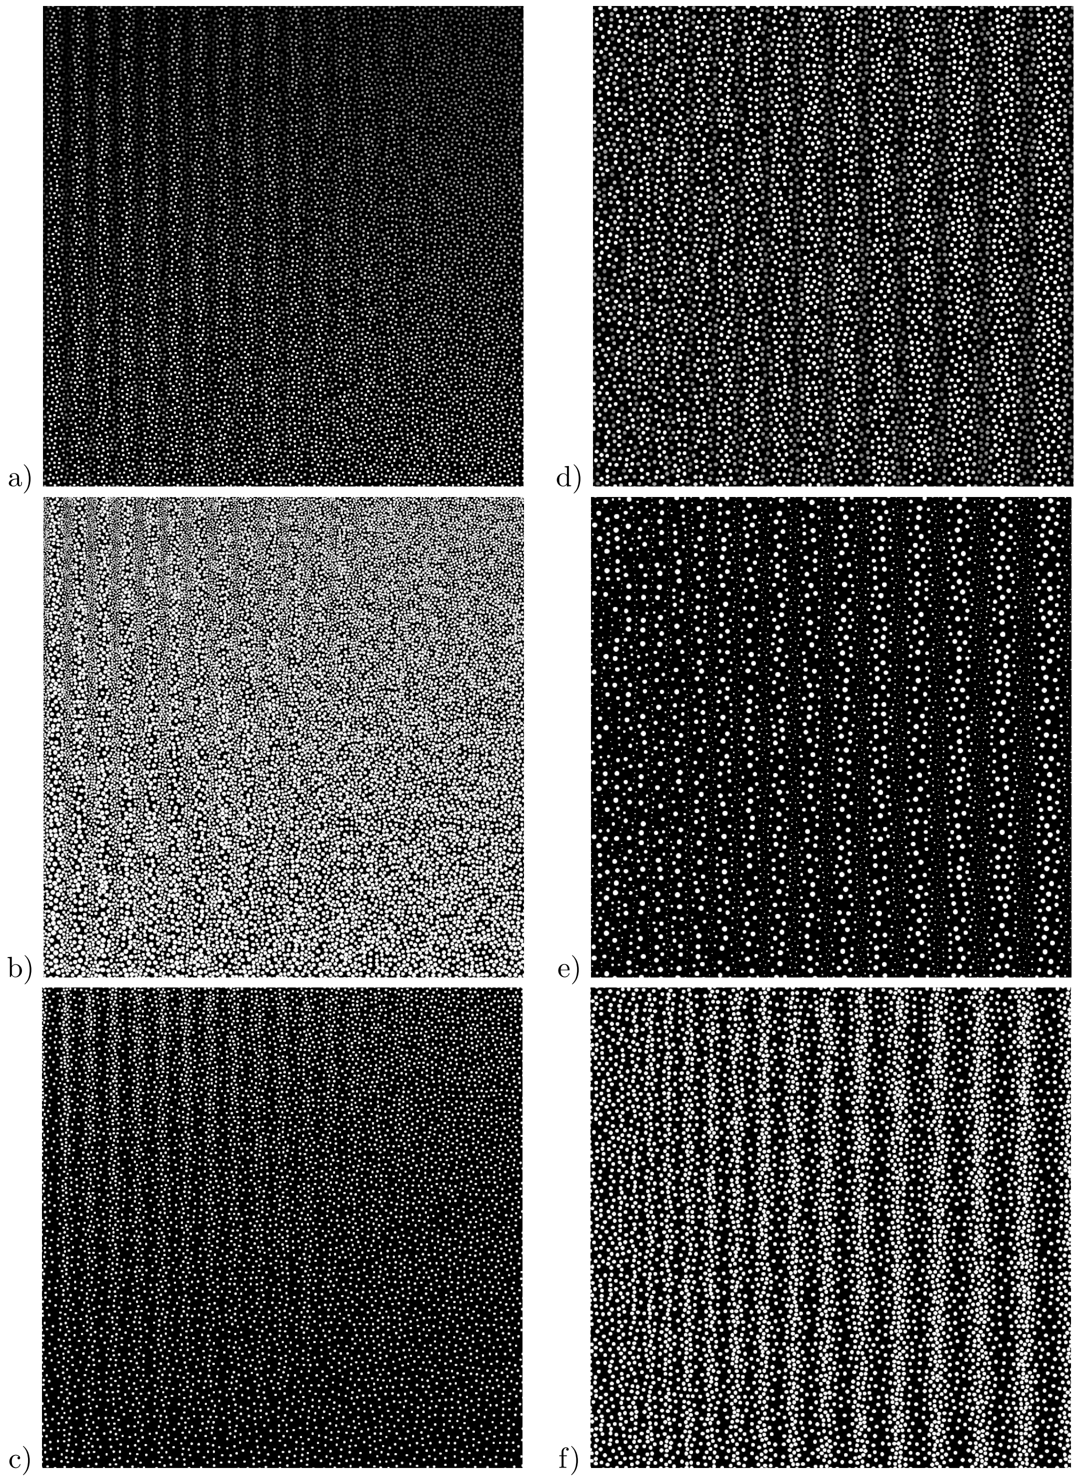

Figure 4.10: Examples of various stimulae used for the experiment. (a), (b) and (c) show data resolution identification stimulae for icon brightness, size and spacing respectively. All of them with $\lambda = 5\%$ and $a = 0.6$. (d), (e) and (f) show spatial feature resolution identification stimulae for the same three visual dimensions.

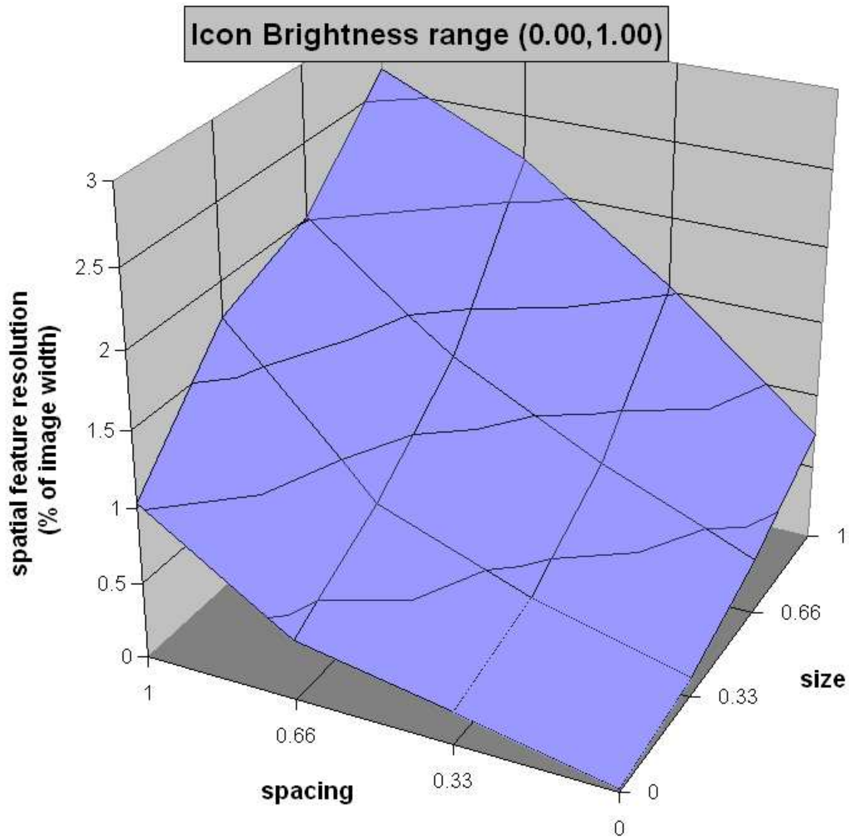

Figure 4.11: Plot of the spatial feature resolution results for icon brightness methods. Note how spacing growth affects spatial feature resolution more than size growth.

