Published in final edited form as:

IEEE Trans Vis Comput Graph. 2012 December ; 18(12): 2411–2420. doi:10.1109/TVCG.2012.214.

Different Strokes for Different Folks: Visual Presentation Design between Disciplines

Steven R. Gomez,

Brown University, steveg@cs.brown.edu

Radu Jianu,

Brown University, jr@cs.brown.edu

Caroline Ziemkiewicz,

Brown University, ciemki@cs.brown.edu

Hua Guo, and

Brown University, huag@cs.brown.edu

David H. Laidlaw

Brown University, dhl@cs.brown.edu

Abstract

We present an ethnographic study of design differences in visual presentations between academic disciplines. Characterizing design conventions between users and data domains is an important step in developing hypotheses, tools, and design guidelines for information visualization. In this paper, disciplines are compared at a coarse scale between four groups of fields: social, natural, and formal sciences; and the humanities. Two commonplace presentation types were analyzed: electronic slideshows and whiteboard “chalk talks”. We found design differences in slideshows using two methods – coding and comparing manually-selected features, like charts and diagrams, and an image-based analysis using PCA called eigenslides. In whiteboard talks with controlled topics, we observed design behaviors, including using representations and formalisms from a participant’s own discipline, that suggest authors might benefit from novel assistive tools for designing presentations. Based on these findings, we discuss opportunities for visualization ethnography and human-centered authoring tools for visual information.

Index Terms

Presentations; information visualization; design; visual analysis

1 Introduction

Visual presentations are ubiquitous information-sharing platforms. Slideshow presentations, for instance, commonly accompany research papers or projects at academic conferences, and are shared online on sites like Prezi and Slideshare. Video-recorded presentations are freely

available on user-upload sites like YouTube and Vimeo, or curated sites like TED.com, which hosts more than 1100 talks for online streaming¹. Despite their common use and function in framing visualization content like data charts, user-created presentations have not been evaluated extensively in the broader context of information visualization (“infovis”). In fact, presentations are themselves visualizations of both topical information and authors’ design conventions and aesthetics, which are intuitively shaped by their disciplines. It is reasonable that architects, for instance, adhere to different standards than computer scientists when presenting information in their respective communities. Little work has been done to extract design principles or form theories of visual thinking from these narrative visualizations. Recent work by Walny et al. [18] examined the types and frequency of hand-drawn visualizations on whiteboards, but it focused on drawings found at one research institution. Therefore, it is difficult to generalize guidelines across many visualization users who include, for instance, mathematicians, archaeologists, and public health researchers.

In this paper, we examine visual presentations in the context of visualization and use them to analyze the connections between author disciplines, visual thinking, and design conventions. We collected 65 slideshows in PDF format from graduate students in various fields at three major American universities. Slideshows were grouped by each author’s field of study into four coarse disciplines: the humanities, social sciences, formal sciences (e.g., mathematics, computer science), and natural and physical sciences. We used this dataset to evaluate a broad hypothesis about presentation design:

- Presentation authors from different academic disciplines use visual representations and narrative features, like building on those representations over time, in ways that are characteristic of their own disciplines.

In other words, disciplines have distinct visual representation and narrative conventions for presentation design.

We report results from analyzing two encodings of presentations. Sets of visual building blocks for slides, called *eigenslides*, in each discipline were computed using principal component analysis (PCA) on bitmaps of slides. Slideshows were also reviewed by human coders and embedded into a space of manually-selected features of visual presentations. Comparisons between eigenslides and between clusters of slideshows indicate discipline-dependent slide design patterns. Finally, we report anecdotal results from a controlled study of video-recorded whiteboard presentations (“chalk talks”), which were also encoded into a space of selected features. Our findings reveal design behaviors that suggest authors might benefit from novel assistive tools for designing visual presentations. The contributions of this work are three-fold:

1. an ethnographic study of slide design conventions between academic disciplines using both manually-selected features and a novel eigenslide analysis;

¹Accessed on 27 March 2012. <http://on.ted.com/23>

2. an anecdotal evaluation of topic-controlled whiteboard presentations that revealed unexpected effects, like “formalizing” or “deformalizing” explanations into an accustomed level of rigor;
3. finally, a summary of opportunities and implications for visualization research and authoring tools.

The paper is organized as follows: in Section 2, we review related research in narrative visualization and analyzing hand-designed visualizations; next, methods and results are presented for our experiments with slideshows (Sec. 3) and whiteboard presentations (Sec. 4); finally, we discuss implications from these experiments for information visualization (Sec. 5) and give conclusions from this work (Sec. 6).

2 Related Work

In this work, we study the forms used in visual presentations for insight into how people with different backgrounds think about information. This builds on cognitive science research on information externalization as well as research that examines user-produced designs to provide insight into visualization users.

2.1 Visualization as Knowledge Externalization

As Liu and Stasko [12] argue, visualization can be seen as a tool to support the formation of mental models about information. Research in cognitive science [10, 9, 17] has shown numerous ways in which the form of an external visual representation can influence or embody mental models. This research suggests that the creation of visual representations can provide important insight into a user’s thinking process. Building on the same principle, Alibali et al. [3] examined the verbal expressions and gestures used by people as they talked through the solution to a word problem that can be solved using multiple mental representations. They found that analyzing the continuity of an individual’s gestures could reliably reveal which representation of the problem would be used in the solution. This body of theory broadly motivates our analysis of visual design elements as a proxy for how a user conceives of information in a communication task.

Further work by Tversky et al. [16] studies how visual representations play a role in explanations, a question that is directly related to our own study of visually-supported explanations. They broadly discuss how words, gestures, and diagrams all play a role in spontaneous explanations. We build on this work by specifically studying how visual elements are used in two different types of explanation scenarios.

