

Rich Screen Reader Experiences for Accessible Data Visualization

Jonathan Zong^{*1} , Crystal Lee^{*1} , Alan Lundgard^{*1} , JiWoong Jang² , Daniel Hajas³ , Arvind Satyanarayan¹ 

¹ Massachusetts Institute of Technology ² Carnegie Mellon University ³ University College London

Abstract

Current web accessibility guidelines ask visualization designers to support screen readers via basic non-visual alternatives like textual descriptions and access to raw data tables. But charts do more than summarize data or reproduce tables; they afford interactive data exploration at varying levels of granularity—from fine-grained datum-by-datum reading to skimming and surfacing high-level trends. In response to the lack of comparable non-visual affordances, we present a set of rich screen reader experiences for accessible data visualization and exploration. Through an iterative co-design process, we identify three key design dimensions for expressive screen reader accessibility: structure, or how chart entities should be organized for a screen reader to traverse; navigation, or the structural, spatial, and targeted operations a user might perform to step through the structure; and, description, or the semantic content, composition, and verbosity of the screen reader's narration. We operationalize these dimensions to prototype screen-reader-accessible visualizations that cover a diverse range of chart types and combinations of our design dimensions. We evaluate a subset of these prototypes in a mixed-methods study with 13 blind and visually impaired readers. Our findings demonstrate that these designs help users conceptualize data spatially, selectively attend to data of interest at different levels of granularity, and experience control and agency over their data analysis process. An accessible HTML version of this paper is available at: <http://vis.csail.mit.edu/pubs/rich-screen-reader-vis-experiences>.

CCS Concepts

• Human-centered computing → Visualization design and evaluation methods; Accessibility design and evaluation methods;

1. Introduction

Despite decades of visualization research and recent legal requirements to make web-based content accessible [W3C18b; Hig19], web-based visualizations remain largely inaccessible to people with visual disabilities. Charts on mainstream publications are often completely invisible to screen readers (an assistive technology that transforms text and visual media into speech) or are rendered as incomprehensible strings of “graphic graphic graphic” [SCWR21; Sar20]. Current accessibility guidelines ask visualization designers to provide textual descriptions of their graphics via alt text (short for alternative text) and link to underlying data tables [GOF08; W3C19]. However, these recommendations do not provide modes of information-seeking comparable to what sighted readers enjoy with interactive visualizations. For instance, well-written alt text can provide a high-level takeaway of what the visualization shows, but it does not allow readers to drill down into the data to explore specific sections. While tables provide readers with the ability to hone in on specific data points, reading data line-by-line quickly becomes tedious and makes it difficult to identify overall trends.

Developing rich non-visual screen reader experiences for data visualizations poses several unique challenges. Although visuomotor interactions (like hovering, pointing, clicking, and dragging) have been core to visualization research [DP20], screen readers redefine what interaction is for visualization. Rather than primarily manipulating aspects of the visualization or its backing data

pipeline [YKSJ07; HS12; DP20], screen readers make reading a visualization an interactive operation as well—users must intentionally perform actions with their input devices in order to cognize visualized elements. Moreover, as screen readers narrate elements one-at-a-time, they explicitly linearize reading a visualization. As a result, in contrast to sighted readers who can choose to selectively attend to specific elements and have access to the entire visualization during the reading process, screen reader users are limited to the linear steps made available by the visualization author and must remember (or note down) prior output conveyed by the screen reader. Despite these modality differences, studies have found that screen reader users share the same information-seeking goals as sighted readers: an initial holistic overview followed by comparing data points [SCWR21], akin to the information-seeking mantra of “overview first, zoom and filter, and details on demand” [Shn03].

In this paper, we begin to bridge this divide by conducting an iterative co-design process (co-author Hajas is a blind researcher with relevant experience) prototyping rich and usable screen reader experiences for web-based visualizations. We identify three design dimensions for enabling an expressive space of experiences: *structure*, or how the different elements of a chart should be organized for a screen reader to traverse; *navigation*, which describes the operations a user may perform to move through this structure; and, *description*, which specifies the semantic content, composition, and verbosity of text conveyed at each step. We demonstrate how to operationalize these design dimensions through diverse accessible reading experiences across a variety of chart types.

* Equal contribution

To evaluate our contribution, we conduct an exploratory mixed-methods study with a subset of our prototypes and 13 blind or visually impaired screen reader users. We identify specific features that make visualizations more useful for screen reader users (e.g., hierarchical and segmented approaches to presenting data, cursors and roadmaps for spatial navigation) and identify behavior patterns that screen reader users follow as they read a visualization (e.g., constant hypothesis testing and validating their mental models).

2. Background and Related Work

Screen Reader Assistive Technology. A screen reader is an assistive technology that conveys digital text or images as synthesized speech or braille output. Screen readers are available as standalone third-party software or can be built-in features of desktop and mobile operating systems. A screen reader allows a user to navigate content linearly with input methods native to a given platform (e.g., touch on smartphones, mouse/keyboard input on desktop). Content authors must generate and attach alt text to their visual content like images or charts in order for them to be accessible to screen reader users. Functionality and user experience differs across platforms and screen readers. In this paper, however, we focus on interacting with web-based visualizations with the most widely used desktop screen readers (JAWS/NVDA for Windows, VoiceOver for Mac).

Web Accessibility Standards. In 2014, the World Wide Web Consortium (W3C) adopted the Web Accessibility Initiative's Accessible Rich Internet Applications protocol (WAI-ARIA) which introduced a range of semantically-meaningful HTML attributes to allow screen readers to better parse HTML elements [MDN21]. In particular, these attributes allow a screen reader to convey the state of dynamic widgets (e.g., autocomplete is available for text entry), alert users to live content updates, and identify common sections of a web page for rapid navigation (e.g., banners or the main content). In 2018, the W3C published the WAI-ARIA Graphics Module [W3C18a] with additional attributes to support marking up structured graphics such as charts, maps, and diagrams. These attributes allow designers to annotate individual and groups of graphical elements as well as surface data values and labels for a screen reader to read aloud.

Accessible Visualization Design. In a recent survey, Kim et al. [KJRK21] describe the rich body of work that has explored multi-sensory approaches to visualization for multiple disabilities [YRB01; HMM*11; KWT99; BK99; WPA*21; LLS19]. Here, we focus on screen reader output native to web-based interfaces for blind users (namely via speech). Sharif et al. [SCWR21] find that many web-based charts are intentionally designed to cause screen readers to skip over them. For charts that a screen reader does detect, blind or visually impaired users nevertheless experience significant difficulties: these users spend 211% more time interacting with the charts and are 61% less accurate in extracting information compared to non-screen-reader users [SCWR21]. Despite the availability of ARIA, alt text and data tables remain the most commonly used and recommended methods for making web-based charts accessible to screen readers [GOF08; W3C19; CJP*19]. However, each of these three approaches comes with its own limitations. Static alt text requires blind readers to accept the author's interpretation of the data; by not affording exploratory and interactive

modes, alt text robs readers of the necessary time and space to interpret the numbers for themselves [LS21]. Recent research also suggests that blind people have nuanced preferences for the kinds of visual semantic content conveyed via text [PGFM21; LS21], and desire more interactive and exploratory representations of pictorial images [MJBC18]. Data tables, on the other hand, undo the benefits of abstraction that visualizations enable—they force readers to step sequentially through data values making it difficult to identify larger-scale patterns or trends, and do not leverage the structure inherent to web-based grammars of graphics [BOH11; SMWH17]. Finally, ARIA labels are not a panacea; even when they are used judiciously—a non-trivial task which often results in careless designs that cause screen readers to simply read out long sequences of numbers without any other identifiable information [Sar20]—they present a fairly low expressive ceiling. The current ARIA specification does not afford rich and nuanced information-seeking opportunities equivalent to those available to sighted readers.

There has been some promising progress for improving support for accessibility within visualization toolkits, and vice-versa for improving native support for charts in screen reader technologies. For instance, Vega-Lite [SMWH17] and Highcharts [Hig21] are beginning to provide ARIA support out-of-the-box. Apple's VoiceOver Data Comprehension feature [DE19] affords more granular screen reader navigation within the chart, beyond textual summaries and data tables, via four categories of selectable interactions for charts appearing in Apple's Stocks or Health apps. These interactions include *Describe Chart*, which describes properties of the chart's construction, such as its encodings, axis labels, and ranges; *Summarize Numerical Data*, which reports min and max data values, and summary statistics like mean and standard deviation; *Describe Data Series*, which reports the rate-of-change/growth of a curve, trends, and outliers; and *Play Audiograph*, which plays a tonal representation of the graph's ascending/descending trend over time [DE19]. While Apple's features are presently limited to single-line charts, SAS' Graphics Accelerator [Acc18] supports a similar featureset (including sonification, textual descriptions, and data tables) but for a broader range of statistical charts including bar charts, box plots, contour plots, and scatter plot matrices. Our work follows in the spirit of these tools but focuses on web-based visualizations rather than standalone- or platform-integrated software. We go beyond what ARIA supports today to enable high-level and fine-grained screen reader interactions, and hope that our work will help inform ongoing discussions on improving web accessibility standards (e.g., via an Accessibility Object Model [BCMS22]).