Visual Linearity Perception Task

In this task subjects are shown visualizations of a linear dataset that progresses from a value of 0 on the left of the image to a value of 1 on the right edge (see Figure 4.9 (b)). They are told that 0 and 1 are at the very edge of image and are asked to place five marks for the values $0.0, 0.25, 0.50, 0.75, 1.0$. The two extremes would indicate regions where they do not perceive a change in the visualization's border regions. A visually linear method would maintain a constant ratio between data changes and visualization changes.

Experimental Setup

We ran a fully randomized within-subjects pilot study where 6 computer science graduate students performed all three tasks on the computer screen. The full study consisted of nine separate sections (3 tasks x 3 elements) with a training subsection and a real trial subsection within each one. Response time was recorded for the real trials. There was

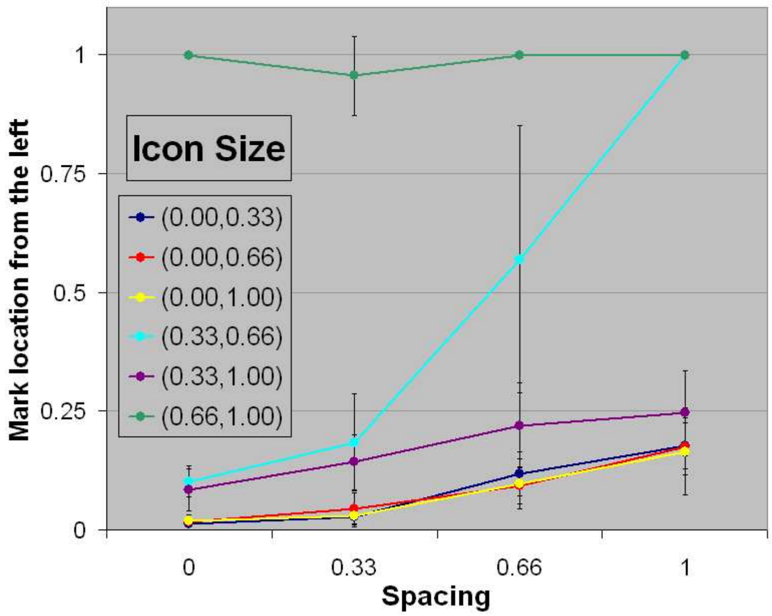

Figure 4.12: Results for all 6 ranges of icon size with respect to the 4 different values of spacing for the spatial feature resolution task. The vertical axis shows the distance from the left border of the display where subjects placed the mark, the closer to zero the lower the spatial feature resolution is (see Figure 4.10(e)). All plots show 95% confidence intervals. Observe how, for the size range (0.66,1.00), almost all subjects reported a "No Pattern" result (plotted as 1 in the graph.) It can be observed that the larger the range the better resolution. For same size ranges, like (0.00,0.33), (0.33,0.66), and (0.66,1.00), the smaller the size the better results. Equivalent plots for spacing versus size show symmetric effects.

no time limit during any part of the study, although subjects were instructed to proceed as quickly and accurately as they could. Subjects took an average of an hour and forty minutes to complete the whole study and were paid for their participation. Subjects were given written instructions before each task. Stimulae for all tasks consisted of images of size 900x900 pixels displayed one by one on an LCD display.

4.3.2 Results and Discussion

Our results successfully characterize the capabilities of each visual dimension, using a variety of value ranges, in combination with potentially interfering dimensions.

Figure 4.11 shows how spatial feature resolution values increase (the actual spatial feature size is measured in percentage of the image width) when icon size and icon spacing values grow larger for the same icon brightness range. The best results, with spatial feature resolutions of around 0.3% of the image width, were obtained for methods that used the smallest icon size (2 pixels) and the smallest spacing (0 pixels). This is reasonable since these methods produce displays that are the closest to simple grayscale displays.

More surprising are the spatial feature resolution values for size and spacing methods. Figure 4.12 shows the results for size. The overall trend is that they have symmetric interaction: larger icon size affects the reading of spacing values in the same way spacing affects the reading of icon size values. The unexpected result is that actual spatial feature resolution values are comparable to icon brightness methods, with the best results being around 0.3% of the image width for small spacing and small size of icons. The explanation for this unexpected good performance of size and spacing comes from the design of our experimental stimulae. Our sine-wave patterns for this task do not change vertically across the image. This produces very strong linear cues that induce subjects to continue perceiving the sine-wave pattern when, locally, there is no clear evidence of it.

Figure 4.13 shows results for data resolution for the same full-range icon brightness visualization methods shown before. The trends are consistent with the results from spatial feature resolution. Maximum data resolution levels are around 21 levels in average, with 95% confidence intervals being ± 5.

The results for size and spacing are particularly interesting since, although they still follow the same trend as expected from the spatial feature resolution data, they are very different in absolute values (see Figure 4.14).

Finally, during the visual linearity task, all subjects reported difficulty completing the task. They easily placed the marks for the extreme values but they could not judge, in general, the 25% intermediate differences we were asking them to indicate, especially for icon brightness methods. Subjects also complained about possible inaccurate gamma calibration of the monitors used. We need to further explore this task and reimplement this portion of the experiment. It is still worth noting that practically all methods, for all three visual dimensions, exhibited clear constant-value areas for the extreme values, sometimes as large as 30% of the image width. This is consistent with our data resolution values where subjects indicate no *jnd*'s for those ranges.

4.3.3 Lessons Learned

With this experiment we solved the high variance issue we had during our pilot study with expert visual designers. The cost for this is two-fold. First, subjects evaluated methods one at a time, avoiding direct comparisons among displays that were possible in the first study. Comparative critique is a very useful tool design educators utilize constantly, but

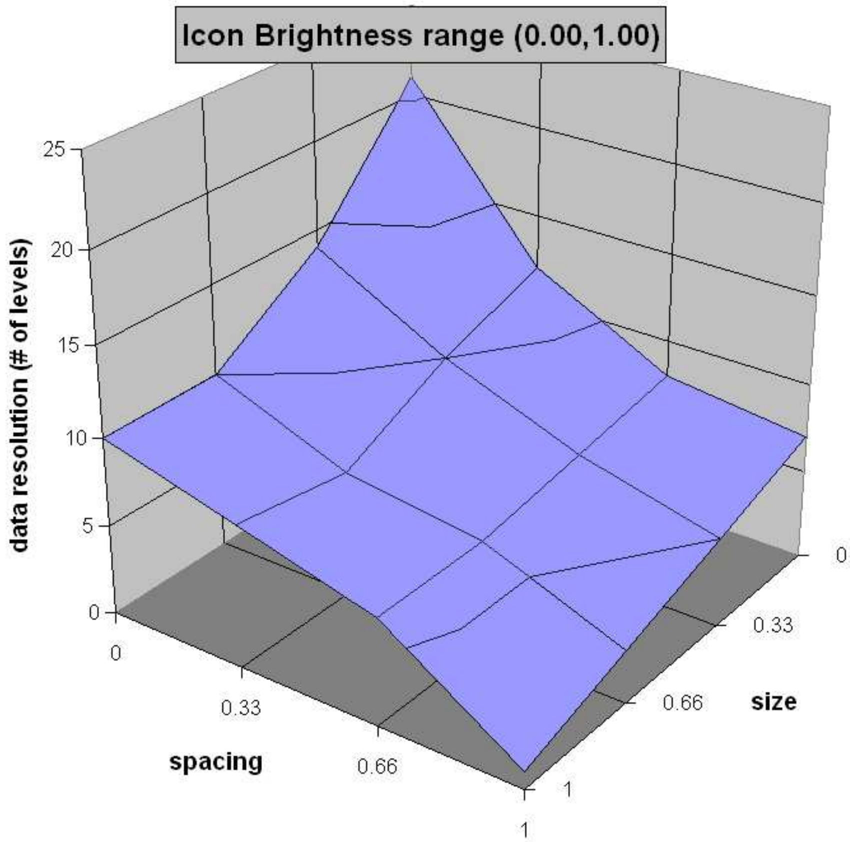

Figure 4.13: Plot of the data resolution results for icon brightness methods. Consistent with the results for spatial feature resolution (the direction of the size and spacing axes is reversed with respect to Figure 4.11), spacing growth affects data resolution more than size growth. This plot corresponds to $\lambda = 5\%$, which gives the highest data resolution values.

