

A Dataset of Alt Texts from HCI Publications

Analyses and Uses Towards Producing More Descriptive Alt Texts of Data Visualizations in Scientific Papers

Sanjana Chintalapati*
sanjanac@cs.washington.edu
University of Washington
Seattle, WA, USA

Jonathan Bragg
jbragg@allenai.org
Allen Institute for AI
Seattle, WA, USA

Lucy Lu Wang
lucyw@allenai.org
Allen Institute for AI; University of
Washington
Seattle, WA, USA

ABSTRACT

Figures in scientific publications contain important information and results, and alt text is needed for blind and low vision readers to engage with their content. We conduct a study to characterize the semantic content of alt text in HCI publications based on a framework introduced by Lundgard and Satyanarayan [30]. Our study focuses on alt text for graphs, charts, and plots extracted from HCI and accessibility publications; we focus on these communities due to the lack of alt text in papers published outside of these disciplines. We find that the capacity of author-written alt text to fulfill blind and low vision user needs is mixed; for example, only 50% of alt texts in our sample contain information about extrema or outliers, and only 31% contain information about major trends or comparisons conveyed by the graph. We release our collected dataset of author-written alt text, and outline possible ways that it can be used to develop tools and models to assist future authors in writing better alt text. Based on our findings, we also discuss recommendations that can be acted upon by publishers and authors to encourage inclusion of more types of semantic content in alt text.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in accessibility**.

KEYWORDS

accessibility, scientific documents, alt text, dataset

ACM Reference Format:

Sanjana Chintalapati, Jonathan Bragg, and Lucy Lu Wang. 2022. A Dataset of Alt Texts from HCI Publications: Analyses and Uses Towards Producing More Descriptive Alt Texts of Data Visualizations in Scientific Papers. In *The 24th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '22)*, October 23–26, 2022, Athens, Greece. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3517428.3544796>

*Work conducted during internship at Allen Institute for AI

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ASSETS '22, October 23–26, 2022, Athens, Greece
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9258-7/22/10.
<https://doi.org/10.1145/3517428.3544796>

1 INTRODUCTION

Alternative text (or “alt text”) describes the content of a visual graphic or image to those who cannot see them. As such, alt text is an important component of accessible design. Most scientific documents use graphics to communicate information alongside text; scientific documents can be especially difficult to make accessible to BLV readers [5], with a large majority of these paper PDFs lacking usable alt text [47]. Mack et al. [31] conducted a study of BLV users and what they need from alt text, and found that graphs and charts are of special importance to these users, as they can be especially important for conveying results. This, coupled with the finding that the vast majority of scientific figures lack alt text altogether, suggests that even if the rest of the text in a scientific paper were accessible to a BLV reader, that a significant portion of the informational content of these works (figures) remain inaccessible, which can negatively impact reader experience.

Though there have been progress and attempts in automatically generating alt text descriptions of images on web and social media platforms [17, 35, 48], these methods do not apply as well to scientific images. From a machine learning perspective, much of the advancement in image recognition and scene understanding in recent years have derived from training neural models on large-scale labeled image datasets (such as ImageNet [42] and Google Open Images [25]), datasets that are primarily composed of natural images, which represent only a small proportion of the types of images found in scientific publications. Figures from scientific papers run the gamut of image types, including but not limited to natural images, medical images, diagrams, schematics, a wide array of graphs and charts, as well as combinations of these types, e.g., a medical image annotated with a histogram. Correspondingly, many established image understanding models cannot be directly or easily adapted for the scientific domain. Hybrid crowdsourcing solutions that integrate human experience with machine functionality may offer useful alternative solutions [17, 19, 39, 43]. For example, Qian et al. [39] advocate for a hybrid approach to image captioning where machines generate caption units and humans perform stitching. To support these solutions for scientific alt text, we need to better understand the current status of alt text content, and develop tools that can support authors in writing more useful descriptions of scientific figures.

Several prior studies have attempted to quantify the availability of alt text in scientific documents [7, 26, 47], though none have investigated the content of author-written alt texts and whether they convey adequate information about figures to blind and low vision

(BLV) readers. In this work, we extract and analyze the content of author-written alt text from papers published by the accessibility and HCI communities, and provide recommendations toward encouraging the inclusion of more types of descriptive information in alt text that may be useful to BLV readers. By extracting realistic author-written alt text, we also provide a useful data resource that can be used to support authors in writing better alt text and study how image understanding models and crowd-authoring techniques can be adapted to more effectively produce figure alt text in the scientific domain.

We process and extract author-written alt text from over 25K publications in the domains of accessibility and HCI, identifying nearly 3.4K pieces of valid alt text from 899 papers. To assess the type of information contained in these alt texts, we use the framework introduced by Lundgard and Satyanarayan [30], which accounts for four different levels of semantic content that may be conveyed by graphical data visualizations. We assess the semantic content present in the alt text corresponding to figures of graphs, charts, and plots (data visualizations), images that are prevalent in scientific papers, and for which the alt text content can be suitably represented using the levels introduced in the Lundgard and Satyanarayan [30] framework. We find that though most alt text contain basic information about the graph type, axes labels, and what is plotted, far fewer contain information beyond this. For example, only 50% of alt text in our sample discuss extrema or outliers in the data, and only 31% discuss trends or comparisons. The lack of this type of semantic content in alt text can make it difficult for a BLV user to understand these kinds of images in the way they desire, as found by Lundgard and Satyanarayan [30].

Our contributions in this work can therefore be summarized as:

- An assessment of the semantic information conveyed by author-written alt text of graph and chart figures extracted from papers published in venues representing work in accessibility, HCI, and related areas. We found that levels of covered content are inadequate, even at accessibility and HCI conferences, which have alt text requirements and writing guidelines.
- A dataset of 3386 author-written alt text from HCI publications, of which 547 have been annotated with semantic levels.¹ The methods used to construct this dataset can be extended to study trends in scientific figure alt text more broadly, and our dataset can be used develop tools and models to support alt text authoring. For example, we experiment with training a classifier that identifies semantic levels in text, which could be used to provide feedback to authors as they are writing alt text. We discuss and explore additional opportunities in Section 5.

2 RELATED WORK

We briefly discuss related work on how to write useful scientific alt text (Section 2.1), resources and methods for scientific figure understanding (Section 2.3), resources and methods for automatic alt text generation (Section 2.4), and other methods for improving figure accessibility (Section 2.5).