Because of this focus, our work is especially relevant to the emerging field of narrative visualization [14, 8]. Researchers in this area have presented taxonomies of strategies for explanation with visualization, but in existing work these design strategies are based on analyses of case studies. However, this area of research would also benefit from information on how and why ordinary users produce different types of visual explanations. In the current work, we take a more focused approach to answer this question for common scenarios that involve visual storytelling to convey information.

2.2 Analyzing User-Produced Visual Representations

There is an existing body of research on the subject of user-produced visual representations of information and how these representations shed light on visualization design. This research has the potential to reveal user preferences and behavior outside the constraints and biases of specific visualization software. By examining how users create visualizations, rather than how they read them, this work provides a valuable complementary view to that provided by traditional user studies.

Agrawala et al. [1] suggest that examining existing hand-designed visualizations gives insight into the design rules that work well with viewers' cognitive and perceptual abilities. This method has successfully been used to improve the automated design of route maps [2], assembly instructions [7], and exploded views of 3D models [11]. We adopt a similar philosophy in analyzing the design of visual narratives. This is a more general task than those examined by Agrawala and colleagues, and so our method is more exploratory and less directed towards developing rules for automation.

Most relevant to our current work are studies of how and why people use visual presentation methods such as whiteboards in conveying information. Cherubini et al. [6] performed an in-depth study of how software developers used whiteboards in daily work, finding that diagrams sketched on whiteboards were most frequently used to communicate preliminary or transient ideas during ad hoc meetings. This gives context to the type of task for which our findings may be most relevant. However, unlike the current work, they were unable to observe any patterns in the type of graphical forms developers used.

In a similar study of whiteboard use in the field, Walny et al. [18] examined more closely how visual and textual representations are used in this context. The authors took snapshots of whiteboards in a number of offices and labs at a research institution and recorded the types of visual elements used on each. The results show a rough estimate of how frequently certain visual forms appear in everyday whiteboard use in this type of environment. This work serves as an important background to our current research questions. However, given the nature of this study, there is little information on what the whiteboards were being used for and who was using them. For this reason, there are limits to how much these results can be generalized beyond the specific institution and scenarios from which they were gathered.

In our work, we adapt and extend some of Walny et al.'s taxonomy of whiteboard marks to focus on information narratives and how they differ based on a participant's background. In addition, we compare presentations across media to examine how general these findings are. In our first study, we examine the rate at which participants use different types of visual representations in slideshow presentations.

3 Slideshow Experiments

Electronic slideshows, usually created with tools like Microsoft PowerPoint or Apple Keynote, are ubiquitous visual artifacts in academic settings. As such, these slideshows are a good starting point for analyzing visual design between groups of presentation authors. In this section, we first describe the construction of a small slideshow dataset, then present

building blocks of slide images, called eigenslides, and a coding scheme for meaningful manually-selected features in presentations. Results are given that demonstrate differences between presentations using these analysis tools.

3.1 Collecting Slideshows

Slideshows were collected via email requests from current graduate students across many fields at Brown University, Harvard University, and the Rhode Island School of Design (RISD). We specifically solicited slideshows from students who identified their fields of study as belonging to one of four coarse disciplines: the *humanities*, *social sciences*, *natural sciences*, and *formal sciences*. We describe these groups as follows: the social sciences and natural sciences primarily use the scientific method to study human behavior and the natural world, respectively; formal sciences and the humanities primarily use non-empirical methods to study formal systems and the human condition, respectively.

In practice, many fields straddle multiple disciplines (e.g., archaeology). We rejected slideshows from fields of study that we felt did not primarily belong to one of these disciplines, such as business or human-computer interaction. Some other well-known conceptual classifications for disciplines exist, like Biglan's dimensions [5], but our categories have the advantage of being few in number and relatively straightforward to classify manually. Additionally, they draw distinctions between methods that seem to employ different information types and visual representations.

Each contributor was allowed to submit one or two slideshows to the dataset, resulting in 65 total slideshows collected from 52 unique contributors. The breakdown of the dataset is shown in Table 1. Slideshows were not restricted to specific functions, such as class presentations or lab talks, though all were created by graduate students.

3.2 Image-Based Analysis with Eigenslides

Performing an image-based analysis of slides is a reasonable first step at visualizing variations in slide designs. Most people could squint or look at thumbnails of slides and sort them into similar designs without actually reading details on any slide. Therefore, we expect methods used in image analysis, like dimensionality reduction, to find some structure in slides.

The first experiment we performed recovers variations in the way that pixels form marks on slides. Before we began, each slideshow was chopped into individual slides to create a “bag of slides” image set for each discipline. Individual slides were treated as images – vectors of pixel brightness values – then converted to gray-scale PNG format and scaled to 320 by 240 pixels, using the ImageMagick library. We did not observe significant distortion in images after scaling them to the 4:3 aspect ratio. In general, images were legible enough for a reader to distinguish between most semantic features of the slide (e.g., diagrams and images, captions, bullet points).

Next, principal component analysis (PCA) was performed on sets of images grouped by discipline. PCA has been used in computer vision to study how natural images, like photos of faces [15], vary across a set. Here, we applied PCA to synthetic images that can include

text, visualizations, and other images, with the goal of gaining insight about slide design. For each discipline, the analysis gives a mean-valued slide image and an ordered list of principal components, which we call *eigenslides*, that are dimensions for a space of slides. Eigenslides can also be thought of as building blocks that combine in different quantities to form slide images. In Section 3.4, we interpret visual differences between the disciplines' eigenslides.

