Toward Supporting Quality Alt Text in Computing Publications

Candace Williams
Carnegie Mellon University
Pittsburgh, USA
candacew@andrew.cmu.edu

Lilian de Greef
Apple Inc.
Cupertino, USA
ldegreef@apple.com

Ed Harris III
Apple Inc.
Cupertino, USA
edward.harris@apple.com

Leah Findlater
Apple Inc.
Cupertino, USA
lfindlater@apple.com

Amy Pavel
Apple Inc.
Cupertino, USA
apavel@apple.com

Cynthia Bennett
Apple Inc.
Cupertino, USA
cynthia_bennett@apple.com

ABSTRACT

While researchers have examined alternative (alt) text for social media and news contexts, few have studied the status and challenges for authoring alt text of figures in computing-related publications. These figures are distinct, often conveying dense visual information, and may necessitate unique accessibility solutions. Accordingly, we explored how to support authors in creating alt text in computing publications—specifically in the field of human-computer interaction (HCI). We conducted two studies: (1) an analysis of 300 recently published figures at a general HCI conference (ACM CHI), and (2) interviews with 10 researchers in HCI and related fields who have varying levels of experience writing alt text. Our findings characterize the prevalence, quality, and patterns of recent figure alt text and captions. We further identify challenges authors encounter, describing their workflow barriers and confusions around how to compose alt text for complex figures. We conclude by outlining a research agenda on process, education, and tooling opportunities to improve alt text in computing-related publications.

KEYWORDS

Accessibility, Alt Text, Blind, Data Representations, Scientific Figures, Screen Readers, Vision Impairment

ACM Reference format:

Candace Williams, Lilian de Greef, Ed Harris III, Leah Findlater, Amy Pavel, Cynthia Bennett. 2022. Toward Supporting Quality Alt Text in Computing Publications. In *Proceedings of the 19th International Web for All Conference (ACM W4A'22), April 25-26, 2022, Lyon, France.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3493612.3520449

W4A'22, April 25–26, 2022, Lyon, France

ACM ISBN 978-1-4503-9170-2/22/04.
https://doi.org/10.1145/3493612.3520449

1 Introduction

Accessibility research has examined the prevalence and tooling for alternative (alt) text in contexts like social media, news outlets, and presentations [5, 24, 25, 32, 33, 38, 43, 44, 46, 47, 52, 53, 59]. However, computing-related publications and their often complex figures is an understudied context. This gap is concerning for web accessibility researchers specifically as these papers are primarily circulated online. Figures common in computing-related publications, such as diagrams, visualizations, and screenshots, convey information differently from other image types; they often contain dense visual information—sometimes even combining multiple images, what we call elements, into a single figure—and are key to communicating data and concepts in research papers. Figures subsequently have distinct needs for their alt text, as reflected by the specialized guidelines and writing tools published for these types of images [8, 9, 18, 19, 28, 37, 41, 54].

While specific guidelines are promising, recent research on the prevalence of alt text in academic papers surfaced low rates [29], [51]. This raises questions around whether existing guidelines are useful, and what supports authors need to compose high-quality alt text. Accordingly, we conducted two studies to characterize the needs and opportunities for supporting authors in creating alt text in computing-related publications, focused on the field of human-computer interaction (HCI). First, we analyzed alt text plus caption prevalence and quality across 300 papers published from 2019 to 2021 at the ACM CHI conference, which we chose as it is a top-tier and general HCI venue that includes a broad range of research contributions (including accessibility), which has also recently increased the number of resources they provide to authors with instructions for writing and incorporating quality alt text [11, 12, 13, 14, 54, 57]. To build on Jung et al [29], we structured an analysis around the properties of scientific figures. Specifically, we (1) broke down the prevalence of different figure types and (2) examine the descriptive quality of figure alt text, caption, and combined alt text and caption. Second, we interviewed 10 authors of HCI and related publications to understand how alt text fits into their current paper writing workflows.

From the analysis of published alt text and captions we found that alt text prevalence at CHI increased over the last three years, correlating with the increasing number of resources for authors. However, score variance suggests there is room for improvement beyond increasing authors' access to guidelines. First, the most common were images, data visualizations, and diagrams, reflecting figure types of focus in scientific alt text guidelines [9, 18, 28, 41, 54, 55]. However, a substantial proportion of figures (38%) were complex composites of multiple elements (e.g., photographs, or visualizations and photographs combined)—formats absent from current guideline instructions. Next, alt text with more words generally received higher scores, contrasting with the common guideline advice for making alt text brief [18, 54]. Finally, caption quality varied greatly. However, a subset of papers scored high when both the combined caption and alt text quality were considered. Yet current guidelines do not advise authors on how to ensure that multiple sources work together to provide an optimal nonvisual understanding of figures.

From the interviews, we learned that while authors spend ample time developing the visual communication of their figures, they lacked tools, knowledge, and support from their research communities to compose high-quality alt text. For example, the design of paper-writing workflows and tools incentivized composing alt text at the last minute. Although some interviewees had read guidelines, they particularly struggled to decide what information belonged in captions versus alt text. They also struggled to describe composite figures and dense and highly visual, domain-specific information.

Combined, our work offers three main contributions: (1) an analysis of description prevalence and quality across a range of figure types, specifically published in the last 3 years at an international HCI conference (CHI); (2) an interview study on publication authors' experiences writing alt text for figures that offers insights into associated opportunities and challenges; and (3) future directions to improve alt text guidelines, process, and tooling support for authors of computing-related publications.

2 Related Work

Prior work offers guidelines for composing alt text, describes its prevalence in different contexts, and shares methods to increase the quantity and quality of alt text. We provide an overview and highlight the work that most closely relates to our study.

2.1 Alt Text Guidelines and Standards

Standards like the widely used Web Content Accessibility Guidelines 2.1 (WCAG) state that alt text should serve an "equivalent purpose" to an image and they include general best practices for writing alt text [55]. However, few alt text guidelines offer instruction specific to writing descriptions of complex figures like data visualizations. One of these exceptions was published by the Benetech Diagram Center [18]. In addition to describing general best practices for alt text, like keeping descriptions concise and ordering information from general to specific, Benetech offers advice specific for STEM-related images like diagrams, flowcharts, graphs, and maps. For example, in their guidance for describing line or bar charts, they recommend including chart titles and axes labels, but to exclude color when it is insignificant (e.g., when a bar chart's color assignments are arbitrary and not referenced elsewhere).

Shari Trewin incorporated Benetech's guidelines into a resource on how to describe figures in ACM publications [54], which publication venues like CHI have encouraged authors to follow. These guidelines adapt alt text best practices to the context of academic publications, such as not only providing data tables for charts but also when, where, and how to offer data points in alt text vs. an appendix. This resource also offers examples for how to describe figure types more specific to ACM publications, like architecture diagrams.

Recent work has also focused on evaluating and extending alt text guidelines specifically to data visualizations, can be unwieldy to navigate with screen readers [50]. Lundgard et al. [31] extended guidelines by introducing a conceptual model for the semantic content of visualization alt text, and further identifying which types of content are most useful according to blind and sighted readers. We build on this work by recommending additional guidance (e.g., composite images and addressing tensions between brief and complete alt text).

2.2 Alt Text Prevalence

Missing or low-quality alt text is a persistent issue. A recent analysis of the top one million home pages in February 2021 [57] revealed that 60.6% of their images are missing alternative text, corroborating other evaluations [25, 29]. Research specific to academic publications also signals low rates. A 2014 analysis of the top 10 academic mathematics journals found that none had an accessibility policy statement, only two referenced making descriptive figure captions to improve access, and none of the sampled articles used alt text for their figures [51]. A more recent analysis focusing on one figure type, data visualizations, published in HCI and visualization venues found that in 2019 and 2020, 0% of IEEE VIS and TVCG, 65% of ACM ASSETS, and 51% of ACM CHI data visualizations had alt text [29]. Our analysis builds on this work by considering both caption and alt text quality for a wider variety of figures (e.g., diagrams, photos, screenshots).

