

An Exploratory Study of Data Sketching for Visual Representation

J. Walny¹, S. Huron^{1,2} and S. Carpendale¹

¹University of Calgary, Canada, ²IRI Centre Pompidou, Paris, France

Abstract

Hand-drawn sketching on napkins or whiteboards is a common, accessible method for generating visual representations. This practice is shared by experts and non-experts and is probably one of the faster and more expressive ways to draft a visual representation of data. In order to better understand the types of and variations in what people produce when sketching data, we conducted a qualitative study. We asked people with varying degrees of visualization expertise, from novices to experts, to manually sketch representations of a small, easily understandable dataset using pencils and paper and to report on what they learned or found interesting about the data. From this study, we extract a data sketching representation continuum from numeracy to abstraction; a data report spectrum from individual data items to speculative data hypothesis; and show the correspondence between the representation types and the data reports from our results set. From these observations we discuss the participants' representations in relation to their data reports, indicating implications for design and potentially fruitful directions for research.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Interfaces and Presentation]: User Interfaces—

1. Introduction

Creating mappings from data to visual representations is a key issue when developing visualizations. Theories of representational mapping arise from perceptual theory and experiments [CM84, War04], from graphical primitives [Ber11], by the visualization pipeline process [CSM99], through use of metaphors [Bla06], and from specific data and tasks [AS04, BM13]. To complement these approaches, a growing body of research has started to study how people create visualizations, sketches, and diagrams for themselves. Previous investigations have included observing how novices use existing software with the aid of a human mediator to develop visualizations [GTS10], observing how people spatially reorganize nodes and links to create meaningful graph layouts [vHR08], observing the residue of visual thinking as left on whiteboards [WCR*11], and examining the evidence of visual communication used to explain ideas [GJZ*12]. These initial investigations have indicated that this is a rich avenue of research that can potentially lend insight into how to best create visualizations and improve our chances of fulfilling the original goal of information visualization – generat-

ing interactive visual tools that amplify cognition [CSM99]. While holding this long-term goal, our practical short-term goals involve questions about data sketching.

We extend this research stream by studying how people, without the aid (or inherent constraints) of software, take data represented as a table of numbers and transform it into visual representations via freehand sketching. The data sketching process is of interest because sketches or drawings can be seen as personal, spontaneous visualizations created to represent internal thought. Using hand-drawn sketches and diagrams is known to be effective in promoting innovation, creativity, and thinking in general [Bux07, GCMB11]. However, the relationship between digital visualizations and people's thinking processes remains elusive [Nor06, SND05, CZGR09]. Grammel et al.'s study of visualization novices [GTS10] suggests that people had difficulties creating visualizations via software even with the aid of a human assistant, yet writings abound about how readily people use sketches during thinking [Kir10, Tve11]. However, the literature that praises people's general facility

with using sketching as part of ideation does not talk about sketching data. This raises questions such as:

- What diversity of representations do people produce when asked to sketch data manually?
- Is there a relationship between sketching data and people's understanding of that data?

Our goal is to develop a deeper understanding of the answers to these questions with a closer look at data sketching.

2. Related Work

Several different senses of the term “sketching” have arisen in previous literature. “Sketchiness” can refer to the rendered visual style [WII*12]; it has also been explored as a potential visual variable [WII*12, BBIF12]. Sketching can also refer to sketch-based interfaces for visualization [CMvdP10, BLC*11, WLJ*12, LKS13]. In the domain of human-computer interaction, sketching is often used as a rapid ideation technique for interface design [Bux07]. More generally, sketching is an accessible and common form of externalization, a reification of internal thought [Kir10].

It is in this last sense of sketching as externalization that we are most interested. While a sketchy visual style is an inevitable by-product of drawing sketches by hand, the visual style is not our focus in this study; rather, we were interested in harnessing the characteristics that make sketching a popular ideation and externalization tool. The nature of sketching for ideation is that it is rapid and has a low barrier to entry: one needs only a writing implement and paper; no artistic skills are necessary [GCMB11]. It also does not restrict the spatial mapping of elements: sketches can be very freeform. This type of sketching is of particular interest in fields involved in the creation of physical objects (such as design [Sch83], architecture [Gol94], and engineering [Fer94]) and those focused on creating digital objects, such as human-computer interaction [Bux07].

Tversky argues that sketches represent thought by revealing an “idea” of what the sketcher has in mind [Tve08]. In this vein, several researchers have turned to studying manually manipulated representations to inform the design of visualizations. These have largely focused on improving graph-drawing aesthetics [vHR08, RPHHP07, PPBP10, DLF*09] or automatically generating visualizations [TAH*06]. Walny et al. studied how hand-drawn representations of information can inform the design of information visualizations [WCR*11] and Huron et al. [HJC14] have deconstructed the visual representation process by observing people creating representations using tangible tokens. Meanwhile, Gomez et al.'s study has revealed that there are differences in the way people from different disciplines create visual presentations of information [GJZ*12]. These human-in-the-loop studies shed light on aspects of representational mapping that can be difficult to see when focusing on the more established data- or task-driven approaches to developing new representations.

Established approaches to representation for visualization have been strongly guided by research into human visual

perception [War04, CM84], based on the data type being represented [Shn96, CM97, TM05], or guided by the tasks that are to be performed to gain an understanding of the data [AS04, BM13, SNHS13]. Recently, more attention has been paid to how representations are not just perceived, but also understood. Kindlmann and Scheidegger proposed an algebraic approach to representation, focusing on whether changes to data have corresponding changes in the representation [KS14]. Visual metaphors have been investigated [Bla06, ZK08]. A greater emphasis on representations that support the broader goal of insight generation has been called for [SND05, Nor06, PFG07, YKSJ08, CZGR09]. In our study we were more concerned with this understanding-based approach than with strict perceptual accuracy, focusing on what information could be read from a representation.