3. Design Dimensions for Rich Screen Reader Experiences

Currently, the most common ways of making a visualization accessible to screen readers include adding a single high-level textual description (via alt text), providing access to low-level data via a table, or tagging visualization elements with ARIA labels to allow screen readers to step through them linearly (e.g., as with Highcharts [Hig21]). While promising, these approaches do not afford rich information-seeking behaviors akin to what sighted readers enjoy with interactive visualizations. To support systematic thinking about accessible visualization design, we introduce three design dimensions that support rich, accessible reading experiences: *struc-*

ture, or how elements of the visualization should be organized for a screen reader to traverse; *navigation*, or the mechanisms by which a screen reader user can move from one element to another; and *description*, or what semantic content the screen reader conveys.

Methods. We began by studying the development of multi-sensory graphical systems, covering work in critical cartography [WG69; Koc12], blind education [AS01; GL15], tactile graphics [FTT18; HMM*11; dGMB21; AC97; BHR*21], and multi-sensory visualization [MTS19; CPR*21; BTOJ10; BMD*16]. Drawing on conventions and literature on crip, reflective, and participatory design [Ham13; SBDK05; Cos20], all authors began an iterative co-design process with Hajas, who is a blind researcher with relevant expertise. Hajas is a screen reader user with a PhD in HCI and accessible science communication, but he is not an expert in visualization research. Co-design — particularly as encapsulated in the disability activism slogan, “*Nothing about us, without us*” [Cos20] — is important because it can eliminate prototypes that replicate existing tools, solve imaginary problems (i.e., by creating disability dongles [Jac19]) or unintentionally produce harmful technology [She20]. To balance engaging disabled users while acknowledging academia’s traditionally extractive relationship with marginalized populations [CJ95], we intentionally acknowledge Hajas as both co-designer and co-author. We believe that the distinction between co-designer — a phrase that often discounts lived experience as insufficiently academic — and researcher is minimal; technical, qualitative, and experiential expertise are all important components of this research. Hajas’ profile is a perfect example of the intersection between lived experience of existing challenges and solutions, academic experience of research procedures, and an interest in the science of visualization. While he does not represent all screen reader users, his academic expertise and lived experience uniquely qualify him to be both researcher and co-designer. Nevertheless, to incorporate a diverse range of perspectives, we recruited additional participants as part of an evaluative study (§ 5).

Our work unfolded over 6 months and yielded 15 prototypes. All authors met weekly for hour-long video conferences. In each session, we would discuss the structure and affordances of the prototypes, often by observing and recording Hajas’ screen as he worked through them. We would also use these meetings to reflect on how the prototypes have evolved, compare their similarities and differences, and whiteboard potential design dimensions to capture these insights. Following these meetings, Hajas wrote memos detailing the motivations for each prototype, tagging its most salient features, summarizing the types of interactions that were available, enumerating questions that the prototype raises, and finally providing high-level feedback about its usefulness and usability. In the following section, we liberally quote these memos to provide evidence and additional context for our design dimensions.

3.1. Structure

We define *structure* to mean an underlying representation of a visualization that organizes its data and visual elements into a format that can be traversed by a screen reader. Through our co-design process, we identified two components important to analyzing accessible structures: their *form*, or the shape they organize information into; and *entities*, or which parts of the visualization specifica-

tion are used to translate a chart into a non-visual structure. Design decisions about form and entities are guided by considerations of *information granularity*, or how many levels comprise the range between a high-level overview and individual data values.

Form. Accessible structures organize information about the visualization into different forms, including lists, tables, and trees. Consider existing best practices and common approaches. A rasterized chart with alt text is represented to a screen reader as a single node. SVG-based visualizations can additionally be tagged with ARIA labels to describe the axes, legends, and individual data points. Despite SVG’s nesting, screen readers linearize these ARIA labels into a list structure so that the user can step through them sequentially. Data tables, on the other hand, provide a grid structure for screen readers to traverse. At each cell of the grid, the screen reader reads out a different textual description, allowing the user to explore a space by traversing the grid spatially (up, down, left, and right) instead of merely linearly. Accessible visualization research has begun to explore the use of tree structures for storing chart metadata [WOH*15], but they remain relatively rare in practice. Our prototypes primarily use trees as their branching and hierarchical organization allows users to browse different components of a visualization and traverse them at different levels of detail.

Entities. Where form refers to how nodes in a structure are arranged, entities instead refers to what aspects of the visualization the nodes represent. These aspects can include:

- *Data*, where nodes in the structure represent individual data values or different slices of the data cube (e.g., by field, bins, categories, or interval ranges). For example, in a data table, every node (i.e. cell) represents a data value designated by the row and column coordinates. Depending on the form, data entities can be presented at different levels of detail. For example, one prototype we explored represents a line chart as a binary tree structure (Fig. 2e): the root node represents the entire x-axis domain, and each left and right child node recursively splits the domain in half. Users can traverse the tree downward to binary search for specific values or understand the data distribution.
- *Encodings*, where nodes in the structure correspond to visual channels (e.g., position, color, size) that data fields map to. For instance, consider Figure 1a which depicts the encoding structure of a Vega-Lite scatterplot. The visualization is specified as mappings from data fields to three visual encoding channels: *x*, *y*, and *color*. Thus, the encoding structure, which here takes the form of a tree, comprises a root node that represents the entire visualization and then branches for each encoding channel as well as the data rectangle (x-y grid). Descending into these branches yields nodes that select different categories or interval regions, determined by the visual affordances of the channel. For instance, descending into axis branches yields nodes for each interval between major ticks; x-y grid nodes represent cells in the data rectangle as determined by intersections of the axes gridlines; and legend nodes reflect the categories or intervals of the encoding channel (i.e., for nominal or quantitative data respectively). Finally, the leaves of these branches represent individual data values that fall within the selected interval or category.
- *Annotations*, where nodes in the structure represent the rhetorical devices a visualization author may use to shape a visual

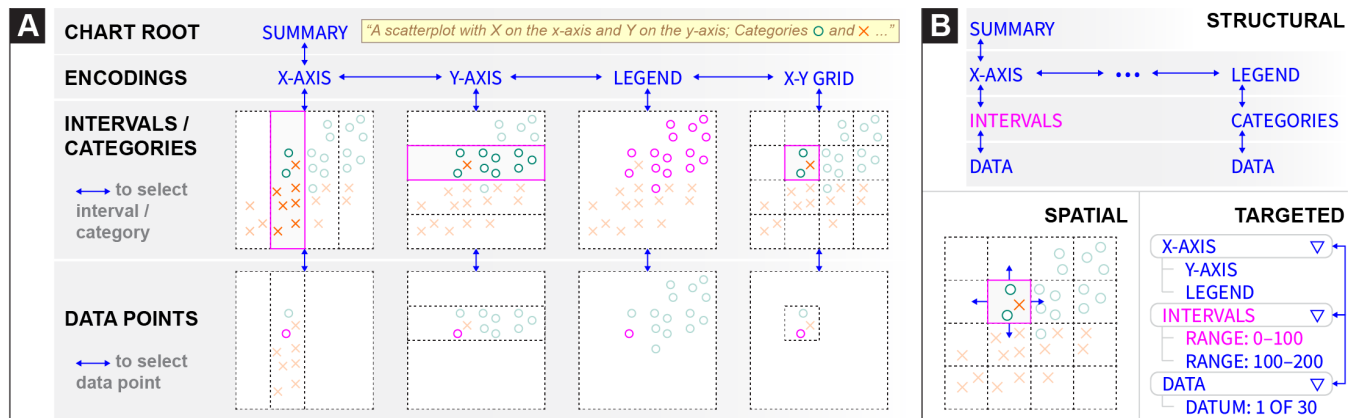


Figure 1: (a) An accessible visualization structure in the form of a tree and comprised of encoding entities. Solid magenta outlines indicate the location of the screen reader cursor. Solid blue arrows between labels indicate available next steps via keyboard navigability (up, down, left, right). (b) Three ways of navigating accessible visualization structures: structural, spatial, and targeted.

narrative or guide reader interpretation of data (e.g., by drawing attention to specific data points or visual regions). Surfacing annotations in the visualization structure allows screen reader users to also benefit from and be guided by the author's narrative intent. For example, Figure 2d illustrates an annotation tree structure derived from an example line chart with two annotations highlighting intervals in the temporal x-axis. The root of the tree has two children representing the two annotated regions. These two annotation nodes have a child node for each data point that is highlighted within the region of interest.

Considerations: Information Granularity. When might users prefer nested structures (i.e. trees) over flat structures (i.e., lists and tables)? Like sighted users, screen reader users seek information by looking for an overview before identifying subsets to view in more detail [SCWR21]. Trees allow users to read summary information at the top of the structure, and traverse deeper into branches to acquire details-on-demand. Kim et al. use the term *information granularity* to refer to the different levels of detail at which an accessible visualization might reveal information [KJRK21]. They organize granularity into three levels: existence, overview, and detail. *Existence* includes information that a chart is present, but no information about underlying data. *Overview* includes summary information about data — e.g. axes, legends, and summary statistics like min, max, or mean — but not individual data points. *Detail* includes information about precise data values.

We use the root node to signal the existence of the tree, and deeper nodes in the tree reflect finer levels of granularity. Branch nodes give an overview summary about the data underneath, providing information scent [PC99], while leaf nodes map to individual data points. In his feedback about the prototype shown in Figure 1, Hajas wrote “considering how difficult reading a scatterplot with a screen reader is due to its sequential reading nature, the tree structure makes the huge number of data points fairly readable”.

