Identifying, Instantiating and Evaluating Design Principles for Visual Communication

Maneesh Agrawala
UC Berkeley
635 Soda Hall #1776
Berkeley, CA 97420

Wilmot Li
Adobe Systems Inc.
601 Towsend St.
San Francisco, CA 94103

Floraine Grabler
UC Berkeley
514 Soda Hall #1776
Berkeley, CA 94720

Introduction

Visual communication via diagrams, sketches, charts, photographs, video, and animation is fundamental to the process of exploring concepts, and disseminating information. The most effective visualizations capitalize on the human facility for processing visual information and thereby improve comprehension, memory, and inference. Such visualizations can help analysts rapidly find patterns lurking within large data sets, and they can help audiences quickly understand complex ideas.

Over the last two decades, a variety of books [8,14,18,34] have gathered examples of effective visual displays. One thing is evident from inspecting these examples – the best visualizations are carefully crafted by skilled human designers. Yet even with the aid of computers, hand-designing effective visualizations is time-consuming and requires considerable human effort. Moreover, the rate at which we generate new data is growing exponentially. Gantz et al. [4] estimate that we produced 161 exabytes of new information in 2006 alone and that the compound growth rate between 2007 and 2011 will be 60% annually. Thus, we will produce 1800 exabytes of information in 2011, 10 times more than the amount of information we produced in 2006. Yet acquiring and storing this data is, by itself, of little value. We must understand the data in order to produce real value and make decisions based on it.

The problem is that human designers do not have the time to hand-design effective visualizations for this wealth of data. Too often, data is either poorly visualized, or it is not visualized at all. Either way, the results can be catastrophic. For example, Tufte [19] has explained how the Morton Thiokol engineers failed to visually communicate the risks of launching the Challenger Space Shuttle to NASA management. While Robison et al. [15] argue that the engineers must not be blamed for the Challenger accident, better communication of the risks may have prevented the disaster.

Skilled visual designers manipulate the perception, cognition, and communicative intent of visualizations by carefully applying principles of good design. These design principles explain how visual techniques can be used to either emphasize important information or de-emphasize irrelevant details in the display. For example, the most important information in a subway map is the sequence of stops along each line and the transfer stops that allow riders to change lines. Most subway passengers do not need to know the true geographic path of each line. Based on this insight, map designer Harry Beck re-designed the map of the London Underground in 1933 using two main principles; he straightened the subway lines and evenly spaced the stops to visually emphasize the sequence of stops and transfer points (Figure 1).

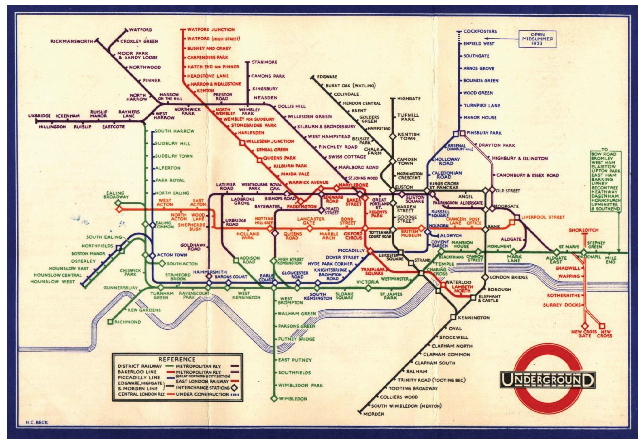

Image credit: Harry Beck © Tfl from the London Transport Museum collection

Figure 1: Harry Beck's map of the London Underground subway from 1933. Beck straightened the subway lines and more evenly spaced the stops to visually emphasize the sequence of stops along each line.

Such *cognitive design principles* connect the visual design of a visualization with the viewer's perception and cognition of the underlying information the visualization is meant to convey. In the field of design, there is a long standing debate regarding the interaction of aesthetic and functional properties of designed artifacts. We do not seek to engage in this debate here. Instead we focus on how particular design choices affect perception and cognition of the visualization rather than the aesthetic style of the visualization. Accordingly, we use the term design principle as a shorthand for guidelines that help improve viewers' comprehension of visually-encoded information.

Design principles usually are not strict rules, but rather "rules of thumb" that may even oppose and contradict one another. For example, Beck did not completely straighten the subway lines. He includes a few turns in the lines to give viewers a sense of the overall spatial layout of the subway