2.1 Guidelines for writing scientific alt text

The Web Content Accessibility Guidelines (WCAG) [9, 13] contain guidance on when alt text should be provided and the suggested content for the alt text. The National Center for Accessible Media (NCAM) has published guidelines including high-level recommendations for writing alt text for graphs, suggesting that a complete description should include text describing (i) the layout of the graph, (ii) the location of variables on the graph, and (iii) for static graphs, the overall trends presented, and for dynamic graphs, summary information such as the range of the axes.² The Benetech Diagram Center also provides image description guidelines with the goal of making it easier, cheaper, and faster to create and use accessible digital images.³ The referenced documentation includes both general best practices concerning aspects such as style and language that apply to every type of image, along with specific considerations for bar graphs, pie graphs, line graphs, and scatter plots such as listing the numbers in a pie graph from smallest to largest, and focusing on the change of concentration in scatter plots. Several academic publishers have also provided research-based guidelines for improving the accessibility of digital media. For example, the Association for Computing Machinery (ACM) strongly encourages authors to provide alt text for images and charts, and includes instructions for authors such as not duplicating the caption text and providing keywords.⁴

Lundgard and Satyanarayan [30] introduced a four-level conceptual model describing the semantic content of information that should be present in alt text descriptions of scientific data visualizations. Level 1 includes construction details such as the type of figure (e.g., bar plot or line plot), and labels of the axes. Level 2 includes statistics about the figure data, such as extremes and correlations. Level 3 includes larger takeaways, such as trends and patterns in the data. Level 4 includes domain-specific insights and societal context for the figure data. The authors conducted studies including BLV users, and found that these users gained the most information from textual descriptions conveying information from semantic levels 1–3 [30]. This finding corresponds to the recommendations of alt text content made by the NCAM. Though this framework is not a guideline document *per se*, we adopt it for this study in order to evaluate the quality of author-written alt text for graphs and charts found in scientific publications. It has been validated through studies with BLV participants, and we are unaware of alternative frameworks.

Towards alt text preferences, Bennett et al. [4] conducted a study on best practices for describing race, gender, and disability status in alt text, and found that people in photographs preferred to be described with the language that they use to talk about themselves, and that descriptions of concrete visual details were more appropriate than language around identities. Though this does not apply directly to scientific figures, alt text written by someone other than the authors of a paper may want to consider how the authors intended for the figure to be understood as well as the language used by the authors in the rest of the publication. After reviewing BLV people's experiences with digital image types such as news

²<https://www.wgbh.org/foundation/ncam/guidelines/accessible-digital-media-guidelines>

³<http://diagramcenter.org/making-images-accessible.html>

⁴<https://authors.acm.org/proceedings/production-information/describing-figures>

¹The dataset and annotations are available at <https://github.com/allenai/hci-alt-texts>.

articles and employment websites, Stangl et al. [46] found that a one-size-fits-all approach for image descriptions is not ideal. Similarly, through interviews with screen reader users, Mack et al. [31] found that different BLV users had varying preferences about the level of detail that they found to be most helpful in alt text, although most users concluded that both brevity and the availability of detailed information were desirable traits. Such findings should be kept in mind when authoring effective alt text.

2.2 Interfaces for authoring alt text

Morash et al. [33] developed interfaces to guide novice web workers in writing descriptions of scientific images. The authors queried workers for information about select image attributes based on the NCAM guidelines, on attributes such as image type, title, and units shown. They found that the templated query method was preferred by the workers and produced better image descriptions.

Mack et al. [31] built a prototype interface for authoring alt text, and measured the quality of alt text on a four-point scale based on three interface variations: the current PowerPoint interface; a free-form interface, where suggestions were presented as a bulleted list; and a template interface, where each prompt was listed separately and included a text box to respond to that prompt. Participants who use screen readers were asked to rank the quality of alt text written under the PowerPoint, free-form, and template interfaces, they found that, in general, the free-form interface encouraged authors to write alt text that is more closely aligned with the preferences of screen reader users.

2.3 Automated methods for scientific figure understanding

Scientific figure understanding tasks such as figure classification, visual question-answering (VQA), or image captioning have received significant attention from the AI community in recent years [10, 11, 20, 21, 23, 24, 27, 32, 34, 38, 40, 45]. In Table 1, we describe a number of datasets that have been introduced to train models and evaluate their performance on these tasks.

Datasets introduced for scientific image classification include FigureSeer [45], DocFigure [21], and SlideImages [34]. All three datasets include realistic images extracted from scientific papers, along with labels to classes such as graph, medical image, or natural image. These datasets have been used to train models that can detect figure type, which is an essential piece of information that should be available in alt text. However, figure type is only one of many pieces of information that BLV users may need to understand the content of an image, and therefore these datasets are of limited use in the alt text generation setting.

Towards more detailed figure understanding, datasets such as FigureQA [24], DVQA [23], and PlotQA [32] have been introduced. These datasets are made up of synthetically generated graphs and charts, along with associated questions and answers about the graph contents. For example, questions may be related to the graph title, axes, the x and y -axes values associated with specific data series, or the names of each data series. In PlotQA [23], many questions also go beyond the structure of the plot and may require data retrieval or additional reasoning (e.g. What is the average difference between two data series?) Models trained on these datasets have

shown improving performance [24, 27], though because the plots in these datasets are generated synthetically, they may not transfer well to graphs found in actual scientific publications, which are significantly more noisy, diverse, and variable than those found in these datasets. Also, though figure understanding through VQA is related to the task of producing alt text, the task of VQA itself does not produce a coherent textual description of the figure, which is the desired outcome for alt text. The FigCAP project [10, 11] attempts to bridge this divide by deriving figure descriptive text from the questions and answers of the FigureQA [24] dataset; though FigCAP refers to itself as a figure captioning dataset, the “captions” they provide are more similar to the notion of alt text, including descriptions about the figure structure, axes labels, data values, etc.

Recent work has aimed to automatically generate captions for both natural images [12, 44] and figures [20, 38]; however, caption generation and alt text generation are not the same task, and could be said to have differing goals. Captions are intended to be consumed by all readers, and contain information that complements the content of the image; while alt text is meant to explain the informational content of the image for users who cannot see it. SciCap [20] introduces a large dataset of graph figures and captions derived from arXiv; captions are those originally written by authors, and are post-processed to remove tokens corresponding to numbers and equations. The ImageClefMed Captioning task released a relatively much smaller dataset focused on captioning of medical images derived from scientific publications [38].

As far as we know, there are no datasets available for studying scientific figure alt text generation with realistic, author-written alt text. Though FigCAP [10, 11] and FigJAM [40] explore text generation in the alt text setting, the images and target texts used are synthetic (derived from FigureQA [24]) and not representative of the analogous task in a realistic setting.

2.4 Automated methods for alt text generation

Office 365 generates alt text for any image or figure pasted into Microsoft PowerPoint.⁵ Though easy to use, the feature shows limited performance on scientific figures, usually only describing the type of the figure. For example, for the figures in Table 2, the corresponding alt texts generated by Office 365 are, respectively: “Chart, bar chart,” “Chart, line chart,” and “Chart.” Though Office 365 usually identifies the correct type of chart, no other information about the figure content is generated, and the resulting alt text is of limited use to the reader.