Entities are not mutually exclusive, and a structure might opt to surface different entities in parallel branches. We prototyped a version of Figure 2d which placed an encoding tree and annotation tree as sibling branches under the root node. Users could descend down a given branch, and switch to the equivalent location in the

other branch at will. These design decisions are motivated by findings in prior work: by placing encodings and annotations as co-equal branches, we produce a structure that preserves the agency of screen reader users either to start with the narrative arc of annotations, or follow it after having the chance to interpret the data for themselves [LS21]. As Hajas confirms “Depending on my task, either the encoding or annotation tree could be more important. If my task involved checking population growth in the last 100 years, I would start with the encodings. If I were to look for sudden changes in population numbers, such war-time mortality effects, I would start exploring the annotations, then tunnel back to the other tree.”

3.2. Navigation

Screen reader users need ways to traverse accessible structures to explore data or locate specific points. When browsing a webpage, screen readers provide a cursor that represents the current location in the page. Users use keyboard commands to step the cursor backward and forward in a sequential list of selectable items on the page, or jump to important locations such as headers and links. Through our prototyping process, we developed three ways of navigating through an accessible structure: *structural navigation*, *spatial navigation*, and *targeted navigation* (Fig. 1b). A key concern across these navigation schemes is reducing a user's cognitive load by affording a sense of the *boundaries* of the structure.

Structural Navigation. Structural navigation refers to ways users move within the accessible structure. We identify two types of structural navigation. *Local navigation* refers to step-by-step movements between adjacent nodes in the structure. This includes moving up and down levels of a hierarchy, or moving side to side between sibling elements. *Lateral navigation* refers to movement between equivalent nodes in adjacent sub-structures. For example, Fig. 2a depicts a multi-view visualization with six facets. When the cursor is on a Y-axis interval for the first facet, directly moving to the same Y-axis interval on the second facet is a lateral move.

Spatial Navigation. Sometimes users want to traverse the visualization according to directions in the screen coordinate system. We refer to this as spatial navigation. For example, when travers-

ing part of an encoding structure that represents the visualization's X-Y grid, a downward structural navigation would go down a level into the currently selected cell of the grid, showing the data points inside the cell. A downward spatial navigation, in contrast, would move to the grid cell below the current one — i.e. towards the bottom of the Y-axis. Spatial navigation is also useful when navigating lists of data points, which may not be sorted by X or Y value in the encoding structure. Where a leftward structural navigation would move to the previous data point in the structure, a leftward spatial navigation would move to the point with the next lowest X value.

Targeted Navigation. Navigating structurally and spatially requires a user to maintain a mental map of where their cursor is relative to where they want to go. If the user has a specific target location in mind, maintaining this mental map in order to find the correct path in the structure to their target can create unnecessary cognitive load. We use targeted navigation to refer to methods that only require the user to specify a target location, without needing to specify a path to get there. For example, the user might open a list of locations in the structure and select one to jump directly there. Screen readers including JAWS and VoiceOver implement an analogous form of navigation within webpages. Instead of manually stepping through the page to find a specific piece of content, users can open a menu with a list of locations in the page. These locations are defined in HTML using ARIA landmark roles, which can designate parts of the DOM as distinct sections when read by a screen reader. When a screen reader user opens the list of landmarks and selects a landmark, their cursor moves directly to that element.

Considerations: Boundaries & Cognitive Load. Screen reader users only read part of the visualization at a time, akin to a sighted user reading a map through a small tube [HMM*11]. How do they keep track of where they are? In our co-design process, we found it easiest for a user to remember their location relative to a known starting point, which is corroborated by literature on developing spatial awareness for blind people [WG69; LKM*19; CPR*21]. Hajas noted the prevalence of the Home and End shortcuts across applications for returning to a known position in a bounded space (e.g. the start/end of a line in a text editor). We also found that grouping data by category or interval was helpful for maintaining position. Hajas noted that exploring data within a bounded region was like entering a room in a house. In his analogy, a house with many smaller rooms with doors is better than a house with one big room and no doors. Bounded spaces alleviate cognitive load by allowing a user to maintain their position relative to entry points.

Comparing navigation techniques, Hajas noted that spatial felt “shallow but broad” while targeted felt “deep but narrow.” While he expressed a personal preference for deep-narrow structures, he nevertheless “would not give up [spatial navigation] because it makes me believe I’m actually interacting with a visualization.” This insight demonstrates the value of offering multiple complementary navigation techniques. Moreover, while targeted navigation facilitates quick searching and doesn’t require the user to maintain a mental map to find specific data points, structural and spatial exploration enable more open-ended data exploration. It also provides a mechanism for establishing common ground with sighted readers (e.g., allowing both blind and sighted readers to understand a line segment as being “above” or “higher” than another).

3.3. Description

When a user navigates to a node in a structure, the screen reader narrates a description associated with that node. For example, when navigating to the chart’s legend, the screen reader output might articulate visual properties of the chart’s encoding: “Category O has color encoding green; X has color encoding orange” (Figure 1). Or, if that visual semantic content isn’t relevant to understanding the data, it might ignore the color: “each datum belongs to either Category O or X.” The content, composition, and verbosity of the description can affect a user’s comprehension of the data. Designers must consider context & customization when describing charts.

Content. Semantic content is the meaningful information conveyed not only through natural language utterances, but also through the visualization (a graphical language [Ber83]). Because graphics convey myriad different kinds of content, the challenge of natural language description is to convey information that is not only commensurate with what the chart expresses via graphical language, but also useful to its readers. Accessible chart description guidelines from WGBH [GOF08], W3C [W3C19], and others [JMK*21] offer prescriptions for conveying specific content for blind readers (such as the chart’s title, axis encodings, and noteworthy trends). Lundgard and Satyanarayan expand the scope of these guidelines with a more general conceptual model of four levels of semantic content: *chart construction properties* (e.g., axes, encodings, marks, title); *statistical concepts and relations* (e.g., outliers, correlations, descriptive statistics); *perceptual and cognitive phenomena* (e.g., complex trends, patterns); and *domain-specific insights* (e.g., socio-political context relevant to the data) [LS21].

Decoupling a chart’s semantic content from its visual representation helps us better understand what data representations afford for different readers. For instance, Lundgard and Satyanarayan find that what blind readers report as most useful in a chart description is not a straightforward translation of the visual data representation. Specifically, simply listing the chart’s encodings is much less useful to blind readers than conveying summary statistics and overall trends in the data [LS21]. As Hajas noted, “I want to see the global trend, which is why sighted people rely on visualization.” For instance, for a stock market chart the reader “might see the overview from first to last data points, and then zoom into an outlier in the middle.” These findings suggest opportunities interleave different kinds of content at different levels of a hierarchical structure to yield richer, more useful screen reader navigation. For example, injecting summary statistics (say, the existence of outliers within a particular subcategory of the data) higher up in the chart’s tree structure (e.g., at the legend encoding node) might afford “scent” for “information foraging” [PC99], or further exploration down a particular branch (data subcategory) of the tree. Or, if navigating in a targeted fashion, the user might be afforded the option to directly navigate to outliers without traversing the tree.

Composition. The usefulness of a description depends not only on the content conveyed by its constituent sentences, but also on its composition: how those sentences are ordered in relation to each other. For example, during our co-design process, Hajas found that when navigating a chart’s tree structure, the screen reader output could quickly become redundant, affecting how quickly and efficiently he could pick out the meaningful information at each node.

For instance, the utterance “*Category: O, Point 3 of 15, x = 5, y = 12*” and the utterance “*x = 5, y = 12, Category: O, Point 3 of 15*” afford significantly different experiences for a user who wishes to quickly scan through individual data points. In the first utterance, the reader immediately receives content that helps to situate them in a broader data context, namely data labeled as “Category: A” at the legend node. In the second utterance, the reader immediately receives datum-specific content that helps to rapidly explore the fine-grained details within that data context. Whether a reader prefers one compositional ordering to another will depend on the task they are attempting to accomplish. As Hajas noted “*I like the label at the beginning of the information, saying at which level of the tree I am at. It is important for knowing where I am. It is also great that this information is only spoken out when I change level, but not when I navigate laterally.*” These compositional choices are highly consequential for readers’ experience, when they must repeatedly read nearly-identical utterances while navigating a structure.

Verbosity. Whereas composition refers to the ordering of content, *verbosity* refers to how much content the screen reader conveys. More content is not always better. As Hajas noted of Apple’s Data Comprehension feature [DE19]: “*It can sometimes be too much information all at once, if it starts reading out all of the data. This is very difficult if you’re interested in some data points that are in the middle. It is very play-or-stop.*” Depending on the screen reader software, a user may be afforded control over how much content is conveyed. For instance, JAWS offers high, medium, and low verbosity levels [Sci21]. At higher verbosity the screen reader announces more structural, wayfinding content (e.g. the start and end of regions). For data tables, verbosity configurations can affect whether the table size is read as part of the description, and whether row and column labels are repeated for every cell. Descriptions of nodes in an encoding structure might analogously include information about the path from the root—for example, by reminding the user that they are reading Y-axis intervals. These repetitions can help users remember their location within a structure, but additional verbosity is less efficient for comprehending the data quickly.

Considerations: Context & Customization. Apart from its constituent parts (content, composition, verbosity), a description’s usefulness also depends on the *context* in which it is read: namely, the reader’s task or intent, and familiarity with the data interface. The same description might be useful in some situations, but relatively useless in others. A reader’s information needs are fundamentally context-sensitive. For example, as Hajas noted, when reading a news article, it may be satisfactory to accept a journalist’s description of the data on good faith. But, when reviewing scientific research, “*I don’t necessarily want to just believe what is said in the text, I want to check and double-check the authors’ claims. Go down to the smallest numbers in the analysis. I want to be able to look at the confusion matrix and see if they made a mistake or not.*” This targeted verification requires a description to afford users with precise look-up capabilities, in contrast to descriptions that may be generated when browsing or exploring the data.