Because PCA is computationally expensive on images (320×240 pixels as features), we ran PCA on 250 randomly selected, unique image samples from each discipline. The mean and top 15 eigenslides for each discipline are shown in Figure 2. Because eigenslides are created from samples of the full dataset, it is important to check whether analyzing 250 samples is enough to capture the pixel variance in the entire dataset. We ran the sampling and eigenslide process three times for each discipline. Cumulative variance accounted for by the eigenslides, computed during PCA, is plotted for all samplings and shown in Figure 3.

3.3 Semantic Feature Analysis

Eigenslides give insight about contrast variations in groups of pixels. However, these variations do not directly correspond to semantic features in slideshows, like visualization constructs, bullet points, or image captions. Two slides with the same design “blueprint” of semantic features might look different simply due to the specific information presented. Analyzing these features gives insight about a more abstract level of design than pixel variations.

We give a set of semantic features that we used to code slideshows into 19-dimensional numerical descriptors that can be compared and visualized. The approach to studying features in the slideshow medium is similar to recent work analyzing whiteboards by Walny et al. [18]. We created a 19-feature coding scheme – described in Table 2 – meant to characterize visual design in presentations. Fewer distinctions are drawn between different visualization classifications than in the whiteboard taxonomy; we did not want features to be too specific to particular fields of study or information types. Other features are added that would have been impossible to count in leftover whiteboards, such as “building on” previous slides or diagrams.

Each of the 65 slideshows was coded manually by a single reviewer, and coding duties were divided between four trained reviewers. After coding, a feature vector $\{f_{S,1}, f_{S,2}, \dots, f_{S,16}\}$ is computed for each slideshow S . The i th feature score is computed by averaging an indicator variable for the feature over all individual slides.

$$f(S, i) = \frac{\sum_{s \in S} I(s, i)}{|S|}, I(s, i) = \begin{cases} 1 & \text{if feature } i \text{ present in } s, \\ 0 & \text{otherwise} \end{cases}$$

Differences in semantic features between groups are visualized using force-directed Multidimensional Scaling (MDS). With MDS, high-dimensional points are plotted onto a plane in a way that tries to preserve their relative distances from one another. We use it as a

blackbox tool to see whether slideshows cluster by discipline using features based on a simple coding scheme. In this case, the 19-dimensional feature vectors are plotted and are colored by discipline. Results are shown in Figure 5.

3.4 Results

3.4.1 Eigenslides—The mean slide and top 15 eigenslides (ranked left to right, top to bottom) in each discipline are shown in Figure 2. These images are generated from one sampling of 250 unique slide images per discipline. Eigenslides generated from repeated samplings of each discipline's set look similar. Furthermore, in Figure 3 we plot the cumulative variance accounted for by the eigenslides generated from three repeated samplings; we do not see major differences in these curves between samplings within each discipline. This is evidence that sampling is sufficient for generating representative eigenslides for each discipline.

Vignetting, or shading around the edges of an image in contrast to the center, appears in both the natural sciences and humanities mean slides. This suggests heavier use of images, which are usually centered in slides, by these disciplines compared to the others. Highly ranked eigenslides in both the formal and social sciences show high-frequency contrast details, like pieces of individual diagrams or horizontal striping that suggests text. By comparison, the top eigenslides in the humanities and natural sciences mostly show low-resolution contrast patterns.

Evidence of these differences is also seen in plots of cumulative variance for each discipline. Generally, we expect a visually diverse set of slides, compared to a visually homogeneous one, will require more eigenslides to account for a high percentage of its variance. As seen in our data, the formal and social sciences need fewer eigenslides than the humanities or natural sciences to account for a given high threshold of variance in their respective slide sets. Therefore, slides from the formal and social sciences might be less visually diverse than the others. This is also consistent with the higher-frequency visual details seen in these disciplines' top 15 eigenslides – the first few eigenslides account for a large amount of variance, leaving the remaining ones to account for high-frequency details.

3.4.2 Presentation Features—A visualization of slideshows with respect to semantic features is shown in Figure 5. Points represent entire slideshows and are colored by discipline. In (a), the MDS visualization of all slideshows reveals that humanities, social sciences, and natural sciences are well separated, but formal sciences are scattered in this space. In (b), the pairwise MDS plots show separation in the plane between all pairs except social and formal sciences. The fact that these are hard to separate is consistent with some visual similarities in eigenslides between these disciplines.

In general, qualitative analysis is most appropriate for analyzing this type of data, since they are based on content analysis rather than a formal experiment. However, we found it useful to statistically analyze the normalized use rate of the semantic features we coded in order to determine which varied most significantly between discipline groups. To help assess the effect of discipline on the appearance rate of semantic features, we performed a multivariate analysis of variance (MANOVA) using Wilks' Lambda as the test statistic, with discipline

as the independent variable and the 19 features as dependent variables. This analysis found a significant main effect of discipline, $F(57, 129) = 3.1, p < .001, \eta^2 = .57$. Therefore, univariate analysis of variance (ANOVA) was used to test for significant differences on specific semantic features. Significant findings are highlighted in Figure 4.

4 Whiteboard Presentation Experiment

In this section, we describe a second common form of visual presentation in academia – whiteboard “chalk talks” that let presenters explain topics using freeform marks on the board as a visual aid. Whiteboard talks are studied for two reasons: 1) analyzing a second form of visual presentations lets us generalize insights more easily; 2) topics can be controlled across all participants, unlike in the slideshow set, where most contributors submitted presentations about their own fields of study. Controlling topics lets us see whether the expression of some features, like labels or using color for organization, is dependent on particular topics or, instead, reflects design choices that stick with participants between topics.