CACM '09

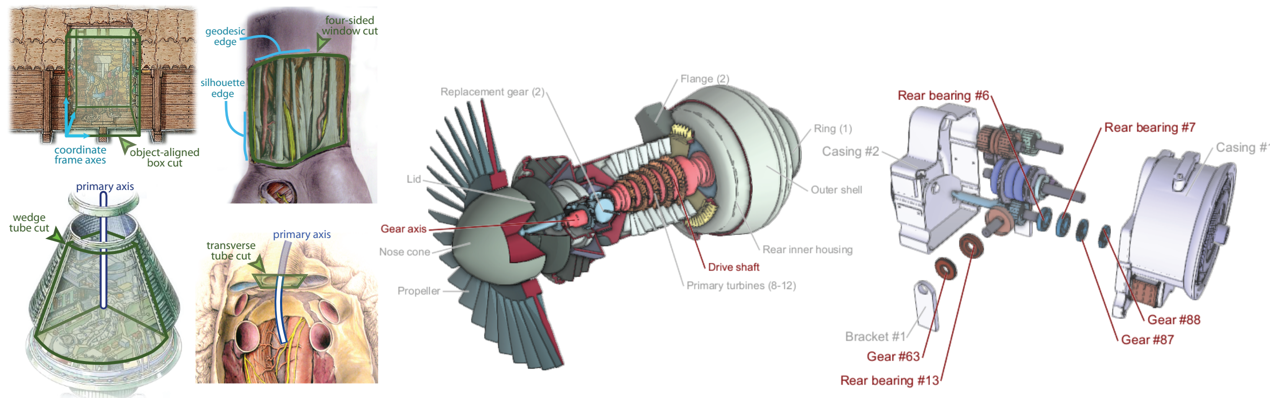

Image credits: (left) Stephen Biesty © Dorling Kindersley; (right) LifeART images

Figure 2: Hand-designed cutaway and exploded view illustrations design the cuts and explosions to emphasize the shape of the missing geometry and the spatial relationships between parts (left). Our system incorporates such principles to generate interactive cutaway and exploded view illustrations (middle, right).

lines. Skilled visual designers implicitly apply the relevant design principles and balance the tradeoffs between them in an iterative process of creating example designs, critiquing the examples, and then improving the designs based on the critiques. Designers usually do not directly apply an explicitly defined set of design principles. The principles are a form of tacit knowledge that designers learn by creating and studying examples. It is far more common for books on visual design to contain examples rather than explicit design principles.

Yet, many of the analysts and end-users that are inundated with data and charged with creating visualizations of it, are not trained designers. Thus, our work is aimed at identifying domain specific design principles and instantiating them within automated visualization design systems that enable non-designers to easily create more effective visual displays. While other researchers have considered specific ways of using cognitive design principles to generate visualizations (see Sidebar), we have been developing a general, three-stage approach for creating visualization design systems.

Stage 1: Identify Design Principles. We identify domain-specific design principles by analyzing the best hand-designed visualizations within a particular information domain. We connect this analysis with research on perception and cognition of visualizations.

Stage 2: Instantiate Design Principles. We encode the design principles into algorithms and interfaces for creating visualizations.

Stage 3: Evaluate Design Principles. We measure improvements in information processing, communication and decision-making that result from our visualizations. Thus, these evaluations also serve to validate the design principles.

We have used this three stage approach to build automated visualization design systems in two domains. In the domain of cartographic visualizations we have developed automated algorithms for creating route maps [1, 3] as well as tourist maps of cities [6]. In the domain of technical illustration we have developed automated techniques for generating assembly instructions of furniture and toys [2, 7] and for creating interactive cutaway and exploded view illustrations of complex mechanical and biological objects [9, 10]. Here, we focus on articulating the techniques we have used to identify and evaluate the design principles for each of these domains. We believe that these techniques will generalize to other domains and that applying our three stage approach will result in a better understanding of the strategies people use to make inferences from visualizations.

Stage 1: Identify Design Principles

Design principles are prescriptive rules that describe how visual techniques affect the perception and cognition of the information in the display. In some cases, these design principles are explicitly outlined in books. For example, books on photography techniques explain a variety of rules for composing pleasing photographs, such as cropping images of people just below the shoulders or near the waist, rather than at the neck or the knees. Researchers have directly applied these principles to build a variety of automated photo manipulation algorithms (see Sidebar for examples).

Yet, our experience is that the design principles are rarely stated so explicitly. Thus, we have developed three strategies for extracting and formulating domain-specific design principles that involve 1) analyzing the best hand-designed visualizations in the domain, 2) examining prior research on the perception and cognition of visualizations, and when necessary, 3) conducting new user studies that investigate how visual techniques affect perception and cognition. We consider each of these strategies in detail.

Analyzing Hand-Designed Visualizations

We have found that a useful first step in identifying design principles is to analyze examples of the best visualizations in the domain. This analysis is designed to find similarities and recurring patterns in the kinds of information that these visualizations highlight as well as the techniques used to emphasize this information

Consider the problem of depicting the internal structure of complex mechanical, anatomical or architectural objects. Illustrators often use cutaways and exploded views to reveal the structure of such objects. They carefully choose the size and shape of cuts, as well as the placement of the parts relative to one another to expose and highlight the internal structure and spatial relationships between parts. We have analyzed a large corpus of cutaways and exploded views to identify the principles and conventions expert illustrators