Most approaches to representation design are geared towards expert information visualization designers developing representations for use on specific datasets or data types and with specific tasks in mind, ideally in consultation with domain experts. However, tools are emerging that empower people to specify their own visualizations. Myers et al.'s tool automatically generates visualizations based on drawn examples [MGG94]. Pretorius and van Wijk's tool for visualizing system traces allows analysts to define schematic diagrams based on their personal conceptualization of the system [PvW08]. Several commercial and publicly available tools allow people to create basic visualizations of their data using a limited set of pre-defined representations, for example ManyEyes [VWvH*07], GraphSketcher [SS09], Tableau [Tab], Tibco Spotfire [Tib], charting tools available in Microsoft Excel [Mic] and Adobe Illustrator [Ado], and Google Charts [Goo]. Most of these tools are restricted in the variety of their representations. Non-experts who wish to create more freeform data representations are currently better supported with the release of popular data visualization toolkits such as D3 [BOH11] and simplified environments built on top of D3, such as Vega [Tri]; however, these environments still require non-trivial programming skills. A promising concept in this regard is a tool demonstrated by Victor for creating freeform data visualizations using an interface similar to a drawing program [Vic]. Despite the strides made in supporting non-experts in creating their own representations, current tools do not match the combination of representational freedom and ease of use that sketching provides.

3. Data Sketching Study

In our qualitative study, we asked people to sketch a visual representation of a small data set and briefly tell us what they learned or found interesting about the data.

Participants. We ran three hour-long sessions with 7, 8, and 7 participants in each, for a total of 22 unique participants, 13 male and 9 female. One participant had completed some university courses; 6 had completed a Bachelor's degree; 13 had completed a Master's degree; and 2 had completed a PhD degree. Eighteen participants had education in computer-related fields such as computer science or software

engineering; of those, 3 reported education in graphics or visualization, and 1 in business. Additionally, we had one participant each in degrees related to design, communication-illustration, kinesiology, and languages, and one did not report this information.

Set-up and Materials. We used a simple classroom with ample tables and chairs and good lighting. We provided blank sheets of standard letter-size paper (participants could use as many as they liked) and a large variety of coloured pencils. Each participant received a printout of the dataset.

3.1. Dataset

We spent considerable effort looking for an appropriate dataset. Our goal was to find a universally interesting and, if possible, entertaining dataset so as to keep participants engaged. It was also important that participants find the data understandable without need of expert knowledge; ideally, participants would feel familiar with the data at least to some extent. Our chosen dataset contains appropriateness ratings of behaviours in different social situations extracted from a 1974 social psychology study performed on university students [PB74]. (This dataset is readily available in the common statistics package SPSS [IBM].) This data could yield interesting insights even to those outside the field of social psychology because it deals with common everyday situations. Thus, we consider our participants experts in understanding this data from a human perspective. In this way, we were able to gather manually sketched data representations from a group of experts who all had an intimate understanding of the data and pertinent domain information.

The dataset was presented as a table containing mean scores of appropriateness of behaviour for 15 behaviours and 15 situations for a total of 225 behaviour-situation pairs. Values ranged from 0 (highly inappropriate) to 9 (highly appropriate). We used the same dataset for all three sessions and did not provide any additional information from the original study, such as summary statistics or standard deviations. The data set was printed on one sheet of letter-size paper.

3.2. Procedure

Participants entered the room during one of three group sessions and were free to sit at whichever seat they chose. They were provided with all necessary materials. We scripted the explanation of the dataset and the data sketching rules to ensure consistency across different sessions. We asked participants to represent the data on the blank paper and told them that how they sketched was entirely up to them and that there was no right or wrong method. We encouraged them to draw the data as they explored it and, in order to give minimal guidance to participants unfamiliar with visualization, pointed out that they might want to think about: connections between different pieces of data; ways to group the data; similarities and differences in the data; interesting patterns; and surprising findings in the data.

We planned answers to common questions to avoid influencing the representations. For instance, if asked how to

draw something, we answered ‘any way that makes sense to you’. Though participants were free to talk, having a prepared script made for a more formal atmosphere, so talk tended to be minimal. When a participant handed in a sketch, or after 45 minutes, we administered a simple post-sketching questionnaire, which asked for basic demographic information, experience in visualization, and one important question with a full page for possible answers. The question was: *Please describe what you learned or found interesting about this data during the session (there are no wrong answers).*

4. Analysis

In our analysis, two of the authors used a qualitative approach to analyze the three types of artefacts we gathered: data sketches, data reports, and dataset printouts. Working first from the data sketches alone, through careful examination, affinity diagramming, several coding passes, and finally working together until agreement was reached, we arranged the sketches on a continuum from numeric to abstract representations. We considered a representation sketch as *numeric* to the degree to which the actual numeric data was directly readable from the sketch, and increasingly *abstract* to the extent to which the data had been manipulated or worked with before being graphically represented in the sketch.

Separately, we divided the data reports into individual statements and analyzed these using an open coding approach [SC98], which resulted in a spectrum of statements ranging from referencing individual data items to statements containing speculative data hypotheses. Here, one author performed open coding; two authors performed focused coding [Cha06] independently, then discussed each coding instance until consensus was reached. Taking note that many participants had made additional markings, notes and annotations on the data sets they had been given as well as on their own representational sketches, we also included these in our data reports analysis.

5. Results

We present the results in terms of a representation continuum, a spectrum of data reports, and augmentation and annotation of the dataset printouts.