4.1 Presentation Prompt Inventory

For this experiment, we constructed a small inventory of 6 explanation prompts written in casual American English. These are shown in Figure 6. The criteria for prompts included the following:

1. No prompt should be fully obvious or axiomatic to the average viewer on YouTube; it should be non-trivial to create a compelling explanation for each prompt;
2. Given up to five minutes of planning time, graduate students studying any field at Brown should be able to construct some explanation (though not necessarily a correct one);
3. The set of all prompts should be diverse in the types of visual representations that can be used to explain them.

We applied each criterion to the best of our abilities. Many more prompts pass these criteria than were possible to include in this study. To ensure (3), prompt candidates were labeled as *high-spatial*, *low-spatial*, and *neutral*. In highly spatial prompts, we expect participants to externalize spatial mental models during explanations. We chose prompts related to route-map drawing and mechanical assembly, as these kinds of hand-drawn visualizations have been previously studied [2, 7]. In low-spatial prompts, we expect participants to reason using representations that do not correspond with spatial or visibly observable phenomena. In [WIFI], we were interested in seeing how abstract ideas like sharing or dividing a rate would be explained by those without specific technical knowledge. The other low-spatial prompt [ROOM] is based on the pigeonhole principle² and is easy to confirm in specific instances but challenging to explain generally. Neutral prompts are neither predominantly spatial or predominantly abstract. The [CARS] and [PHOTO] prompts involve phenomena that are difficult to visualize directly but act on spatial entities. These prompts were designed so that two prompts fell into each of the spatial categories.

²Based on an example at http://www.cut-the-knot.org/do_you_know/pigeon.shtml

4.2 Controlled Presentations

We recruited four participants from each discipline we identified in Section 3. Participants included 15 PhD students and one Masters student at Brown University; all were right-handed and were fluent English speakers. Participants were asked to give a video-recorded, explanatory presentation for each prompt described in Figure 6; in total, 96 presentations were recorded ($96 = 6 \text{ prompts} \times 4 \text{ disciplines} \times 4 \text{ participants/discipline}$). These presentations were controlled in two ways that the slideshows we collected were not: 1) the content was controlled by having each participant give a presentation for each of our prompts; 2) the purpose/audience of the presentation was explicitly given beforehand. Presenters were told to be convincing as possible to a Web audience of YouTube viewers.

Each participant was given a whiteboard space approximately 1.5m wide by 1.0m tall, and red, blue, and black markers and an eraser. Prompts were given to each participant one by one, with a short questionnaire completed between prompts. To account for possible priming effects by low-spatial and high-spatial prompts, participants were split evenly and given two different orderings of the six prompts: [*neutral, low-spatial, high-spatial*] and [*neutral, high-spatial, low-spatial*]. The order of prompts within each spatial category was fixed (see rows in Figure 6). For the [MAP] prompt, all participants were familiar with the local landmarks mentioned. For each prompt, participants were given up to five minutes of planning time, then up to five minutes for the final presentation; during both periods, they were allowed to use the whiteboard.

4.3 Semantic Feature Analysis

Whiteboard talks were video-recorded, and features were manually coded by one reviewer using the presentation coding scheme shown in Table 2, which we also used to code slideshows. The ‘Building on previous slide’ feature was changed to ‘Building on a whiteboard drawing’, and was counted if a presenter drew some visual element, stopped drawing and explained it verbally, then extended it with more drawing. Additional features added for the controlled case included total presentation time and time spent drawing. We also noted presentations in which participants gave unnecessary or irrelevant details. These extra features were impossible to evaluate in slideshows without more information about each author or deeper knowledge of each slideshow’s domain content. In each presentation, individual features were coded either as 0 (not present) or 1 (present) except the time features, which were counted in seconds.

4.4 Results

Figure 7 shows the frequency of each feature per presentation, averaged over all participants and prompts for each discipline. A bar chart also shows average drawing time and total time per presentation. As in Section 3.4, qualitative analysis is most appropriate for this data, given the small study size and that features were manually counted from video. A MANOVA on the aggregated whiteboard frequencies was not significant, and there were no signs of significant differences in individual features.

Participants spent on average 20-25% of presentation time drawing on the whiteboard. About 29% of presentations (28 out of 96) contained contextual details that we judged as

unnecessary or irrelevant to the prompt. Of these, some individual presenters were more eager than others: 5 participants gave 3+ such presentations; 6 participants gave 1 or 2; and the remaining 5 participants gave none.

5 Discussion

Results from both our slideshow analyses and whiteboard experiment are discussed below. We found evidence for between-discipline differences in both visual representations and narrative elements. We evaluate these differences with respect to two abstract presentation dimensions we have identified that seem to correspond with design differences. The first is *visualness*, which reflects the amount and diversity of visual marks used. The second is *systematicness*, which reflects the level of formality or focus on proof in information narratives. Intuitively, we expect to see high presentation systematicness in disciplines whose research methods are concerned with proofs and formulas, like the formal and natural sciences, and lower systematicness in disciplines that use critical or interpretative methods, like the humanities.

For both representations and narrative features, we discuss design implications for visualization, and then describe other application opportunities, like intelligent user interfaces and improved visual information retrieval systems.

5.1 Representations Between Disciplines

5.1.1 Slideshows—Representations used in slideshows reflected the domain content of those presentations. More than half of all formal and social science slides contained bullet-point text elements, whereas only around 35% of natural science slides and 8% of humanities slides contained them. This may be due to the explanation of systems, models, and statistics that appear in fields like computer science and economics, but are less prevalent in fields that primarily work with images or diagrams, like art history or biology. In the humanities, around 80% of slides contained some piece of art or photograph. These differences are visible in the top 15 eigenslides we generated for each discipline. Formal and social sciences slides show high-frequency ‘striping’, indicative of text; humanities and natural sciences slides show lower frequencies (e.g., where details of diagrams or text are indistinguishable). We also identified vignetting – shading that radiates from the image center – in the mean humanities slide and social sciences slide, which indicates high visualness in the center of the slide canvas.