This context-sensitivity reveals an important aspect of usability: a user’s familiarity (or lack thereof) with the data interface. Wayfinding content (e.g., “*Legend. Category: O.*”) can help a user remember their location in a structure, and may be useful while they

assemble a mental map of the visualization. But, as they become accustomed to the interface and visualization, such descriptions may prove cumbersome. Because user needs depend on their task, preferences, and familiarity, interfaces might afford personalization and customization to facilitate context-sensitive description.

4. Example Gallery

Our co-design process yielded prototypes that demonstrate a breadth of ways to operationalize our design dimensions. Figure 2 excerpts some of our highest-fidelity prototypes, implemented on top of Vega-Lite [SMWH17]. As deeply nested structures and dynamic content are not well-supported by ARIA, we implemented our designs as in-memory data structures. Event listeners update the user’s position in the structure on keypress, and write text descriptions to an ARIA-Live region (an ARIA role typically used for temporary notifications). To establish common ground with sighted users, we also render the visualization graphically. The user’s position in the tree drives a Vega-Lite selection that highlights points when the screen reader user is attending to them.

For every prototype, the up, down, left, and right arrow keys enable structural navigation (moving up or down a level, or stepping through siblings respectively). For example, within the facet level of Fig. 2(a), the user can press left or right keys to move between the six subplots of the multiview chart. On charts that contain a node representing the x-y grid, users can also use the WASD keys to spatially navigate the grid and data points within that branch (mimicking an interaction found in video games).

These prototypes highlight different compositions of structures and navigation schemes. Fig. 2(a) includes shift+left and shift+right for lateral navigation across facets: pressing these keys at any node within a facet branch will navigate to the same location under an adjacent branch (subplot). With the choropleth (Fig. 2(b)), we group data in the encoding structure by U.S. state; users can then drill down into counties across either this branch or the legend one. Fig. 2(c) offers two different paths for drilling down: month first, or weather first. Fig. 2(d) structures the tree by annotations rather than encoding: users can descend into the time intervals designated by the orange and blue rectangles, and view points within those intervals. Finally, Fig. 2(e) organizes its tree in terms of data, offering a binary search structure through the years.

5. Evaluation

To evaluate our contribution, we conducted 90-minute Zoom studies with 13 blind and visually impaired participants. Participants were asked to explore three prototype accessible screen reader experiences, shown one after another each with a different dataset. The goal of our evaluation was not to determine which particular combination of design elements was “best,” but rather to be exploratory—to compare the relative strengths and advantages of instantiations of our design dimensions, and understand how they afford different modes of accessible data exploration.

5.1. Evaluation Setup & Design

Following Frøkjær and Hornbæk’s Cooperative Usability Testing (CUT) method [FH05], Zong and Lee conducted each session by

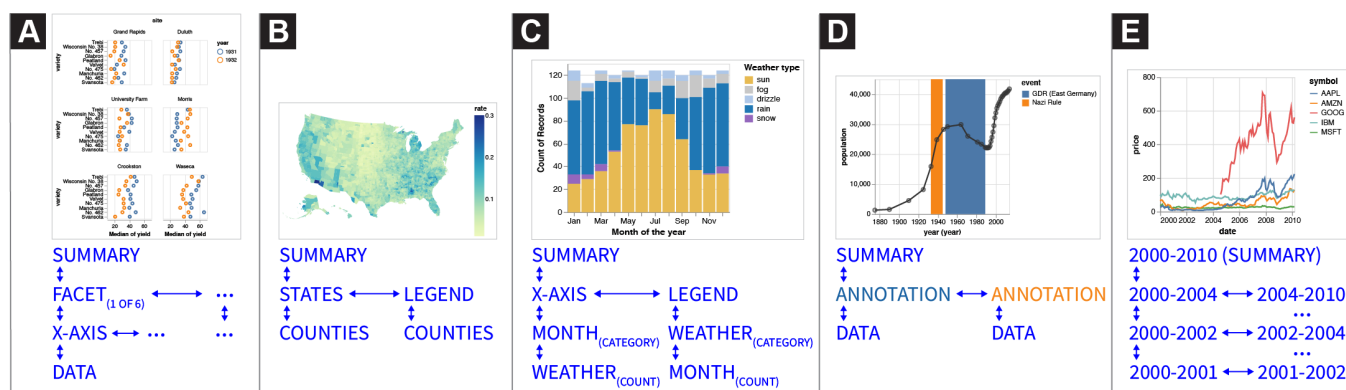


Figure 2: Example structural and navigational schemes generated as part of our co-design process, and applied to diverse chart types.

alternating between the role of guide (i.e., talking to the user and explaining the prototype) and logger (i.e., keeping track of potential usability problems, interpreting the data to prepare for becoming the guide). We began each session with a semi-structured interview to understand participants' current experiences with data and the methods they use to make inaccessible forms of data representation usable (script included in supplementary material). The rest of the session focused on each of the three prototypes in turn, with each condition split into two phases: interaction and interpretation. In the interaction phase, Zong or Lee guided participants through the prototypes and asked participants to use them and comment on their process, in the style of Hutchinson et al.'s technology probes [HMW*03]. Then, the authors switched roles and began a cooperative interpretation phase, where the authors and participants engaged in a constructive dialogue to jointly interpret the usability problems and brainstorm possible alternatives to the current prototype. In this method, participants influence data interpretation, allowing for more rapid analysis than traditional think-aloud studies as some analysis is built into each evaluation session with instant feedback or explanation from participants [FH05].

Prototypes. The in-depth nature of our cooperative interpretation sessions required us to balance the total number of prototypes evaluated (so that participants would have time to thoroughly learn and interact with each one) with a time duration appropriate for a Zoom session (limited to 90 minutes to avoid exhausting participants). Accordingly, we selected the following three prototypes, each representing a different aspect of our design dimensions:

- **TABLE:** An accessible HTML data table with all rows and three columns from the classic Cars dataset, in order to compare our prototypes with existing accessibility best practice.
- **MULTI-VIEW:** Becker's barley yield trellis display [BCS96] as shown in Fig. 2a. This prototype features local and lateral structural navigation via the arrow keys and with the shift modifier respectively, as well as spatial navigation via WASD.
- **TARGET:** A single-view scatterplot, illustrated in Fig. 1, depicting the Palmer Penguins dataset [HHG20]. In addition to structural and spatial navigation, targeted navigation is available via three dropdown menus corresponding to the structural levels.

TABLE is our control condition, as it follows existing best practice for making data accessible to screen readers. MULTI-VIEW enables us to study how users move between levels of detail, and

whether they could navigate and compare small multiple charts. Finally, TARGET allows us to compare how and when our participants use the three different styles of navigation (structural, spatial, and targeted). We presented the prototypes in this sequence to all participants to introduce new features incrementally.

Participants. We recruited 13 blind and visually impaired participants through our collaborators in the blind community and through a public call on Twitter. Each participant received \$50 for a 90-minute Zoom session. We provide aggregate participant data following ethnographic practice to protect privacy and not reduce participants to their demographics [SKK15]. Half of our participants were totally blind ($n=7$), while others were almost totally blind with some light perception ($n=4$) or low vision ($n=2$). Half of them have been blind since birth ($n=7$). Participants were split evenly between Windows/Chrome ($n=7$) and Mac/Safari ($n=6$). Windows users were also split evenly between the two major screen readers (JAWS, $n=3$; NVDA, $n=4$), while all Mac participants used Apple VoiceOver. These figures are consistent with recent surveys conducted by WebAIM which indicate that JAWS, NVDA, and VoiceOver are the three most commonly used screen readers [Web21]. Demographically, 70% of our participants use he/him pronouns ($n=9$) and the rest use she/her pronouns ($n=4$). One participant was based in the UK while the rest were spread across eight US states. Participants self-reported their ethnicities (Caucasian/white, Asian, and Black/African, Hispanic/Latinx), represented a diverse range of ages (20–50+) and had a variety of educational backgrounds (high school through to undergraduate, graduate, and trade school). Nine participants self-reported as slightly or moderately familiar with statistical concepts and data visualization methods, two as expertly familiar, and one as not at all familiar. Five participants described data analysis and visualization tools as an important component in their professional workflows, and 8 interacted with data or visualizations more than 1–2 times/week.

5.2. Quantitative Results

To supplement the cooperative interpretation sessions, participants rated each prototype using a series of Likert questions. We designed a questionnaire with six prompts measuring a subset of Brehmer and Munzner's multi-level typology of abstract visualization tasks [BM13]. This framework, however, required some adaptation for non-visual modes of search. In particular, searching with

Table 1: Rating scores for each prototype (Table, Multi-view, Targeted) on a five point Likert scale where 1 = Very Difficult (Very Unenjoyable) and 5 = Very Easy (Very Enjoyable). Median scores are shown in boldface, averages in brackets, standard deviations in parentheses.