5.1. Representation Continuum

We collected at least one data sketch from each of our 22 participants, resulting in 35 representations in total, with as many as 4 coming from a single participant. Uniqueness and diversity of representation types were an important factor, however, although each data sketch had unique qualities, there were commonalities by which we were able to group them.

Despite the fact that 41% (9/22) of participants reported no previous experience with visualization, many of the sketched representations could be classified as known types, as follows: The most frequent representation we received was the bar chart – 12 in total coming from 9 different participants. There were also 3 countable representations; 5 dot




Figure 1: Sample representations along a continuum of numeracy to abstractness. Tokens indicate how many representations of each type were collected. An extended version of this figure is available online: <http://bit.ly/datasketching>.

plots and matrices; 3 types of line graphs or parallel coordinates; 3 ranked lists; 3 graph-like representations, 2 that relate best to Venn diagrams and 4 that were mostly pictorial. Representations did not always fall neatly into each category; for instance, one graph-like representation (seen in Figure 3A) bore some similarity to a Venn diagram. For clarity, we report it here in the graph-like category.

Through our analysis we derived a continuum of representations from numeracy to abstractness, defined as in Section 4: the more fully the raw data could be read from the representation, the more *numeric* the representation was; conversely, the more the data had been manipulated, transformed, or otherwise abstracted, the more *abstract* the representation was. For instance, token-filled matrices (Figure 1), which allow the viewer to retrieve the cardinality of a given datum, are placed in the numeric range, on the left. In contrast, a diagram that shows binned over-lapping sets of data (Figure 3) is closer to the abstract end of the continuum. If a representation of our dataset conveys information at the numeric level, one can read the numerical values of the appropriateness levels of specific behaviour-situation pairs with some accuracy. However, note that some representations had the actual numbers written on them, performing a tooltip-like function of typical visualizations. We treated these as labels rather than numerical representations.

We illustrate this continuum with samples and descriptions of several points along the continuum (Figure 1). Note that while the samples are discrete and grouped into similarity clusters, this is a continuum because a representation can always be created between any given two representations.

Countable Representations (3) The most numeric representations were simple tallies and countable tables: here, the dataset printout was reproduced quite faithfully with countable points or tokens used instead of numbers. In some cases, smaller points were used to indicate decimal amounts.

Dot/Scatter Plots and Matrices (5) The dot / scatter plot variations were quite numeric representations. Although the numbers here are not read as directly as in countable variations, the reading method is relatively well established. One representation we collected used one categorical axis and one numerical axis, making it more similar to a dot plot than a scatterplot (Figure 2A). In this particular sketch, many data aspects are coded in multiple ways. The position of the

points allows for a fairly precise numerical reading (considering that the representation is hand-drawn). The size of the points double-encodes the numerical value and the colour of the background groups the behaviours into categories of high, medium, or low appropriateness.



Figure 2: Parts of a dot plot and two matrices collected during the study.

Bar Charts (12) We saw 12 instances of various bar charts, including stacked bar charts, ones grouped according to situation or activity type, and ones with embedded descriptive icons in the bars. Bar charts are generally numeric, however several of the bar charts we collected showed aggregates and groupings of data points rather than showing the raw data directly, moving them towards the right along the continuum.

Line Charts (1) and Coordinate Plots (2) We saw only one line graph; this low number makes sense due to the non-serial nature of the data set. In fact, the line chart could easily have been turned into a parallel coordinates plot by adding vertical bars through points in a category. There were two coordinate plots, one parallel and one radial (a star plot). Both coordinate plots were organized similarly, by situation. The parallel plot also included small pictorial symbols as labels – a smiling face for appropriate behaviours, and a frowning one to indicate inappropriate behaviours.

Ranked Lists (3) Ranked lists incorporated data binning, grouping and ranking. They made simple use of colours and a considerable amount of labeling.

Graphs (3) Graphs or graph-like representations all used some form of grouping or binning of the depicted values, providing compact overviews of entire data dimensions (situations or behaviours). One graph linked behaviours with situations below some appropriateness threshold. Another linked situations to colour-coded behaviours arranged by appropriateness level. A “graph-like” representation, shown in Figure 3A, linked similar behaviours across situations.




Figure 3: Graph-like representations that place appropriate behaviours in orbits around a situation and link behaviours across situations.




Figure 4: A data sketch of a “decision helper” tool

Venn diagrams (2) Venn diagram-like representations showed relations between particular behaviours or situations as a whole. One was clearly approximate, showing the number of shared “appropriate” behaviours per situation. The other was a hybrid of bar chart and Venn diagram: plotting one value as a horizontal bar and one value as an overlapping vertical bar, the diagram depicted a relationship between two values. In small-multiple format, this showed interesting parallels between two different situations (dating and restroom). The participant used an aggregate value derived from grouping of behaviours as ‘more’ or ‘less’ physically active.

Pictorial Representations (4) At the far right end of the continuum, we placed pictorial or storytelling representations. These used pictorial icons such as stick figures and line drawings to present the data as a story. These representations depicted approximations of the data, mostly by summarizing dimensions of the data. For instance, one participant drew icons depicting select situations, then wrote short summaries, e.g. “NO. Most behaviours highly inappropriate”. One participant sketched a “decision helper” representation (see Figure 4). It depicts a person moving through various situations during the day and uses the data suggest the most appropriate behaviours for each situation.

5.2. Spectrum of Data Reports

Participants’ data reports provided a glimpse of the knowledge that they had discovered about the data. Some participants filled the page, most wrote about one-third to one-half of a page, and some also annotated their visualizations and data sheets with data reports. Most statements discussed the

data. We categorized the data reports across these responses, developing a spectrum with direct readings of individual data values at one end and relatively high-level conjectures and fledgling hypotheses on the other end (see Figure 5).