Qian et al. [40] created synthetic datasets for figure alt text generation by adapting the FigureQA [24] and DVQA [23] datasets for figure VQA. Figures in both datasets are synthetic (not from scientific papers), and alt text units are derived from the data and information used to construct each figure. The authors then trained the FigJAM model, which generates alt text descriptions based on the multimodal inputs of the raw figure image and figure metadata [40]. Though performance on synthetic data was shown to be good,

⁵<https://www.microsoft.com/en-us/microsoft-365/powerpoint>

Dataset	Domain	Realistic/ Synthetic	Task	Size of dataset
FigureSEER [45]	Scientific figures	Realistic	Image classification (6 classes)	60K figures
DocFigure [21]	Scientific figures	Realistic	Image classification (28 classes)	33K images
SlideImages [34]	Educational illustrations	Realistic	Image classification (8 classes)	3K images
FigureQA [24]	Scientific graphs and charts	Synthetic	VQA	100K images
DVQA [23]	Bar charts	Synthetic	VQA	300K images
PlotQA [32]	Plots	Synthetic plots; Realistic data	VQA	224K images
SciCap [20]	Scientific figures	Realistic	Image captioning	416K figures
ImageClefMed Caption [38]	Medical images	Realistic	Image captioning	5K images
FigCAP [10, 11]	Bar charts, pie charts, line plots	Synthetic	alt text generation	110K figures

Table 1: Datasets for scientific figure understanding

the model may not generalize to realistic figures found in scientific papers, which exhibit significantly more variability than those evaluated by Qian et al. [40].⁶

Researchers have developed tools for generating alt text information for images on non-scientific social platforms. Gleason et al. [17] addressed the accessibility barrier on Twitter by creating a browser extension which adds alt text to Twitter using six methods, such as reverse image searching and automatic image captioning. However, this project focused on Twitter images, which differ from realistic scientific figures. Additionally, Wu et al. [48] deployed an automatic alt text system to identify faces, objects, and themes for photos on Facebook in order to make them more accessible to screen reader users. The domain of natural images on social media again significantly differs from our domain of scientific graphs and charts.

2.5 Other methods for improving figure accessibility

Researchers have also developed alternatives to alt text for improving figure accessibility. ChartSense [22] and PlotDigitizer⁷ are chart data extraction methods which convert chart images into structured data tables. However, charts of the same type are too diverse in style to apply a single extraction algorithm, algorithms have trouble interpreting overlapped visual entities, and there is no text-region detection algorithm for chart images with sufficient accuracy [22]. Auditory graphs, tactile graphs, and various multimedia approaches have also been introduced to improve the accessibility of graphs and charts [15, 36].

Crowdsourcing, in which a group of non-experts completes a task that is currently infeasible to accomplish via automated methods, has also been used to good effect for making images more accessible. Salisbury et al. [43] explored a novel approach in which

crowdworkers were paired together to create image descriptions. Crowdworkers were asked questions to extract desired details about an image such as the location of the image and what emotions the image evoked. Platforms such as VizWiz [6] or Be My Eyes⁸ connect BLV users to sighted crowdworkers and volunteers via an app for assistance with questions and daily tasks. Building upon the successes of the VizWiz Grand Challenge [19], a similar solution could be created to connect BLV researchers with volunteers that can provide suitable descriptions or answers to questions about scientific figures. Mack et al. [31] built user interfaces for both authoring alt text and providing feedback on automatic alt text, which could be theoretically connected to a crowdsourcing platform and adapted to collect alt text for scientific figures based on user demand.

Lastly, researchers have developed methods for surfacing alt text that do not involve manual effort from humans or automated alt text generation. Guinness et al. [18] found that many images appear in several places across the web, and used this insight to develop Caption Crawler. This system uses reverse image search to find alt text from similar images available on the web and surfaces these alternate descriptions to the user. This method works quite well in the domain of natural images, where many similar images of places or things can be found on the open internet.

3 METHODS

We extract and analyze the presence and content of alt text for graphs, charts, and plots in scientific papers. These figures are typically used to visualize data and results, are of special importance to BLV readers [31], and are the types of images for which Lundgard and Satyanarayan [30] have defined semantic levels. The Lundgard and Satyanarayan [30] framework describes four levels of semantic content, organized by increasing complexity:

Level 1: enumerating visualization construction details (e.g., type, marks, and encodings)

⁶The authors do not release their dataset or pretrained models, so we are unable to perform a comparative analysis on realistic figures derived from scientific papers.

⁷<http://plotdigitizer.sourceforge.net/>

⁸<https://www.bemyeyes.com/>

Level 2: identifying statistical concepts and relations (e.g., extremes and correlations)

Level 3: characterizing perceptual and cognitive phenomena (e.g., trends and patterns)

Level 4: articulating domain-specific insights or societal context.

Lundgard and Satyanarayan [30] found that levels 1–3 were reported most useful by blind and low vision readers. Level 4, which incorporates significantly more subjective information, was found to be less essential to understanding; in fact, a majority of blind readers in their study (63%, $n=19$) believed that figure alt text should not contain level 4 content.

3.1 Sampling papers and extracting author-written alt text

Our goal is to construct a dataset of author-written alt text by automatically sampling and extracting alt text from papers. We start with the set of papers from two conferences: ACM CHI Conference on Human Factors in Computing Systems (CHI) and ACM SIGACCESS Conference on Computers and Accessibility (ASSETS) published in the years 2010–2020. We identify these papers using the ACM’s reported DOIs, and link these to PDFs in the Semantic Scholar corpus [1]. This yields 5218 PDFs. We then iteratively extend our paper sample to include all papers written by authors who have published in CHI and/or ASSETS. We prioritize authors based on the frequency of their publications in CHI and ASSETS, and retrieve these authors’ other publications using the Semantic Scholar API.⁹ From these queries, we assemble a further sample of 20000 paper PDFs to process and extract alt text.

CHI and ASSETS are premiere conferences on human-computer interaction and accessible computing, and both conferences have a history of soliciting accessible paper submissions¹⁰ and requiring authors to provide alt text for figures. We therefore expect a much higher percentage of papers published at these venues to contain valid alt text. Similarly, our rationale for sampling additional papers published at other venues by authors who have published at CHI and ASSETS is based on our hypothesis that such authors are more likely to write alt text in general, even if another publishing venue may not require alt text for submission. Initially, we intended to extract alt text from a stratified random sample of papers representing all fields of study, but a pilot attempt showed that random sampling would yield virtually no alt text. We arrived at this conclusion after processing a stratified sample of 5000 PDFs from 2010 to 2020 and extracting only a single piece of descriptive alt text. Given the time and expense of processing PDFs to extract alt text at scale, we decided that the stratified random sample strategy was untenable.