Natural sciences slides used the most visualization features between the disciplines. These slides contained more labels than slides in other disciplines; in most cases, labels were used to specify diagram elements. This contrasts with humanities slides, which also have high visualness (e.g., significantly more photographs and visual art than other disciplines) but use labels the least. It is possible that the images used in humanities slides are more straightforward and obviate labeling. Another hypothesis is that humanities slideshows tend to be less systematic and that formally specifying images or parts of images is less conventional in these fields.

Some representations, like infovis constructs for network relations and numeric equations, were predominantly used by the natural and formal sciences, as expected.

5.1.2 Whiteboard Talks—We found evidence that in topic-controlled presentations, some participants chose familiar representations from their own fields of study. Sometimes this occurred even when those representations seemed to us to be less appropriate or obvious than other choices. Two observations during the whiteboard experiment support this hypothesis:

- Some representations were only used by authors of some disciplines. For example, only authors from formal and natural sciences used numeric equations (in [WIFI], [MAP], [ROOM], and [CARS] prompts). These were used to explain concepts like bandwidth, time as a function of walking speed and distance, graph properties, and mechanical work in powering a vehicle.
- Some prompts elicited a variety of representations across participants. Visual representations used during the [MAP] prompt (Figure 8) ranged from stylized maps – what we considered the most obvious visual aid – to a map overlaid with a geometric diagram, to a set of distance/time equations. In some cases, presentations made similar arguments, but used different representations to organize or reason about similar information. Two representations used during the [ROOM] prompt (Figure 9) show how participants selected different kinds of network data visualizations to explain the prompt. It is possible that matrix representations are more familiar than node-link diagrams to the cognitive scientists who sketched (a), and vice versa for the computer scientist who sketched (b).

Another hypothesis we formed is that the visual medium can affect how representations were used. For instance, no participants wrote captions on the whiteboard; however, humanities presenters, who used captions in around 44% of slides, used labels on the whiteboards more frequently than others, despite using labels in slideshows least frequently among disciplines. It is possible that labels, which require less writing than captions, assumed some functions of captions during whiteboard presentations. A reason for this change could be to reduce the amount of handwriting on boards, which can be illegible or challenging to produce during timed, videotaped presentations.

5.1.3 Implications—The foremost implication is that different disciplines use different kinds of visual representations; in cases where presenters are given topics that require only general knowledge or knowledge outside their domain, familiar domain representations are sometimes used to reason about and explain these topics. Our hypothesis is that presenters choose effective representations with respect their own understanding, not just the target audience. We imagine a “golden rule” for presenting information: Design information representations for others as you would like them presented to you. Some presenters in our whiteboard experiment seemed to follow this rule, despite talks being specifically targeted for a general audience.

Even though we analyzed presentations with a taxonomy of visual features that generalizes to both slideshows and whiteboard talks, some features (e.g, labels, captions) seemed to be

used more or less frequently in each medium due to specific affordances by the medium. This is consistent with other studies of whiteboards [18, 6] that suggest users take advantage of the ability with that medium to make freeform marks or customize representations. Therefore, understanding differences between presentation media is critical when trying to generalize about design conventions from individual visual artifacts. We have also seen that different semantic features can serve similar abstract roles. For instance, both labels and captions augment charts or pictures with words to help viewers interpret visual information. Understanding how and why different representations can be functionally similar could indicate underutilized areas of the design space.

5.2 Narrative Conventions

5.2.1 Slideshows—Systematicness of disciplines appears related to differences in narrative development and conventions, like the use of color to organize marks or how presentations build on representations over time. Students in more proof-based disciplines, like the formal and natural sciences, seemed to structure information in presentations in more organized ways. Science disciplines were more likely to build on slides than the humanities, though these differences were not significant. Outside of images, titles were apparent in over 75% of slides from the formal, natural, and social sciences, but appeared in just 27% of humanities slides. One possible explanation is that titling every slide is not aesthetically pleasing, and that more systematic disciplines are less likely to sacrifice organizational elements for aesthetics.

Narrative features that illustrate ideas – like similes, metaphors, and examples – might be especially important when presentations are not very visual. Formal and social science slideshows, whose eigenslides and feature analysis suggest they are less visual than slideshows from the other disciplines, used these elements significantly more than others. It is possible that these features help capture audience attention or give supporting context, which might otherwise be done by images and diagrams.

5.2.2 Whiteboard Talks—Like we observed in slideshows, whiteboard talks showed evidence that discipline systematicness is reflected in presentation narratives. The first observation is that authors from more systematic fields spent more overall time on presentations. In some cases, these participants ran out of their 5 minutes of presentation time while “proving” the prompt for the viewer. In the [ROOM] prompt, which could be explained by enumerating a set of scenarios, participants from more systematic fields like computer science were aware that proving the prompt required them to explain all scenarios. On the other hand, participants from less systematic fields usually gave just a handful of supporting scenarios. It is unclear whether fast presentations are the result of conscious or subconscious satisficing – giving just enough of an argument to be convincing – or just being parsimonious with time.

Participants from science disciplines also built on whiteboard drawings more frequently than humanities participants, which is consistent with our slideshow findings. One hypothesis is that building on drawings is a systematic process, like constructing a mathematical proof, and that science disciplines have more experience with this style of argument development.