Low alt text prevalence may stem partly from challenges that authors encounter when writing alt text. For example, in a study with 20 Twitter users, the most common reasons for leaving out alt text were the time required to write it or simply forgetting to do so [25]. Even when authors make the effort to write alt text, multiple studies highlighted critical differences in understanding between sighted authors and screen reader users of what information is important to include [31, 32, 37]. This mismatch in intuition points to the importance of guidelines. However, there is also evidence that people do not implement guidelines even when they are explicitly provided. For example, in a study where crowd workers were given a text box and guidelines for authoring alt text, they frequently left out key information covered in the guidelines [37]. Similarly, an analysis of alt text for data visualizations in publication venues that offered guidelines found the data visualization alt text lacked some information requested

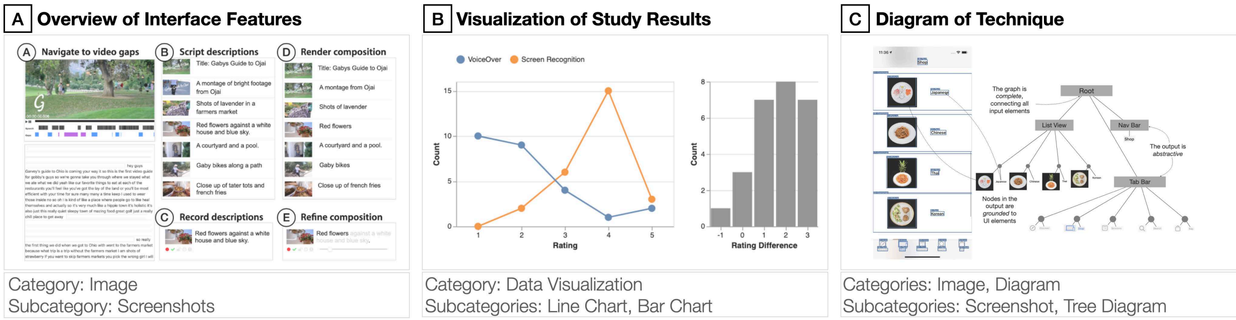

Figure 1. Three examples of common types of figures (top) with their categories and subcategories (bottom): (A) "Overview of Interface Figures" includes 5 images (screenshots) with labels, (B) "Visualization of Study Results" includes two data visualizations (a line chart and bar chart), and (C) "Diagram of Technique" includes an image and diagram (a screenshot and a tree diagram).

in guidelines [29]. Our interview study furthers these works by examining and reporting authors' first-hand experiences in the context of academic publications.

2.3 Improving Alt Text Prevalence and Quality

One approach to improve alt text prevalence and quality is to automatically generate alt text using artificial intelligence (AI) approaches [2, 34, 46, 58]. However, alt text for research figures, like data visualizations, is more specialized than general image descriptions, rendering most commercial automation unsuitable for the near future [31]. Efforts to automatically describe data visualizations are still in research stages, working to either generate descriptions directly from data sources [6, 17, 21, 35, 36, 42], or from raster images [7, 10, 15, 45]. Rather than offer just one high-level description, some other methods offer additional chart details broken into smaller components, like individual data points on a chart for screen readers to navigate over [3, 20, 48] — approaches that may help with vector visualizations in applications or webpages, but that would not work yet for rasterized visualizations in scientific publications.

Other efforts to improve alt text quality, which we build on, have focused on authoring assistance tools [8, 19, 27, 32, 37]. Benetech's Poet Training Tool [19], for example, lets authors upload an image and select the image's type from a list (e.g., Venn diagram, pie graph); the tool then shows alt text guidelines for the image type and a textbox for authors to write alt text. Additional scaffolding then further improved figure alt text. For example, Morash et al. [37] asked authors to write answers to a series of questions based on best practices for describing charts, which a template-based system automatically assembled into descriptions that were often superior to ones written in free response. Our suggestions build on this work by recommending sourcing feedback from research communities and potential technical advancements to ease alt text composition.

3 Methods

We assembled a team with expertise in digital accessibility and HCI methods, and one of the lead mentors is a blind researcher. To evaluate the current state of alt text in computing research figures, we conducted an analysis of image types, alt text, and captions on a sample of 300 figures from the past three years of ACM CHI. To understand authors' experiences with creating that alt text, we also conducted an interview study with 10 authors of publications in HCI and related fields.

3.1 Existing Alt Text Analysis

For our analysis of figures, alt text, and captions in recent HCI papers, we focused on the ACM CHI conference both because it includes a wide range of general HCI research as opposed to specifically focusing on accessibility, and because CHI recently increased the number of accessibility resources they provide to authors, including on writing and implementing quality alt text into their publications [11, 12, 13, 14, 54, 57]. We generated a dataset of 300 figures for analysis by randomly sampling 100 papers with figures from each of the CHI 2019, 2020, and 2021 technical paper proceedings, and randomly sampling one figure per paper. Our sample was scoped small since our analysis required figures to be graded manually.

To conduct our analysis, Williams and Harris, who are sighted, categorized each figure's types and graded the quality of its alt text and caption. For categorization, they labeled each figure with one or more of the four following high-level image types: (1) **data visualizations** like box plots or line charts, (2) **diagrams** depicting relationships between concepts such as flowcharts, circuits, or quadrants, (3) **tables** and **text blocks** which were formatted as an image rather than as screen reader-accessible content, and (4) other **images** such as illustrations, photographs, or screenshots (Fig. 1A). Graders then coded whether a figure with multiple elements repeated the same high-level image type (Fig. 1B), and whether it included multiple high-level types (Fig. 1C).

When grading image description quality, the two researchers assigned scores to four information sources: (1) the caption and all alt text combined, (2) the caption only, without any alt text, (3) the "short field" of the alternative text only, and (4) the "long field" of the alternative text only (see Appendix A for more information about "short" and "long" alternative text fields). We excluded body text from the analysis to keep manual grading feasible, but we recognize useful information may be conveyed there. Graders also noted clear errors in alt text (e.g., textual data points that differed from visible data points).

To assess alt text and caption quality, we adapted Gleason et al.'s alt text quality rubric [25] to include detail from guidelines for scientific figures [29, 54]. Our rubric includes the following codes:

0 Blank Description: There was no information in the source (e.g., the "long" alt text field).

1 Not Descriptive: The information source contained a description, but the description did not include any information about the figure content (e.g., *"Figure 1.")*. This is similar to Gleason et al.'s [25] and Jung et al.'s [29] criteria for "Irrelevant" or "ignorable" alt text, respectively.

2 Somewhat Descriptive: The information source included a high-level summary of the figure but no other information about the visual content or, some information about the visual content but no high-level summary. Trewin's guidelines requested a high-level summary [54] (e.g., *"3. Make the first sentence a 'title' less than 125 characters long"*, and *"4. Go from general to more specific details"*). This code is similar to Gleason et al.'s "Somewhat Relevant" [25]. An example of a summary sentence from Trewin's guidelines is: *"Boxplots of user response times in milliseconds for the six experimental conditions."* The second part of this code (some information but no high-level summary) captures cases where authors use descriptions only to supplement the visuals, such that the description is relevant but not necessarily complete (e.g., *"The red circles indicate the stopping time"*).

3 Descriptive: The information source included a high-level summary and *at least one* additional piece of information about the visual content. This is similar to Gleason et al.'s "Good" [25]. From Trewin's guidelines, an example of a high-level description followed by two details about the visual content of a flow chart is: *"Flowchart labelled '(a)' State Diagram with eight elements connected by action and flow links. The start state is 'Email Event (New arrival or Interaction)'. The end state is 'Email is finished'."* [54]. However, this description does not include additional details about the other six elements or connections on the flow chart.

4 Very Descriptive: The information source included a high-level summary and a description of much of the visual content in the figure. This rubric code is similar to Gleason et al.'s "Great" [25], but a description may also receive "Very Descriptive" if it achieves a score of "3" and points to further detail in supplemental material in line with Trewin's guidelines [54]. For example, the following description includes a high-level description of the figure along with all figure details: *"Schema graph headed by type 'company'. One type ('key_people') and four attributes ('key', 'cik', 'name' and 'founded') are linked to 'company'. 'key_people' has two further linked attributes: 'title' and 'name'. Examples of 'key' values are "IBM" and "IBM Corp" … [continues for two sentences to describe elements on schema graph]"* [54].