Twelve participants added statements about how to read their visualizations and about their process in creating them. This included comments about the type of visualization participants were sketching, explanations of colour schemes, and assertions that people’s sketches would differ depending on their interests. We call these *process statements* and have not placed them in the data report spectrum. We now explain each portion of the data report spectrum in detail.

5.2.1. Information Intrinsic to Dataset (A, B, C)

Statements about *individual data items* (Figure 5-A1 and A2) addressed individual values in the data set, for example, “fighting in church is inappropriate”. No statements included precise numeric values, but some included relative values (e.g. “> 0”) or general values (“OK”).

Nineteen participants *summarized individual rows or columns* (Figure 5-B1) with statements that described individual situations or behaviours as a whole, for example, “there aren’t many behaviours appropriate in church.” Again, no statements included precise mean values, although annotations on dataset printouts showed that some participants had calculated such values.

Some data reports *compared two rows or columns* (Figure 5-C1), making specific pairwise comparisons between situations or behaviours. For example, “Date and own room have the similar rating for ‘kiss’. Other ratings in these two situations are close to each other.”

5.2.2. Comparisons and Trends Within Dataset (D)

Several data reports *compared three or more rows or columns*, noting trends (Figure 5-D1). They also *classified* rows or columns (Figure 5-D2), *grouped* items according to values (D3), and made *global comparisons* (Figure 5-D4). For example, “Several situations which have a similar rate for one specific behavior tend to be similar for other behaviors.” Some statements indicated that values, rows, or columns had been classified according to some scheme based on the information in the data set. For instance, several people binned values into levels of appropriateness (“completely appropriate”, “somewhat appropriate”, “highly inappropriate”); one participant grouped situations based on overall similarities in their values.

5.2.3. Including Extrinsic Information (E)

In 15 participants’ data reports, the data was in some way related to external information, for example through *classification* (Figure 5-E1). These statements classified values, rows, or columns similarly to that in Figure 5-D2, but with an added judgment on the meaning of the classification. For example, high appropriateness of behaviour was related with comfort, safety, or aggressiveness; one participant characterized some situations as “work-related”. Others related external concepts through *comparison to expectations* (Figure 5-E2). Their statements indicated that the participants

	A. Individual data points	B. Low level summary	C. Low-level comparison	D. Comparisons and Trends in the dataset	E. Including Extrinsic Information	F. Statements with Analytic Potential
1	Refer to specific individual data values	Summarize individual rows or columns.	Compare two rows or columns.	Compare 3 or more rows or columns; note trends.	Classify as named type.	Fledgling hypothesis.
2	Refer to approximate individual data values.			Classify values, rows or columns.	Compare against expectations.	Conjecture about reasons for values.
3				Group items according to value.	Explain in domain context.	
4				Make global comparisons.		

Figure 5: The spectrum of data report statements.

had used their experience to set expectations for the data, and then compared those to the actual data, for instance, “no surprises” or “mumbling + talking diverged more than expected.” Another way of relating the data to external concepts was by adding *explanations in the domain context* (Figure 5-E3), such as explaining the data in terms of social behaviour. For example, “people care a lot in job interviews” was used to explain the overall low appropriateness ratings in that situation.

5.2.4. Analytic Potential (F)

Several participants made hypotheses or conjectures about the reasons behind the values in the dataset (Figure 5-F1 and F2). One participant hypothesized that the similarity between *park* and *own room* might be due to the relative anonymity of both situations. Another speculated that there were more females than males in the original survey because the acceptability of talking in bathrooms was quite high.

5.3. Augmentation and Annotation

Almost all participants (19) added various types of annotations to their dataset printouts while developing their visual representations. We categorized these, first by amount of activity and then by the type or purpose of the activity. We characterised 11 printouts as lightly annotated, 6 as moderately annotated, and 2 as heavily annotated. Individuals would often employ more than one type of annotation. We found people using picking (11), grouping (10), adding symbols (9), filtering (2), and more generally adding words, calculations, notes, labels, and values (12). One participant even used physical manipulation – folding – to juxtapose different parts of the data.

6. Integration of Results

We consider the representation continuum and the data report spectrum in combination. We note the apparent relationship between the continuum and the spectrum amongst our participants and consider the variances in representation in light of the data reports. From this comparison we derive a specification of the levels of data description that a representation and a data report can encompass (shown in Figure 7).

6.1. Considering the Representation Continuum and Data Report Spectrum in Combination

The relation of participants’ data reports to the position of their data sketches along the representation continuum can

be seen in Figure 6, ordered from those that submitted the most abstract sketches to the most numeric (in the case of multiple submitted sketches, the most abstract was chosen for the ordering). From this juxtaposition, we can see that in our sample, the participants who submitted the most abstract sketches were among the participants whose data reports tended to be in categories E3 (including extrinsic information) and F (statements with analytic potential). Note that in this figure, columns A1 and A2 were collapsed into A due to their similarity and the fact that nearly every participant provided data reports in these categories; F1 and F2 were also collapsed due to their similarity.

The apparent correspondence between the data reports and the representation continuum, while inconclusive, does suggest an approach to considering a more formal specification of the variation in representations based on the level of data description they contain. We define these levels based on what we saw in the data reports and representations.

6.1.1. Levels of Data Description

We describe the levels of data description on the basis of a numerical data set D of data points d , each of which is a tuple $(v_1, v_2, \dots, v_n, x)$ where every v_i is a dimension and x is a numerical value. In our dataset, each data point is a triple (*behaviour*, *situation*, *appropriateness rating*), where *appropriateness rating* is a numerical value between 0 and 9. We define the *level of data description* of a representation in terms of the information that can be retrieved about the data points in the dataset. This information can be conveyed about a data set at the level of *values* (1), *dimensions* (2), *globally* (3), or in terms of *external information* (4). Figure 7 shows examples of these levels on a two-dimensional dataset. Next we discuss each one of these levels separately.