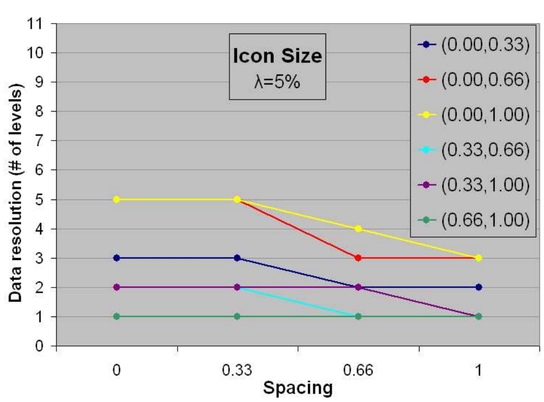

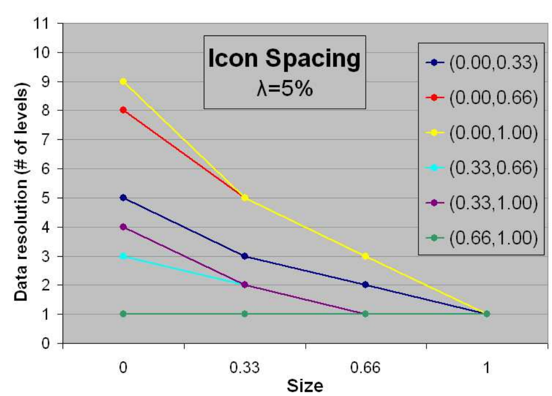

Figure 4.14: Results of the data resolution task for size and spacing. Both plots correspond to $\lambda = 5\%$. 95% confidence intervals are around ± 2. Observe how increased values for the distractor variable decrease the data resolution results consistently across different ranges.

one that we had to sacrifice to improve the quality and quantity of data obtained. Secondly, we did not use expert visual designers as subjects, so we could not expect feedback on why a method performs as it does for a given task. Our tasks now are more perceptual than conceptual and the low variance of the data, along with consistent trends, validates our choice of non-expert subjects.

Chapter 5

Proposed Work and Expected Results

Our first priorities after obtaining these encouraging results will be to increase the number of visual elements involved, including color saturation and orientation, and to perform measurements of relative saliency among all five dimensions. This will require moving to a between-subject design to avoid fatigue when running the experiments. Once we have more data about how the different dimensions interact, we will begin defining hypotheses for higher order combinations that we will evaluate using expert visual designers again.

Finally, and after defining our optimization strategy, a final study will evaluate how, utilizing real datasets, our model is able to search and present effective combinations of visual dimensions to visualize multiple data variables simultaneously.

5.1 The Next Experiments

5.1.1 Experiment 4: Icon Orientation and Icon Color Saturation

We will complete the previous study by including icon orientation and color saturation. The tasks will be the same as explained before, and we will use the values shown in Table 5.1.

We will not evaluate visual linearity explicitly this time. During our experiment with expert visual designers, orientation proved to be the most consistently evaluated dimension of all. In particular, the results showed a remarkable perceived linearity of the visual displays. Saturation, on the other hand, did not have the low variance results orientation did, but we can still *fix* its visual non-linearity based on previous psychophysical studies.

Also, for the spatial feature resolution task, we will limit the stimulae to only a horizontal

	values			
	0.00	0.33	0.66	1.00
size (pixels)	2	5	7	10
spacing (pixels)	0	3	6	10
orientation (degrees)	0	45	90	180
brightness	0.00	0.33	0.66	1.00
saturation	0.00	0.33	0.66	1.00

Table 5.1: Values used for each of our visual dimensions. Note that we non-linearly map the real values for each dimension to the range $(0.00, 1.00)$.

band of the full display shown in the previous study. This will both limit the over-estimation issue encountered before and ease the *apparent flow* effect experts described when looking at orientation displays. Our subjects noted that, even though they were able to abstract from it and concentrate on the scalar changes indicated by the oriented icons, the appearance of flow distracted them enough to rate the method poorly in general.

Finally, these two new dimensions will only be evaluated against icon size and spacing, the same way icon brightness was. For the cases when orientation would be a constant, our previous results should hold true, since it would amount to a different shape being used. Meanwhile, icon saturation should not be used in combination with icon brightness on the same layer, so we are not evaluating those cases either.

This is expected to be a fairly straight forward continuation of the previous study, with the new data obtained being used to complete our baseline for the rest of the project.

5.1.2 Experiment 5: Saliency Evaluation

For this study we will measure saliency differences between pairs of methods. Each visualization method will display a linear dataset as shown in Figure 5.1. Subjects will indicate which dataset they perceive first through a continuous scale, also shown in Figure 5.1. With this very simple setup, subjects should be able to go through many different pairs of methods in a reasonable amount of time, which will help ease the fatigue problems we had in previous experiments.

Apart from the five visual dimensions we are concerned with, layering and shape will play an important role in this experiment. We will evaluate how saliency is perceived when

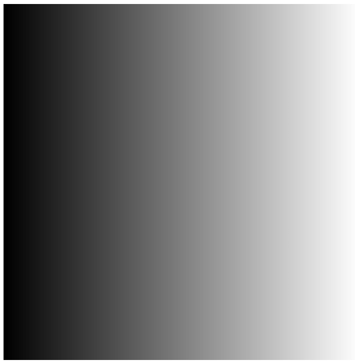

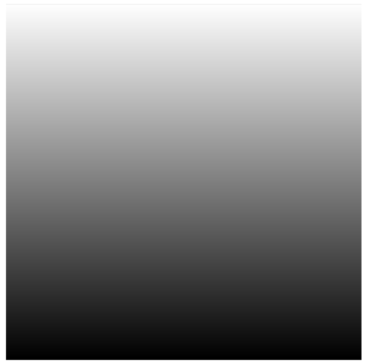

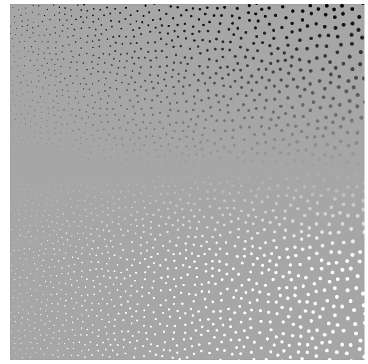

(a) (b) (c)

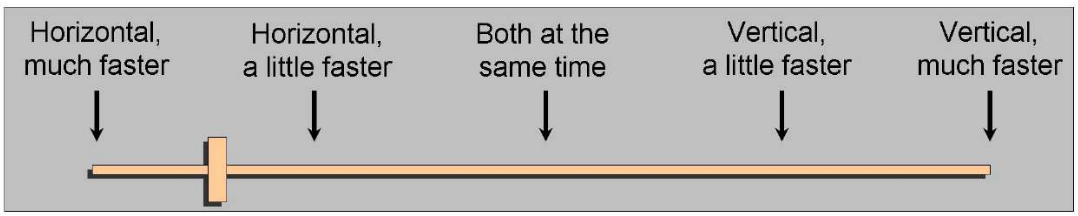

Figure 5.1: Images (a) and (b) show the two datasets used for this saliency evaluation experiment. Image (c) shows an example stimuli generated using icon size and brightness. Subjects will answer which dataset they perceived first using the scale shown.

two layers of icons are present. This allows us to also evaluate saliency within a single visual dimension by both using a different shape icon or changing the effective spatial resolution for both layers. Our previous experiments will inform the capabilities of each layer, while this one will inform the visibility and preattentiveness characteristics for each combination.