Another hypothesis is that images in these disciplines typically serve different supporting roles in narratives. Many images in science presentations are diagrams that the author has created. In many cases, these diagrams or visualizations are themselves the focus of the presentation and do not merely illustrate a broader concept. In history or literature slideshows, images like art and photos are often works created by others and are not typically modified by the presenter. As such, for humanities participants, whiteboard drawings may serve a more familiar role as static illustrations rather than dynamic parts of the presentation narrative.

Whiteboard organization seems to reflect the systematicness of participants' disciplines. Formal science participants used color for organization on boards frequently. Furthermore, the variety of whiteboard organization suggests that not all participants found the medium equally 'forgiving' [18]. Some participants expressed concern over how viewers would interpret the marks they drew. One natural scientist apologized to the camera during his [MAP] presentation because he could not draw a map to scale and ultimately suggested to viewers that "it just comes down to looking at a map." On the other hand, one formal scientist relied on the ease of erasing and redrawing on the board during his [DOOR] presentation to illustrate keyframes of a door swinging open.

Informally, we observed two additional factors in the content-controlled whiteboard presentations – *eagerness*, or giving extraneous details, and *formalizing*, a term we use when participants seem to fit a presentation topic to an accustomed level of systematicness.

Eagerness: We noted that some participants gave unnecessary or irrelevant details in presentations, perhaps to sound more convincing to YouTube viewers. This is related to Bernstein et al.'s characterization of "eager beavers" among Amazon Mechanical Turk workers [4], who go beyond the scope of a task's instructions. Sometimes these details were used to help set up an example. In one [PHOTO] talk, a participant explained that cameras in the nineteenth century had much slower shutter speeds, which is why smiling is rare in old photos – subjects had trouble holding the smile while the film was exposed, causing blurring in the photograph. In other cases, extra details were not linked to elaborate examples. In [MAP] talks, for instance, some participants named extra landmarks without providing their relative locations. In one [CARS] talk, a physicist spent minutes giving equations before relating them back to the car scenario, and eventually ran out of time. Afterwards, he remarked that he elaborated so much because he "just graded a [college physics] homework set with almost the exact same problem". This suggests that some eager beavers simply forgot about or ignored the instructions to be convincing to average YouTube viewers.

Formalizing: A related second factor we observed is about formalizing or de-formalizing a prompt to one's accustomed level of systematicness. Often in these cases, information is embellished or subtly changed. In Figure 10(a), a computer scientist presented a geometric proof based on the Pythagorean Theorem for the [MAP] prompt. In order to make the proof work, the presenter disregarded one key fact – the actual roads are not straight and do not perfectly form a right triangle. In 10(b), a public health student used de-formalizing to treat the [ROOM] prompt as a statement about social networks in a community. During the

presentation, the participant argued that “chances are” the prompt was true, indicating relatively low concern about proving the prompt in certain terms to the audience.

5.2.3 Implications—Systematicness appears to be an important factor in visualization design choices. Sharing information between disciplines using the “golden rule” might be less effective than structuring narratives more in-line with the target audience. Visualization designers could also benefit from seeing how others in their own discipline create visualizations. Some systematic design principles, like organizing all information in one drawing, might compete with other systematic principles, like building on drawings over time. Examining visualizations from like-minded domains could illuminate the design space.

Finally, eagerness and formalizing reflect on ways in which information narratives are embellished by authors. Both authors and the audience could benefit from knowing when and how formalizing happens. Eagerness can sometimes improve visualizations or narratives by providing context; other times it can detract by diluting critical information or resources (e.g., wasting on-screen “ink”). Support systems for authoring information visualizations might benefit by enhanced context-gathering and evaluation tools. Bringing humans into the loop with crowdsourcing tools like Mechanical Turk could provide a mechanism to identify and evaluate these embellishments.

5.3 Opportunities

We studied user-created visual presentations to gain insights about visualization use and visual thinking in user domains. This work has applications in assistive design tools and opens the door to important ethnographic studies of visualization and design.

5.3.1 Visualization Tools—Our experiments demonstrate some feasibility for gathering useful design statistics, like distributions of semantic features, from different populations of authors. Novel authoring tools could use these statistics to provide content-aware design recommendations, helping users build presentations more quickly or effectively. The analysis and applications extend easily to other artifacts, like posters or papers with graphics. Furthermore, design models could be trained on artifacts from specific venues like *IEEE InfoVis* to better target design for these audiences.

A related opportunity lies in inferring metadata about presentations or visualizations, like discipline or content area, which might improve tools like image search engines that have indexed visualizations or presentations on the Web. Our finding that slideshows tended to group by discipline in MDS plots (see Figure 5) suggests that a classifier could have some predictive power in labeling the discipline of an unknown presentation based on its features. Some visual and text-based features, like charts or bullet points, themselves might be classified automatically to obviate the need for manual coding of features. Recent work in this direction includes ReVision by Savva et al. [13], which can classify the type of simple information visualizations before extracting quantitative information.

5.3.2 Ethnography—Characterizing how individuals create visualizations and apply them in settings like presentations is an important step in understanding patterns of visual communication and developing assistive tools for these users. In this paper, we presented an

ethnographic study of visual presentation design between groups of users. The results of studies like this and Walny et al.'s [18] can be helpful for generating hypotheses that lead to applicable design guidelines or design-space exploration for visualization.