(1) Value Level of Data Description. Representations at the value level convey information about: *individual raw values* if a set of marks in the representation can be interpreted as a specific data point $d = (v_1, \dots, v_n, x)$; *pairwise comparisons* between raw values if a set of marks in the representation can be interpreted as showing two data points d_1 and d_2 and the relationship between them; and *trends* (comparisons between three or more raw values) if a set of marks in the representation can be interpreted as showing a set of data points d_1, \dots, d_n (where $n > 2$) and the relationships between them.

This level of data description is derived from category A of the data report spectrum. A representation conveying in-

Data Report Spectrum													Representation Continuum			Others		
A	B	C	D1	D2	D3	D4	E1	E2	E3	F	Max					P	A	X
													P	P				
													P	P				
													P	P				
													P	P				
													V	B	V	C		
													V	B	V			
													G	G	G			
													G	G				
													R	R	R	R		
													L	D	D	B	L	
													L	L				
													B	B				
													B	B				
													B	B	B	B		
													B	B				
													B	B				
													B	C	B	B		
													D	D				
													D	D				
													D	D				
													D	D				
													C	C				

Figure 6: Summary of artifacts collected for each participant. Columns show (L-R): Statements on the data report spectrum (A-F), corresponding to Fig. 5; Participant's most abstract sketch along the representation continuum (Max); Classification of submitted representations – letters correspond to those in Fig. 1; Presence of process comments (P); Presence of full-sentence annotation on data sketches, in yellow (A); Self-reported experience in visualization (X) – darkest denotes most experience. Rows are ordered by Max, from most abstract to most numeric. The abstract representations tend to correspond to data reports in columns E3 and F.

formation at the value level would be at the numeric end of the representation continuum, e.g. a countable matrix.

(2) Dimension Level of Data Description. At the dimension level, a representation conveys information about *individual dimensions* if a set of marks in the representation can be interpreted as a summary description of a set of data points $D_{i,k} \in D$ for which the value of dimension v_i is fixed to k . For a two-dimensional data set, the set of marks describes one entire row or column in the data set. For instance, in our dataset we might represent the average appropriateness of 'talking' across all situations by representing the average value in the set of points with dimension *behaviour* fixed to the value 'talk'. A representation conveys information about *groups of dimensions* if a set of marks in the representation can be interpreted as a summary description of a group of individual dimensions, i.e. a set of data points $D_i \in D$ for which the value of dimension $v_i \in \{k_1, \dots, k_n\}$; and *dimensional trends* within three or more dimensions or groups of dimensions if a set of marks in the representation can be interpreted as summaries of three or more dimensions (or groups of dimensions) and the relationships between them.

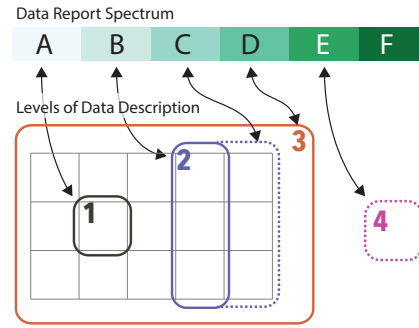


Figure 7: Relation of data report spectrum to the levels of data description that a representation can convey about a two-dimensional dataset: (1) value level, (2) dimension level, (3) global level, and (4) external information.

sional pairwise comparisons between individual or groups of dimensions if a set of marks in the representation can be interpreted as summaries of two dimensions (or groups of dimensions) and the relationship between them; and *dimensional trends* within three or more dimensions or groups of dimensions if a set of marks in the representation can be interpreted as summaries of three or more dimensions (or groups of dimensions) and the relationships between them.

The dimension level of information stems from columns B and C of the data reports spectrum, describing individual rows or columns and comparing pairs of rows or columns.

(3) Global Level of Data Description. When a set of marks in a representation can be interpreted as an aggregate description of the entire set of data points in the data set D , they can be said to offer a *global* level of description. While we did not collect any representations that described the entire data set in aggregate (unsurprising given the amount of manual work that would have required), we did receive some data reports (column D4) that corresponded to the global level of information. One possible example of a representation that would convey information at the global level for our dataset might be a matrix that orders behaviours and situations by similarity, allowing one to observe, for example, the set of behaviours that are inappropriate in most situations.

(4) External Information Level of Data Description. The data sketches and data reports we collected did not always include exclusively information contained in the data set; occasionally they included contextual or domain information. Thus we define an additional, *external* information level of data description, in which either the value, dimension, or global levels are related to external concepts or external information. For example, one participant ordered bars by level of privacy, a dimension not included in the dataset.

Hypothesis Generation Our data report spectrum contains an additional column, F, which includes hypotheses or conjectures about the data. We do not include this as a level of data description because, rather than describing the data, these signify the beginnings of analysis.

Precision of Data Description Regardless of the level of data description, a representation can convey information at various levels of precision. In the case of individual values, it can convey: precise, actual values (e.g. “7.37”); fairly accurate binned values (e.g. “between 7 and 9”); or approximate values (e.g. “high”); symbols and pictures, e.g. a happy face symbol, can also depict approximate or imprecise values.

7. Discussion

We discuss our results in light of previous work on characterizing the information that can be read from representations and the possibility of insight.

