UNIVERSITY OF CALIFORNIA,
IRVINE

The Who and How of Image Accessibility in Practice: Studies of Alt Text Production Processes and Roles in Differing Contexts

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Informatics

by

Emory James Edwards

Dissertation Committee:
Associate Professor Stacy Branham, Chair
Associate Professor Anne Marie Piper
Professor Katie Salen Tekinbaş

2023

DEDICATION

To

My colleagues, friends, and mentors:
Who have helped more than I can say.

For

My mother:
My favorite person on earth,
For whom bringing me into the world was the least of it
Because without everything she's done for me since, I would do and be nothing.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to thank my advisor, Dr. Stacy Branham, for going above and beyond in mentoring me over the past four years. But especially for teaching me that in this digital and COVID era we all have more balls in the air than we can possibly juggle. Everyone is constantly deciding which ones are glass and which ones are rubber. The nice thing about the rubber ones is that they bounce back.

I would like to thank my committee members, Dr. Katie Salen and Dr. Anne Marie Piper, whose genuine encouragement and thoughtful consideration made the terrifying process of facing down a dissertation that much easier.

In addition, a thank you to Dr. Tom Boellstorff, whose love for Anthropology is catching. You made theory an exciting destination, even if it was not my final home.

Thank you to the ARCS Foundation, UCI's Graduate Division, and the Informatics Department, for the fellowships and scholarships that made this dissertation possible.

And an extra thank you to my therapist, Tracey, who has accompanied me down so many paths and led me up so many mountains. May every grad student, every queer or trans person, every neurodivergent dude, have as wonderful a therapist as you.

The text of this dissertation includes some reprinted material that originally appeared in ASSETS '21: Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility and ACM Transactions on Accessible Computing. The co-authors listed in those publications are Emily Blank, Michael Gilbert, Kyle Lewis Polster, Isabel Tuason, and Stacy Branham. Stacy Branham directed and supervised the research which forms the basis for the dissertation.

VITA

Emory James Edwards

2015-17	Undergraduate Researcher EmCOMP (Emerging Capacities of Mass Participation) and SoMe (Social Media) Lab University of Washington
2016	Undergraduate Summer Researcher CHIMPS (Computer Human Interaction: Mobility, Privacy, Security) Carnegie Mellon University
2017	B.A. in English (Cum Laude with Honors), University of Washington B.S. in Informatics (Magna Cum Laude), University of Washington
2018	Teaching Assistant, Department of Informatics University of California, Irvine
2018-19	Graduate Researcher and Founding Lab Member CATS (Critical Approaches to Technology and the Social) Lab University of California, Irvine
2019	Teaching Assistant, School of Information and Computer Science University of California, Irvine
2019-23	Graduate Research Fellow INsite (Inclusive Studio for Innovative Technology and Education) Lab Department of Informatics, University of California, Irvine
2020	Virtual Teaching Assistant, School of Information and Computer Science University of California, Irvine
2020	M.S. in Informatics, University of California, Irvine
2021-22	Intel-funded Research Fellow Steckler Center for Responsible Ethical and Accessible Technology (CREATE) University of California, Irvine
2022	Research Intern, Ability Group, Microsoft Corporation
2022	Accessible Computing UX Research Intern Research and Experience Definition (RED) Team, Client Computing Group Intel Corporation
2023	Ph.D. in Informatics, University of California, Irvine

FIELD OF STUDY

Accessible Computing and Human-Computer Interaction

PUBLICATIONS

"Could this be true? I think so! Expressed uncertainty in online rumoring." Kate Starbird, Emma Spiro, Isabelle Edwards, Kaitlyn Zhou, Jim Maddock, and Sindhuja Narasimhan. May 2016. Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.

"Centralized, parallel, and distributed information processing during collective sensemaking." Peter Krafft, Kaitlyn Zhou, Isabelle Edwards, Kate Starbird, and Emma S. Spiro. May 2017. Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems.

"Three Tensions Between Personas and Complex Disability Identities." Emory James Edwards, Cella Monet Sum, and Stacy M. Branham. May 2020. Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems.

"Towards More Universal Wayfinding Technologies: Navigation Preferences Across Disabilities." Maya Gupta, Ali Abdolrahmani, Emory James Edwards, Mayra Cortez, Andrew Tumang, Yasmin Majali, Marc Lazaga, Samhitha Tarra, Prasad Patil, and Ravi Kuber. May 2020. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.

"Migration, non-use, and the 'Tumblrpocalypse': Towards a unified theory of digital exodus." Emory James Edwards and Tom Boellstorff. Apr 2021. Media, Culture & Society.