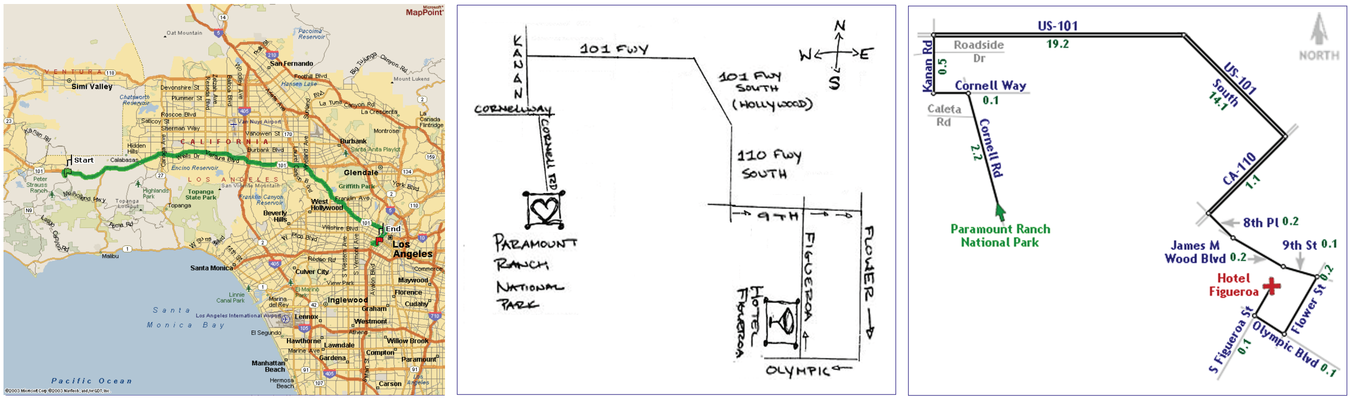

Image credit: MapPoint screen shot reprinted with permission from Microsoft Corp.

Figure 3: A computer-generated route map rendered at a fixed scale does not depict all the turns necessary for navigation (left). A hand-designed map emphasizes the turning points by exaggerating the lengths of short roads and simplifying the shape of the roads (middle). Our LineDrive system incorporates these design principles into an automated map design algorithm (right).

commonly use to generate these images [9,10]. Our process for identifying these design principles was based on three main objectives:

Style-Independence: In order to identify a general set of principles that we could apply to a variety of complex 3D objects, we looked for visual techniques that were common across different artistic styles and types of objects.

Generative Rules: To ensure that we could apply our principles in a generative manner to create cutaways or exploded views, we formed explicit, well-defined rules describing when and how each principle should be applied. The rules were designed to be as general as possible but also remain consistent with the evidence from the example illustrations.

Perceptual/Cognitive Rationale: We motivated each principle by hypothesizing a perceptual or cognitive rationale that explains how the convention helps viewers better understand the structure of the depicted 3D object.

Through this analysis, we identified a set of general, perceptually motivated design principles for creating cutaways and exploded views. For example, the size and shape of cuts in a cutaway illustration are often designed to not only reveal internal parts, but also help viewers mentally reconstruct any occluding geometry that has been removed. Thus, illustrators cut radially-symmetric objects with wedge-shaped cutaways that emphasize the object's cylindrical structure. Similarly, rectangular objects are cut with object-aligned cutting planes (i.e., box cuts), skin-like covering surfaces are cut using window cuts, and long tubular structures are cut using transverse tube cuts. After identifying the design principles, we implemented them algorithmically as part of an interactive cutaway and exploded view system (Figure 2).

We have applied a similar approach to identify the design principles for depicting driving directions from one location to another. We analyzed a variety of hand-drawn route maps and found that they are usually far more useful than computer-generated route maps (as available at sites like local.live.com and maps.google.com) because they emphasize roads, turning points, and local landmarks [1,3]. These maps significantly distort the distance, angle, and shape of roads while eliminating many details that would only serve to clutter the map. Tufte [18] points out that triptiks and subway maps similarly distort the shape of routes and eliminate unnecessary details. Many of these maps use multi-perspective rendering in which the roads are drawn in top-down plan view while important landmark buildings are drawn from a side view so that their facades are visible.

Although analyzing hand-designed visualizations is often a good initial approach for identifying design principles, there are some limitations to this strategy. In some cases it may be tempting to form generative rules that are too specific and do not apply outside the range of analyzed examples. In other cases the rules may be so general that it is unclear how to apply them to specific examples. Such difficulties often arise when the perceptual or cognitive rationale behind a particular visual technique is not clear. In the context of route maps for example, although our analysis revealed that mapmakers often distort road length, angle and shape it was not immediately obvious how such distortions improve perception and cognition of route maps.

Similarly, we have found that one of the challenges in analyzing hand-designed visualizations is to factor out differences due to artistic style. Designers may choose visual attributes such as font type, color palette, line width, etc. for aesthetic reasons; one font looks nicer to the designer than another. Although such aesthetic design choices are important considerations, the goal of our analysis is determine how the design choices improve perception and cognition of the information rather than how these choices improve aesthetics. The difficulty is that these design choices often affect both the aesthetics of the display and the perception and cognition of the information and it is not always clear how to separate these two effects.