Williams and Harris each graded a sample of 30 figures. They then arbitrated disagreements with the larger team including the blind researcher, and adapted the rubric to increase clarity. For example, "2 Somewhat Descriptive" included only a high-level summary until the team noticed a common case of authors adding visual detail without a high-level summary (e.g., the caption only described the meaning of a color on a graph with no further information). Williams and Harris then graded 30 more figures and computed the inter-rater reliability on the resulting 60 figures. The Weighted Kappa for the overall description scores was 0.91, indicating "Excellent" agreement [23]. They each then independently graded 120 figures to complete the remaining 240 ungraded figures. With Pavel supervising the analysis, the team computed alt text, caption, and their combined descriptive quality. We additionally compared quality scores to figure type, description length, and publication year.

3.2 Interview Study

Upon obtaining approval from our organization's user studies review process, Williams and Bennett interviewed 10 employees who voluntarily gave informed consent to participate. To be eligible, participants had to be at least 18 years old and to have helped create a figure for a computing-related research publication in the past five years. We explicitly stated in the recruitment that participants did not need to have experience writing alt text.

During the 60-minute interview, we asked participants about their research backgrounds, what types of figures they create, and their experience composing alt text. We also developed an alt text writing activity as a probe [56] to concretize the discussion. Before their interview, each participant shared example papers that they co-authored, from which we selected two figures for use during their study. Our sample of 20 figures represented diverse image types, including data visualizations, diagrams, images of system designs, and figures with more than one element. During the activity, participants first read through Trewin's alt text guidelines [54], which were also made available to CHI 2020 and 2021 authors. Participants then received five minutes per sample figure (10 minutes total) to compose alt text while thinking out loud. We assured participants that they did not have to finish composing alt text within the short time span and that we would not evaluate their alt text; their drafted alt text was instead meant to spark discussion. In a post-activity reflection, participants shared thoughts on the guidelines, the quality of their alt text, and what support might help them with writing alt text for their publications in the future.

Williams and Bennett took notes during each interview, and human transcribers produced text transcripts of each interview's audio recording. They then conducted a thematic analysis [6] on the interview transcripts, and they discussed and strengthened the themes with the larger team's input.

4 Results

We present first the results from the figure description analysis followed by the interview findings.

4.1 Existing Figures, Alt Text, and Captions

4.1.1 What types of figures are authors creating? First, we overview figure type frequency (Fig. 2). Note that because figures could have multiple elements, the percentages sum to more than

100. Authors published a variety of figure types, the most prevalent being images, which appeared in 47% (N=142) of all figures. The most frequent image subtypes were photos, illustrations, and screenshots. The second most frequent figure type was data visualizations, appearing in 33% (N=98) of figures. The most common visualization subtypes were bar charts, box plots, line charts, scatter plots, and heat maps. Next, 17% (N=50) of figures contained a diagram where the most frequent subtypes were flowcharts, system diagrams, quadrant diagrams, and circuit diagrams. This variance is reflected in guidelines specific to figures [9, 18, 28, 41, 54] and underscores the need for figure alt text research to advise on several content types.

However, when considering figures holistically, 38% (N=115) of them consisted of multiple elements—a format that guidelines currently offer little instruction for. Of these figures, 84% (N=97) contained repeating instances of an element type (e.g., a series of 4 bar charts to represent data taken in 4 different weeks, or a set of result renderings from 3 different algorithms) (Fig. 1A). The other 26% (N=30) of these figures included more than one element type (e.g., a set of prototype photos with line charts alongside each photo depicting the prototype's performance) (Fig. 1C). Some figures contained both repeating elements and multiple figure types. For example, one figure included the following 7 elements: a series of 4 photos depicting a movement, a series of 2 3D renderings depicting the same movement, and a line graph depicting the force per distance of the movement. Finally, to distinguish elements from one another, figures often included visual annotations for each element (N=52). For instance, in the prior example, the photo series and the 3D rendering series were labeled "Physical Manipulation" and "Virtual Manipulation", respectively. These results show a mismatch between published figure compositions and example figures shown in guidelines, which advise on writing alt text for just one element.

4.1.2 How do authors describe figure content? As we illustrate in Figure 2, almost all (91%, N=274) of the figures included a caption and 81% (N=243) included alt text in the short and/or long fields. After removing uninformative captions (e.g., only the figure number, as in "Figure 1"), 90% (N=269) of figures still had captions. Figures with uninformative alt text were more common. Of the 81% (N=243) of figures with alt text, 20% (N=48) of alt text did not include any information about the visual content in the figure (e.g., "Figure 1"), such that only 65% (N=195) of figures included informative alt text[1].

Considering the entire figure description, that is, including *both* the caption and alt text combined: 38% (N=114) of figures received a description score of 2 (either a summary or a single piece of information about the visual content with no summary), 30% (N=89) of figures received a description score of 3, and only 23% (N=68) of figures received a description score of 4 for describing most of the visual content needed to understand the image. For example, a system architecture diagram received a score of 2 with the caption *"Figure 3: Architecture of the platform."* and alt text of *"An architecture diagram describing the data*

[1] Our finding that 65% of figures contained informative alternative text (for CHI 2019-2021) indicates a number higher than Jung et. al.'s finding that 51% of figures contained alt text (they only considered CHI 2019-2020).

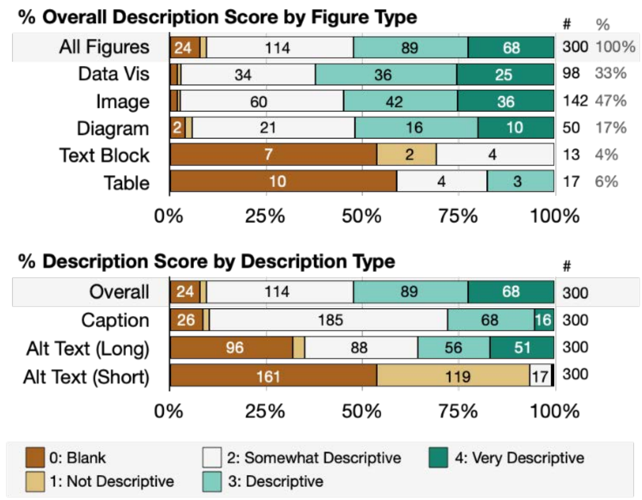

Figure 2. Two stacked bar charts depicting the percentage of overall scores by figure type (top) and the percent description scores by description type (bottom). Each figure may have more than one figure type (e.g., Data Visualization, and Image). N=300

architecture of the platform". While both the caption and alt text included a summary of the visual content, both lacked any detail about the system architecture itself, which showed Calendar and Collection APIs feeding into a Documentation Module, then a Visualization Module that a library patron can view. In contrast, a figure with flowchart diagrams received a score of 4 as it featured a caption with an overview of the figure content (*"Figure 1: Interaction Trajectories among various Identity Claims."*) and alt text providing detail about each element of the chart (e.g., "*Schema with four flowcharts placed in a row and labelled as'(a),' '(b),' '(c),' and '(d).' Flowchart labelled (a) has two elements connected by descriptor and flow links. The start state is 'I have sense of responsibility (SoR).' The end state is 'I am an activist/advocate.' …*"). The variance in alt text quality scores, and some redundancy observed between some captions and alt text (such as the score 2 example) demonstrate there is room for improving figure alt text quality.

When considering only the captions and not the alt text, 69% (N=185) of informative captions (i.e., that had any information about the figure) received a score of 2 indicating the caption included only a summary or a single piece of visual information. Some 2-scoring captions included, *"Figure 3: Left: Flow chart for iterative training process. Right: Network architecture."* and *"Figure 3: Worker accuracy (top) and task time (bottom) across crowd tasks when using the web and voice assistant."* These captions supplemented the figure's visual content with guidance on how it should be interpreted. Less frequently, captions included more detail about the visual content of the image: 23% (N=68) scored 3 for including a summary and at least one additional piece of information, and 5% (N=16) received a score of 4 for including a