In this study the exponential nature of our exploration of the space of visualization methods finally catches up with us. If we wanted to explore the full set of pairwise combinations during this study, including two-layer combinations, the total number of stimulae needed would be $53,748$, which would take a little more than 74 hours to evaluate assuming an average of 5 seconds per image. This is obviously out of the question, so we will divide the experiment into two parts.

The first part involves all the possible combinations for single-layer visualization displays. This excludes combinations of icon brightness and icon saturation, since those are dimensions nearly impossible to distinguish even for the trained eye. The number of stimulae for this first part of the experiment is only $1,262$, which makes it doable in less than 2 hours.

The second part of the experiment will evaluate two-layer combinations, but reducing the number of values each dimension can take. From Table 5.1 we will evaluate only $(b_i, e_i) =$

$\{(0.00, 0.66), (0.00, 1.00), (0.33, 1.00)\}$ for mapped dimensions and $c_i = \{0.00, 0.66, 1.00\}$ for constant dimensions. This generates a total of $5,022$ stimulae which should take approximately 8 hours to evaluate. For this part, the ordering of the layers matters when making saliency judgments, so the actual number of stimulae is actually doubled to $10,044$.

We will design the whole experiment as a between-subjects study, with a minimum of 12 subjects needed to fully evaluate all stimulae at least once (assuming 1.5 hours per subject). We will pilot this design to test for timing and fatigue issues, and whether this reduction of combinations is enough for a doable experiment.

5.1.3 Experiment 6: Two-Variable Datasets

The goal of this experiment is to bring back expert visual designers and ask them to critique visualization displays with two scalar fields present. We will sample among the more than $53,000$ combinations mentioned before and, based on the results obtained so far, hypothesize the evaluations for our design factors. We will compare these hypotheses with the subjective critiques offered by the experts but, more importantly, we will gather quantitative information about how the combinations might be improved by modifying the dimensions' parameterizations.

These modifications will be recorded through a combination of video from the critique sessions and a set of interactive sessions where experts will have direct control over the actual visualization methods. One of the issues with extracting numerical data out of verbal explanations is the lack of a specific protocol that could be applied and filled out as the subjects explain their decisions. The interactive sessions will help generating the guidelines for a protocol that we can use for analysis of the videotapes.

Examples of the type of images we will use are shown in Figure 5.2. For this study we will use recognizable datasets, obtained from grayscale photographs of every day life scenes. We will match the methods used to the characteristics of the datasets (spatial resolution and data resolution). Recognition of the scenes and time taken to recognize them will be factors recorded during the study.

In our first experiment with visual design educators they complained about the lack of knowledge of what it was they were supposed to be looking at. We showed them general, low frequency scalar fields, but they did not know whether the hills and valleys they perceived were supposed to be there or not. Using this database of natural scenes we are trying to correct that and improve the subjects' confidence. Nevertheless, these datasets present issues of their own due to subjective bias or attraction to certain kinds of images that is

Figure 5.2: Samples of the two-variable visualizations used in the study. All three images show the same two datasets. The datasets used for the study are obtained from grayscale photographs of recognizable scenes. The boat dataset has a higher average spatial frequency than the softball scene, hence its practical invisibility on the saturation and spacing visualization on the left. Meanwhile, in the last two images, the saliency of icon brightness, combined with smaller icon size, allows for faster recognition of the two datasets by flipping the mapping between brightness and spacing.

very difficult to measure and control. To avoid this we will also experiment with abstract height fields similar to the ones from experiment 2 but with spatial feature sizes closer to the expressive ranges of our visual dimensions. To present the ground truth we can use grayscale images since, in this case, we are not using that dimension in our experiments and it will not produce any learning effects.

Expected Results

During their interactive transformation of the methods presented, we will give the expert visual designers the goal of making both datasets as visible as possible. The goal will basically be to fix the saliency of both methods so a user could perceive both datasets as best as possible. We will use the original method parameterizations presented, along with the path towards equal saliency, to evaluate our initial hypotheses based on the results of our previous studies.

This new quantitative evaluation will provide us with the values for each dimension that are best to represent a certain type of dataset, described in terms of data and spatial resolution. Since these values will be obtained when a second dataset is present, we expect them to be quite different from the first results we obtained in Experiment 3 (Section 4.3). In that study we obtained spatial and data resolution characterizations for a constant distractor, while here the distractor is changing based on a second data variable.

The types of displays generated here are closer to what real dataset visualizations will be like. We are already moving away from the previous psychophysical study type of stimulae and presenting subjects with more complex situations to evaluate. We believe the expertise of our illustration educators will be critical at this point.

5.2 Modeling Strategy and Optimization

The previous experiments are meant to provide us with a solid baseline of the expressiveness of many different visualization methods and their combinations.

For each pair of methods shown in Experiment 6, the final result indicates a parameterization of those methods that shows both variables effectively (data resolution,w_0, and spatial resolution,w_1) and with equal saliency (w_2). Let us name the two resulting methods v and v', which are displaying data variables d_0 and d_1, respectively. The following table shows the full results from that experiment explained:

$w_0(v\|v')$	Data resolution value for d_0 given the v' distractor
$w_0(v'\|v)$	Data resolution value for d_1 given the v distractor
$w_1(v\|v')$	Spatial feature resolution value for d_0 given the v' distractor
$w_1(v'\|v)$	Spatial feature resolution value for d_1 given the v distractor
$w_2(v', v) = 0$	Relative saliency (zero, given the task of the experiment)

Table 5.2: These are the results we will obtain from Experiment 6.

From Experiments 3 and 4 we have the following results, for the same methods, v and v'.

$w_0(v\|v'[constant])$	Data resolution value for d_0 given the v' distractor is constant
$w_0(v'\|v[constant])$	Data resolution value for d_1 given the v distractor is constant
$w_1(v\|v'[constant])$	Spatial feature resolution value for d_0 given the v' distractor is constant
$w_1(v'\|v[constant])$	Spatial feature resolution value for d_1 given the v distractor is constant

Table 5.3: These are the results we will obtain from experiments 3 and 4.

The mapping of the distractor method in Experiment 6 uses a range (b_i, e_i) where both the minimum and maximum value has been evaluated as constant distractors in Experiments 3 and 4. We can fit a mathematical model to our results from Experiment 6 as a function of the results from Experiments 3 and 4 in the following way:

$$w_i(v|v') = f(w_i(v|v'[constant]))$$

For example, if we choose to model this as a linear combination, given that we have four different constant values, our model would look as follows:

$$w_i(v|v') = \alpha_0 w_i(v|v'[c_0]) + \alpha_1 w_i(v|v'[c_1]) + \alpha_2 w_i(v|v'[c_2]) + \alpha_3 w_i(v|v'[c_3])$$

We can fit the values of $\alpha_0, \alpha_1, \alpha_2, \alpha_3$ to our data. More complex models are possible, but a linear combination is a good initial step.