In light of these limitations and challenges, we have found that it is often useful to connect our observations and hypotheses from the analysis of hand-designed examples with relevant work from perception and cognitive psychology. These connections serve to clarify the perceptual or cognitive rationale for design principles.

Connecting to Prior Work in Perception & Cognition

In some cases, prior research in perception and cognition suggests or formalizes the appropriate design principles. For example, cognitive psychologists have shown that people

Image credit: © Unique Media

Figure 4: A hand-designed tourist map of San Francisco emphasizes semantically, visually and structurally important landmarks, paths, districts nodes and edges. They use multi-perspective rendering to ensure that the front facades of buildings are visible(left). Our tourist map design system is based on these principles and similarly emphasizes the information that is most important for tourists in this map of San Francisco (right).

think of routes as a sequence of turns [20] and that when following a route the exact length of a road is far less important than properly executing the turns. In other words, the topology of the route is more important than its absolute geometry. This helps explain why hand-drawn maps often distort geometry – the distance, angle and shape of roads – to ensure that all roads and turning points are visible but almost never modify the topology of the route.

In this case, the prior research confirmed and formalized the perceptual/cognitive rationale for the visual techniques we first noticed when analyzing hand-drawn route maps. Based on the resulting design principles, we developed LineDrive, a fully automated system for rendering route maps in the style of hand-drawn maps [3]. LineDrive has been publicly accessible[1] since Oct. 2000, and in surveys we found that for navigation tasks users strongly prefer LineDrive maps to computer-generated maps drawn at a fixed scale (Figure 3).

More recently, we applied a similar approach to build an automated system for generating maps for tourists visiting a new city [6]. Prior work on mental representations of cities [12] has shown that people consider five main elements in a city; *landmarks, paths, districts, nodes* and *edges*. However, a map containing every instance of these elements would be cluttered with excessive detail. The most effective tourist maps include only those elements that are semantically meaningful (e.g., the home of a well-known writer), visually distinctive (e.g., an oddly shaped or colored building), or placed in a structurally important location (e.g., a building located at a prominent intersection) [17]. After choosing the elements to include in the map, mapmakers usually apply a variety of cartographic generalization techniques including simplification, displacement, deformation and selection. Cognitive psychologists and cartographers studying the cognition of maps have shown that such generalizations improve the clarity of the map because they emphasize the most important map elements while preserving spatial relationships between these elements [13].

Our tourist map design system is based on these design principles. The input consists of a geometric model of a city, including streets, bodies of water, parks and buildings (with textures). Our system automatically determines the importance of map elements using top-down web-based information extraction techniques to extract semantic importance and bottom-up vision-based image/geometry analysis to extract visual and structural importance. It then generates a map that emphasizes the most important map elements, using a combination of multi-perspective rendering and cartographic generalization to highlight the important landmarks, paths, districts, nodes and edges and de-emphasize less important elements (Figure 4).

Conducting Experiments on Perception & Cognition

In some domains, new perception and cognition research is required to provide the rationale for the design principles. Working with cognitive psychologist Barbara Tversky, we have developed a methodology for conducting human-subject experiments to understand how people think about and communicate the information within a domain. We first applied this methodology to identify the design principles for creating assembly instructions for everyday objects, such as furniture and toys [2, 7]. The experiments are conducted in three phases:

Production: Participants create visualizations for a given domain. In the context of assembly instructions, participants assembled a TV-stand without instructions, using only a photograph of the assembled stand as a guide. They then drew a set of instructions showing how to assemble the stand.

Preference: Participants rate the effectiveness of the visualizations. In the assembly instructions project a new set of participants assembled the TV-stand, without instructions. They then rated the quality of the instructions created by the first set of participants, redrawn to control for clarity, legibility and aesthetics.

Comprehension: Participants use the ranked visualizations, and we test for improvements in learning, comprehension, and decision-making. In the assembly instructions project yet another set of participants assembled the TV-stand, this time using the instructions rated in the prefer-

[1] http://vis.berkeley.edu/LineDrive

ence phase. Tests showed that the highly rated instructions were easier to use and follow; participants spent less time assembling the TV-stand and made fewer errors.

Once we have conducted these three experiments, we look for commonalities in the highly rated visualizations to identify the design principles. In the context of assembly instructions we identified three main principles: 1) use a step-by-step sequence of diagrams that show one primary action in each diagram; 2) use guidelines and arrows to depict the actions required to fit parts together; and 3) ensure that the parts added in each step are visible. Our automated assembly instruction design system is based on these principles (Figure 5). Tversky and Lee [20] have studied mental representations of maps using a similar methodology where subjects first drew maps to familiar locations and then other subjects rated the effectiveness of the maps.

Stage 2: Instantiate Design Principles

Designing a visualization usually requires choosing visual properties or attributes for each element in the display. For example, to create a route map, the designer must choose attributes including position, size and orientation for each road, landmark and label that appears in the map. Similarly, to create a cutaway illustration the designer has to choose how and where to cut each structure that occludes the target part. Because there are many possible choices for each attribute the design space of possible visualizations is usually large. To build automated visualization design systems we treat the relevant design principles as guidelines for making these design decisions. Thus, we use the principles to navigate through the design space and obtain an effective design.