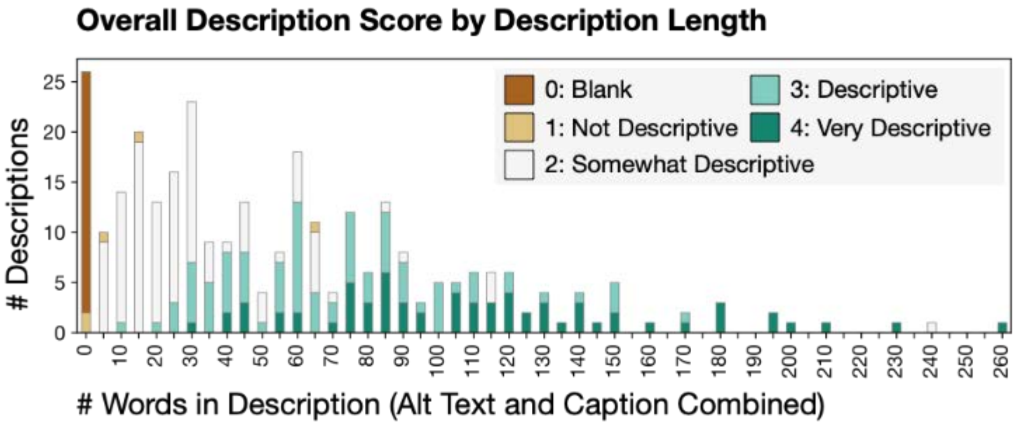

Figure 3. A histogram depicting the description length and overall description score. Descriptions with higher word counts (combining the alt text and captions) achieved higher scores than descriptions with shorter word counts.

summary and most of the information necessary to understand the visual content of the figure. Figures with only a caption and no alt text were unlikely to describe most of the visual content required to understand the image. Although unsurprising, as captions are not meant to serve as alt text, the caption scores indicate that they contain some visual information that may supplement the alt text, yet guidelines do not instruct authors to ensure that captions and alt text complement one another.

In contrast to the captions alone, the alt text contained more detail about visual content. For this alt text analysis we examined the highest score across both the short and long alt text fields for each figure. Of figures with at least one informative alt text field, 45% had a highest alt text score of 2 across both the short and long alt text fields, 29% (N=56) received a highest score of 3, and 26% (N=51) received a highest score of 4. When using both alt text and captions, authors occasionally complemented a 2-scoring caption, with more detailed (3-scoring or 4-scoring) alt text. For example, a figure with one of the 2-scoring captions shared previously (*"Figure 3: Worker accuracy (top) and task time (bottom) across crowd tasks when using the web and voice assistant."*) was complemented by 151 words of alt text (in the long field only) that received a score of 4 for describing the content of the top and bottom parts of the figure: *"The figure contains two subplots. Top - A boxplots of participant accuracy across 6 tasks and two conditions, web interface and voice assistants. Interquartile ranges are typically 0.2 - 0.3, with Comprehension task in Voice assistant condition and Text Moderation task in both conditions having broader ranges of 0.4. …"*. Interestingly, we found discrepancies in alt text and visuals in 14 figures; among these errors were incorrect labels and calling out an incorrect number of elements. The nearly half of 1 and 2-scoring alt text points to areas for improvement which go beyond developing guidelines.

Often when the alt text fields received a highest score of 2, so did the caption (72% of the time). When no field (either alt text fields or the caption field) received a score higher than 2, they often contained similar, but slightly different summaries of the visual content. For instance, one figure included the caption *"Figure 1: Sound Forest as seen from above."* and the alt text *"Map of Sound Forest room."* both of which did not describe the actual layout of the Sound Forest room that was depicted in the figure, a control room behind a mirror wall to an exhibition room with several small and large platforms inside. These examples potentially indicate author confusion in the purpose of the different fields, which are not delineated in current guidelines.

4.1.3 What factors impacted the level of description? We examined the potential relationship between the level of visual description with other factors, including figure type, length of description, and year of publication.

The level of visual description differed based on the type of figure (Fig. 2). Data visualizations, images, and diagrams were most likely to have high quality alt text, with 20% of diagrams (N=10/50) to 26% of visualizations (25/98) receiving a score of 4. In contrast, the majority of text blocks and tables, captured as inaccessible images (rather than screen reader navigable elements, a phenomenon also noted by Pareddy et. al. [43]), had blank alt text. These differences may point to a need for more figure type-specific examples in alt text guidelines and in particular, education on what makes text screen reader-accessible (i.e., text rendered in images must be written in alt text).

Higher word counts also achieved better scores (Fig. 3). That is, 4-scoring descriptions (captions and alt text combined) included more words (μ=34, σ=30 words) than 3-scoring descriptions (μ=75, σ=49 words), which in turn included more words than 2-scoring descriptions (μ=117, σ=54 words). Length also varied by figure type—82 words (σ=61 words) for data visualizations vs. 66 (σ=32) for images and 61 (σ=45) for diagrams. Increasing length did not always increase quality. Data visualizations' longer descriptions often provided a summary and attempted to describe part of the visual content, but still left off some important content, often receiving a score of 3 (Figure 2, top). Take the following example: *"A bar chart of quality of alt text generated by each method. Crowdsourcing, original alt text, and text recognition have the highest percentage of "Great" alt text. Scene description has high "Good" alt text but fails to meet the same bar. Other methods have more low-quality alt text."* It described a visualization that compared several methods by including a summary and relative comparison for a few of the methods, but it omitted other methods and lacks absolute values (e.g., did the top method perform with 80% accuracy or 20% accuracy?), resulting in a score of 3. Alt text guidelines emphasize brevity [18, 54]. However, complex figures may require longer descriptions to be comprehensible.

Finally, each subsequent year of CHI gave authors an increasing number of resources to write alt text. Figure description scores, combined from the caption and alt text, increased from CHI 2019 (μ=2.28, σ=1.33) to CHI 2020 (μ=2.65, σ=1.08) and again from CHI 2020 to CHI 2021 (μ=2.81, σ=0.90). This increase indicates that the resources and education may be leading to more and higher-quality alt text over time. But the increase could result from other reasons such as greater recognition of accessibility research and disabled scholars. Further, in CHI 2021 fewer than 30% of figures scored 4 for having a description that provided most of the information required to understand the figure content.

4.2 Interview Study Results

In this section, we overview findings from the interviews with 10 researchers who publish in HCI and related computing fields. We first briefly describe their backgrounds and experience followed by their processes for writing alt text both in the past and during the study itself. To preserve anonymity and because the alt text writing activity was used as a probe, we do not present or evaluate the alt text they wrote during the study. Instead, we share our observations about how they responded to the guidelines they read during the study. We finally overview challenges interviewees encountered around developing content for and structuring their alt text.

4.2.1 Interviewee backgrounds. Interviewees self-reported that they published in one or more of the following domains: HCI (8 participants), applied machine learning (4), visualization (3), design (2), software engineering (2), and one each in accessibility, computer science education, and learning sciences. For these publications, interviewees created a variety of figure types, including screenshots or images of technical systems (9 participants), data representations like charts and graphs (8), diagrams or flowcharts (6), and photos from user studies (1). Interviewees had between 3 and 21 years (median 9) of experience publishing, yet none of them had written alt text for all of those years; two participants had never written alt text for a publication though they were generally familiar with the concept, while the remaining eight interviewees had from 1 to 7 years of experience writing alt text for publications (median 5 years). Finally, two participants had learned what alt text was because it was required for an ACM publication submission.

4.2.2 Figure and alt text creation workflow. All interviewees shared their workflows for creating figures, and the eight who had alt text experience also shared their workflows for alt text. By comparing these two processes, we report differences in how authors conceptualized and implemented figures versus alt text.

The eight interviewees who had experience writing alt text reported spending much more time iterating on their figures than they did composing alt text. Two participants (P2 and P10) kept a file with alt text during the figure creation process but the other six participants consistently waited until the day the paper was due to write it. These discrepancies in time and effort between creating figures and creating alt text were primarily due to a perception that alt text need not be written until the figure is complete, and a lack of alt text creation support from their larger research communities and/or the tools used for paper-writing.

Interviewees conveyed tensions on when to create alt text. For example, they generally agreed alt text should be written earlier in the process rather than right before deadlines, but as prior mentioned, in practice they tended to wait until the last minute. One reason they reported proactivity to be infeasible was the iterative nature of their figure creation process, meaning early drafts of alt text would quickly become inaccurate. P6 explained differences in their figure and alt text creation workflows, "[Figure creation] comes very early in my process. It helps me think. The alt text doesn't evolve with the figure but more so it's when the figure is done." P6 and others considered figure creation as part of the paper's story whereas they treated alt text as metadata. P10 had practical reasons for waiting to write alt text. They stopped keeping early-written alt text (as mentioned before) and transitioned to writing it last-minute because their research team did not have expertise to quickly verify the correct alt text was in their final versions, "The alt text is not something people are looking at as part of the document. We noticed at the last minute that the figure has actually changed, but the alt text is still the old alt text, and we had to fix it at the last minute." P8 corroborated the difficulty of verifying whether alt text was accurate and functional, "The conferences we submitted to didn't give very good guidelines for us to actually check whether the alt text was saved correctly." As such, interviewees' workflows responded to the general invisibility of alt text: because it did not show up visually during drafting, authors could easily forget that it needed updating. As a result, seven opted to start writing alt text right before a deadline.