We focused on academic disciplines of users and how discipline might influence design choices; exploring other user factors, like design experience, working in industry versus academia, and culture could provide more insights about how users think about design and visualization. In addition to examining different user groups, another opportunity lies in refining the set of semantic features used to encode presentations or other visualizations. The set of features we used to code presentations (see Table 2) was intentionally general to fit multiple presentation types from many disciplines. One disadvantage of this approach is that these features might not be detailed or specific enough to discriminate between presentations from closely related user groups, like pure mathematicians and computer scientists. Additionally, it is possible that general features are more likely to be coded inconsistently by humans than very specific ones. Other feature sets could have a better trade-off between ease of coding and discriminative power for experiments like ours.

Finally, research into what makes information visualization convincing, interesting, or memorable could produce important insights about visualization design. In this work, we focused on how different design conventions are used between groups, rather than assessing the effectiveness of these conventions individually. Following this work up with controlled studies about design features will inform guidelines about when and how to use features beyond what is simply conventional in specific domains. For instance, for the whiteboard prompts in Section 4, we might like to know whether equations or diagrams are more or less convincing for most YouTube viewers than simple metaphors or examples.

6 Conclusion

We reported the results of an ethnographic study of visual presentation design that suggests opportunities for research and tool development in information visualization. We examined design differences between four coarse academic disciplines (social, natural, and formal sciences; and the humanities) in two settings: electronic slideshows and prompted whiteboard presentations. The study characterized differences in visual representation design and narrative features, like building on slides or whiteboard drawings over time, between author disciplines. Two representations of slides were used: eigenslides, which illustrate differences in *visualness* in slides between disciplines, and a space of hand-selected features in which slideshows clustered by discipline.

A second more controlled study of whiteboard presentations showed that some participants used representations and argument styles from their own domains. One behavior we observed was formalizing, or reframing of topic explanations into an accustomed level of formality; for instance, a computer scientist turned an explanation about walking routes into a geometric proof that ignored details like traffic or curved streets. This suggests a dimension of presentations we called *systematicness* that seems to be orthogonal to visualness. While the amount of visual representations used might indicate visualness, the

type of representations (e.g., recognized information-visualization constructs versus stylized imagery) or the way they are used indicates systematicness.

In general, we found evidence for a space of academic disciplines that is linked to patterns in visual-presentation design conventions. Several implications for information visualization follow:

- Narrative visualization designers can learn the representations familiar to specific domains by doing observational studies of visual artifacts in those domains.
- Metaphors, similes, and examples can illustrate abstract ideas alongside or in place of conventional diagrams or visualizations.
- Infovis systems that support design activities, including tools like Microsoft PowerPoint, could incorporate domain-specific design models for templates or recommender systems.
- Infovis systems that support finding and sharing extra context might be preferred by “eager beavers” over other systems. Additionally, a crowd-powered interface could help identify and evaluate extra context or narrative embellishments.
- Infovis systems that support retargeting visualizations to different formality levels are likely to make visualizations more effective for a diverse audience.
- Showing the stages of diagram development could be an effective construct for explanatory visualizations, especially in highly systematic domains, like mathematics or natural sciences.

While this work focused on presentations, we believe studying other hand-created visualizations presents an opportunity to learn new principles and generate new hypotheses for visualization. Studies like these can give insight about the human factors and data factors, like scale or uncertainty, that affect design choices for visualization. Possible applications include novel user interfaces that make it easier to author and publish visualizations effectively to a wide audience.

Acknowledgments

This work was supported in part by NSF award IIS-10-16623 and NIH award R01-EB004155. All opinions, findings, conclusions, or recommendations expressed in this document are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

References

- [1]. Agrawala M, Li W, Berthouzoz F. Design principles for visual communication. *Commun. ACM*. Apr.2011 54:60–69. [PubMed: 21984822]
- [2]. Agrawala, M.; Stolte, C. Rendering effective route maps: Improving usability through generalization. *Proceedings of SIGGRAPH*; 2001. p. 241-250.
- [3]. Alibali MW, Bassok M, Solomon KO, Syc SE, Goldin-Meadows S. Illumination mental representations through speech and gesture. *Psychological Science*. 1999; 10(4):327–333.
- [4]. Bernstein, MS.; Little, G.; Miller, RC.; Hartmann, B.; Ackerman, MS.; Karger, DR.; Crowell, D.; Panovich, K. Soylent: a word processor with a crowd inside. *Proceedings of the 23rd annual ACM symposium on User interface software and technology, UIST '10*; New York, NY, USA: ACM; 2010. p. 313-322.

- [5]. Biglan A. The characteristics of subject matter in different academic areas. *Journal of Applied Psychology*. 1973; 57(3):195–203.
- [6]. Cherubini, M.; Venolia, G.; DeLine, R.; Ko, A.J. Let's go to the whiteboard: how and why software developers use drawings. *Proceedings of the SIGCHI conference on Human factors in computing systems, CHI '07*; New York, NY, USA: ACM; 2007. p. 557-566.
- [7]. Heiser, J.; Phan, D.; Agrawala, M.; Tversky, B.; Hanrahan, P. Identification and validation of cognitive design principles for automated generation of assembly instructions. *Proceedings of Advanced Visual Interfaces*; 2004. p. 311-319.
- [8]. Hullman J, Diakopoulos N. Visualization rhetoric: Framing effects in narrative visualization. *IEEE Transactions on Visualization and Computer Graphics*. 2011; 17(12):2231–2240. [PubMed: 22034342]
- [9]. Kirsh D. Thinking with external representations. *AI Soc*. Nov.2010 25:441–454.
- [10]. Larkin J, Simon HA. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*. 1987; 11:65–99.
- [11]. Li W, Agrawala M, Curless B, Salesin D. Automated generation of interactive exploded-view diagrams. *ACM Transactions on Graphics*. 2008; 27(3):101:1–101:11.
- [12]. Liu Z, Stasko JT. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE Transactions on Visualization and Computer Graphics*. 2010; 16(6): 999–1008. [PubMed: 20975137]
- [13]. Savva M, Kong N, Chhajta A, Fei-Fei L, Agrawala M, Heer J. Re-Vision: Automated classification, analysis and redesign of chart images. *ACM User Interface Software & Technology (UIST)*. 2011
- [14]. Segel E, Heer J. Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics*. 2010; 16(6):1139–1148. [PubMed: 20975152]
- [15]. Turk M, Pentland A. Eigenfaces for recognition. *J. Cognitive Neuroscience*. Jan; 1991 3(1):71–86.
- [16]. Tversky, B.; Heiser, J.; Lee, P.; Daniel, M-P. Explanations in gesture, diagram, and word. In: Coventry, KR.; Tenbrink, T.; Bateman, J., editors. *Spatial Language and Dialogue*. Oxford University Press; 2009.
- [17]. Tversky, B.; Suwa, M. Thinking with sketches. In: Markman, A., editor. *Tools for Innovation*. Oxford University Press; 2009.
- [18]. Walny J, Carpendale S, Henry Riche N, Venolia G, Fawcett P. Visual thinking in action: Visualizations as used on whiteboards. *IEEE Transactions on Visualization and Computer Graphics*. Dec.2011 17:2508–2517. [PubMed: 22034372]