With this model in place, we can go back to the initial combinations shown in Experiment 6 and obtain the expected values for w_0 and w_1. Since we have also the final results of the interactive session, we know how a change in the visual dimension mapping affects all three design factors. Following our original nomenclature, we have:

$$\frac{d\mathbb{W}(v,v')}{d(v,v')} = (\frac{dw_0(v|v')}{d(v,v')}, \frac{dw_0(v'|v)}{d(v,v')}\frac{dw_1(v|v')}{d(v,v')}, \frac{dw_1(v'|v)}{d(v,v')}, \frac{dw_2(v,v')}{d(v,v')})$$

Optimization Process

Our goal is to be able to, given a set of requirements for the visualization of multiple scalar variables, search for the most appropriate set of methods that fulfill those requirements. For any given number of variables, we can proceed in the following way:

1. For every pair of variables, our previous model will provide us with a set of methods that will fulfill our requirements pair-wise.

2. The derivatives described above can now guide a search algorithm to areas of the space where all design goals are fulfilled as best as possible.

3. Given enough cases of multi-valued visualization methods for more than two variables, we can begin building a similar model to the one we have for pair-wise interactions

4. Apart from the the requirements on our three design factors, users can provide weights for each constraint. We will use those weights to control how far from the actual requirements we can go when not all of the requirements can be fulfilled.

Effectiveness calculations can be done as the process continues, so that a final report will contain a ranking of possible visualization methods and their performance. This will, in effect, create an effectiveness map for our space visualization methods. It is also possible to annotate the resulting combined methods with instructions on how modifying each dimension will affect the overall expected effectiveness.

5.3 Evaluation on Practical Applications

We believe that computational tools must be directed toward driving application areas, which serve as a source of requirements and as evaluation platforms [Brooks, 1996]. To that end, our visualization design factors are targeted at the specific application area of remote sensing data visualization. Additionally, our experience with 2D computational fluid dynamics visualization [Kirby et al., 1999; Laidlaw et al., 2001, 2005] will enable us to use fluid simulation data to test our model during its development. Both areas share a need to display many data values at each point in a 2D plane.

Our driving application area is multi-valued geographic remote sensing using Synthetic Aperture Radar (SAR) imaging. This technique is used to gather information about how large areas of the Earth are used for agricultural and forestry purposes. Scientists using SAR face the challenge of incorporating and understanding the relationships among many different wave parameters; conventional visualization methods do not capture the multi-valued characteristics and complex relationships characteristic of this type of data [Woodhouse et al., 2002]. This scientific domain will provide explicit design goals for generating visualizations and a validation ground to test the results from our optimization process.

5.3.1 Remote Sensing

Remote sensing provides an unparalleled means of obtaining information over extensive areas of the Earth, often providing information of a type unobtainable by other means. Currently, remote sensing is used for environmental monitoring and resource management by a spectrum of disciplines including meteorology, oceanography, geology, vulcanology, agriculture, forestry and disaster management.

Recent advances in the development of airborne and spaceborne SAR imaging systems have permitted the measurement of surface characteristics that were previously unobtainable. However, because of their high data acquisition rate and unique ability to measure many wave parameters simultaneously, SAR instruments can generate large volumes of data during a single overpass of the test site; these data are impossible to represent simultaneously using conventional visualization methods. To exploit such data fully, and to permit efficient extraction of relevant information, requires the development of suitable analysis techniques that are both informative and straightforward. Existing commercial software packages designed for the analysis of remotely sensed imagery are focused around red-green-blue color composites or hue-saturation-value images, in both cases the information content being limited to three channels.

Many model-based analysis methods have been developed to provide automated land cover classification and target identification. However, this method tends to be prescriptive rather than descriptive, as they effectively interpret the data within the limits of the model used [Cloude and Pottier, 1996]] and lose information in the process. Additionally, the results are rarely presented in such a way as to make them straightforward to visualize - because these methods have concentrated on the physics of polarimetry they have only used conventional display methods or classified images with discrete color schemes.

We will utilize image data available at the University of Edinburgh. This comprises fully polarimetric C and L-band data from the Danish EMISAR airborne instrument over test sites in Sweden, L-band polarimetric and interferometric data from the German E-SAR instrument over a Scottish test site, and NASA-JPL AirSAR data over woodland savanna in Belize. The latter data set was acquired in 2004 and includes image data at three wavelengths, two of which (L and P) are fully polarimetric, and the third is interferometric (C-band, giving information on ground and canopy heights). Ground truth data is available for all of the test sites, including results from two intensive field campaigns in Belize in 2004 and 2005 that recorded 3-D attributes of $>$ 1000 trees over a 1km transect [Woodhouse et al., 1997; Cameron et al., 2005]

The development of new visualization techniques for polarimetric data is particularly timely: the recent launch of ESA's Envisat satellite saw the first permanent dual-polarimetric SAR system in orbit, and by 2006 the Canadian Radarsat 2, the German TerraSAR-X and the Japanese ALOS-PALSAR should be in orbit, all of which will carry SAR systems capable of acquiring fully polarimetric data.

There are several important tasks that scientists need to perform with this type of data:

- Visualize the multi-valued output from polarimetric decomposition models for single and dual-wavelength data with a view to optimizing the usefulness of these analytical methods. The ability to study the correlations amongst these parameters and across wavelength bands, along with the ground data, will provide new insight into the value of these analytical methods for determining useful properties over vegetated areas.

- Identify large-scale variations due to land cover changes in multi-valued images of vegetated landscapes. This will include the identification of spatial patterns within different vegetated areas. Separation of land cover into common parameters such as tree species, tree number densities and species has not proven entirely effective using the simple analytical methods, but some success has been demonstrated when the full polarimetric data is used [Hajnsek et al., 2003; Lee et al., 2004]. Identifying patterns

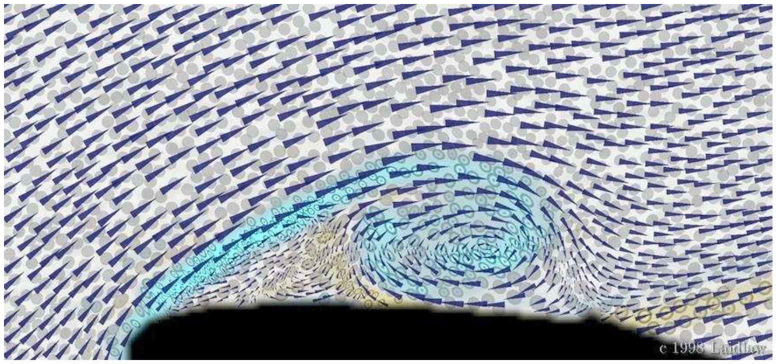

Figure 5.3: 2D flow field visualization. The image simultaneously displays the velocity (two values), the vorticity (one value), and the deformation-rate tensor (three values). The values are encoded, respectively, in the size and orientation of a layer of wedges; in a color base layer; and in the size, shape, and texture of a layer of ellipses [Kirby et al., 1999].

in the multi-valued data will help disambiguate which polarimetric parameters are most indicative of these variations in land cover.

- Identify patterns in high-resolution multi-valued polarimetric data of individual trees (or small groups of trees) to improve understanding of the 3-D scattering properties throughout the height of the tree and to aid species determination from radar data [Fortuny and Sieber, 1999].
- Identify patterns in high resolution multi-valued interferometric data over distributed forest areas to understand the location of the scattering phase centers for tree height determination and thus refine current tree height retrieval algorithms. Using the results of the previous task on individual trees, this task will identify correlations between the multi-valued polarimetric data and the interferometric coherence values.