For each paper in our sample, we followed a three step process to extract alt text. First, we processed the PDF using Adobe Acrobat Pro¹¹ to convert the PDF to HTML using the Adobe Acrobat Pro Action Wizard. We used Adobe Acrobat Pro rather than an alternate programmatic approach because we were unable to identify

⁹<https://api.semanticscholar.org/>

¹⁰CHI directs authors to the SIGCHI guidelines for an accessible submission: <https://sigchi.org/conferences/author-resources/accessibility-guide/> and ASSETS provides these instructions: https://assets21.sigaccess.org/creating_accessible_pdfs.html

¹¹<https://www.adobe.com/acrobat/acrobat-pro.html>

a PDF processing library capable of extracting alt text.¹² Second, we extracted alt text from the converted HTML document. Third, we filtered the extracted alt text according to a set of heuristics. Our pilot attempts at extracting alt text revealed that most of the extracted alt text consisted of uninformative short descriptions like “Image” or file paths like “C:\path_to\figure1.jpeg,” so we defined a set of filtering criteria to remove these. We also filtered out alt text shorter than 80 characters since many shorter alt text fall under the category of uninformative short descriptions. To determine this 80-character threshold, we analyzed a sample of 100 extracted alt text to find a limit that maximizes recall without sacrificing precision. We refer to the alt text that pass these filtering criteria as “valid alt text.”

For all valid alt texts, we identify those that are likely to correspond to graphs, charts, and plots. We iteratively defined another set of heuristics: a list of words and phrases that correspond to graphs and charts (e.g., “graph”, “chart”, “error bar”; the full list of terms are provided in Appendix A). For each figure, we search for token matches in the alt text and image caption against this list of terms, and retain only figures and alt text matching at least one term. The alt text of the matching figures are then annotated with the Lundgard and Satyanarayan [30] semantic levels.

3.2 Annotation of alt text semantic levels

To study what types of content are present in author-written alt text, we ask annotators to assess the semantic content levels present in each sentence of each piece of alt text. We split alt text into sentences using the scispaCy NLP library [37]. Annotators are shown these sentences along with the corresponding figure caption and a link to view the figure.

Six label options were provided for each sentence:

- Level 1: Figure logistics
- Level 2: Statistical properties and comparisons
- Level 3: Complex trends and patterns in data
- Level 4: Domain-specific insights or societal concepts to help explain Level 3 trends
- This alt text contains no levels of content
- This image is not a graph or chart

If a figure is not a graph or chart, the annotator is instructed to select the last option. Otherwise, up to three semantic levels could be selected for each sentence. These label options were adapted from the level descriptions given by Lundgard and Satyanarayan [30] and shortened and simplified to make them easier for annotators to understand. Annotators were given an additional instruction document that provides more detail on each of the label options, along with examples of sentences corresponding to each label option.

The annotators were instructed to exhaustively label each sentence with all of the levels it contains. We recruited two annotators through the UpWork platform.¹³ The annotators had undergraduate-level education in math, statistics, and materials science, and had

¹²We experiment with several other widely available PDF libraries and conversion tools, including PDFTOHTML (<http://pdfhtml.sourceforge.net/>), PDFMiner (<https://github.com/euske/pdfminer>), pdf2xml (<https://sourceforge.net/projects/pdf2xml/>) etc., and found that none of these tools allowed access to the embedded alt text. Scaling in Adobe Acrobat Pro was the only working solution we were able to identify.

¹³<https://www.upwork.com/>

previous experience reading graphs, plots, and other scientific figures. The alt text retained after the filtering steps described in the previous section were split into sets of 100 alt texts each for annotation. Two individuals annotated an initial sample of 100 alt texts (298 sentences) to refine the task and ensure high annotator agreement. The inter-annotator agreement computed over this sample was 87.6%, with $\alpha = 0.80$, indicating very good agreement. We further clarified the instructions following a discussion of disagreements. Given the high agreement level, a single annotator annotated the remaining alt text. For the final analysis, the first annotator’s annotations are used for the sample that was doubly annotated. Examples of extracted alt text and the corresponding annotated semantic levels are provided in Table 2.

4 RESULTS

We analyze the semantic content of the extracted alt text, and attempt to answer the following research questions:

- **RQ1: What is the distribution of semantic content in author-written alt text?**

We want to determine the proportion of alt text containing level 1, 2, 3, and 4 content. Of these, what proportion contains levels 1–3 content, which satisfies most BLV user needs? Correspondingly, which semantic levels are most often missing?

- **RQ2: How does the distribution of semantic content in alt text change over time?** We want to determine whether the presence of levels 1–3 semantic content is increasing over time, and by how much. We expect that with improvements in alt text awareness and workflows over time, that the amount of content available in alt text should correspondingly increase.

- **RQ3: How does length of alt text correlate with semantic levels?**

Mack et al. [31] find that there is tension between detail and brevity in alt text. The ideal alt text may vary based on user needs, but should balance length and completeness of semantic content. We determine the relationship between length and presence of semantic levels using our data.

4.1 Descriptive statistics

We process 25218 total paper PDFs to extract alt text. Of these, 19500 (77.3%) are successfully converted to HTML by Adobe Acrobat Pro. Only 897 (4.6%) of these converted documents contain at least one piece of valid alt text. Around 2048 pieces of alt text corresponded to file paths and 2545 alt texts did not meet our length criteria, and were filtered out; all other alt texts that were filtered out did not have any content besides “image.” After this filtering, the 897 papers contain 3386 valid author-written alt texts. Using our keyword heuristics, we determine that 1085 of these alt texts are likely to correspond to graphs or charts. Based on a cursory examination, the alt texts that were filtered out using these keyword heuristics consist primarily of natural images and diagrams.

We ask annotators to assess the semantic levels present in each sentence of the remaining 1085 alt text. Of these, 547 figure alt texts (consisting of 2127 sentences) are labeled as belonging to graphs, charts, and plots by our annotators, indicating that our keyword

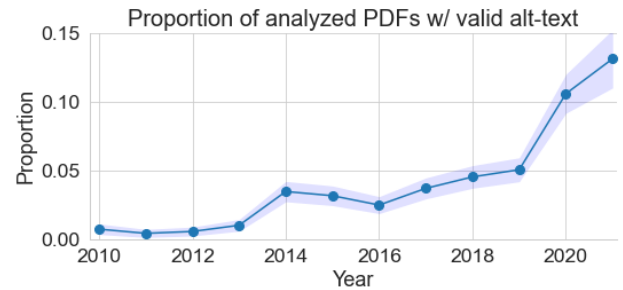


Figure 1: The proportion of PDFs in our analyzed sample which contain valid alt text over time, with 95% confidence intervals computed through bootstrap resampling.

heuristics have approximately 50% precision. Alt texts of figures corresponding to these data visualizations are further annotated for semantic content. Several examples of author-written alt text and semantic level annotations are given in Table 2. In Table 3, we provide the numbers of PDFs processed and alt text retained after each filtering step.