Yet, most design principles are stated as qualitative guidelines, rather than procedures that we can directly instantiate in an automated design algorithm. The challenge is to transform such high-level principles into implementable algorithms.

Design principles generally fall into two categories: 1) *design rules* or 2) *evaluation criteria*. Design rules separate the design space into regions that contain effective designs and those that contain inviable designs. They are essentially hard constraints in the design space. In creating route maps for example, designers commonly adjust the turn angle to emphasize the orientation of the turn, to the left or right. However, it is unacceptable to adjust the turn angle so much that a left turn appears to be a right turn or vice-versa. This design rule puts a hard constraint on how much designers can adjust the the turn angle.

Evaluation criteria quantify the effectiveness of some aspect of the visualization. We can assess the overall effectiveness of a visualization by considering a set of evaluation criteria that covers all major aspects of the visual design. In creating exploded views for instance, designers must balance two such criteria. A good exploded view must separate the parts so that all of them are visible, yet the visualization must also remain compact and maintain a roughly square aspect ratio to make the best use of the available screen space. To quantify the overall effectiveness of an exploded view we measure the visibility of each part as well as the compactness of the overall visualization. Similarly in designing route maps, designers must ensure that all the roads are visible. To quantify this criterion we compute the length of each road in the map and check that the length is greater than a minimum visibility threshold. The number of roads that are longer than the threshold length is a quantitative measure of the effectiveness of the map with respect to this criterion.

Given a set of design rules and quantitative evaluation criteria, we can use procedural techniques to build an automated visualization design system. For example our system for designing cutaways and exploded views is driven exclusively by procedural techniques. In this case we encode the design rules as a decision tree that describes how to cut or explode away occluding parts based on their geometry. Another approach is to consider visualization design as an energy-minimizing optimization problem. In this case we treat the design rules as hard constraints that define the boundaries of the design space and we treat the evaluation criteria as soft constraints that guide the system to the optimal visualization. While this optimization-based approach is general, we have found that it is essential to develop a set of design rules and evaluation criteria and sufficiently limit the space so that it is feasible to complete the optimization. Both LineDrive and our assembly instruction design system use such an energy-minimizing optimization.

Stage 3: Evaluate Design Principles

The final stage of our approach is to measure the usefulness of the visualizations produced by our automated design systems. We consider several such measures including feedback from users in the form of qualitative interviews and quantitative usage statistics. In some cases we have also conducted more formal user studies to check how well the visualizations improve information processing, communication and decision-making.

User Feedback: We have found that it is critical to involve users early on and conduct qualitative interviews and surveys to check users' overall impressions of the visualizations produced by our systems. Such feedback is essential for quickly identifying problems and it helps to ensure that our design principles and the visualizations converge on effective designs. The interviews and surveys provide high-level checks of the effectiveness of our design principles and allow us to tweak the principles when they are not quite right. For example, early on in building LineDrive, we asked users to rate hand-crafted prototype route map designs and found that 79 out of 90 respondents preferred the distorted LineDrive prototypes to maps drawn to scale [1]. The feedback confirmed that users found the distorted maps to be useful. We believe continual feedback and evaluation yields more effective algorithms and tools.

Another approach is to release the visualization on the world wide web and then check usage statistics. For example, at its peak LineDrive served over 750,000 maps per day and became the default route mapping style for MapBlast, an early web-based provider of route maps. We believe that such public feedback is a strong test of effectiveness as ineffective solutions will be quickly rejected. Yet, we recognize that usage statistics are at best an indirect measure of effectiveness. Many excellent solutions remain little-used due to a wide variety of external forces that have little to do with the usefulness or effectiveness of the system.