5.3.2 Computational Fluid Dynamics (CFD)

2D fluid flow provides an excellent scientific area to which we can apply our visualization framework. Fluid flow fields exhibit a continuous set of data points, each of which contains many values for a variety of variables describing the flow. Scientists who study these flow fields want to understand the characteristics of the flow by studying individual flow variables and the complex relationships among them [Urness et al., 2003].

For example, the visualization describing the flow around an airfoil in Figure 5.3 simultaneously shows six values over three variables at each point in the flow: velocity (two values), vorticity (one value), and deformation-rate tensor (three values). The authors [Kirby et al., 1999] found that simultaneously viewing these six values allowed scientists to understand visually the underlying mathematical relationships between the variables describing the flow.

Chapter 6

Discussion and Conclusion

In this dissertation we are developing a framework for the modeling of the effectiveness of visualization methods based upon the quantification of perceptual interactions among visual dimensions. To facilitate this process we also utilize visual design experts, who perform subjective critiques of visualization displays. We have designed a series of experiments that combine psychophysical and visual design components, and they lead up to the practical evaluation of a methodology to generate effective visualization methods for multi-valued remote sensing data.

The encouraging results from the first three experiments described in this document motivates us to continue forward and complete this framework. We are proposing three new experiments to complete this dissertation. The scope to which we plan to evaluate our interactions is one that, although constrained to very few dimensions and to 2D visualization methods, has not been proposed or completed before, to the best of our knowledge. The difficulty and high dimensionality of the problem at hand inevitably makes this thesis a proof of concept for the methodology utilized, but it will provide a solid experimental and computational base for the effectiveness quantification of other visualization methods.

Our contributions include new experimental and computational techniques to evaluate how different visualization methods perform when displaying multi-valued scientific datasets in 2D. This thesis is built around what we consider the four basic elements of visualization: a visualization problem, a set of visual dimensions used to create the solution, a set of design factors that characterize the expressiveness of those dimensions, and a methodology to evaluate the effectiveness of that solution.

Our target visualizations are exploratory scientific visualizations of multi-valued scalar datasets in 2D. Exploration of a dataset means confirming the expected and discovering the unexpected, and this is the visualization problem we are dealing with. At this stage, researchers are interested in gaining an overall understanding of the data as they search for clues that will help them pose the next hypotheses and more specific questions.

For the building blocks of our solutions, we concentrate in the evaluation of perceptual interactions among five visuals dimensions: icon size, icon spacing, icon orientation, icon brightness and icon color saturation. With the data collected in our experiments we build up a model of how the expressiveness of each of these dimensions is affected when they are combined to show multiple scalar variables simultaneously. We use a set of design factors for our visualizations to describe the expressiveness of the visualization methods.

The type of design factors that define the exploratory nature of the visualizations, and constitute the basis of our effectiveness evaluation, include the spatial characteristics of the data variables, the number of different levels of the variables researchers are interested in exploring, and the visual dominance or saliency each data variable should have in the overall visualization display.

6.1 Summary of our Research Plan

Our guiding principle in designing our research plan was to complete our model from the ground up: we must understand and quantify how the individual dimensions that form our visualizations work before trying to approach multi-valued scenarios where perceptual interactions will play a significant role. In following this principle we gather perceptual interaction data as exhaustively as possible for lower order combinations and try to extrapolate and hypothesize results for higher order cases.

For practical reasons, we cannot get to explore even a full set of pairwise combinations, let alone cases where three data variables are represented in the same visualization display. This limitation clearly presents a high risk of having too little data to meaningfully characterize more practical cases where four or more data variables need to be visualized. The use of real remote sensing datasets, and having field experts evaluate the effectiveness of such higher order cases, will provide us with some validation about how well our model works.

A key contribution of this dissertation is the use of visual design experts during our experiments. Our hypothesis is that their critiques of our visualization methods help us model the effectiveness of complex methods better. They accomplish this by providing us with reasons for certain visual dimension interactions, and by indicating how a change in these dimensions would further affect the reading of the data variables.

The capture of those critiques in a quantitative way is clearly a challenge for the project and we have developed techniques to resolve this. For example, interactive sessions in which experts modify visualization methods to accomplish a given design goal help us gather derivative information for our model, i.e. how a change in a visual dimensions produces a change in the expression of one or more design factors.

6.2 Impact of this Dissertation and Future Directions of Enquiry

This dissertation will provide scientists with a first step towards quantifiably effective designs for visualizations of multi-valued datasets given specific design goals. This will allow researchers to concentrate on data analysis instead of visualization creation. Although we will focus on remote sensing data, the generality of both the data and the scientific problems in this field will allow the proposed model to have a broader impact on many other disciplines, such as meteorology and oceanography, and will eventually extend to the space of 3D visualizations.

We recognize that this project is just scratching the surface of the complex and not clearly defined problem of effective visualization design. Our hope is that new lines of research, involving collaboration with visual design and art experts, as well as perceptual psychologists, will develop around the basis we will form with these initial results. In particular, these are some of the directions we would like to follow after completing this thesis:

- *Double Mapping:* The practical limits of experimental design constrained our investigations to a handful of visual dimensions and limited discretizations of the continuous axes defined by each one. Although we have taken into account and explored the different implications of single versus multiple layers of icons, we could not study the case where more than one visual dimensions is mapped to the same data variable. It is recognized in the perceptual literature that synergistic relationships might occur when multiple visual cues are combined, forming emergent features. We would like to explore how multiple simultaneous mappings would affect the expressiveness of high order visualization methods. Indeed, in cases where some visual dimensions in a layer are *free*, maybe mapping them to already mapped data variables would enhance the overall effectiveness of the visualization.

- *Genetic Algorithms:* Given the high dimensionality of the space of visualization methods we are working with, an efficient search strategy is difficult to design. In particular, since our data is based on low order interactions and we are trying to look for effective solutions for higher order problems, the guiding of that search will not be efficient. Our original idea for this dissertation was to implement a genetic algorithm approach for this problem. Each visualization method would be an individual in our population, and the genome would be built with the different visual dimensions we are interested in using. The definition of an evaluation function that could select surviving individuals for each generation led to the current dissertation. We still believe, as do some other researchers in the field [House and Ware, 2002; House et al., 2006], that a GA approach will provide good results and fast exploration of the vast extent of the visualization space. The evaluation function would be based on the interactions and model for them defined in this dissertation.

- *Dataset Contrast:* Given that our datasets are scalar fields in 2D, we could analyze the local contrast between every pair of data variables by using a grayscale representation. Given a set of design goals, we could use local measurements of contrast to see how to optimally fit a single mapping (data to visual dimension) that would maintain the salience requirement across the image. One hypothesis that would need to be evaluated would be that changing the mapping across the image to favor perception would not affect the reading of the data values. In other words, if we know that a certain values of brightness and spacing conflict with each other be decreasing the data resolution of the brightness, we could tweak the brightness mappings in those areas to perceptually maintain data resolution. This is a very complex proposition since it basically means that we are creating a mathematically incorrect (the mapping is not constant across the display) visualization that is perceived correctly.

These are some possible lines of future research that could be followed. Since this dissertation proposes a framework for the initial evaluation of effectiveness, completing and realizing that framework for more visual dimensions would be required in order to pursue these other extensions.

6.3 Conclusion

Acquiring and using expert visual design knowledge and perceptual interaction data at the scope proposed will be intellectually challenging, and will be relevant to various knowledge domains. The new framework we present in this thesis will advance the state of the art in scientific visualization synthesis, as well as in the application areas our results apply to, since data analysis will be improved by the optimized visualizations.