Figure 1 shows the proportion of PDFs in our sample that contain any valid alt text, and how this proportion changes in our sample over the last decade. There is a slight increase in alt text coverage in 2014; this is the same year that CHI specified that alt text is required in submissions. We also observe that the proportion of papers with valid alt text has improved over time, especially in the past few years, although the overall proportion is still quite low (below 15%). We note that Figure 1 does not indicate the actual proportions of papers from these years that have valid alt text (we do not process the version of record for all papers in our sample due to the difficulty in ascertaining these versions and copyright challenges in obtaining them; we also cannot guarantee that the pipeline we use succeeds in extracting alt text for all of the papers we process). Rather, the figure describes our success rate in creating this dataset, and provides some sense of the trend towards more valid alt text in recent years.

4.2 Analysis of alt text semantic content

Distribution of semantic content in alt text (RQ1). In Figure 2 (left), we present the maximum level of content found in each figure alt text in our sample. We observe a fairly even distribution between max levels 1, 2, and 3. Recall that in Lundgard and Satyanarayan [30], BLV users found levels 1–3 content to be most useful. Although some figures in our sample have alt text with level 3 content, over two-thirds of the figure alt texts that we examined do not contain level 3 content (level 3 content describes trends and patterns). At the same time, over one-third of the figure alt texts lack both level 2 and level 3 content, meaning that there is no information on extrema and outliers in addition to trends and patterns. The lack of such content can make it more difficult for BLV users to acquire the information that they need from these graphs and charts.

In Figure 2 (right), we show the total number of levels of content present in all figure alt text we analyzed. We find that the vast majority of alt text in our sample only contain one or two levels. This

Figure & Source

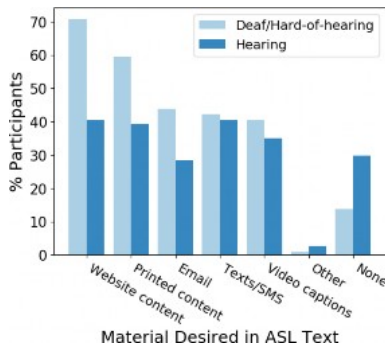


Figure 4 reproduced from Bragg et al. [8]

Author-written alt text with annotations

“Figure 4: Materials participants reported wanting to read in ASL text (L1). This figure presents a bar chart, with separate bars for DHH (light blue) and hearing (dark blue) populations (L1). Y-axis is % participants, ranging from 0-70 (L1). X-axis is Material Desired in ASL Text (L1), sorted by DHH popularity (most popular first): Website content, Printed content, Email, Texts/SMS, Video captions, Other, None (L1).”

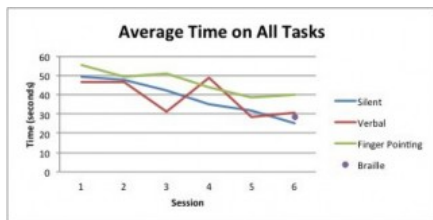
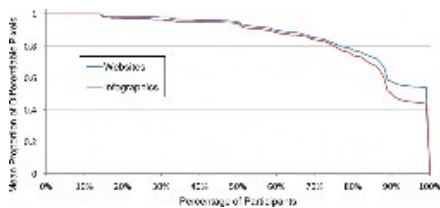


Figure 4 reproduced from Baker et al. [2]

“A line chart showing the average time it took participants on all tasks (L1). The y-axis of the chart is time in seconds (ranges from 0 to 60), the x-axis of the chart is session number (ranges from 1 to 6) (L1). There is a line for the three modes: Silent, Verbal and Finger Pointing (L1). They all appear to be going down, but there is a big spike in the Verbal mode line at session 4 (L3). In general, the Finger Pointing mode is the highest (takes the most time), the Silent mode is next and the Verbal takes the least amount of time, although in the fourth and sixth sessions, the Verbal line is above the Silent one (L3). There is a dot corresponding to the Braille mode at Session 6, it is between the Verbal and Silent modes (L1).”



(Figure 7 reproduced from Reinecke et al. [41])

“Plot of mean proportion of image pixels differentiable (independent) for 0% - 100% of the population (dependent) for websites and infographics (L1). Increasing from 0% of the population, both plots start at 100% differentiable and gradually fall to 80% differentiable at 75% of the population (L2, L3). Plots begin to diverge here as they both fall off more quickly until a discontinuity plateau is reached at 88% of the population (websites = 60% differentiable, infographics = 50% differentiable) (L2, L3). Plateau gradually declines to 99% of population (websites = 55% differentiable, infographics = 45% differentiable), and then both plots fall to 0% differentiable for 100% of the population (L2, L3).”

Table 2: Example figures and author-written alt text with annotated semantic levels added to the end of each sentence, with the prefix L, in parentheses and colored (blue).

means that though the maximum levels of content are somewhat evenly distributed between levels 1, 2, and 3, only one or two of these levels are present in most figure alt texts.

Semantic content in alt text over time (RQ2). Figure 3 shows the proportion of alt text from each year that contain text of each of the semantic levels. Though the vast majority of alt text contain level 1 information, a much lower proportion contain level 2 and 3 information. Over time, there have not been significant changes to the proportion of alt text that contains level 2 and 3 information.

Relationship between alt text length and content (RQ3). In Figure 4, we show the relationship between the length of alt text and the number of levels of content that it contains. Alt text containing more levels of information tend to be longer, though alt text of comparable length can have different numbers of levels present (as indicated by the overlapping boxes across levels). This suggests that alt text does not have to be longer in order to have more levels of content. In regards to the balance between brevity and detail, authors may want to optimize for the amount of information included in alt text without resorting to writing something overly long.

Processing step	Count
Total PDFs processed	25218
PDFs successfully converted to HTML	19500 (77.3%)
Papers with at least one valid alt text	897 (4.6%)
Pieces of valid alt text	3386
Number of figure alt text annotated (heuristically filtered to likely be graphs or charts)	1085
Number of annotated figure alt text that correspond to graphs or charts	547

Table 3: The numbers of papers and figure alt text that remain after each filtering step.