Our results confirm and extend Bertin’s [Ber11] concept of *levels of reading*. Bertin explains his “elementary”, “intermediate”, and “overall” levels of reading a representation on the basis of a simple one-dimensional dataset. Bertin’s “elementary level” is the reading of single values, the “intermediate level” refers to reading groups of values along a dimension, and his “overall level” refers to reading all items along a dimension at once. Bertin’s levels of reading do not map exactly to our levels of data description; they are based upon a less complex dataset. Our *value* level of data representation corresponds to Bertin’s elementary level or intermediate level when referring to multiple values at once; our *dimension* level can be associated with Bertin’s overall level. We additionally include the *global* and *external* levels.

There is a clear correspondence between what can be read from a representation and what can be understood from the data: the insights that can be gained. Ample literature defines “insights” in various ways, some of which overlap with our levels of data description. Saraiya et al. [SND05] consider insights as “units of discovery” that fall into categories of overview (corresponding to *global*), patterns and groups (roughly corresponding to *dimensions*), and details (corresponding to *values*). Yi et al. [YKSJ08] describe processes that provide insights rather than describing insights themselves, but we see some parallels: their “provide overview” aligns with the *global* level of data description; their “adjust” refers to viewing data at various levels of abstraction. Perhaps the closest correspondence to our collected data reports is the taxonomy of analytic events observed by Liu and Heer [LH14]. Correspondences to our results can be found in their observation (*value*), and generalization (*dimension*) events. Liu and Heer also observed that hypotheses include “prior knowledge or personal experience brought into working memory”, aligning with our observations of extrinsic information in the data reports. North [Nor06] also brings up extrinsic information, arguing that insights are “relevant” in the context of a domain, which gives the data meaning.

Tversky et al.’s [TMB02] Congruence Principle states that “the structure and content of the external representation should correspond to the desired structure and content of the internal representation,” suggesting that a representation affects how we think about data. Chang et al. [CZGR09] argue that many insights (“knowledge-building insights”) come from a foundation of knowledge about data. These obser-

vations account for, in part, the degree to which our participants incorporated their external knowledge into their representations. In this light, it is intriguing that it was the participants with the most abstract representations and with data reports drawing upon external knowledge who also tended to provide fledgling hypotheses and conjectures about the data.

8. Limitations

This small study provided us with a rich, thought-provoking, yet non-generalizable sample of representations and associated data reports. It was highly exploratory in nature, pointing to intriguing new research directions, but it does not allow us to make generalizations.

Due to the exploratory nature of this study, it is impossible to tell whether the depth of thought shown by participants was a result of the representations they drew, the process of drawing, their engagement with this particular dataset, their own analytic skills, or some combination of these. The study was not videotaped, nor were participants interviewed. Thus our results reflect what could be reasonably interpreted from the collected artifacts, but we cannot comment about whether participants analyzed the data before, during, or after drawing their data representations. (It should be noted, however, that questionnaires with the open-ended question were distributed at the end of the session). Nevertheless, these observations bear further investigation.

The question eliciting the data reports was open-ended; the absence of particular statements does not necessarily mean an absence of that type of thought. However, we can surmise that statements thought to be highly important by our participants were more likely to have been reported. Additionally, our study used one dataset, so our spectrum of data reports may not be complete for other dataset types.

The diversity, type, and completeness of representations was influenced by the tools provided - paper and coloured pencils. The manual nature of the tools made it difficult for participants to try many variations or to represent the complete dataset, and also to represent it precisely. Conversely, the freeform nature likely enabled the diversity of representations. Note that not all representations we received were immediately readable to us. In creating their own representations, people personalized the data according to what interested them. It is not clear how readable the representations would be for others, particularly given the lack of precision afforded by the tools used. However, again we note that the pictorial representations are immediately understandable at a meta-level, though the raw values are not readable.

9. Implications for Research and Design

Our study suggests several directions for research and design:

Potential use of levels of data description. Viewing a representation in terms of its levels of data description may be a useful tool for describing the representation and for assessing whether there are additional variations of representations that may offer different levels of data description. While a

representation may not include a particular level of data description, it may have validity at a different one.

Understanding the advantages of sketching in terms of representation and data understanding. We observed that participants who provided more abstract representations (e.g. pictorial sketches) tended to provide data reports with analytic potential. This raises several questions for further research: Is this relationship repeatable? Is it causal? Can it be inverted? Does the process of sketching a representation, rather than only viewing it, affect the analysis outcome? Is there more analytic activity involved in the process of creating an abstract representation than a numeric one?

Diversity of representations. We collected a diverse set of representations, perhaps due to the inherent expressivity of pencil and paper sketching, which allowed participants with various levels of visualization expertise to express their visual interpretation of the data. Some of these representations are less studied in the information visualization community, for example abstract pictorial representations. This raises questions about potential advantages of expanding the diversity of representations we study.

Informing the design of InfoVis authoring tools. There has been considerable interest recently in creating visualization authoring tools accessible to novices [GTS10, HJC14], with some taking a sketching-oriented approach [LKS13, Vic]. An emerging issue is the trade-off between expressivity and steepness of learning curve. While we did not study the comparative ease with which participants sketched data, we did observe that each participant created at least one sketch. This raises questions such as: Where does data sketching fall in terms of the ease of representation authoring? How can software support data sketching freedom?

10. Conclusion and Future Work

We have presented a qualitative analysis of 35 manually sketched representations of a small dataset and the resulting reports about the data. From our initial analysis we developed a visual representation continuum from numeracy to abstraction and a data report spectrum from statements about individual data values to conjectures and fledgling hypotheses. From these we derived a specification of the level of data description a representation or data report can have and compared it to several related specifications.