PROMPT : When using this prototype ...	TASK [BM13]	TABLE	MULTI-VIEW	TARGETED
How enjoyable was it to interact with the data?	enjoy	3 [3.31] (0.95)	4 [3.77] (1.01)	4 [3.54] (0.97)
How easy was it to generate and answer questions?	discover	4 [3.15] (1.34)	3 [3.00] (1.08)	3 [3.23] (1.17)
If you already knew what information you were trying to find, how easy would it be to look up or locate those data?	lookup-locate	3 [3.31] (1.32)	4 [3.77] (1.17)	4 [3.38] (1.19)
If you didn't already know which information you were trying to find, how easy would it be to browse or explore the data?	browse-explore	2 [3.00] (1.68)	2 [2.69] (1.11)	3 [3.00] (1.29)
PROMPT : When using this prototype ...	USE [KS99]	TABLE	MULTI-VIEW	TARGETED
How easy was it to learn to use?	ease-of-use	4 [4.15] (0.99)	3 [2.69] (0.75)	3 [3.15] (1.34)
How useful would it be to have access to this interaction style for engaging with data?	perceived usefulness	4 [4.15] (0.80)	4 [4.00] (0.82)	4 [4.15] (1.07)

a screen reader requires a sequential approach to data that is at odds with the “at-a-glance” approach sighted readers take to browsing and exploring data. As our prototypes focus on navigation through charts, we collapsed the *location* dimension of Brehmer and Munzner’s search decomposition resulting in two prompts that jointly measure *lookup-locate* and *browse-explore*. We formulated additional questions to measure Brehmer and Munzner’s *discover* and *enjoy* tasks as well as more traditional aspects of technology acceptance including *ease-of-use* and *perceived usefulness* [KS99]. Participants responded on a five point scale where 1 = Very Difficult/Unenjoyable and 5 = Very Easy/Enjoyable.

Table 1 displays the questionnaire prompts, their corresponding tasks, and statistics summarizing the participants’ ratings. A Friedman test found a significant rating difference for the ease-of-use of the prototypes $\chi^2(2, N = 13) = 15.05, p < 0.01$, with a large effect size (Kendall’s $W = 0.58$). Follow-up Nemenyi tests revealed that MULTI-VIEW was more difficult to use than TABLE with statistical significance ($p < 0.01$), but TARGET was not. Additional tests for the other prompts found neither statistically significant differences, nor large effect sizes, between the prototypes. However, median scores (which are more robust to outliers than means [MJBC18]) suggest that participants generally *enjoy* interacting with MULTI-VIEW and TARGET more, and found them easier to *lookup* or *locate* data with. Moreover, TARGET had the highest median score for affording *browse* or *explore* capabilities. Conversely TABLE was easiest to learn to use, and generally made it easy to *discover*, or ask and answer questions about the data. Notably, in response to the question “How useful would it be to have access to this interaction style for engaging with data?” participants on average ranked all prototypes as more-than-useful ($med = 4, \mu \geq 4$). These statistics provide only a partial picture of participants’ experiences with the prototypes [Bag07]. Thus, we elucidate and contextualize reasons behind their scores through qualitative analysis.

5.3. Qualitative Results

After the interviews, we qualitatively coded the notes taken by the logger with a grounded theory approach [Cha06]. We performed open coding in parallel with the interviews (i.e., coding Monday’s interviews after finishing Tuesday’s interviews). We then synthesized the codes into memos, from which we derived these themes.

Tables are familiar, tedious, but necessary. Every participant noted that tables were their primary way of accessing data and visualizations. While tables are an important accessible option, par-

ticipants overwhelmingly reported the same problems: they are ill-suited for processing large amounts of data and impose high cognitive load as users must remember previous lines of the table in order to contextualize subsequent values. As P2 reported, “if I’m trying to get a general sense of the table, I’ll just scroll through and see what values there are. But there’s 393 rows, so I’ll never scroll through all of it...I can’t really get a snapshot.” P11 said that “Finding relationships can be tricky if you’re in a table, because you’ve got to either have a really good memory or just get really lucky. [...] If I didn’t know what I was looking for, forget it.” At most, participants tabbed through 20–30 rows during our sessions, but did so only because of the questions we posed (e.g., “is there a relationship between horsepower and mileage?”) and noted that if they encountered this table outside of the study, they would tab past a few rows to check for summary statistics and then move on.

While it is not enjoyable to explore or build a mental model of data with static tables, participants still emphasized their necessity because of the format’s familiarity: “in terms of accessibility, tables are infinitely more useful because there is a standard way of navigating them in whatever your preferred screen reader is. With different representations, a blind person may not be trained to interpret it” (P2). This builds on prior literature [SCWR21] and echoes testimony from participants who had some difficulty with the new prototypes; they reported that they lacked expertise and therefore found it difficult to work with non-tabular data (P8, 10). In other words, to maximize accessibility, it is crucial to include a table view of the data *in addition* to other forms of novel interaction.

Prior exposure to data analysis and representations increases the efficacy of spatial representations. Participants who had experience conducting data analysis or reading tactile graphs/maps were able to easily develop a spatial understanding of how each prototype worked. Five participants (P2–4, 11, 13) made direct connections between the MULTI-VIEW and TARGET prototypes, and the tactile graphs they encountered in school. Three participants (P2, 11, 12) found their software engineering experience made it easier to understand and navigate the prototypes’ hierarchical structure. Previous literature on tactile mapping has also shown how developing tactile graphical literacy is crucial for building spatial knowledge, but they emphasize that it is not a sufficient for being able to conduct and understand data analysis. [HMM*11; GL15] Since our participants already had an existing spatial framework, it became easier to explain how a prototype might work using their prior experience as a benchmark, which has been corroborated by

similar studies in tactile cartography. [WG69; AS01; SA01] Importantly, our participants were able to find specific origin points that they could return to in order to navigate the different branches of the tree, which would be further aided with help menus and mini-tutorials to understand the keyboard shortcuts (P2). Being able to shift between origin points is especially important for switching between graphs or between variables. By contrast, participants who had more difficulty with the prototypes (P8, 10) pointed to their lack of experience working with non-tabular data. P10 reported that being able to mentally visualize data points within a grid was a specific challenge. *"I suspect that this might be understandable to someone who's done this before,"* he said, *"I don't do well with these charts unless they're converted back into tables."*

Structure: Hierarchical representations make it possible to effectively convey insights with minimal cognitive load. While static tables are the most common accessible option to interactive visualizations, eight of our participants (P2–5, 7, 10, 11, 13) expressed a desire to filter and sort the data so that they could begin to explore possible trends without wading line by line. Sorting and filtering a table is one way to look for trends but, to get a summary view of the data quickly, a system must provide snapshots in smaller intervals so that users can easily construct a larger picture or choose specific slices of the data to explore further (i.e., "details on demand"). [Shn03; KJRK21] With MULTI-VIEW and TARGET P4 said, *"I always want more layers and details, but some charts had too much...This was a happy medium between having the information I wanted and presenting it in a way that I can keep up with."* P5 also noted that he liked *"having the ability to scroll through at a higher level and then drill down deeper if that's of interest."* By giving users a way to quickly skip through the data across specific axes, they are able to rapidly generate a broader mental image of each graph and drill down further to collect more details. *"When I was working with the table, I [started building] a table in my head,"* P2 shared. *"I had a rough representation of it as a scatter plot. But here, I know how to drill down and up between different layers of data grids, so that I can get the overall picture... [It gives me] different ways of thinking."* Being able to control the parts of the data that were most important to them was also an issue of trust, as it also provided a way for users to reach conclusions for themselves rather than rely on the interpretation of others: *"It's hard to mix...doing your own analysis and be given a text description that you have to just trust"* (P12). In their own workflows, these participants reported downloading static tables to further examine and manipulate with Excel, which they would use to create summary statistics or intervals to move more quickly through the data.

Navigation: Reading a visualization with a screen reader entails constant hypothesis testing and pattern-making. Since screen reader users parse data iteratively, nine of our participants (P1–5, 7, 8, 11, 13) described reading a visualization as a process of slowly building up a mental model and constantly testing it to see where the patterns may no longer hold. *"I'm going row by row, not memorizing exact numbers but building a pattern in my head, and looking at the other rows to test my theory,"* reported P3. In other words, our participants engaged in a continuous state of updating and validating [Mun09] their mental images as new data challenged the existing patterns they have pieced together. MULTI-VIEW and TARGET accelerated this process, as participants were able to more

rapidly identify specific components that they wanted to test. For example, P2 intentionally moved quickly across each level of the structure hoping to find its "edges," or the minimum and maximum limits of each axis and grid. *"Visually, it might look like I'm doing a lot of jumping around,"* he said, *"[but it's] because I'm trying to build the picture in a way that makes sense for me."* Similarly, P5 started building his mental model of the visualization by drilling up and down the grid to create a spatial image of the data: *"I'm thinking more in spatial terms just because [this] is a new method of navigating to me. [...] I'm moving through the grid...I'm thinking of drilling down into that square to get more information."*

TARGET made it especially easy for participants to test their hypotheses by giving them direct access to components that might break their hypotheses. P5 reported that it allowed him to *"navigate to areas...that I'm interested in, skipping over stuff that's not of interest,"* and P4 likened it to *"[being] able to go directly to what you want in a grocery inventory rather than going through each item one by one."* The ability to use structural, spatial, and target navigation in both MULTI-VIEW and TARGET respectively facilitated the hypothesis-testing and pattern-making behaviors that our participants were accustomed to with static tables, and gave them an additional mental model for working with the data. As P1 noted, these prototypes gave her a richer understanding of the data by helping her piece together *"both the picture and the mathematical pattern,"* whereas TABLE afforded only the latter.