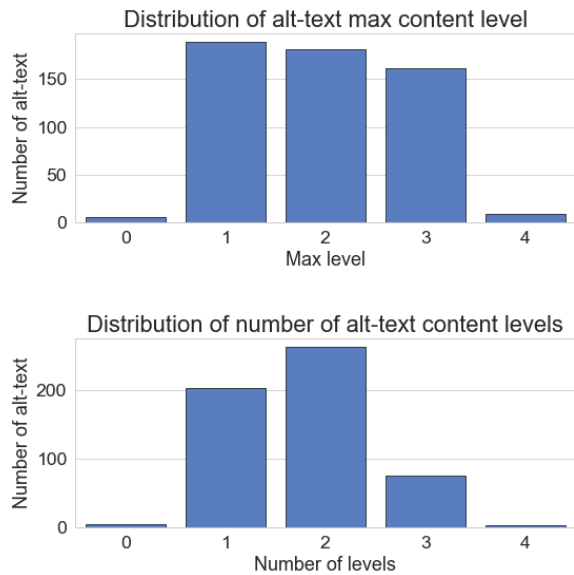


Figure 2: Distribution of the maximum level (top) and total number of levels (bottom) of semantic content found in the sample of annotated author-written alt text.

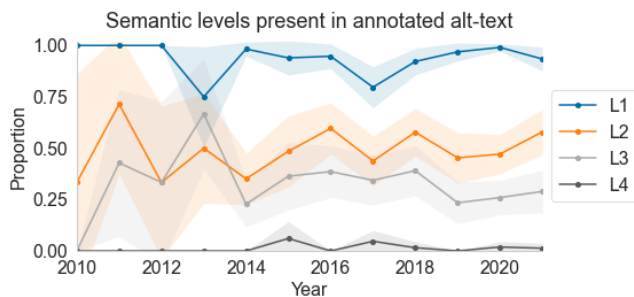


Figure 3: Proportion of alt text containing text of each semantic level over time with 95% confidence intervals computed using bootstrap resampling.

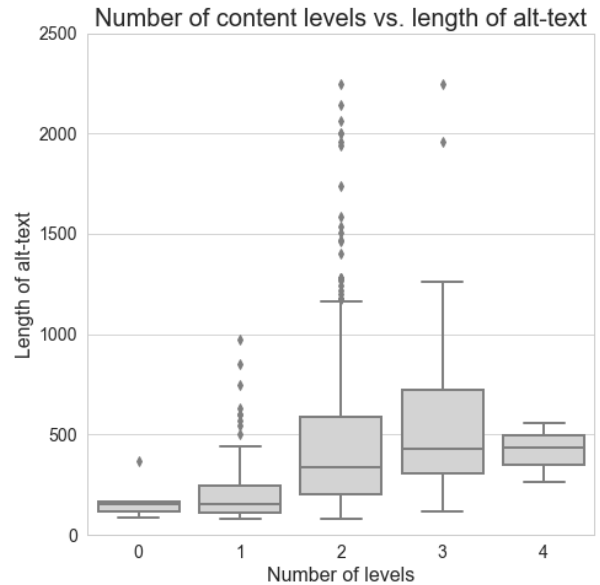


Figure 4: Relationship between length of alt text (character count) and the number of semantic levels of information represented. Though length correlates with the number of levels, there are many longer alt texts that do not necessarily contain more levels of semantic content.

5 DATASET USES

Alongside our analysis, we release our dataset of 3386 alt text collected from 897 HCI publications.¹⁴ Of these, the contents of 547 alt texts (2127 sentences) are annotated with the semantic levels introduced by [30]. This dataset can be used to develop tools to support authors in writing better alt text, or to develop authoring tools or models for producing alt text where none are available. We discuss some of these potential applications below.

5.1 Improving author-written alt text & supporting reading interfaces

To improve authoring of alt text, we can develop tools to help identify potentially missing content in alt text and prompt authors to add such information during the authoring and editing process. For graphs and charts, the semantic levels serve as a proxy for content. Based on the findings by Lundgard and Satyanarayan [30], BLV users found levels 1–3 information most beneficial, and a tool aimed to improve alt text for graphs and charts could assess alt text quality based on the presence or absence of information at these levels. In other words, we could train a classifier based on the semantic level annotations in our dataset to detect which levels are present and build a tool that prompts authors for the missing levels, similar to interfaces that have been built in the past for other tasks like providing peer feedback [16]. Such a classifier could also be used to automate trend monitoring for alt text content, enabling

¹⁴Data and analyses are available at <https://github.com/allenai/hci-alt-texts>.

Model	Accuracy	Micro-F1
Random Forest (<i>tf-idf</i>)	0.689 (0.021)	0.782 (0.011)
BERT-base	0.912 (0.016)	0.824 (0.032)
SciBERT-base	0.910 (0.010)	0.819 (0.021)

Table 4: Model performance averaged over five folds, shown with standard deviations in parentheses.

expanded and continually updating versions of our analysis with reduced human labor.

To test out the viability of this theory, we train several classification models using our collected annotations. We approach this as a multi-class, multi-label classification problem. The classes are the four semantic levels, and up to four labels can be assigned to a single piece of text. The input to the model is a single sentence of alt text, and the output is a distribution of labels over the four classes. We experiment with a Random Forest classifier using *tf-idf* word representations, as well as classifiers based on BERT [14] and SciBERT [3]. We use 5-fold cross-validation to train and evaluate all models. The training data is split into folds preserving each alt text as a unit (547 instances), while text and labels are provided to the model at the sentence level (2127 sentences). We report mean accuracy and F1 over all five folds for each model in Table 4.

Baseline performance of these models suggests that they are reasonably good at identifying the correct semantic levels present in an alt text sentence. Performance could be improved further through additional model tuning or data annotation. The outputs of these models can feasibly be used to provide feedback to authors who are writing alt text, to indicate when content of certain semantic levels may be lacking.

In addition to helping to improve author-written alt text, such a classifier could be used to support improved reading experiences of existing alt text. For example, semantic level classification could enable users to make informed decisions about whether and how to read alt text by filtering for semantic levels that may be more relevant to their needs. Morash et al. [33] described similar types of personalized reading experiences, which could be enabled for alt text that has been written using standardized templates. We leave the implementation of these authoring and reading interfaces, as well as explorations on user interface design, to future work.

5.2 Training and evaluating NLP models for alt text generation

Recent developments in multimodal image-language pretraining [28, 29, 49] hold promise towards the eventual automatability of scientific figure alt text generation. Currently, figure alt text generation is hampered by the lack of realistic training and evaluation data. Though the size of our dataset is small and insufficient for training neural models, it may still be useful to help scale the collection of training data. Alt text from this dataset can be used to provide high-quality examples to annotators. Additionally, a classifier trained to predict alt text semantic levels such as the one introduced in Section 5.1 could be used to offer feedback to annotators during the annotation process, e.g., by indicating when information of a

certain content level is missing in the annotation. Instruction specific to the missing level could be provided to the annotator to elicit further description information, as in the techniques employed by Morash et al. [33]. Several works propose a hybrid approach that combines machine learning model output with human writers to create better image descriptions [19, 31, 39]. Our dataset and classifiers could be used in collaboration with human writers to produce more descriptive alt text.