User Studies: To quantitatively assess the effectiveness

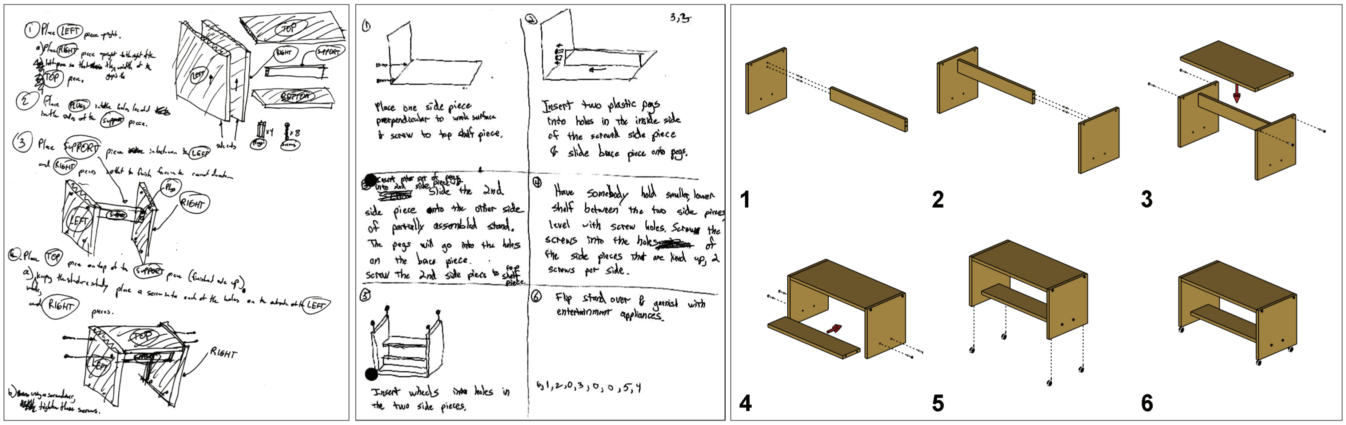

Figure 5: We asked subjects to assemble a TV stand and then create instructions for a novice, explaining how to assemble it (left, middle). Analyzing the hand-drawn instructions we found that diagrammatic, step-by-step instructions that use guidelines and arrows to indicate the actions required for assembly and that provide good visibility for the attached parts are easiest to use and follow. Our system automatically generates assembly instructions based on these principles (right).

of a visualization, we must conduct user studies that compare visualizations created with our design algorithms to the best hand-designed visualizations in the domain. For example, we have compared our computer-designed instructions to factory-produced instructions and hand-drawn instructions for assembling a TV stand. We found that users completed the assembly task about 35% faster and made 50% fewer errors using our instructions. In addition to completion time and error rates, it is also possible to use eye-tracking to determine how the visualization affects the way people scan and process information [5, 39]. Such eye tracking studies can help us evaluate the effectiveness of low-level design choices in creating visualizations. Rigorous user studies are especially important because they also serve to validate the effectiveness of the design principles that the visualizations are based on.

Yet, it is not always clear how to design such quantitative studies. One challenge is that in some domains it is unclear how to compare one visualization to another. For example, in the domain of anatomical illustrations it is not obvious how to compare our cutaway illustrations to hand-designed illustrations. What task should we ask users to perform using the two different illustrations? One approach might be to measure how quickly and accurately viewers can locate a particular organ in the body. However, if the task is to learn the location of the organ, then both illustrations would label the organ and with labels it is unlikely that there would be a significant differences in speed and accuracy. Our cutaways and exploded views are also designed to convey the layering relationship between parts. So an alternative task might be to ask viewers to indicate the layering relationships between parts. But, how can we ask users to do this without leading them to an answer? For many domains, like anatomical illustrations, we believe that it will be necessary to develop new methodology for evaluating the effectiveness of the visualizations and validating the underlying design principles.

Conclusions and Future Work

The approach we have outlined for identifying, instantiating and evaluating design principles for visual communication is a general methodology for combining findings about human perception and cognition with automated design algorithms. The systems we have built for generating route maps, tourist maps, and technical illustrations demonstrate that this methodology can be used to develop effective automated visualization design systems. However, we believe there is much room for extending our proposed approach, and we hope that researchers will improve upon the methods we have described here. There are several directions for future work.

Many other information domains could benefit from a deeper understanding of the ways in which visual display techniques affect the perception and cognition of the information. We commonly encounter a wide variety of different types of information, including cooking recipes, budgets and financial data, dance steps, tutorials on using software, explanations of strategies and plays in sports, and political polling numbers. More effective visualizations of such everyday information could empower citizens to make better decisions.

While we have presented three strategies for identifying design principles, there may be other strategies as well. The strategies we have presented all require significant human effort to look for commonalities in hand-designed visualization, to find and synthesize the relevant prior studies in perception and cognition or to conduct such studies. Moreover, the Internet has made a great deal of visual content publicly available online. Within any information domain, there are often thousands of example visualizations. Thus, a viable alternative strategy for identifying design principles may be to *learn* them from a large collection of examples using statistical machine learning techniques. We have taken an initial step in this direction, with a project designed to learn how to label diagrams from a few examples [21]. One advantage of this approach is that skilled designers often find it easier to create example visualizations than to explicitly describe design principles.

Techniques for evaluating the effectiveness of visualizations and validating the design principles could also be improved. Design principles are essentially models that predict how visual techniques affect perception and cognition. Yet, as we have noted, it is not always clear how to check the effectiveness of a visualization. More sophisticated evaluation methodology could provide stronger evidence for these models and thereby experimentally validate the design principles.

Acknowledgments

We would like to thank Brian Curless, Pat Hanrahan, John Haymaker, Julie Heiser, Jeff Klingner, Mark Pauly, Doantam Phan, Lincoln Ritter, David Salesin, Chris Stolte, Robert Sumner, and Barbara Tversky for their contributions to this research. Jeff Heer, Takeo Igarashi and Tamara Munzner provided excellent suggestions and feedback on early drafts of this manuscript.

1. REFERENCES

[1] M. Agrawala. *Visualizing Route Maps*. PhD thesis, Stanford University, 2002.