Description: Cursors and roadmaps are important for understanding where you are. Being able to capture both a high-level overview of the information while preserving the ability to drill down into the data is a crucial component to accessing interactive visualizations [SCWR21]. To navigate between these two levels, however, our participants emphasized the importance of markers to help them understand where they could move. TARGET addressed this with dropdown menus that allowed participants to navigate to any part of the visualization, explore, and then return to where they had started. In the words of P4, *"[This] mode is freedom for the user. Being able to jump around and move in real time as you would with your hand gives you a new way of exploring the information."* MULTI-VIEW approached this issue by allowing participants to move throughout the grid. *"With the table, I was trying to hold the numbers in my head and I wasn't trying to visualize it or anything,"* said P3. *"With [MULTI-VIEW], I can sort of think about it more like a visualization since I can move up and down, left and right. Even though I can use the arrows in the table, it just doesn't feel the same. I'm still feeling around and seeing what I can find."* Without these navigation tools, P7 noted that *"It's too easy to get lost ...I don't know how to backtrack."* To orient herself, P13 would first test to see if she was at the corner cells in the visualizations (e.g., *"Am I in the upper left or the bottom right cell here?"*) so that she could contextualize her position within the visualization and return to a point of origin. *"I know that I must be at the bottom left cell here because I can't go to the left,"* P13 said, *"but being able to know where that is beforehand would be very helpful."*

6. Discussion and Future Work

In this paper, we explore how structure, navigation, and description compose together to yield richer screen reader experiences for data

visualizations than are possible via alt text, data tables, or the current ARIA specification. Our results suggest promising next steps about accessible interaction and representation for visualizations.

6.1. Enabling Richer Screen Reader Experiences Today

Although our design dimensions highlight a diverse landscape of screen reader experiences for data visualizations, our study participants attested to the value of following existing best practices. Namely, alt text and data tables provide a good baseline for making visualizations accessible. Thus, visualization authors should consider adopting our design dimensions to enable more granular information access patterns only after these initial pieces are in place.

Existing visualization authoring methods, however, are likely insufficient for instantiating our design dimensions or producing usable experiences for screen reader users. In particular, it currently falls entirely on visualization authors to handcraft appropriate structures, navigational techniques, and description schemes on a per-visualization basis. As a result, besides being a time-consuming endeavor, idiosyncratic implementations can introduce friction to the reading process. For instance, per-visualization approaches might not account for an individual user's preferences in terms of verbosity, speed, or order of narrated output—three properties which varied widely among our study participant in ways that did not correlate with education level or experience with data. Thus, to scale and standardize this process, some responsibility for making visualizations screen reader accessible must be shared by *toolkits* as well. For example, our prototypes suggest a strategy for translating visualization specifications into hierarchical encoding structures (i.e., encoding channels as individual branches, and using visual affordances such as axis ticks and grid lines to populate the hierarchy levels). If toolkits provide default experiences out-of-the-box, visualization authors can instead focus on customizing them to be more meaningful for their specific visualization, and screen reader users have a stronger guarantee that the resultant experiences will be more usable and respectful of their individual preferences.

Current web accessibility standards also present limitations for realizing our design dimensions. For instance, there is no standard way to determine which element the screen reader cursor is selecting. Where ARIA has thus far focused on annotating documents with the semantics of a pre-defined palette of widgets, future web standards might instead express how elements respond to the *interaction affordances* of screen readers. For example, ARIA could offer explicit support for overview/detail hierarchies and different levels of description detail that can be progressively read according to user preferences.

6.2. Studying and Refining the Design Dimensions

Our conversations with study participants also helped highlight that design considerations can differ substantially for users who are totally blind compared to those who have low-vision. For example, partially-sighted participants used screen magnifiers alongside screen readers. As a result, they preferred verbose written descriptions alongside more terse verbal narration. Magnifier users also wished for in situ tooltips, which would eliminate the need to scroll back and forth between points and axes to understand data values.

However, promisingly, we found that using a screen reader and magnifier together affords unique benefits: *"I would have missed this point visually if I solely relied on the magnifier because the point is hidden behind another point"* (P12). Future work should more deeply explore how accommodations might complement and conflict when designing for different kinds of visual disability.

Similarly, in scoping our focus to screen readers and, thus, text-to-speech narration, we refrained from considering multi-sensory modalities in our design dimensions. Yet, we found that most participants had previous experience with multi-sensory visualization, including sonification (P5, 7, 9, 13), tactile statistical charts (P2–4, 10, 11, 13), and haptic graphics (P3, 4, 11, 13). Some participants reported that a *combination* of modalities would further enhance their experience—for example, getting a sonic overview of a line chart before reading more detailed text descriptions. Other participants, however, cautioned that adding multiple modalities can create additional confusion. For example, P7 noted that *"There's often a lack of explanation about how to map between sound and text."* Based on this testimony, it is unlikely that "sensory modalities" are merely an additional, independent dimension within our framework. Rather, future work must unpack the affordances of individual modalities, how they interact with one another, and how they impact the design of structure, navigation, and description.

6.3. What are Accessible Interactions for Data Visualizations?

In visualization research, we typically distinguish between static and interactive visualizations, where the latter allows readers to actively manipulate visualized elements or the backing data pipeline [YKSJ07; HS12]. Screen readers, however, complicate this model: reading is no longer a process that occurs purely "in the head" but rather becomes an embodied and interactive experience, as screen reader users must intentionally perform actions with their input devices in order to step through the visualization structure. While some aspects of this dichotomy may still hold, it is unclear how to cleanly separate *static reading* from *interactive manipulation* in the context of screen reader accessible visualizations, if these notions are conceptually separable at all. For instance, Hajas likened the navigation our prototypes afforded to *"shifting eye gaze, shifting focus of perceptual attention. When I navigate a visualization, naturally I would say 'I'm looking at this figure' and not that 'I'm interacting with this figure'."* Analogously, recent results in graphical perception find that sighted readers do not simply "see" visualizations in a single glance but rather perform active visual filtering operations [BMF21]. However, when using the binary tree prototype (Fig. 2e), Hajas noted a more distinct shift from reading to interacting. He said, *"it gave me the impression that I'm not just looking selectively, but I focus and zoom into the data,"* analogous to zoom interactions that change the viewport for sighted readers. Better characterizing the shift that occurs with this prototype, and exploring accessible manipulations of visualizations that allow screen reader users to meaningfully conduct data analysis, are compelling opportunities for future work.

Acknowledgements

We thank Matthew Blanco, Evan Peck, the NYU Digital Theory Lab, and the MIT Accessibility Office. This work was supported by NSF awards #1942659, #1122374, and #1941577.