The alt texts in our dataset could also serve as part of a viable evaluation corpus. Though not all alt text in the dataset contain semantic level annotations, the texts themselves are written by the original paper authors, and are therefore more likely to be faithful to the original intent and content of the paper. We release all 3386 valid alt text extracted from our sample of papers, which includes alt text in addition to the 547 alt texts belonging to data visualizations which we annotate for semantic content. These alt texts can be used to assess pieces of information that authors found important enough to include in the image description. The output of a general-purpose scientific alt text generation model can be evaluated against the information contained in the original author-written alt text associated with these figures.

6 DISCUSSION & CONCLUSION

In regards to scientific alt text, availability is still the primary issue. However, in the alt text we were able to extract from HCI publications, we observed that for papers where authors have taken the time to write alt text, the content and level of detail available in these alt text is also worth considering. What does it mean to write useful alt text? What does it mean to include enough detail such that the content of an image can be understood by a BLV user? For graphs and charts, we propose that authors leverage the framework introduced in Lundgard and Satyanarayan [30] to ensure that some basic semantic information is provided, enough such that BLV users can understand the structure of the graph, its extrema and outliers, as well as the obvious trends and comparisons that can be made. In our current analysis, we find that many author-written alt text are not yet meeting these thresholds.

We recognize the limitations of our techniques. The alt text and figures included in our dataset make up a biased sample, containing only papers from CHI and ASSETS and from the authors publishing in these venues. They are not representative of figures in all scholarly documents. Though we would have liked to construct a more representative sample of scientific figures, the overwhelming lack of figure alt text in scholarly publications prevents us from doing so. As more authors from other disciplines begin including alt text and the barriers to adding alt text to scientific figures decreases, we hope that it will become easier to create such a dataset.

Additionally, our analysis and annotations are limited to figures containing data visualizations (graphs, charts, and plots), which are only one of many types of images present in scientific publications. Our results on the suitability or missingness of semantic content in author-written alt text cannot generalize beyond these image types. Going beyond graphs, frameworks other than that introduced in Lundgard and Satyanarayan [30] may be needed to capture the availability and distribution of informational content. We emphasize that the lack of certain semantic levels in alt text is

not equivalent to an assessment about the alt text's quality. Different users may want different information out of alt text, and there is no one-size-fits-all solution [31]. Rather, we use the framework as a proxy for content availability, which can be used to elicit different kinds of descriptive information that may be missing in an author's original alt text. Finally, in Lundgard and Satyanarayan [30], the authors assessed whether the levels were useful for BLV users, but not whether all of levels 1–3 were necessary for an alt text to be considered complete. Further study is needed to determine the appropriate balance of depth of information, completeness, and brevity in relation to the usefulness of alt text.

We propose in Section 5 several uses of this dataset towards improving author and publisher workflows around writing figure alt text. A classifier trained to detect semantic levels can be used to provide feedback to authors who are writing alt text, and can be used to elicit alt text containing more levels of information. The dataset can also be used to help develop machine learning models that can generate better alt text based on the image itself. We believe that coupling machine learning models with crowdsourced image descriptions may provide a reasonable solution to problems around alt text availability, and we plan to explore such solutions in future work.

Policy clearly matters. We faced significant challenges when attempting to extract alt text from a broad swathe of scientific publications, and had to limit ourselves ultimately to HCI publishing venues such as ASSETS and CHI. There is no doubt in our minds that the alt text we were able to extract are only there because of the efforts of the accessibility and HCI research community and the importance that members have placed on digital accessibility. Significant work remains to encourage researchers outside of these communities to participate in making their work accessible. Within the community, there are also ways that we can improve figure accessibility, by providing information that are described by BLV users as being more relevant or more important towards interpreting these images.

ACKNOWLEDGMENTS

We thank Donal Fitzpatrick, Doug Downey, the reviewers, and members of the Semantic Scholar research team for their valuable feedback. We thank Bailey Kuehl and Ihsan Allah Rakha for their assistance in the annotation process.