Participants who related their responses to concepts extrinsic to the dataset and those who drew pictorial representations of the data tended to have the deepest observations and questions about the data. Our findings from this exploratory study raise intriguing questions for further study, including: questions about the relationship between the sketching process, abstractness of representation, and data understanding; questions about individuals' ability to create representations given appropriate authoring tools; and questions about the benefits of diversifying the space of representations by considering representations at differing levels of data description, or abstraction.

We also offer additional factors for further exploration:

Our observations of participants' annotations on their data sheets points to interesting research possibilities in the area of active reading of visualizations, while the richness of data thought behind story-like representations reinforces the visualization community's interest in narrative visualization. Finally, the study findings point to personal creation of visualizations as a potential driver of deeper engagement with data.

11. Acknowledgements

This research was supported in part by NSERC, AITF, Surfnet, GRAND, and SMART Technologies.

References

- [Ado] ADOBE INC.: Adobe Illustrator. www.adobe.com. 2
- [AS04] AMAR R. A., STASKO J. T.: A knowledge task-based framework for design and evaluation of information visualizations. *IEEE Symposium on InfoVis* (2004), 143–150. 1, 2
- [BBIF12] BOUKHELIFA N., BEZERIANOS A., ISENBERG T., FEKETE J.: Evaluating sketchiness as a visual variable for the depiction of qualitative uncertainty. *IEEE Trans Vis Comput Graph* 18, 12 (Dec 2012), 2769–2778. 2
- [Ber11] BERTIN J.: *Semiology of graphics: Diagrams, networks, maps*. ESRI Press, 2011. 1, 8
- [Bla06] BLACKWELL A. F.: The reification of metaphor as a design tool. *ACM Trans. Comput.-Hum. Interact.* 13, 4 (Dec 2006), 490–530. 1, 2
- [BLC*11] BROWNE J., LEE B., CARPENDALE S., RICKE N., SHERWOOD T.: Data analysis on interactive whiteboards through sketch-based interaction. In *ACM ITS* (2011). 2
- [BM13] BREHMER M., MUNZNER T.: A multi-level typology of abstract visualization tasks. *IEEE Trans Vis Comput Graph* 19, 12 (Dec 2013). 1, 2
- [BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D3: Data-driven documents. *IEEE Trans Vis Comput Graph* 17, 12 (Dec 2011). 2
- [Bux07] BUXTON B.: *Sketching user experiences: Getting the design right and the right design*. Morgan Kaufmann, 2007. 1, 2
- [Cha06] CHARMAZ K.: *Constructing grounded theory*. Sage, 2006. 3
- [CM84] CLEVELAND W. S., MCGILL R.: Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association* (1984). 1, 2
- [CM97] CARD S. K., MACKINLAY J.: The structure of the information visualization design space. In *Information Visualization, 1997. Proceedings., IEEE Symposium on* (1997). 2
- [CMvdP10] CHAO W. O., MUNZNER T., VAN DE PANNE M.: Poster: Rapid pen-centric authoring of improvisational visualizations with NapkinVis. *Posters Compendium InfoVis* (2010). 2
- [CSM99] CARD S., SHNEIDERMAN B., MACKINLAY J.: *Readings in information visualization*. Turtleback, 1999. 1
- [CZGR09] CHANG R., ZIEMKIEWICZ C., GREEN T. M., RIBARSKY W.: Defining insight for visual analytics. *Computer Graphics and Applications, IEEE* 29, 2 (March 2009). 1, 2, 8
- [DLF*09] DWYER T., LEE B., FISHER D., QUINN K. I., ISENBERG P., ROBERTSON G., NORTH C.: A comparison of user-generated and automatic graph layouts. *IEEE Trans Vis Comput Graph* 15, 6 (Nov 2009). 2