References

- [AC97] AMICK, NANCY and CORCORAN, JANE. *Guidelines for Design of Tactile Graphics*. Tech. rep. American Printing House for the Blind, 1997. URL: <https://www.aph.org/research/guides/> 3.
- [Acc18] ACCELERATOR, SAS GRAPHICS. *SAS Graphics Accelerator Customer Product Page*. 2018. URL: <https://support.sas.com/software/products/graphics-accelerator/index.html#s1=12>.
- [AS01] ALDRICH, FRANCES K. and SHEPPARD, LINDA. "Tactile Graphics In School Education: Perspectives From Pupils". *British Journal of Visual Impairment* (2001). DOI: 10.1177/026461960101900204. URL: <https://doi.org/10.1177/026461960101900204> 3, 9.
- [Bag07] BAGOZZI, RICHARD P. "The Legacy of the Technology Acceptance Model and a Proposal for a Paradigm Shift." *Journal of the Association for Information Systems* (2007). DOI: 10.17705/1jais.00122. URL: <https://aisel.aisnet.org/jais/vol8/iss4/128>.
- [BCMS22] BOXHALL, ALICE, CRAIG, JAMES, MAZZONI, DOMINIC, and SURKOV, ALEXANDER. *The Accessibility Object Model (AOM)*. 2022. URL: <https://wicg.github.io/aom/> 2.
- [BCS96] BECKER, RICHARD A., CLEVELAND, WILLIAM S., and SHYU, MING-JEN. "The Visual Design and Control of Trellis Display". *Journal of Computational and Graphical Statistics* (1996). DOI: 10.1080/10618600.1996.10474701. URL: <http://www.tandfonline.com/doi/abs/10.1080/10618600.1996.10474701> 7.
- [Ber83] BERTIN, JACQUES. *Semiology of Graphics*. University of Wisconsin Press, 1983. URL: <https://dl.acm.org/doi/10.5555/10955975>.
- [BHR*21] BUTLER, MATTHEW, HOLLOWAY, LEONA M, REINDERS, SAMUEL, et al. "Technology Developments in Touch-Based Accessible Graphics: A Systematic Review of Research 2010-2020". *ACM Conference on Human Factors in Computing Systems (CHI)*. 2021. URL: <https://doi.org/10.1145/3411764.3445207> 3.
- [BK99] BARRASS, STEPHEN and KRAMER, GREGORY. "Using Sonification". *Multimedia Systems* (1999). DOI: 10.1007/s005300050108. URL: <https://doi.org/10.1007/s005300050108> 2.
- [BM13] BREHMER, MATTHEW and MUNZNER, TAMARA. "A Multi-Level Typology of Abstract Visualization Tasks". *IEEE Transactions on Visualization & Computer Graphics (Proc. IEEE VIS)*. 2013. DOI: 10.1109/TVCG.2013.124. URL: <http://ieeexplore.ieee.org/document/6634168/> 7, 8.
- [BMD*16] BAKER, CATHERINE M., MILNE, LAUREN R., DRAPEAU, RYAN, et al. "Tactile Graphics with a Voice". *ACM Transactions on Accessible Computing (TACCESS)* (2016). DOI: 10.1145/2854005. URL: <http://doi.org/10.1145/2854005> 3.
- [BMF21] BOGER, TAL, MOST, STEVEN B., and FRANCONERI, STEVEN L. "Jurassic Mark: Inattentional Blindness for a Datasaurus Reveals that Visualizations are Explored, not Seen". *IEEE Transactions on Visualization & Computer Graphics (Proc. IEEE VIS)*. 2021. DOI: 10.1109/VIS49827.2021.9623273. URL: <https://ieeexplore.ieee.org/abstract/document/9623273> 10.
- [BOH11] BOSTOCK, MICHAEL, OGIEVETSKY, VADIM, and HEER, JEFFREY. "D3: Data-Driven Documents". *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2011). DOI: 10.1109/TVCG.2011.185. URL: <http://vis.stanford.edu/papers/d3> 2.
- [BTOJ10] BROCK, ANKE, TRUILLET, PHILIPPE, ORIOLA, BERNARD, and JOUFFRAIS, CHRISTOPHE. "Usage Of Multimodal Maps For Blind People: Why And How". *ACM International Conference on Interactive Tabletops and Surfaces (ISS)*. 2010. DOI: 10.1145/1936652.1936699. URL: <https://doi.org/10.1145/1936652.1936699> 3.
- [Cha06] CHARMAZ, KATHY. *Constructing Grounded Theory*. Sage Publications, 2006. URL: <https://uk.sagepub.com/en-gb/eur/constructing-grounded-theory/book2359608>.
- [CJ95] CORNWALL, ANDREA and JEWKES, RACHEL. "What Is Participatory Research?": *Social Science & Medicine* (1995). DOI: 10.1016/0277-9536(95)00127-S. URL: <https://www.sciencedirect.com/science/article/pii/S0277953695001273>.
- [CJP*19] CHOI, JINHO, JUNG, SANGHUN, PARK, DEOK GUN, et al. "Visualizing for the Non-Visual: Enabling the Visually Impaired to Use Visualization". *Computer Graphics Forum (EuroVis)* (June 2019). DOI: 10.1111/cgf.13686. URL: <https://onlinelibrary.wiley.com/doi/10.1111/cgf.13686> 2.
- [Cos20] COSTANZA-CHOCK, SASHA. *Design Justice: Towards an Intersectional Feminist Framework for Design Theory and Practice*. MIT Press, 2020. URL: <https://papers.ssrn.com/abstract=3189696> 3.
- [CPR*21] CHUNDURY, PRAMOD, PATNAIK, BISWAKSEN, REYAZUD-DIN, YASMIN, et al. "Towards Understanding Sensory Substitution for Accessible Visualization: An Interview Study". *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS)* (2021). DOI: 10.1109/TVCG.2021.3114829. URL: <https://ieeexplore.ieee.org/document/9552177/> 3, 5.
- [DE19] DAVER, SCOTT and EDITORIAL TEAM, APPLEVIS. *What's New In iOS 13 Accessibility For Individuals Who Are Blind or Deaf-Blind*. 2019. URL: <https://www.applevis.com/blog/whats-new-ios-13-accessibility-individuals-who-are-blind-or-deaf-blind> 2, 6.
- [dGMB21] De GREEF, LILIAN, MORITZ, DOMINIK, and BENNETT, CYNTHIA. "Interdependent Variables: Remotely Designing Tactile Graphics for an Accessible Workflow". *ACM Conference on Computers and Accessibility (SIGACCESS)*. Association for Computing Machinery, 2021. DOI: 10.1145/3441852.3476468. URL: <https://doi.org/10.1145/3441852.3476468> 3.
- [DP20] DIMARA, EVANTHIA and PERIN, CHARLES. "What is Interaction for Data Visualization?": *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS)* (2020). DOI: 10.1109/TVCG.2019.2934283. URL: <https://hal.archives-ouvertes.fr/hal-02197062> 1.
- [FFT18] FUJIYOSHI, MAMORU, FUJIYOSHI, AKIO, TANAKA, HIROSHI, and ISHIDA, TORU. "Universal Design Tactile Graphics Production System BPL0T4 for Blind Teachers and Blind Staffs to Produce Tactile Graphics and Ink Print Graphics of High Quality". *Computers Helping People with Special Needs*. Ed. by MIESENBERGER, KLAUS and KOUROUPETROGLOU, GEORGIOS. Springer International Publishing, 2018. URL: http://link.springer.com/10.1007/978-3-319-94274-2_23 3.
- [FH05] FRØKJÆR, ERIK and HORNBEK, KASPER. "Cooperative Usability Testing: Complementing Usability Tests With User-Supported Interpretation Sessions". *ACM Extended Abstracts on Human Factors in Computing Systems (CHI)*. 2005. DOI: 10.1145/1056808.1056922. URL: <https://doi.org/10.1145/1056808.1056922> 6, 7.
- [GL15] GODFREY, A. JONATHAN R. and LOOTS, M. THEODOR. "Advice From Blind Teachers on How to Teach Statistics to Blind Students". *Journal of Statistics Education* (2015). DOI: 10.1080/10691898.2015.11889746. URL: <https://doi.org/10.1080/10691898.2015.11889746> 3, 8.
- [GOF08] GOULD, BRYAN, O'CONNELL, TRISHA, and FREED, GEOFFREY. *Effective Practices for Description of Science Content within Digital Talking Books*. Tech. rep. The WGBH National Center for Accessible Media, 2008. URL: <https://www.wgbh.org/foundation/ncam/guidelines/effective-practices-for-description-of-science-content-within-digital-talking-books> 1, 2, 5.