REFERENCES

- [1] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu A. Ha, Rodney Michael Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler C. Murray, Hsu-Han Ooi, Matthew E. Peters, Joanna L. Power, Sam Skjonsberg, Lucy Lu Wang, Christopher Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. 2018. Construction of the Literature Graph in Semantic Scholar. In *NAACL*.
- [2] Catherine M. Baker, Lauren R. Milne, Jeffrey Scofield, Cynthia L. Bennett, and Richard E. Ladner. 2014. Tactile graphics with a voice: using QR codes to access text in tactile graphics. *Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility* (2014).
- [3] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. In *EMNLP*.
- [4] Cynthia L. Bennett, Cole Gleason, Morgan Klaus Scheuerman, Jeffrey P. Bigham, Anhong Guo, and Alexandra To. 2021. "It's Complicated": Negotiating Accessibility and (Mis)Representation in Image Descriptions of Race, Gender, and Disability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (*CHI '21*). Association for Computing Machinery, New York, NY, USA, Article 375, 19 pages. <https://doi.org/10.1145/3411764.3445498>
- [5] Jeffrey P. Bigham, E. Brady, Cole Gleason, Anhong Guo, and D. Shamma. 2016. An Uninteresting Tour Through Why Our Research Papers Aren't Accessible. *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (2016).
- [6] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Rob Miller, Rob Miller, Aubrey Tatarowicz, Brandyn Allen White, Samuel White, and Tom Yeh. 2010. VizWiz: nearly real-time answers to visual questions. *Proceedings of the 23rd annual ACM symposium on User interface software and technology* (2010).
- [7] E. Brady, Y. Zhong, and Jeffrey P. Bigham. 2015. Creating accessible PDFs for conference proceedings. *Proceedings of the 12th Web for All Conference* (2015).
- [8] Danielle Bragg, Raja S. Kushalnagar, and Richard E. Ladner. 2018. Designing an Animated Character System for American Sign Language. *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (2018).
- [9] B. Caldwell, M. Cooper, Loretta Guarino Reid, and G. Vanderheiden. 2008. Web Content Accessibility Guidelines (WCAG) 2.0.
- [10] Chen Chen, Ruiyi Zhang, Sungchul Kim, Scott D. Cohen, Tong Yu, Ryan A. Rossi, and Razvan C. Bunescu. 2019. Neural caption generation over figures. *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers* (2019).
- [11] Charles C. Chen, Ruiyi Zhang, Eunye Koh, Sungchul Kim, Scott D. Cohen, Tong Yu, Razvan C. Bunescu, and Razvan C. Bunescu. 2019. Figure Captioning with Reasoning and Sequence-Level Training. *ArXiv abs/1906.02850* (2019).
- [12] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. 2015. Microsoft COCO Captions: Data Collection and Evaluation Server. *ArXiv abs/1504.00325* (2015).
- [13] W. Chisholm, G. Vanderheiden, and Ian Jacobs. 2001. Web content accessibility guidelines 1.0. *Interactions* 8 (2001), 35–54.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
- [15] Christin Engel and Gerhard Weber. 2017. Improve the Accessibility of Tactile Charts. In *INTERACT*.
- [16] C. Ailie Fraser, Tricia J. Ngoon, Ariel S. Weingarten, Mira Dontcheva, and Scott Klemmer. 2017. CritiqueKit: A Mixed-Initiative, Real-Time Interface For Improving Feedback. In *Adjunct Publication of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) (*UIST '17*). Association for Computing Machinery, New York, NY, USA, 7–9. <https://doi.org/10.1145/3131785.3131791>
- [17] 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* (2020).
- [18] 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* (2018).
- [19] Danna Gurari, Qing Li, Abigale Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. VizWiz Grand Challenge: Answering Visual Questions from Blind People. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), 3608–3617.
- [20] Ting-Yao Hsu, C. Lee Giles, and Ting-Hao Kenneth Huang. 2021. SciCap: Generating Captions for Scientific Figures. In *EMNLP*.
- [21] K. V. Jobin, Ajoy Mondal, and C. V. Jawahar. 2019. DocFigure: A Dataset for Scientific Document Figure Classification. *2019 International Conference on Document Analysis and Recognition Workshops (ICDARW) 1* (2019), 74–79.
- [22] Daekyoung Jung, Wonjae Kim, Hyunjoon Song, Jeongin Hwang, Bongshin Lee, Bo Hyoung Kim, and Jinwook Seo. 2017. ChartSense: Interactive Data Extraction from Chart Images. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017).
- [23] Kushal Kafle, Scott D. Cohen, Brian L. Price, and Christopher Kanan. 2018. DVQA: Understanding Data Visualizations via Question Answering. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), 5648–5656.
- [24] Samira Ebrahimi Kahou, Adam Atkinson, Vincent Michalski, Ákos Kádár, Adam Trischler, and Yoshua Bengio. 2018. FigureQA: An Annotated Figure Dataset for Visual Reasoning. *ArXiv abs/1710.07300* (2018).
- [25] Alina Kuznetsova, Hassan Rom, Neil Gordon Alldrin, Jasper R. R. Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. 2020. The Open Images Dataset V4. *International Journal of Computer Vision* 128 (2020), 1956–1981.
- [26] J. Lazar, E. Churchill, T. Grossman, G. V. D. Veer, Philippe A. Palanque, J. Morris, and Jennifer Mankoff. 2017. Making the field of computing more inclusive. *Commun. ACM* 60 (2017), 50–59.
- [27] Matan Levy, Rami Ben-Ari, and Dani Lischinski. 2021. Classification-Regression for Chart Comprehension.

- [28] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. VisualBERT: A Simple and Performant Baseline for Vision and Language. *ArXiv abs/1908.03557* (2019).
- [29] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *NeurIPS*.
- [30] Alan Lundgard and Arvind Satyanarayan. 2022. Accessible Visualization via Natural Language Descriptions: A Four-Level Model of Semantic Content. *IEEE Transactions on Visualization and Computer Graphics* 28 (2022), 1073–1083.
- [31] Kelly M. 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* (2021).
- [32] Nitesh Methani, Pritha Ganguly, Mitesh M. Khapra, and Pratyush Kumar. 2020. PlotQA: Reasoning over Scientific Plots. *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)* (2020), 1516–1525.
- [33] 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 Transactions on Accessible Computing (TACCESS)* 7 (2015), 1–21.
- [34] David Morris, Eric Müller-Budack, and Ralph Ewerth. 2020. SlidImages: A Dataset for Educational Image Classification. *Advances in Information Retrieval* 12036 (2020), 289 – 296.
- [35] 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* (2018).
- [36] Azadeh Nazemi and Iain Murray. 2013. A Method to Provide Accessibility for Visual Components to Vision Impaired.
- [37] Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing. *ArXiv abs/1902.07669* (2019).
- [38] Aaron Nicolson, Jason Dowling, and Bevan Koopman. 2021. AEHRC CSIRO at ImageCLEFmed Caption 2021. In *CLEF*.
- [39] Xin Qian, Eunyee Koh, Fan Du, Sungchul Kim, and Joel Chan. 2020. A Formative Study on Designing Accurate and Natural Figure Captioning Systems. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3382946>
- [40] Xin Qian, Eunyee Koh, Fan Du, Sungchul Kim, Joel Chan, Ryan A. Rossi, Sana Malik, and Tak Yeon Lee. 2021. Generating Accurate Caption Units for Figure Captioning. *Proceedings of the Web Conference 2021* (2021).
- [41] Katharina Reinecke, David R. Flatla, and Christopher Brooks. 2016. Enabling Designers to Foresee Which Colors Users Cannot See. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016).
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115 (2015), 211–252.
- [43] Elliot Salisbury, Ece Kamar, and Meredith Ringel Morris. 2017. Toward Scalable Social Alt Text: Conversational Crowdsourcing as a Tool for Refining Vision-to-Language Technology for the Blind. In *HCOMP*.
- [44] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *ACL*.
- [45] Noah Siegel, Zachary Horvitz, Roie Levin, Santosh Kumar Divvala, and Ali Farhadi. 2016. FigureSeer: Parsing Result-Figures in Research Papers. In *ECCV*.
- [46] 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>
- [47] Lucy Lu Wang, Isabel Cachola, Jonathan Bragg, Evie (Yu-Yen) Cheng, Chelsea Hess Haupt, Matt Latzke, Bailey Kuehl, Madeleine van Zuylen, Linda M. Wagner, and Daniel S. Weld. 2021. Improving the Accessibility of Scientific Documents: Current State, User Needs, and a System Solution to Enhance Scientific PDF Accessibility for Blind and Low Vision Users. *ArXiv abs/2105.00076* (2021).
- [48] 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. *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (2017).
- [49] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J. Corso, and Jianfeng Gao. 2020. Unified Vision-Language Pre-Training for Image Captioning and VQA. *ArXiv abs/1909.11059* (2020).

A HEURISTICS FOR IDENTIFYING GRAPHS AND CHARTS

The full list of words and phrases used to identify figures corresponding to graphs and charts include:

- graph
- chart
- plot
- scatter plot
- scatter
- distribution
- data
- points
- error
- error bar
- trial
- trials
- bar plot
- bar
- venn
- mean
- average