[2] M. Agrawala, D. Phan, J. Heiser, J. Haymaker, and J. Klingner. Designing effective step-by-step assembly instructions. In *Proceedings of SIGGRAPH 03*, pages 828–837, 2003.

[3] M. Agrawala and C. Stolte. Rendering effective route maps: Improving usability through generalization. In *Proceedings of SIGGRAPH 01*, pages 241–250, 2001.

[4] J. Gantz, C. Chute, A. Manfrediz, S. Minton, D. Reinsel, W. Schlichting, and A. Toncheva. The diverse and exploding digital universe: An updated forecast of worldwide information growth through 2011. Technical report, IDC White Paper, March 2008. http://www.emc.com/about/destination/digital_universe/.

[5] J. Goldberg, M. Stimson, M. Lewenstein, N. Scott, and A. Wichansky. Eye tracking in web search tasks: design implications. *Proceedings of the Symposium on ETRA 2002: Eye tracking Research and Applications Symposium*, pages 51–58, 2002.

[6] F. Grabler, M. Agrawala, R. W. Sumner, and M. Pauly. Automatic generation of tourist maps. *ACM Transactions on Graphics (TOG)*, 27(3):100:1–100:11, 2008.

[7] J. Heiser, D. Phan, M. Agrawala, B. Tversky, and P. Hanrahan. Identification and validation of cognitive design principles for automated generation of assembly instructions. In *Proceedings of Advanced Visual Interfaces 04*, pages 311–319, 2004.

[8] E. Hodges. *The Guild Handbook of Scientific Illustration*. Van Nostrand Reinhold, 1989.

[9] W. Li, M. Agrawala, B. Curless, and D. Salesin. Automated generation of interactive exploded view diagrams. *ACM Transactions on Graphics (TOG)*, 27(3):101:1–101:11, 2008.

[10] W. Li, L. Ritter, M. Agrawala, B. Curless, and D. Salesin. Interactive cutaway illustrations of complex 3D models. *ACM Transactions on Graphics (TOG)*, 26(3):31:1–31:11, 2007.

[11] B. London and J. Upton. *Photography*. Addison-Wesley, 1997.

[12] K. Lynch. *The image of the city*. The MIT Press, 1960.

[13] A. M. MacEachren. *How Maps Work*. The Guilford Press, 1995.

[14] P. Mijksenaar and P. Westendorp. *Open Here: The Art of Instructional Design*. Joost Elffers Books, New York, 1999.

[15] W. Robison, R. Boisjoly, D. Hoeker, and S. Young. Representation and misrepresentation: Tufte and the Morton Thiokol engineers on the Challenger. *Science and Engineering Ethics*, 8(1):59–81, 2002.

[16] A. Santella, M. Agrawala, D. DeCarlo, D. Salesin, and M. Cohen. Gaze-based interaction for semi-automatic photo cropping. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 771–780, 2006.

[17] M. Sorrows and S. Hirtle. The nature of landmarks for real and electronic spaces. *In Proc. COSIT '99*, pages 37–50, 1999.

[18] E. Tufte. *Envisioning Information*. Graphics Press, 1990.

[19] E. Tufte. *Visual Explanations*. Graphics Press, 1997.

[20] B. Tversky and P. Lee. Pictorial and verbal tools for conveying routes. In C. Freksa and D. M. Mark, editors, *Proceedings of COSIT 99*, pages 51–64, 1999.

[21] I. Vollick, D. Vogel, M. Agrawala, and A. Hertzmann. Specifying label layout by example. In *Proc. UIST*, pages 221–230, 2007.

Sidebar: Visualization Research Based on Cognitive Design Principles

Designing visualization systems based on cognitive design principles is not a new methodology. Visualization researchers have applied similar strategies in a variety of data domains. As in our work, these researchers borrow results from the perception and cognition community. In some cases they conduct new user studies to identify the appropriate design principles and to evaluate the effectiveness of the resulting visual representations. Here we consider several notable examples of research that has followed this methodology.

In the domain of photography, books on photographic techniques have outlined a wide variety of design principles for creating good composition and effective lighting [34, 37]. These books explain that photographs are most pleasing when the subject is placed according to geometric criteria such as centering, the rule of thirds, or fifths, and the golden ratio. They describe how to position light sources to emphasize shape and material properties of objects. Recent techniques for automated photo cropping [39, 42] and image relighting [22, 29] are based directly on such principles.

In domains where the design principles are not stated so directly, one common approach is to identify the relevant prior research in perception and cognition and then synthesize a set of design principles based on this earlier research. While we used this strategy to develop our route map and tourist map visualization systems, others have applied this approach in different domains. For example, researchers have applied research on human perception of shape from image cues [32, 43] such as texture, shading and lighting to develop new non-photorealistic rendering techniques that emphasize the shape of a 3D object via texture [31], suggestive contours [28], or exaggerated shading [38]. The key contribution in this style of research is to connect the relevant studies on perception and cognition of visual displays with the algorithmic techniques for generating such displays.