Better visualizations have the potential to advance science more quickly by improving our understanding of physical and biological phenomena, applied science, and engineering. This dissertation will enhance channels of collaboration in education and research among the disciplines of cognitive science, visual design, art and scientific visualization. It will advance our understanding about the areas in which each discipline influences the visualization design process and the quality of the final product: effective scientific visualizations.

Bibliography

Acevedo, D., Jackson, C., Laidlaw, D. H., and Drury, F. (2005). Using visual design expertise to characterize the effectiveness of 2D scientific visualization methods. In *Proceedings IEEE Visualization, Poster Compendium (BEST POSTER AWARD)*.

Acevedo, D. and Laidlaw, D. H. (2006). Subjective quantification of perceptual interactions among some 2D scientific visualization methods. *IEEE Transactions on Visualization and Computer Graphics (Proceedings Visualization / Information Visualization)*, 12(5). In Press.

Acevedo, D., Vote, E., Laidlaw, D. H., and Joukowsky, M. (2001). Archaeological data visualization in VR: Analysis of lamp finds at the Great Temple of Petra, a case study. In *Proceedings IEEE Visualization 2001 (BEST CASE-STUDY AWARD)*, pages 493–496.

Acevedo, D., Zhang, S., Laidlaw, D. H., and Bull, C. (2004). Color rapid prototyping for diffusion tensor MRI visualization. In *Proceedings of MICCAI 2004 Short Papers*.

Andrienko, G. and Andrienko, N. (1999). Data characterization schema for intelligent support in visual data analysis. In *Lecture Notes in Computer Science*, pages 349–366. Springer-verlag.

Bergen, J. R. (1991). Theories of visual texture perception. In *Spatial Vision*. CRC, Boca Raton, FL.

Bertin, J. (1983). *Semiology of Graphics*. University of Wisconsin Press.

Bokinsky, A. A. (2003). *Multivariate Data Visualization with Data-Driven Spots*. PhD thesis, University of North Carolina at Chapel Hill.

Brooks, F. P. (1996). The computer scientist as a toolsmith II. *Communications of the ACM*, 39(3):61–68.

Callaghan, T. C. (1984). Dimensional interaction of hue and brightness in preattentive field segregation. *Perception and Psychopysics*, 36(1):25–34.

Callaghan, T. C. (1989). Interference and dominance in texture segregation: Hue, geometric form, and line orientation. *Perception and Psychopysics*, 46(4):299–311.

Cameron, I., Wallington, E., Viergever, K., Moss, D., Woodhouse, I., and Stuart, N. (2005). Synthetic aperture radar for neo-tropical botanical inventory and vegetation height retrieval. In *Proceedings of the EARSeL Symposium*, Porto, Portugal.

Card, S. K. and Mackinlay, J. (1997). The structure of the information visualization design space. In *Proceedings of IEEE Symposium on Information Visualization*, pages 92–99.

Carswell, C. M. and Wickens, C. (1990). The perceptual interaction of graphical attributes: Configurality, stimulus homogeneity, and object integration. *Perception and Psychopysics*, 47:157–168.

Casner, S. M. (1991). A task-analytic approach to the automated design of graphic presentations. *ACM Transactions on Graphics*, 10(2):111–151.

Chi, E. H. (2000). A taxonomy of visualization techniques using the data state reference model. In *Proceedings of InfoVis 2000*, pages 69–76.

Cleveland, W. S. and McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association*, 79(387):531–554.

Cloude, S. R. and Pottier, E. (1996). A review of target decomposition theorems in radar polarimetry. *IEEE Transactions on Geoscience and Remote Sensing*, 34(2):498–518.

Dastani, M. (2002). The role of visual perception in data visualization. *Journal of Visual Languages and Computing*, 13(6):601–622.

Demiralp, C., Jackson, C., Karelitz, D., Zhang, S., and Laidlaw, D. H. (2006). Cave and fishtank virtual-reality displays: A qualitative and quantitative comparison. *IEEE Transactions on Visualization and Computer Graphics*, 12(3):323–330.

Eick, S. G. (1995). Engineering perceptually effective visualizations for abstract data. In *Scientific Visualization Overviews, Methodologies and Techniques*, pages 191–210. IEEE Computer Science Press.

Ellis, W. D. (1939). *A source book of gestalt psychology.* Harcourt, Brace, and company, New York, NY.

Fortuny, J. and Sieber, A. (1999). Three-dimensional synthetic aperture radar imaging of a fir tree: First results. *IEEE Transactions on Geoscience and Remote Sensing*, 37:1006–1014.

Hajnsek, I., Pottier, E., and Cloude, S. (2003). Inversion of surface parameters from polarimetric SAR. *IEEE Transactions on Geoscience and Remote Sensing*, 41:727–744.

Hanrahan, P. (2005). Teaching visualization. *Computer Graphics*, 39(1):4–5.

Healey, C. G. (1996). Choosing effective colours for data visualization. In *Proceedings of IEEE Visualization'96*, San Francisco, CA.

Healey, C. G., Amant, R. S., and Elhaddad, M. S. (1999). Via: A perceptual visualization assistant. In *28th Workshop on Advanced Imagery Pattern Recognition.*

Healey, C. G., Booth, K. S., and Enns, J. T. (1993). Harnessing preattentive processes for multivariate data visualization. In *Proceedings of Graphics Interface*, pages 107–117, Toronto, Canada.

Healey, C. G., Booth, K. S., and Enns, J. T. (1996). High-speed visual estimation using preattentive processing. *ACM Transactions on Human-Computer Interaction*, 3(2):107–135.

Healey, C. G., Enns, J. T., Tateosian, L., and Rempel, M. (2004). Perceptually-based brush strokes for nonphotorealistic visualization. *Transactions on Graphics*, 23(1).

Hibbard, B. (2004). The top five problems that motivated my work. *IEEE Computer Graphics and Applications*, 24(6):9–13.

House, D. and Ware, C. (2002). A method for perceptual optimization of complex visualizations. In *Proceedings of Advanced Visual Interface.*

House, D. H., Bair, A. S., and Ware, C. (2006). An approach to the perceptual optimization of complex visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):509–521.

Interrante, V., Fuchs, H., and Pizer, S. M. (1997). Conveying shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):98–117.

Jackson, C., Acevedo, D., Laidlaw, D. H., Drury, F., Vote, E., and Keefe, D. (2003). Designer-critiqued comparison of 2D vector visualization methods: A pilot study. In *SIGGRAPH 2003 Sketches and Applications*. ACM SIGGRAPH.

Jankun-Kelly, T. (2003). *Visualizing Visualization: A Model and Framework for Visualization Exploration*. PhD thesis, Center for Image Processing and Integrated Computing, University of California, Davis.

Johnson, C. (2004). Top scientific visualization research problems. *IEEE Computer Graphics and Applications*, 24(4):13–17.

Keefe, D., Karelitz, D., Vote, E., and Laidlaw, D. H. (2005). Artistic collaboration in designing VR visualizations. *IEEE Computer Graphics and Applications*, 25(2):18–23.

Kirby, M., Keefe, D., and Laidlaw, D. H. (2004). Painting and visualization. In *Visualization Handbook*. Academic Press.

Kirby, M., Marmanis, H., and Laidlaw, D. H. (1999). Visualizing multivalued data from 2D incompressible flows using concepts from painting. In *Proceedings of IEEE Visualization 1999*, pages 333–340.

Kosara, R., Healey, C. G., Interrante, V., Laidlaw, D. H., and Ware, C. (2003). User studies: Why, how, and when. *Computer Graphics and Applications*, 23(4):20–25.