Interviewees also illuminated challenges incorporating alt text into their papers. For example, the ACM's custom publication pipeline (TAPS) includes a LaTeX "Description" tag for figures, the content of which does not export directly to PDFs but provides alt text in the pipeline's HTML version of the publication. The distinction created inconsistent behavior for the nine interviewees who wrote papers in LaTeX. This paper-writing workflow necessitated most authors to add accessibility tags (including alt text) to their PDFs with Acrobat Pro, which requires a paid subscription and for the source file to be complete.

P2, an accessibility expert, was the only participant whose workflow resonated with the most recent (CHI 2021) alt text guidance [13, 14]. However, even with expertise, according to interviewees, using Acrobat Pro was impractical until final preparations, as explained by P10, "If you change even a little word or something [in the source file], it's a very clerical process because you have to go back through the entire document and add all the accessibility alt text [again]." P10 and their research teams needed to be confident no further changes were needed as to not repeat the manual labor of making their PDFs accessible.

4.2.3 Alt text support and helpful resources. Participants reported that practice and past experiences were the most helpful for writing alt text. While two interviewees had consulted technical writers to assist with alt text and one participant had gained experience by describing technical diagrams to a blind colleague, the other interviewees relied on support from publication venues. P5 and others noted that an expectation to include alt text from the venue or tools they used incentivized them to do so, "That's [what has] largely driven whether we've written it or not, for better or worse." P6 agreed that explicit reminders were necessary, "Because [alt text] is optional [in LaTeX], you could publish it and no [warning] is ever raised." In some cases, these external pressures were crucial. For example, P1 only discovered alt text when it was required for a publication, and thus relied on the instructions, "I just did what the venue told me to do." All interviewees except P2 communicated feeling under-prepared to write quality alt text, even if they had read, or tried to follow, specific instructions.

Easing alt text implementation helped some participants. P2 and others noted that the burden for authors to prepare papers according to the venue's formatting requirements interfered with their ability to focus on writing high-quality alt text, and P2 and others expressed optimism that steps such as allowing authors to submit alt text without having to format it with specialized tools could relieve some of this burden. For example, P4, who was new to writing alt text, praised recent online paper submission forms for including an edit box for them to write the alt text, which removed the labor of adding the alt text as an accessibility tag in their PDF. P4 trusted their alt text would be incorporated into the final submission but they weren't sure of the process and did not confirm whether their alt text made it to the final copy. Finally, despite feeling unprepared, all 10 participants expressed that the guidelines presented to them during the interview were helpful; we observed that they were easily able to implement the instructions to begin with a summary of the figure and to follow with important details. After, interviewees remarked that they would like easier access to guidelines, such as having the guidelines directly within authoring tools, as they did not always find linked resources.

4.2.4 Challenges with alt text structure and content. Even with the guidelines we provided, and even when they had experience writing alt text for their computing papers, interviewees reported challenges, which we also observed during the alt text writing activity. These challenges included (1) where to write what information (e.g., caption or alt text), (2) knowing when interpretive descriptions were necessary, and (3) how to structure descriptions of dense visual information and figures with more than one element.

First, participants were confused by what to include in the alt text versus the body of the paper or figure caption. This concern arose as the guidelines we showed interviewees instruct that alt text should not repeat information presented in the figure caption or body text, and that alt text should be brief. Taken together, participants encountered challenges complying with both guidelines while drafting alt text during the study. First, while no one challenged the value of alt text, some interviewees argued that quality captions could provide most of the information necessary to understand the figure, and they struggled to know what additional detail was necessary to include in the alt text. As P1 put it, "The figure caption describes what's in a figure, right? So I was wondering how to not repeat the text if the caption is already telling what you see in the image." P4 relayed a similar concern: "Alt text is supposed to be brief but I have no idea what are the key elements I should talk about because my figure caption already describes the figure, so what's missing that is in the picture that I need to describe?" Some interviewees were confident their captions were near comprehensive, so they were confused as to where the "dividing line" (P9) was between a quality caption and alt text. They didn't have the nonvisual communication skills to distill missing elements appropriate for the alt text.

While some interviewees attempted to not repeat information reported elsewhere (e.g., paper body) while composing their alt text, others argued that the complexity of their figures and effort required to consolidate information about the figure from different sources made some information overlap appropriate. P5 reflected this concern, "You don't know if the user can get back and forth between the two things [paper body and figure alt text] easily to be able to put all of the [information] together." Interviewees also pointed out the value of perusing the paper figure-by-figure as a quick way to understand main points of a paper. Current guidelines assume there are clean separations between paper body text, figure captions, and alt text, and assume these information sources are consumed together. However, interviewees' experiences contrasted, leading them to request more examples of key differences between the content types and situations when it may be appropriate to repeat information in multiple places (e.g., writing an overview of a figure in both the paper body and figure caption).

Next, in some cases, interviewees were unsure of their readers' scientific and technical knowledge, which could impact what is appropriate to include in the alt text. Specifically, they struggled to know whether visual phenomena popular in their research specializations would be familiar and comprehensible to blind readers, even if the blind readers were area experts. P4, who builds interfaces, explained, "I'm relying on all this knowledge of, oh, somebody definitely knows what an image editor looks like. And maybe they don't." P4's introductory knowledge of screen reader navigation caused them to wonder if blind readers may not have the knowledge about spatial layouts to be able to quickly discern the novelty of their visual interface research. Interviewees also referred to canonical visuals in their field that they often used to convey examples, such as the computer graphics community's Stanford bunny (P1).[2] While an abstracted alt text such as "the Stanford bunny" would be brief, it assumes background knowledge that is not only domain specific but also highly visual. Conversely, a description of the bunny's physical features could also mislead unfamiliar readers; as with prior research on cultural imagery (e.g., memes) [27]) some underlying knowledge seemed potentially appropriate to communicate through alt text in some cases even if it was not an explicit visual feature of the image.

[2] https://graphics.stanford.edu/software/scanview/models/bunny.html

Finally, some interviewees struggled to describe figures with dense visual information, such as composite figures with more than one element. P4 researches developer tool user interfaces, and described the goal of one of their densely formatted figures, "I want to give people this impression: we're taking this overwhelming amount of information and making it manageable by how we organize this screen. But if you can't see it, I don't know how to convey that there's a structure to the screen so it's not overwhelming." For figures with multiple elements, some interviewees pointed out that relational qualities within the figures were particularly difficult to describe. P6 explained how their figures depicted how action on one UI would change the visualization displayed on another, "One big challenge is a lot of the data visualization research is interactive. So, that uses buttons and sliders and toggles, hover over, mouse over, clicking." Relatedly, P8 publishes figures that convey user study and fabrication processes. They reflected on tensions between describing elements within their figures and their relationship to one another, "We graphically showed all the various steps that one might take. This was an interesting challenge when it came to alt text because we could use alt text for each of the individual images. But then when it came to the entire layout, it wasn't clear to us how to more generally orientate someone that this is a collage." Interviewees wanted advice on structure and content to communicate both dense visual information of individual elements and those elements' relationships to others.

5 Discussion

In this paper, we inquired into the current state of alt text in computing-related publications. Though we narrowed our scope to figures in HCI papers and to authors of HCI-and related research, we believe our contributions extend to any publication presenting similar figures (e.g., data representations, diagrams, images, screenshots) and composites thereof. First, summarizing our analysis of existing alt text for 300 recently published figures, our results confirm and expand on prior research [29, 51] by characterizing the prevalence of different types of figures including composites, analyzing the alt text for those figures using a rubric specifically designed for this context, and considering the value of visual descriptions found in figure captions, the alt text itself, and the combination of the two. Even though alt text instructions had been available to authors of the figures we analyzed, like Jung *et al.* [29], we found that alt text prevalence and quality was still highly variable, and near half of figure descriptions had just one or fewer helpful pieces of information. Further, figure types varied widely, underscoring that guidance needs to address disparate content types. Additionally, some complex types, like the 38% of figures composed of multiple elements, are not represented in examples in current alt text guidelines [9, 18, 28, 41, 54]. Finally, longer alt text tended to receive a higher quality score, which contrasts with the brevity guideline common in alt text resources [18, 54] such as, "Keep the description as short as possible" [54].