- [Ham13] HAMRAIE, AIMI. "Designing Collective Access: A Feminist Disability Theory of Universal Design". *Disability Studies Quarterly* (2013). DOI: 10.18061/dsq.v33i4.3871. URL: <http://dsq-sds.org/article/view/38713>.
- [HHG20] HORST, ALLISON M., HILL, ALISON PRESMANES, and GORMAN, KRISTEN B. *Palmerpenguins: Palmer Archipelago (Antarctica) Penguin Data*. 2020. URL: <https://github.com/allisonhorst/palmerpenguins7>.
- [Hig19] HIGGINS, TUCKER. "Supreme Court Hands Victory To Blind Man Who Sued Domino's Over Site Accessibility". *CNBC* (2019). URL: <https://www.cnbc.com/2019/10/07/dominos-supreme-court.html>.
- [Hig21] HIGHCHARTS. *Accessibility Module*. 2021. URL: <https://highcharts.com/docs/accessibility/accessibility-module2>.
- [HMM*11] HASTY, LUCIA, MILBURY, JANET, MILLER, IRENE, et al. *Guidelines and Standards for Tactile Graphics*. Tech. rep. Braille Authority of North America, 2011. URL: <http://www.brailleauthority.org/tg/2,3,5,8>.
- [HMM*03] HUTCHINSON, HILARY, MACKAY, WENDY, WESTERLUND, BO, et al. "Technology Probes: Inspiring Design for and with Families". *ACM Conference on Human Factors in Computing Systems (CHI)*. 2003. DOI: 10.1145/642611.642616. URL: <http://doi.acm.org/10.1145/642611.642616>.
- [HS12] HEER, JEFFREY and SHNEIDERMAN, BEN. "Interactive Dynamics for Visual Analysis". *Communications of the ACM* (2012). DOI: 10.1145/2133806.2133821. URL: <https://dl.acm.org/doi/10.1145/2133806.2133821>, 10.
- [Jac19] JACKSON, LIZ. *Disability dangle: A well-intended, elegant, yet useless solution to a problem we never knew we had. Disability dangles are most often conceived of and created in design schools and at IDEO*. Tweet. 2019. URL: <https://twitter.com/elizejackson/status/11106298182348185703>.
- [JMK*21] JUNG, CRESCENTIA, MEHTA, SHUBHAM, KULKARNI, ATHARVA, et al. "Communicating Visualizations without Visuals: Investigation of Visualization Alternative Text for People with Visual Impairments". *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS)*. 2021. DOI: 10.1109/TVCG.2021.3114846. URL: <https://ieeexplore.ieee.org/abstract/document/95529385>.
- [KJRK21] KIM, N. W., JOYNER, S. C., RIEGELHUTH, A., and KIM, Y. "Accessible Visualization: Design Space, Opportunities, and Challenges". *Computer Graphics Forum (EuroVis)* (2021). DOI: 10.1111/cgf.14298. URL: <https://onlinelibrary.wiley.com/doi/10.1111/cgf.14298>, 4, 9.
- [Koc12] KOCH, WOLF GÜNTHER. "State of the Art of Tactile Maps for Visually Impaired People". *True-3D in Cartography: Autostereoscopic and Solid Visualisation of Geodata*. Ed. by BUCHROITHNER, MANFRED. Springer, 2012. DOI: 10.1007/978-3-642-12272-9_9. URL: https://doi.org/10.1007/978-3-642-12272-9_9.
- [KS99] KARAHANNA, ELENA and STRAUB, DETMAR W. "The Psychological Origins of Perceived Usefulness and Ease-of-use". *Information and Management* (1999). DOI: 10.1016/S0378-7206(98)00096-2. URL: [https://doi.org/10.1016/S0378-7206\(98\)00096-2](https://doi.org/10.1016/S0378-7206(98)00096-2), 8.
- [KWT99] KAPER, H.G., WIEBEL, E., and TIPEI, S. "Data Sonification And Sound Visualization". *Computing in Science Engineering* (1999). DOI: 10.1109/5992.774840 2.
- [LKM*19] LI, JINGYI, KIM, SON, MIELE, JOSHUA A., et al. "Editing Spatial Layouts through Tactile Templates for People with Visual Impairments". *ACM Conference on Human Factors in Computing Systems (CHI)*. 2019. DOI: 10.1145/3290605.3300436. URL: <http://dl.acm.org/citation.cfm?doid=3290605.3300436>.
- [LLS19] LUNDGARD, ALAN, LEE, CRYSTAL, and SATYANARAYAN, ARVIND. "Sociotechnical Considerations for Accessible Visualization Design". *IEEE Transactions on Visualization & Computer Graphics (Proc. IEEE VIS)*. 2019. DOI: 10.1109/VISUAL.2019.8933762. URL: <https://ieeexplore.ieee.org/abstract/document/8933762>.
- [LS21] LUNDGARD, ALAN and SATYANARAYAN, ARVIND. "Accessible Visualization via Natural Language Descriptions: A Four-Level Model of Semantic Content". *IEEE Transactions on Visualization & Computer Graphics (Proc. IEEE VIS)*. 2021. DOI: 10.1109/TVCG.2021.3114770. URL: <http://vis.csail.mit.edu/pubs/vis-text-model2,4,5>.
- [MDN21] MDN CONTRIBUTORS. *ARIA - Accessibility*. Nov. 2021. URL: <https://developer.mozilla.org/en-US/docs/Web/Accessibility/ARIA2>.
- [MJBC18] MORRIS, MEREDITH RINGEL, JOHNSON, JAZETTE, BENNETT, CYNTHIA L., and CUTRELL, EDWARD. "Rich Representations of Visual Content for Screen Reader Users". *ACM Conference on Human Factors in Computing Systems (CHI)*. 2018. DOI: 10.1145/3173574.3173633. URL: <https://dl.acm.org/doi/10.1145/3173574.3173633>, 2, 8.
- [MTS19] MARTÍNEZ, RUBÉN ALCARAZ, TURRÓ, MIREIA RIBERA, and SALTIVERI, TONI GRANOLLERS. "Accessible Statistical Charts For People With Low Vision And Colour Vision Deficiency". *ACM International Conference on Human Computer Interaction*. June 2019. DOI: 10.1145/3335595.3335618. URL: <https://doi.org/10.1145/3335595.3335618>.
- [Mun09] MUNZNER, TAMARA. "A Nested Model for Visualization Design and Validation". *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS)* (2009). Conference Name: IEEE Transactions on Visualization and Computer Graphics. DOI: 10.1109/TVCG.2009.111. URL: <https://ieeexplore.ieee.org/abstract/document/52906959>.
- [PC99] PIROLLO, PETER and CARD, STUART. "Information Foraging". *Psychological Review* (1999). DOI: 10.1037/0033-295X.106.4.643 4, 5.
- [PGFM21] POTLURI, VENKATESH, GRINDELAND, TADASHI E, FROELICH, JON E., and MANKOFF, JENNIFER. "Examining Visual Semantic Understanding in Blind and Low-Vision Technology Users". *ACM Conference on Human Factors in Computing Systems (CHI)*. 2021. DOI: 10.1145/3411764.3445040. URL: <https://dl.acm.org/doi/abs/10.1145/3411764.3445040>.
- [SA01] SHEPPARD, LINDA and ALDRICH, FRANCES K. "Tactile Graphics In School Education: Perspectives From Pupils". *British Journal of Visual Impairment* (2001). DOI: 10.1177/026461960101900303. URL: <https://doi.org/10.1177/026461960101900303>.
- [Sar20] SARAH L. FOSSHEIM. *How (not) to make accessible data visualizations, illustrated by the US presidential election*. 2020. URL: <https://fossheim.io/writing/posts/accessible-dataviz-us-elections/1,2>.
- [SBDK05] SENGERS, PHOEBE, BOEHNER, KIRSTEN, DAVID, SHAY, and KAYE, JOSEPH 'JOFISH'. "Reflective Design". *ACM Conference on Critical Computing: Between Sense and Sensibility*. 2005. DOI: 10.1145/1094562.1094569. URL: <http://doi.acm.org/10.1145/1094562.1094569>.
- [Sci21] SCIENTIFIC, FREEDOM. *JAWS Web Verbosity*. 2021. URL: https://www.freedomscientific.com/SurfsUp/Web_Verbosity.htm6.
- [SCWR21] SHARIF, ATHER, CHINTALAPATI, SANJANA SHIVANI, WOBROCK, JACOB O., and REINECKE, KATHARINA. "Understanding Screen-Reader Users' Experiences with Online Data Visualizations". *ACM Conference on Computers and Accessibility (SIGACCESS)*. 2021. DOI: 10.1145/3441852.3471202. URL: <http://doi.org/10.1145/3441852.3471202>, 1, 2, 4, 8, 9.

- [She20] SHEW, ASHLEY. "Ableism, Technoableism, and Future AI". *IEEE Technology and Society Magazine* (Mar. 2020). DOI: [10.1109/MTS.2020.2967492](https://doi.org/10.1109/MTS.2020.2967492) 3.
- [Shn03] SHNEIDERMAN, BEN. "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations". *The Craft of Information Visualization*. Ed. by BEDERSON, BENJAMIN B. and SHNEIDERMAN, BEN. Morgan Kaufmann, 2003. DOI: [10.1016/B978-155860915-0/50046-9](https://doi.org/10.1016/B978-155860915-0/50046-9). URL: <https://www.sciencedirect.com/science/article/pii/B97815586091505004691>, 9.
- [SKK15] SAUNDERS, BENJAMIN, KITZINGER, JENNY, and KITZINGER, CELIA. "Anonymising Interview Data: Challenges And Compromise In Practice". *Qualitative Research* (2015). DOI: [10.1177/1468794114550439](https://doi.org/10.1177/1468794114550439). URL: <https://doi.org/10.1177/1468794114550439> 7.
- [SMWH17] SATYANARAYAN, ARVIND, MORITZ, DOMINIK, WONG-SUPHASAWAT, KANIT, and HEER, JEFFREY. "Vega-Lite: A Grammar of Interactive Graphics". *IEEE Transactions on Visualization & Computer Graphics (Proc. IEEE VIS)*. 2017. DOI: [10.1109/TVCG.2016.2599030](https://doi.org/10.1109/TVCG.2016.2599030). URL: <https://ieeexplore.ieee.org/abstract/document/7539624> 2, 6.
- [W3C18a] W3C. *WAI-ARIA Graphics Module*. 2018. URL: <https://www.w3.org/TR/graphics-aria-1.0/> 2.
- [W3C18b] W3C. *Web Accessibility Laws & Policies*. 2018. URL: <https://www.w3.org/WAI/policies/> 1.
- [W3C19] W3C. *WAI Web Accessibility Tutorials: Complex Images*. 2019. URL: <https://www.w3.org/WAI/tutorials/images/complex/> 1, 2, 5.
- [Web21] WEBAIM. *Screen Reader User Survey #9 Results*. 2021. URL: <https://webaim.org/projects/screenreadersurvey9/> 7.
- [WG69] WIEDEL, JOSEPH W. and GROVES, PAUL A. *Tactual Mapping: Design, Reproduction, Reading and Interpretation*. Tech. rep. Department of Health, Education, and Welfare, 1969. URL: https://archive.org/details/tactualmappingde00jose_03, 5, 9.
- [WOH*15] WENINGER, MARKUS, ORTNER, GERALD, HAHN, TOBIAS, et al. "ASVG Accessible Scalable Vector Graphics: Intention Trees To Make Charts More Accessible And Usable". *Journal of Assistive Technologies* (2015). DOI: [10.1108/JAT-10-2015-0124](https://doi.org/10.1108/JAT-10-2015-0124) 3.
- [WPA*21] WU, KEKE, PETERSEN, EMMA, AHMAD, TAHMINA, et al. "Understanding Data Accessibility for People with Intellectual and Developmental Disabilities". *ACM Conference on Human Factors in Computing Systems (CHI)*. 2021. DOI: [10.1145/3411764.3445743](https://doi.org/10.1145/3411764.3445743). URL: <https://dl.acm.org/doi/abs/10.1145/3411764.3445743> 2.
- [YKSJ07] YI, JI SOO, KANG, YOUN AH, STASKO, JOHN, and JACKO, J.A. "Toward a Deeper Understanding of the Role of Interaction in Information Visualization". *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS)* (2007). Conference Name: IEEE Transactions on Visualization and Computer Graphics. DOI: [10.1109/TVCG.2007.70515](https://doi.org/10.1109/TVCG.2007.70515). URL: <https://ieeexplore.ieee.org/abstract/document/4376144> 1, 10.
- [YRB01] YU, WAI, RAMLOLL, RAMESH, and BREWSTER, STEPHEN. "Haptic Graphs For Blind Computer Users". *Lecture Notes in Computer Science* (2001). Ed. by BREWSTER, STEPHEN and MURRAY-SMITH, RODERICK. DOI: [10.1007/3-540-44589-7_5](https://doi.org/10.1007/3-540-44589-7_5). URL: http://link.springer.com/10.1007/3-540-44589-7_5 2.