Another approach to generating good visualizations is to conduct studies comparing the effectiveness of two or more visual representations for the same information. In the domain of information visualization for example, Bertin [23] developed the theoretical foundation for encoding data using visual variables such as position, length, angle, color, and shape. Subsequent human-subject experiments on *graphical perception* have rigorously tested the effectiveness of these variables [24]. For example, studies comparing bar charts versus pie charts have generally found that the length judgments required in bar charts are faster and more accurate than the angle judgments required in pie charts [25, 40, 41]. Other studies have investigated shape discrimination of scatter plot symbols [33, 44] and the tradeoff between the size and resolution in time-series visualizations [30]. Based on such experiments, Mackinlay [35] rank ordered the effectiveness of the visual variables (i.e. length is more effective than angle for encoding quantitative data). His APT [35] and ShowMe [36] systems for automatically designing effective charts and graphs choose the appropriate visual encoding for nominal, ordinal and quantitative data based on this rank ordering. Ware [45] has collected many of these findings in his textbook on Information Visualization.

Cole et al. [26, 27] have applied a similar approach to producing line drawings that best convey the shape of a 3D object. They asked artists to manually draw lines to convey the shape of 3D objects and then analyzed how often the artists drew similar sets of lines. Based on this analysis they suggest algorithmic techniques for producing effective line drawings [26]. More recently they have conducted an evaluation checking how well people perceive shape from line drawings, thereby closing the loop and validating the effectiveness of algorithmic line drawing techniques for conveying shape [27]. The primary challenge in this style of work is to develop experiments that will yield useful design principles for creating effective visual representations.

2. REFERENCES

[22] D. Akers, F. Losasso, J. Klingner, M. Agrawala, J. Rick, and P. Hanrahan. Conveying shape and features with image-based relighting. In *Proceedings of IEEE Visualization 03*, pages 349–354, 2003.

[23] J. Bertin. *Semiology of graphics*. University of Wisconsin Press, 1983. translated by W.J. Berg.

[24] W. Cleveland. *Visualizing Data*. Hobart Press, 1993.

[25] W. Cleveland and R. McGill. Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. *Journal of the American Statistical Association*, 79(387):531–554, 1984.

[26] F. Cole, A. Golovinskiy, A. Limpaecher, H. S. Barros, A. Finkelstein, T. Funkhouser, and S. Rusinkiewicz. Where do people draw lines? *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 27(3), 2008.

[27] F. Cole, K. Sanik, D. DeCarlo, A. Finkelstein, T. Funkhouser, S. Rusinkiewicz, and M. Singh. How well do line drawings depict shape? In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, volume 28, Aug. 2009.

[28] D. DeCarlo, A. Finkelstein, S. Rusinkiewicz, and A. Santella. Suggestive contours for conveying shape. In *Proceedings of SIGGRAPH 03*, pages 848–855, 2003.

[29] R. Fattal, M. Agrawala, and S. Rusinkiewicz. Multiscale shape and detail enhancement from multi-light image collections. *ACM Transactions on Graphics (TOG)*, 26(3):51:1–51:10, 2007.

[30] J. Heer, N. Kong, and M. Agrawala. Sizing the horizon: the effects of chart size and layering on the graphical perception of time series visualizations. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 1303–1312. ACM New York, NY, USA, 2009.

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

[32] J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13(3):321–30, 1984.

[33] S. Lewandowsky and I. Spence. Discriminating strata in scatterplots. *Journal of the American Statistical Association*, pages 682–688, 1989.

[34] B. London and J. Upton. *Photography*. Addison-Wesley, 1997.

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

[36] J. Mackinlay, P. Hanrahan, and C. Stolte. Show Me: Automatic presentation for visual analysis. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1137–1144, 2007.

[37] B. Peterson. *Learning to See Creatively: Design, Color and Composition in Photography*. Amphoto Books, 2003.

[38] S. Rusinkiewicz, M. Burns, and D. DeCarlo. Exaggerated shading for depicting shape and detail. *ACM Transactions on Graphics (TOG)*, 25(3):1199–1205, 2006.

[39] A. Santella, M. Agrawala, D. DeCarlo, D. Salesin, and M. Cohen. Gaze-based interaction for semi-automatic photo cropping. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 771–780, 2006.

[40] D. Simkin and R. Hastie. An information-processing analysis of graph perception. *Journal of the American Statistical Association*, 82(398):454–465, 1987.

[41] I. Spence and S. Lewandowsky. Displaying proportions and percentages. *Applied Cognitive Psychology*, 5(1), 1991.

[42] B. Suh, H. Ling, B. B. Bederson, and D. W. Jacobs. Automatic thumbnail cropping and its effectiveness. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 95–104, 2003.

[43] J. Todd. The visual perception of 3D shape. *TRENDS in Cognitive Sciences*, 8(3), 2004.

[44] L. Tremmel. The Visual Separability of Plotting Symbols in Scatterplots. *Journal of Computational and Graphical Statistics*, pages 101–112, 1995.

[45] C. Ware. *Information visualization: Perception for design*. Morgan Kaufmann, 2004.