During our interviews, we learned that authors encountered several challenges in creating alt text. First, as with social media alt text research [25], common paper-writing workflows and tools under-supported alt text and incentivized interviewees to wait until the last minute. For example, if an interviewee added alt text with Acrobat Pro and decided to change something in their LaTeX source file, they had to restart adding alt text. Further, just as our analysis uncovered 14 inaccuracies in existing alt text, interviewees reported that the invisibility of alt text and frequent iteration on figures increased the likelihood that incorrect alt text would ship. Next, interviewees largely felt under-aware of how to write alt text, even when they had years of experience writing alt text or were familiar with the guidelines given to authors. For example, interviewees struggled to know where to present what information about figures (e.g., the paper body text, the caption, or the alt text), whether and how to describe highly visual and domain-specific visual concepts, and how to structure alt text of dense visual information and composite figures. In the rest of this section, we distill these results into recommendations to better support authors to write high-quality alt text.

5.1 Future Policy and Guideline Directions

Interviewees candidly acknowledged that accessibility policies and resources motivated them to add alt text to their papers. We also noticed a small increase in alt text prevalence from 2019 to 2021 as the number of alt text resources made available to CHI authors increased. While we agree that policies and resources will not necessarily lead to universal high-quality alt text, our results showing variable alt text quality and author perspectives that they were under-prepared indicate that additional alt text education and information dissemination may be useful.

Our interviews surfaced potential improvements to existing guidelines. For example, several interviewees appreciated reading good examples but wanted bad examples of each figure type to help them avoid common mistakes. Interviewee confusion also suggests that guidelines should educate users on content differences between captions and alt text, instruct authors to write complete descriptions before editing for brevity, and advise on describing composite figures. The poorly-scoring figures with tables and text blocks suggest that guidelines could also exemplify when content should be formatted differently, such as actually creating a table or formatting code examples as text while making it distinct from body text (e.g., italics and indentations which are already commonly used to delineate block quotes). For example, the guidelines might begin with an explanation of the differences between figure types, followed by high and low-quality examples of captions and alt text for each type, and composites of repeating and different types. These examples would not only include written text but also annotations on how particular aspects of the caption or alt text contribute to the overall text being high or low quality. For instance, examples could demonstrate when it is appropriate to write a longer description and contrast it with a wordy description that does not surface the most pertinent information first. Other examples could offer structure for introducing all composites in a multi-element figure followed by brief descriptions of each individual composite. Guidelines could also remind authors to overview the figure's purpose in its caption, even if it is detailed in body text, in case blind readers are perusing the paper by navigating figure-by-figure. Finally, such guidelines would conclude with tips for updating and getting feedback on alt text, and for verifying that alt text is readable with different screen readers.

Similar to alt text scaffolding tools [8, 19, 27, 32, 37], additional instruction in multi-level descriptions could assist authors in providing background information on visual phenomena. For example, if an author needs to describe a visual user interface or canonical imagery like the Stanford bunny, instructions may prompt them to reflect on how the message is communicated, and if it is done so primarily through visual means (photos, data representations), they could include background information after the figure description so readers familiar with the concept are not overburdened. In line with proposed systems for enriching descriptions of images that may have visual cultural significance (such as images of media characters or memes [27, 38]), scaffolding could support alt text to provide background information when appropriate.

Even with improved guidelines, however, some uncertainty will likely persist when authors attempt to write alt text for any specific figure. Work in other domains (e.g., writing, product design, and graphic design) indicates that applying general guidelines to a specific content creation task is challenging [1]. To help authors adapt the general guidelines to their own figures, we recommend that publication venues offer feedback mechanisms, a successful approach in domains, such as visual design, that use guidelines [4, 16]. For example, authors could include their alt text in an edit box on the submission form and request feedback on the quality of the text or how to incorporate it into the publication files. A feedback loop may result in authors more confidently writing higher-quality alt text for future publications.

5.2 Toward Improved Tools and Techniques

Similar to tooling shortcomings revealed in other contexts [25], the accessibility researchers we interviewed revealed a tedious implementation process due to the narrow set of tools required for making papers accessible. For example, interviewees noted that alt text was invisible, making it hard for them to remember to write alt text initially and update it if the figure changed. As alt text research in other contexts such as professional work [32] is leveraging popular tools to increase alt text prevalence and quality, design and technical improvements could augment paper authoring tools to be more useful during alt text creation. For example, paper text and image editors could make alt text visible, along with captions, alert authors to write alt text if none exists, and alert authors to update their alt text when the figure file changes (e.g., in Overleaf). Currently, authors often write alt text in a separate program (e.g., LaTeX, Word) from the programs they use to create or edit their figures (e.g., Sketch, Illustrator, Excel, scripting). An opportunity exists to allow authors to create alt text alongside their figure in image editing software, and to create file formats that support saving alt text with the figure, such that the alt text would appear when the file is added to a text editor.

The process of creating alt text is time-consuming and error-prone for complex figures that may be frequently updated. Interviewees noted that they left alt text to the end to avoid updating it, and 14 alt text fields in our analysis contained inaccuracies. Future tools could allow either human-generated or automatically generated links between the figure text, the alt text and the captions, to better support regular figure updates and reduce errors. For example, a system could detect when authors change a label in a figure and automatically update or flag the corresponding label in the alt text and captions. In addition, domain-specific image recognition tools (e.g., a tool to recognize groups in vector illustrations [22, 30]) could surface inconsistencies between the image and alt text (e.g., "Three" in "Three architecture diagrams").

Additionally, automated tools could help authors better follow image description guidelines relevant to their specific figure. For example, an automated tool could step authors through creating alt text for dense visual information by detecting and indicating which figure elements have not yet been described (e.g., by applying techniques from presentation prompts [40]), and by recognizing when a longer description may be appropriate. Another idea is to help authors recognize and remove redundancies, perhaps applying techniques from audio description shortening [39]. And, while pre-populating fields with automatically generated alt text may lower the quality of human edits [32], a more promising approach may be to automatically recognize the figure type (e.g., [49]) and retrieve a relevant template (e.g., from existing guidelines [8, 19, 37]), or generate a custom template by extracting components in a composite figure. Finally, future tools could explore ways to enrich alt text reading experiences. For example, adding the capability for alt text to embed links would allow authors to easily transition among information sources by linking a reader to the point of the paper where the figure is referenced, or to outside resources [27].

6 Conclusion

Though foundational to web accessibility, alt text prevalence remains low and content creators struggle to apply guidelines in their specific contexts. We focused on computing-related research papers, an under-researched area where accurate and quality alt text is crucial for information comprehension. An analysis of 300 figures published at an HCI venue that offered authors resources for composing and implementing alt text showed that alt text's prevalence and quality varied greatly, meaning there is still opportunity to better support authors. Our interviewees—figure authors with varying experiences writing alt text—corroborated this variability: some lacked awareness of resources and even experienced alt text authors struggled to turn their dense visual information and multi-image figures into efficient and comprehensible alt text. Our results contribute to recommendations for expanding existing figure alt text guidelines, for greater support from publication venues, and for improvements to paper and alt text authoring tools that may ease the alt text composition and implementation process.

ACKNOWLEDGMENTS

We thank Cole Gleason for feedback and tool support.

REFERENCES

[1] Maneesh Agrawala, Wilmot Li, and Floraine Berthouzoz. "Design principles for visual communication." Communications of the ACM 54.4 (2011): 60-69.

[2] Apple. Use VoiceOver Recognition on your iPhone or iPad. Retrieved 1/6/2022 from https://support.apple.com/en-sa/HT211899

[3] Apple Developer Documentation. Audio Graphs. Retrieved 1/6/2022 from https://developer.apple.com/documentation/accessibility/audio_graphs

[4] Terry Barrett. "A comparison of the goals of studio professors conducting critiques and art education goals for teaching criticism." Studies in art education 30.1 (1988): 22-27.