Fig. 1.

Principal components of electronic slides, called *eigenslides*, are shown in (a). Contrast patterns seen in eigenslides computed from different academic disciplines indicate some slide-design conventions in these disciplines, like using primarily images or text. Some discipline-related design conventions, like building on representations over time and using visualization constructs, are used by authors even when topics were controlled in video-recorded whiteboard presentations (b).

Fig. 2.

Eigenslides created by running PCA on image samples from each slideshow discipline. The left of each subfigure shows the mean slide for each set; the right side shows thumbnails of the top 15 principal components (ranked left to right, top to bottom).

Fig. 3. Cumulative variance curves from eigenslide analysis. Three curves for each discipline are plotted showing repeated sampling of the slide sets. Curves within each discipline are similar to one another and distinct from other disciplines' curves.

Fig. 4.

Semantic features across slideshows in each discipline. Features with significant differences between at least one pair of disciplines are marked as follows: * $p < .05$, ** $p < .001$.

Fig. 5. Visualization of differences between slideshows using MDS on coded features. MDS is performed on slideshows from all disciplines (left) and pairwise (right).

low-spatial
↑
high-spatial

[WIFI]: "When many customers' laptops are connected to the public WiFi at a coffee shop, each internet connection is slower, on average, than when only one laptop is connected."	[ROOM]: "In a room with four people, who may or may not be acquaintances with others in the room, at least two people in the room have the same number of total acquaintances. (Note that if person A is an acquaintance of person B, then B is also an acquaintance of A.)"
[CARS]: "If three similar cars are driving single file, all at the same speed, the leading car will use more gas than the other two. There will not be much difference in gas use between the two rear cars."	[PHOTO]: "To take a non-blurry photograph of a car driving on the highway, one's camera must use a faster shutter speed than when photographing a slow-moving car."
[MAP]: "From the Brown University main green, it takes less time for the average person to walk to Kennedy Plaza than to walk to the Providence Place Mall."	[DOOR]: "When pulling on a typical household doorknob, a closed door will not open until the knob is turned. Once opened, however, the door can be shut closed without turning the knob."

Fig. 6.
Inventory of explanation prompts. The bold labels and spatial categorization were not shown to participants in our study.

Fig. 7.

Semantic features across presentations in each discipline. Features that were not used by participants during any whiteboard talks, like titles and paragraphs, are not shown. Differences between features were not significant in our 6×4×4 study design.

Fig. 8.
Different approaches to the [MAP] prompt from participants in humanities (a), social sciences (b), natural sciences (c), and formal sciences (d).

Fig. 9.

Familiar representations. From the [ROOM] prompt, (a) shows an adjacency matrix created by a cognitive scientist (social) and (b) is a exhaustive set of node-link diagrams created by a computer scientist (formal).

Fig. 10. Formalizing in the [MAP] prompt by a computer science student (a), and de-formalizing in the [ROOM] prompt by a public health student (b).

Table 1

Composition of the slideshow dataset.

Discipline	# Slide-shows	Total # slides	Fields of Study
Humanities	15	299	Philosophy, Italian Studies, American Studies, Architecture, Landscape Arch., Art History, Middle East Studies, Near Eastern Language and Civilization, Film and Visual Studies
Social Sci	15	472	Sociology, Economics, Cognitive Science, Public Health, Psychology
Natural Sci	20	601	Environmental Studies, Geological Sciences, Neuroscience, Biomechanics, Biology, Chemistry, Physics
Formal Sci	15	638	Computer Science, Biostatistics, Applied Mathematics, Mathematics

Table 2

Presentation coding scheme.

Type	Feature Description
Diagrammatic	Infovis construct for numerical data (e.g., charts, matrices)
	Infovis construct for relations (e.g., trees)
	Diagram showing clarified structure of natural scene
	Other visualization
	Photograph or visual art
	Mark used for emphasis
	Color used for emphasis
	Color used for organization
Textual	Numeric equation or expression
	Qualitative equation or expression (e.g., abstract operators for non-numeric operands)
	Caption that describes an image directly
	Caption that interprets an image (e.g., commentary about meaning or significance)
	Label on part of a visual element
	Ellipses (...) and etc. (e.g., suggesting additional information or list items)
	Bullet-point style text
	Paragraphs
	Title
Narrative	Building on or repeating a previous slide
	Simile, metaphor, or example