Laidlaw, D. H. (2001). Loose, artistic "textures" for visualization. *IEEE Computer Graphics and Applications*, 21(2):6–9.

Laidlaw, D. H., davidkremers, Toga, A., Drury, F., and Jacobs, R. E. (2004). Applying lessons from visual art to exploration of the brain. Panel in Thirty-Seventh Annual Winter Conference on Brain Research.

Laidlaw, D. H., Kirby, M., Davidson, J. S., Miller, T., DaSilva, M., Warren, W., and Tarr, M. (2001). Quantitative comparative evaluation of 2D vector field visualization methods. In *Proceedings of IEEE Visualization 2001*, pages 143–150. IEEE.

Laidlaw, D. H., Kirby, M., Jackson, C., Davidson, J. S., Miller, T., DaSilva, M., Warren, W., and Tarr, M. (2005). Comparing 2D vector field visualization methods: A user study. *Transactions on Visualization and Computer Graphics*, 11(1):59–70.

Landy, M. and Movshon, J. (1991). *Computational Models of Visual Processing*. MIT Press, Cambridge, MA.

Landy, M. S. and Bergen, J. R. (1991). Texture segregation and orientation gradient. *Vision Research*, 31(4):679–691.

Lange, S., Schumann, H., Muller, W., and Kromker, D. (1995). Problem-oriented visualisation of multi-dimensional data sets. In *Proceedings of the International Symposium and Scientific Visualization*, pages 1–15.

Laper, N. (1995). *Mix and Match: A Construction Kit for Scientific Visualization*. PhD thesis, University of California Santa Cruz.

Lee, J.-S., Grunes, M., Pottier, E., and Ferro-Famil, L. (2004). Unsupervised terrain classification preserving polarimetric scattering characteristics. *IEEE Transactions on Geoscience and Remote Sensing*, 42:722–731.

MacEachren, A. and Kraak, M.-J. (1997). Exploratory cartographic visualization: Advancing the agenda. *Computers and Geosciences*, 23(4):335–343.

Mackinlay, J. (1986). Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2):110–141.

McCleary Jr., G. F. (1983). An effective graphic vocabulary. *IEEE Computer Graphics and Applications*, 3(2):46–53.

Miceli, K. D. (1992). A framework for the design of effective graphics for scientific visualization. Technical Report RNR-92-035, NASA Ames Research Center.

Nagappan, R. (2001). A compositional model for multidimensional data visualisation. In *In Proceedings of SPIE, Visual Data Exploration and Analysis VIII*, pages 156–167.

Norberg, U. M. and Rayner, J. M. V. (1987). Ecological morphology and flight in bats. wing adaptations, flight performance, foraging strategy and echolocation. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 316:337–419.

North, C. (2006). Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9.

Nowell, L. T. (1997). *Graphical Encoding for Information Visualization: Using Icon Color, Shape, and Size to Convey Nominal and Quantitative Data*. PhD thesis, Virginia Polytechnic Institute and State University, Blacksburg, Virginia.

Owen, S., Budgen, D., and Brereton, P. (2006). Protocol analysis: A neglected practice. *Communications of the ACM*, 49(2):117–122.

Rhyne, T.-M. (2003). Does the difference between information and scientific visualization really matter? *IEEE Computer Graphics and Applications*, 23(3):6–8.

Robertson, P. K. (1991). A methodology for choosing data representations. *IEEE Computer Graphics and Applications*, 11(3):56–67.

Salisbury, L. D. P. (2001). *Automatic Visual Display Design and Creation*. PhD thesis, Department of Computer Science and Engineering, University of Washington.

Saraiya, P., North, C., and Duca, K. (2005). An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):443–456.

Sayre, H. M. (1995). *Writing about Art*. Prentice Hall, 2nd edition.

Senay, H. and Ignatius, E. (1994). A knowledge-based system for visualization design. *IEEE Computer Graphics and Applications*, 14(6):36–47.

Shneiderman, B. (1996). The eyes have it. a task by data type taxonomy for information visualizations. In *Proceedings of IEEE Symposium on Visual Languages*, pages 336–343, Boulder, CO.

Sobel, J., Forsberg, A., Laidlaw, D. H., Zeleznik, R., Keefe, D., Pivkin, I., Karniadakis, G., Richardson, P., and Swartz, S. (2004). Particle flurries: Synoptic 3D pulsatile flow visualization. *IEEE Computer Graphics and Applications*, 24(2):76–85.

Springmeyer, R. R., Blattner, M. R., and Max, N. L. (1992). A characterization of the scientific data analysis process. In *Proceedings of the Second IEEE Visualization*, pages 235–242.

Swan, J. E., Interrante, V., Laidlaw, D. H., Rhyne, T.-M., and Munzner, T. (1999). Visualization needs more visual design! Sensory design issues as a driving problem for visualization research. In *Proceedings of IEEE Visualization Conference*, pages 485–490, San Francisco, California.

Taylor II, R. M. (2002). Visualizing multiple scalar fields on the same surface. *IEEE Computer Graphics and Applications*, 22(2):6–10.

Tory, M. and Moller, T. (2004). Human factors in visualization research. *IEEE Transactions on Visualization and Computer Graphics*, 10(1):72–84.

Tufte, E. (1983). *The Visual Display of Quantitative Information.* Graphics Press.

Tufte, E. (1990). *Envisioning Information.* Graphics Press.

Tufte, E. (1997). *Visual Explanations.* Graphics Press.

Urness, T., Interrante, V., Marusic, I., Longmire, E., and Ganapathisubramani, B. (2003). Effectively visualizing multi-valued flow data using color and texture. In *Proceedings of IEEE Visualization.*

Vote, E., Acevedo, D., Jackson, C., Sobel, J., and Laidlaw, D. H. (2003). Design-by-example: A schema for designing visualizations using examples from art. In *SIGGRAPH 2003 Sketches and Applications.* ACM SIGGRAPH.

Vote, E., Acevedo, D., Laidlaw, D. H., and Joukowsky, M. (2002). Discovering Petra: Archaeological analysis in VR. *IEEE Computer Graphics and Applications*, 22(5):38–50.

Wallschlaeger, C. and Busic-Snyder, C. (1992). *Basic Visual Concepts and Principles for Artists, Architects and Designers.* McGraw Hill.

Ware, C. (2004). *Information Visualization. Perception for Design.* Elsevier, 2nd edition.

Watanabe, T. and Cavanagh, P. (1996). Texture laciness: the texture equivalent of transparency? *Perception*, 25(3):293–303.

Watson, B. (2006). Browadening our collaboration with design. *IEEE Computer Graphics and Applications*, 26(5):18–21.

Weigle, C., Emigh, W., Liu, G., II, R. M. T., Enns, J. T., and Healey, C. G. (2000). Oriented sliver textures: A technique for local value estimation of multiple scalar fields. In *Proceedings of Graphics Interface.*

Wolfe, J. M. (1998). Visual search. In *Attention.* University College London Press, London, UK.

Woodhouse, I., Turner, D., and Laidlaw, D. H. (2002). Improving the visualization of polarimetric response in SAR images: from pixels to images. In *Proceedings of IEEE IGARSS.*

Woodhouse, I., van Oevelen, P., and Hoekman, D. (1997). Backscatter modelling in the NOPEX area: Boreal forests and agricultural fields. In *Workshop Proceedings: EMAC 94/95 Final Results, ESA-ESTEC WPP-136 (ISSN 1022-6656*.

Zhang, S., Demiralp, C., and Laidlaw, D. H. (2003). Visualizing diffusion tensor MR images using streamtubes and streamsurfaces. *IEEE Transactions on Visualization and Computer Graphics*, 9(4):454–462.