[5] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samual White, and Tom Yeh. 2010. VizWiz: nearly real-time answers to visual questions. In Proceedings of the 23rd annual ACM symposium on User interface software and technology (UIST '10). Association for Computing Machinery, New York, NY, USA, 333–342. DOI: https://doi.org/10.1145/1866029.1866080

[6] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. Qualitative Research in Psychology 3, 2 (January 2006), 77–101. DOI:https://doi.org/10.1191/1478088706qp063oa

[7] Richard Burns, Sandra Carberry, and Stephanie Elzer. 2010. Visual and Spatial Factors in a Bayesian Reasoning Framework for the Recognition of Intended Messages in Grouped Bar Charts. In Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI Workshops '10). Available at https://www.aaai.org/ocs/index.php/WS/AAAIW10/paper/view/2017/2428

[8] Amy Cesal. 2020. Writing Alt Text for Data Visualization. (July 2020). Retrieved December 2, 2021 from https://medium.com/nightingale/writing-alt-text-for-data-visualization-2a218ef43f81

[9] CFPB Design System. 2021. Data visualization guidelines. (December 2021). Retreived December 2, 2021 from https://cfpb.github.io/design-system/guidelines/data-visualization-guidelines

[10] Charles Chen, Ruiyi Zhang, Eunyee Koh, Sungchul Kim, Scott Cohen, Ryan Rossi. 2020. Figure Captioning with Relation Maps for Reasoning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV '20). 1537-1545.

[11] CHI 2019. Guide to an Accessible Submission. Retrieved 1/6/2022 from https://chi2019.acm.org/authors/papers/guide-to-an-accessible-submission/

[12] CHI 2020. Guide to an Accessible Submission. Retrieved 1/6/2022 from https://chi2020.acm.org/authors/papers/guide-to-an-accessible-submission/

[13] CHI 2021. Guide to an Accessible Submission. Retrieved 1/6/22 from https://chi2021.acm.org/for-authors/presenting/papers/guide-to-an-accessible-submission

[14] CHI 2021. Making Your Final PDF Accessible. Retrieved 1/6/2022 from https://chi2021.acm.org/for-authors/presenting/papers/making-your-final-pdf-accessible

[15] Jinho Choi, Sanghun Jung, Deok Gun Park, Jaegul Choo, Niklas Elmqvist. 2019. Visualizing for the Non-Visual: Enabling the Visually Impaired to Use Visualization. In Proceedings of the Annual Visualization Conference organized by the Eurographics Working Group on Data Visualization (EuroVis '19). Computer Graphics Forum, 38: 249-260. DOI:https://doi.org/10.1111/cgf.13686

[16] Deanna P. Dannels, and Kelly Norris Martin. "Critiquing critiques: A genre analysis of feedback across novice to expert design studios." Journal of Business and Technical Communication 22.2 (2008): 135-159.

[17] Seniz Demir, Sandra Carberry, Kathleen F. McCoy. 2012. Summarizing Information Graphics Textually. In Computational Linguistics 2012. 38 (3): 527–574. DOI: https://doi.org/10.1162/COLI_a_00091

[18] The Diagram Center. Image Description Guidelines. Retrieved December 1, 2021 from http://diagramcenter.org/table-of-contents-2.html

[19] The Diagram Center. Poet Image Description Tool. Retrieved 1/6/2022 from http://diagramcenter.org/poet.html

[20] Leo Ferres, Petro Verkhogliad, Gitte Lindgaard, Louis Boucher, Antoine Chretien, and Martin Lachance. 2007. Improving accessibility to statistical graphs: the iGraph-Lite system. In Proceedings of the 9th international ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '07). Association for Computing Machinery, New York, NY, USA, 67–74. DOI:https://doi.org/10.1145/1296843.1296857

[21] Leo Ferres, Gitte Lindgaard, Livia Sumegi, and Bruce Tsuji. 2013. Evaluating a Tool for Improving Accessibility to Charts and Graphs. ACM Trans. Comput.-Hum. Interact. 20, 5, Article 28 (November 2013), 32 pages. DOI:https://doi.org/10.1145/2533682.2533683

[22] Matthew Fisher, Vineet Agarwal, and Tarun Beri. 2021. Automatic Hierarchical Arrangement of Vector Designs. (2021).

[23] Joseph Fleiss. "Statistical methods for rates and proportions." Statistical methods for rates and proportions. 1981. 321-321.

[24] Darren Guinness, Edward Cutrell, and Meredith Ringel Morris. 2018. Caption Crawler: Enabling Reusable Alternative Text Descriptions using Reverse Image Search. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Paper 518, 1–11. DOI:https://doi.org/10.1145/3173574.3174092

[25] Cole Gleason, Patrick Carrington, Cameron Cassidy, Meredith Ringel Morris, Kris M. Kitani, and Jefrey P. Bigham. 2019. "It's Almost like They're Trying to Hide It": How User-Provided Image Descriptions Have Failed to Make Twitter Accessible. In The World Wide Web Conference (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 549–559. https://doi.org/10.1145/3308558.3313605

[26] Cole Gleason, Amy Pavel, Xingyu Liu, Patrick Carrington, Lydia B. Chilton, and Jeffrey P. Bigham. 2019. Making Memes Accessible. In The 21st International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '19). Association for Computing Machinery, New York, NY, USA, 367–376. DOI:https://doi.org/10.1145/3308561.3353792

[27] Cole Gleason, Amy Pavel, Emma McCamey, Christina Low, Patrick Carrington, Kris M. Kitani, and Jeffrey P. Bigham. 2020. Twitter A11y: A Browser Extension to Make Twitter Images Accessible. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, 1–12. DOI:https://doi.org/10.1145/3313831.3376728

[28] Bryan Gould, Trisha O'Connell and Geoff Freed. 2008. Effective Practices for Description of Science Content within Digital Talking Book (December 2008). In The WGBH National Center for Accessible Media (NCAM). Retrieved December 2, 2021 from http://ncamftp.wgbh.org/ncam-old-site/experience_learn/educational_media/stemdx.html

[29] Crescentia Jung, Shubham Mehta, Atharva Kulkarni, Yuhang Zhao, and Yea-Seul Kim. 2021. Communicating Visualizations without Visuals: Investigation of Visualization Alternative Text for People with Visual Impairments. In IEEE Transactions on Visualization and Computer Graphics (VIS '21). DOI:https://doi.org/10.1109/TVCG.2021.3114846

[30] David Lindlbauer, Michael Haller, Mark Hancock, Stacey D. Scott, and Wolfgang Stuerzlinger. 2013. Perceptual grouping: selection assistance for digital sketching. In Proceedings of the 2013 ACM international conference on Interactive tabletops and surfaces (ITS '13). Association for Computing Machinery, New York, NY, USA, 51–60. DOI:https://doi.org/10.1145/2512349.2512801

[31] Alan Lundgard and Arvind Satyanarayan. 2021. Accessible Visualization via Natural Language Descriptions: A Four-Level Model of Semantic Content. In IEEE Transactions on Visualization and Computer Graphics (VIS '21). DOI:https://doi.org/10.1109/TVCG.2021.3114770

[32] Kelly Mack, Edward Cutrell, Bongshin Lee, and Meredith Ringel Morris. 2021. Designing Tools for High-Quality Alt Text Authoring. The 23rd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21). Association for Computing Machinery, New York, NY, USA, Article 23, 1–14. DOI:https://doi.org/10.1145/3441852.3471207

[33] Haley MacLeod, Cynthia L. Bennett, Meredith Ringel Morris, and Edward Cutrell. 2017. Understanding Blind People's Experiences with Computer-Generated Captions of Social Media Images. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17). Association for Computing Machinery, New York, NY, USA, 5988–5999. DOI:https://doi.org/10.1145/3025453.3025814

[34] Microsoft. Seeing AI App from Microsoft. Retrieved 1/6/2022 from https://www.microsoft.com/en-us/ai/seeing-ai

[35] Priscilla S. Moraes, Sandra Carberry, and Kathleen McCoy. 2013. Providing access to the high-level content of line graphs from online popular media. In Proceedings of the 10th International Cross-Disciplinary Conference on Web Accessibility (W4A '13). Association for Computing Machinery, New York, NY, USA, Article 11, 1–10. DOI:https://doi.org/10.1145/2461121.2461123