- [Fer94] FERGUSON E. S.: *Engineering and the mind's eye*. The MIT Press, 1994. 2
- [GCMB11] GREENBERG S., CARPENDALE S., MARQUARDT N., BUXTON B.: *Sketching user experiences: The workbook*. Elsevier, 2011. 1, 2
- [GJZ*12] GOMEZ S. R., JIANU R., ZIEMKIEWICZ C., GUO H., LAIDLAW D. H.: Different strokes for different folks: Visual presentation design between disciplines. *IEEE Trans Vis Comput Graph* 18, 12 (Dec 2012). 1, 2
- [Gol94] GOLDSCHMIDT G.: On visual design thinking: The vis kids of architecture. *Design Studies* 15, 2 (1994). 2
- [Goo] GOOGLE INC.: Google Charts — Google Developers. URL: developers.google.com/chart/. 2
- [GTS10] GRAMMEL L., TORY M., STOREY M.-A.: How information visualization novices construct visualizations. *IEEE Trans Vis Comput Graph* 16, 6 (Nov 2010). 1, 9
- [HJC14] HURON S., JANSEN Y., CARPENDALE S.: Constructing visual representations: Investigating the use of tangible tokens. *IEEE Trans Vis Comput Graph* 20, 12 (Dec 2014), 2102–2111. 2, 9
- [IBM] IBM: IBM SPSS software. URL: www.ibm.com/software/analytics/spss/. 3
- [Kir10] KIRSH D.: Thinking with external representations. *AI & Society* 25, 4 (Nov 2010). 1, 2
- [KS14] KINDLMANN G. L., SCHEIDEGGER C. E.: An algebraic process for visualization design. *IEEE Trans Vis Comput Graph* 20, 12 (Dec 2014). 2
- [LH14] LIU Z., HEER J.: The effects of interactive latency on exploratory visual analysis. *IEEE Trans Vis Comput Graph* 20, 12 (2014). 8
- [LKS13] LEE B., KAZI R. H., SMITH G.: SketchStory: Telling more engaging stories with data through freeform sketching. *IEEE Trans Vis Comput Graph* 19, 12 (2013). 2, 9
- [MGG94] MYERS B. A., GOLDSTEIN J., GOLDBERG M. A.: *Creating charts by demonstration*. ACM SIGCHI, 1994. 2
- [Mic] MICROSOFT: Microsoft Excel. office.microsoft.com. 2
- [Nor06] NORTH C.: Toward measuring visualization insight. *Computer Graphics and Applications, IEEE* 26, 3 (2006). 1, 2, 8
- [PB74] PRICE R. H., BOUFFARD D. L.: Behavioral appropriateness and situational constraint as dimensions of social behavior. *Journal of Personality and Social Psychology* (1974). 3
- [PFG07] PLAISANT C., FEKETE J. D., GRINSTEIN G.: Promoting insight-based evaluation of visualizations: From contest to benchmark repository. *IEEE Trans Vis Comput Graph* 14, 1 (2007). 2
- [PPBP10] PURCHASE H. C., PLIMMER B., BAKER R., PILCHER C.: Graph drawing aesthetics in user-sketches graph layouts. In *Australasian Conference on User Interface* (2010). 2
- [PvW08] PRETORIUS A. J., VAN WIJK J. J.: Multiple views on system traces. In *2008 IEEE Pacific Visualization Symposium* (2008). 2
- [RPHHP07] REID P., PURCHASE H. C., HALLETT-HOOK F., PLIMMER B.: Applying layout algorithms to hand-drawn graphs. In *Proceedings of the 19th Australasian conference on Computer-Human Interaction: Entertaining User Interfaces* (2007). 2
- [SC98] STRAUSS A. L., CORBIN J.: *Basics of qualitative research: Grounded theory procedures and techniques*. Sage publications, 1998. 3
- [Sch83] SCHÖN D.: *The reflective practitioner: How professionals think in action*. Basic Books, Inc., New York, NY, 1983. 2
- [Shn96] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *Visual Languages, 1996 Proceedings, IEEE Symposium* (1996). 2
- [SND05] SARAIYA P., NORTH C., DUCA K.: An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Trans Vis Comput Graph* 11, 4 (2005). 1, 2, 8
- [SNHS13] SCHULZ H. J., NOCKE T., HEITZLER M., SCHUMANN H.: A design space of visualization tasks. *IEEE Trans Vis Comput Graph* 19, 12 (2013). 2
- [SS09] STEWART R., SCHRAEFEL M. C.: *Graph sketcher: Extending illustration to quantitative graphs*. ACM, 2009. 2
- [Tab] TABLEAU SOFTWARE: Tableau. tableausoftware.com. 2
- [TAH*06] TVERSKY B., AGRAWALA M., HEISER J., LEE P. U., HANRAHAN P., PHAN D., STOLTE C., DANIELE M. P.: Cognitive design principles: From cognitive models to computer models. *Model-based reasoning in science and engineering* (2006). 2
- [Tib] TIBCO: TIBCO Spotfire - Business Intelligence Analytics Software & Data Visualization. URL: spotfire.tibco.com. 2
- [TM05] TORY M., MOLLER T.: Rethinking visualization: A high-level taxonomy. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on* (2005). 2
- [TMB02] TVERSKY B., MORRISON J. B., BETRANCOURT M.: Animation: Can it facilitate? *International Journal of Human-Computer Studies* (2002). 8
- [Tri] TRIFACTA: Vega: A visualization grammar. URL: trifecta.github.io/vega/. 2
- [Tve08] TVERSKY B.: Making thought visible. In *Proceedings of the International Workshop on Studying Design Creativity*. Springer, (2008). 2
- [Tve11] TVERSKY B.: Visualizing thought. *Topics in Cognitive Science* 3, 3 (2011). 1
- [vHR08] VAN HAM F., ROGOWITZ B.: Perceptual organization in user-generated graph layouts. *IEEE Trans Vis Comput Graph* (2008). 1, 2
- [Vic] VICTOR B.: Drawing dynamic visualizations. URL: worrydream.com/#!/DrawingDynamicVisualizationsTalk. 2, 9
- [VWvH*07] VIEGAS F. B., WATTENBERG M., VAN HAM F., KRISS J., MCKEON M.: ManyEyes: A site for visualization at internet scale. *IEEE Trans Vis Comput Graph* (2007). 2
- [War04] WARE C.: *Information Visualization: Perception for Design*. Morgan Kaufmann, 2004. 1, 2
- [WCR*11] WALNY J., CARPENDALE S., RICKE N. H., VENOLIA G., FAWCETT P.: Visual thinking in action: Visualizations as used on whiteboards. *IEEE Trans Vis Comput Graph* 17, 12 (2011). 1, 2
- [WII*12] WOOD J., ISENBERG P., ISENBERG T., DYKES J., BOUKHELIFA N., SLINGSBY A.: Sketchy Rendering for Information Visualization. *IEEE Trans Vis Comput Graph* 18, 12 (2012). 2
- [WLJ*12] WALNY J., LEE B., JOHNS P., HENRY RICKE N., CARPENDALE S.: Understanding pen and touch interaction for data exploration on interactive whiteboards. *IEEE Trans Vis Comput Graph* 18, 12 (2012). 2
- [YKSJ08] YI J. S., KANG Y.-A., STASKO J. T., JACKO J. A.: Understanding and characterizing insights: How do people gain insights using information visualization? In *BELIV '08* (2008), ACM. 2, 8
- [ZK08] ZIEMKIEWICZ C., KOSARA R.: The shaping of information by visual metaphors. *IEEE Trans Vis Comput Graph* 14, 6 (2008). 2