[36] Priscilla Moraes, Gabriel Sina, Kathleen McCoy, and Sandra Carberry. 2014. Evaluating the accessibility of line graphs through textual summaries for visually impaired users. In Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility (ASSETS '14). Association for Computing Machinery, New York, NY, USA, 83–90. DOI: https://doi.org/10.1145/2661334.2661368

[37] Valerie S. Morash, Yue-Ting Siu, Joshua A. Miele, Lucia Hasty, and Steven Landau. 2015. Guiding Novice Web Workers in Making Image Descriptions Using Templates. ACM Trans. Access. Comput. 7, 4, Article 12 (November 2015), 21 pages. DOI:https://doi.org/10.1145/2764916

[38] Meredith Ringel Morris, Jazette Johnson, Cynthia L. Bennett, and Edward Cutrell. 2018. Rich Representations of Visual Content for Screen Reader Users. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Paper 59, 1–11. DOI:https://doi.org/10.1145/3173574.3173633

[39] Amy Pavel, Gabriel Reyes, and Jeffrey P. Bigham. 2020. Rescribe: Authoring and Automatically Editing Audio Descriptions. Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology. Association for Computing Machinery, New York, NY, USA, 747–759. DOI:https://doi.org/10.1145/3379337.3415864

[40] Yi-Hao Peng, Jeffrey P Bigham, and Amy Pavel. 2021. Slidecho: Flexible Non-Visual Exploration of Presentation Videos. The 23rd International ACM

SIGACCESS Conference on Computers and Accessibility. Association for Computing Machinery, New York, NY, USA, Article 24, 1–12. DOI:https://doi.org/10.1145/3441852.3471234
[41] Pennsylvania State University. 2021. Accessibility and Usability at Penn State: Charts & Accessibility. Retrieved December 2, 2021 from https://accessibility.psu.edu/images/charts/
[42] Jason Obeid, Enamel Hoque. 2020. Chart-to-Text: Generating Natural Language Descriptions for Charts by Adapting the Transformer Model. In Proceedings of the 13th International Conference on Natural Language Generation (INLG '20). Association for Computational Linguistics, 138-147
[43] Sujeath Pareddy, Anhong Guo, and Jeffrey P. Bigham. "X-Ray: Screenshot Accessibility via Embedded Metadata." The 21st International ACM SIGACCESS Conference on Computers and Accessibility. 2019.
[44] Helen Petrie, Chandra Harrison, and Sundeep Dev. 2005. Describing images on the web: a survey of current practice and prospects for the future. In Proceedings of Human Computer Interaction International (HCII '05) 71
[45] Jorge Poco, Jeffrey Heer. 2017. Reverse-Engineering Visualizations: Recovering Visual Encodings from Chart Images. In Proceedings of the Annual Visualization Conference organized by the Eurographics Working Group on Data Visualization (EuroVis '17). Computer Graphics Forum, 36: 353-363. DOI:https://doi.org/10.1111/cgf.13193
[46] Emma Sadjo, Leah Findlater, and Abigale Stangl. 2021. Landscape Analysis of Commercial Visual Assistance Technologies. In Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21). Association for Computing Machinery, New York, NY, USA, Article 76, 1–4. DOI:https://doi.org/10.1145/3441852.3476521
[47] Elliot Salisbury, Ece Kamar, and Meredith Morris. 2017. Toward Scalable Social Alt Text: Conversational Crowdsourcing as a Tool for Refining Vision-to-Language Technology for the Blind. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing. 5, 1 (AAAI '17), 147-156.
[48] SAS. SAS Graphics Accelerator. Retrieved 1/6/2022 from https://chrome.google.com/webstore/detail/sas-graphics-accelerator/ockmipfaiiahknplinepcaogdillgoko?hl=en
[49] Manolis Savva, Nicholas Kong, Arti Chhajta, Li Fei-Fei, Maneesh Agrawala, and Jeffrey Heer. 2011. ReVision: automated classification, analysis and redesign of chart images. In Proceedings of the 24th annual ACM symposium on User interface software and technology (UIST '11). Association for Computing Machinery, New York, NY, USA, 393–402. DOI:https://doi.org/10.1145/2047196.2047247
[50] Ather Sharif, Sanjana Shivani Chintalapati, Jacob O. Wobbrock, and Katharina Reinecke. 2021. Understanding Screen-Reader Users' Experiences with Online Data Visualizations. In Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21). Association for Computing Machinery, New York, NY, USA, Article 14, 1–16. DOI:https://doi.org/10.1145/3441852.3471202
[51] Bruno Splendiani, Mireia Ribera, and Miquel Centelles Velilla. 2014. Are figures accessible in mathematics academic journals?. Comunicació presentada a: 29th Annual International Technology and Persons with Disabilities Conference (San Diego, CA, USA) (CSUN '14).
[52] Abigale Stangl, Meredith Ringel Morris, and Danna Gurari. 2020. ''Person, Shoes, Tree. Is the Person Naked?'' What People with Vision Impairments Want in Image Descriptions. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376404
[53] Abigale Stangl, Nitin Verma, Kenneth R. Fleischmann, Meredith Ringel Morris, and Danna Gurari. 2021. Going Beyond One-Size-Fits-All Image Descriptions to Satisfy the Information Wants of People Who are Blind or Have Low Vision. The 23rd International ACM SIGACCESS Conference on Computers and Accessibility. Association for Computing Machinery, New York, NY, USA, Article 16, 1–15. DOI:https://doi.org/10.1145/3441852.3471233
[54] Shari Trewin. 2019. Describing Figures. In Special Intrest Group on Accessible Computing (SIGACCESS) Resources. Retrieved on December 21, 2021 from https://www.sigaccess.org/welcome-to-sigaccess/resources/describing-figures/
[55] W3C Web Accessibility Initiative. 2019. Web Accessibility Tutorials: Guidance on how to create websites that meet WCAG (July 2019). Retrieved December 2, 2021 from https://www.w3.org/WAI/tutorials/images/complex/
[56] Jayne Wallace, John McCarthy, Peter C. Wright, and Patrick Olivier. 2013. Making design probes work. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13). Association for Computing Machinery, New York, NY, USA, 3441–3450. DOI:https://doi.org/10.1145/2470654.2466473
[57] Web Accessibility in Mind (WebAIM). Alternative Text. Retrieved 1/6/2022 from https://webaim.org/techniques/alttext/
[58] Web Accessibility in Mind (WebAIM). 2021. The WebAIM Million: An annual accessibility analysis of the top 1,000,000 home pages. (April 2021). Retrieved on December 2, 2021 from https://webaim.org/projects/million/
[59] Shaomei Wu, Jeffrey Wieland, Omid Farivar, and Julie Schiller. 2017. Automatic Alt-text: Computer-generated Image Descriptions for Blind Users on a Social Network Service. In Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17). Association for Computing Machinery, New York, NY, USA, 1180–1192. DOI:https://doi.org/10.1145/2998181.2998364

APPENDIX A

This appendix specifies the source and format of the alternative text in our analysis sample. Since 2019 CHI authors have published their papers by submitting a LaTeX or Word file to a central service (TAPS) that produces both a PDF version and an HTML version of the same submitted document.[3] For our analysis, we examined the captions and alternative text present in the HTML version of the paper as it contains any author-submitted alternative text by default (on the other hand, the TAPS service may require authors to re-add their alternative text to the PDF version of the paper). In the HTML version of the paper, each image features two types of alternative text that we call "short field" and "long field". For all years (2019-2021), the "short field" of alternative text is specified in the `alt` attribute of the image (`<img>`) HTML tag. For papers published in 2019 and 2020, the "long field" appears in the `longdesc` attribute of the image HTML tag. For papers published in 2021, the `aria-describedby` attribute of the image HTML tag provides an ID that identifies a hidden paragraph (`<p>`) that contains the "long field" alternative text. Importantly, paper authors added their alternative text to their original LaTeX or Word file, but they did not directly control the formatting of their alternative text in the resulting HTML version. In LaTeX, authors specified the "short field" and "long field" for each figure using a "`\Description[<short field>]{<long field>}`" command. In Word, authors specified the "short field" and "long field" for each figure using the Alt Text's "Title" field and "Description" field, respectively. In Word, the interface for adding Alt Text does not always include both fields (e.g., available through Format Picture > Alt Text).

[3] https://chi2021.acm.org/for-authors/chi-publication-formats