"'That's in the eye of the beholder': Layers of Interpretation in Image Descriptions for Fictional Representations of People with Disabilities." Emory James Edwards, Kyle Lewis Polster, Isabel Tuason, Michael Gilbert, Emily Blank, and Stacy M. Branham. Oct 2021. Proceedings of The 23rd International ACM SIGACCESS Conference on Computers and Accessibility, ACM, Virtual Event.

"Putting the disability in DEI through inclusive imagery." Emory James Edwards. Jul 2022. XRDS: Crossroads, The ACM Magazine for Students.

"How the Alt Text Gets Made: What Roles and Processes of Alt Text Creation Can Teach Us About Inclusive Imagery." Emory James Edwards, Michael Gilbert, Emily Blank, Stacy M. Branham. Mar 2023. ACM Transactions of Accessible Computing (TACCESS).

ABSTRACT OF THE DISSERTATION

The Who and How of Image Accessibility in Practice:
Studies of Alt Text Production Processes and Roles in Differing Contexts
by

Emory James Edwards

Doctor of Philosophy in Informatics

University of California, Irvine, 2023

Associate Professor Stacy Branham, Chair

Alt text, short for alternative text, is text associated with a digital image that describes its visual content for screen reader users. The lack of alt text on many major websites, particularly those with high levels of user generated content, has been an ongoing area of study within the accessible computing space. However, most studies on alt text have surveyed current coverage, defined appropriate alt text, or introduced an intervention to deliver more or better alt text to users. This dissertation instead takes the view that understanding existing practices by which alt text is made can provide a better basis from which to develop our understandings and solutions. Therefore, I present three studies of alt text creation in practice, with a focus on inclusive and accessible imagery. First, I present a qualitative interview study in which I co-designed inclusive alt text with users with disabilities. This study led to a description of the different roles and layers of interpretation at play in the creation and reception of imagery and alt text. Second, I describe a set of interviews with artists and accessibility practitioners at Google investigating how imagery and alt text is created in a technology company. This study provided the first known models of alt text creation processes and described the advantages and disadvantages of

different approaches. And lastly, I document a digital ethnography of an online accessibility community that collaboratively supports image accessibility on social media through volunteer-based crowdsourcing. This ethnography found that community-based accessibility is dedicated to more than just alt text, but to the creation of broadly applicable accessibility practices and a shift in social media platform culture. From these studies I develop recommendations for future accessibility research, implications for future social media design, and advice for future practice of accessibility by professionals and activists.

CHAPTER 1: Introduction

There are billions of new images entering the digital sphere every day [80]. They are the clipart, stock photos or interface mockups your coworker includes in their PowerPoints, and the drawings of characters or family vacations from your friends, and the infographics of Covid-19 cases, or screenshots of tweets you see on social media. They contain everything from urgent news and data for decision making, to mundane social and interpersonal information about friends, family, and colleagues that keeps our inboxes and social feeds full every day. Images are used online to convey vast amounts of information, emotion, creativity, personality, and more.

1.1 Motivation

However, for the 2.2 billion people worldwide who are blind or have low vision (BLV) [267], the ubiquity of digital images is a disabling force. Blind computer users generally "see" digital images by using screen reading software, a synthesized voice which reads aloud the text elements on a digital screen or interface. When it comes to images, screen readers are only helpful if the images have alt text (alternative text) or image descriptions associated with them. Alt text and image descriptions are written descriptions of visual content that are either displayed near the image or accompany it as descriptive metadata.

Alt text and image descriptions don't just appear whenever an image comes into being. They have to be created, either manually by human authors or automatically through algorithmic and computer vision-related means. The huge number of digital images therefore corresponds to an equally huge accessibility problem. Through various studies

over the past decade of accessible computing research, alt text has been estimated to accompany between 83% [100] and less than 1% [95] of images, depending on which sites are being assessed. If you consider digital images on the web, popular sites, including government websites, often have better coverage, whereas sites like Twitter with large amounts of user-generated content have alt text for as little as 0.1% of their hosted images [95]. As more users join social networking sites [264] and as some of those sites become more visual over time [167], the need surges for mass methods for creating alt text.

Even understanding what makes good alt text, when it is supplied, is a complicated issue. Human generated alt text has historically been the default solution for digital content, and it can be completed in many different ways, from reusing metadata gathered from other sources [100] to social media-based microvolunteering [31]. It's assumed that human generated alt text is going to be high quality, although actually defining what makes for good alt text is complicated, with different contexts, different users, and different personal preferences impacting what should be included in alt text [18,166,227]. It requires time investment to teach people how to write helpful and thorough descriptions [165], and with more content creators posting more images all the time, the amount of work required to describe, or teach people to describe, them all is enormous.

This has led researchers and companies to pursue automatic approaches to alt text creation,[1] which unfortunately often result in lower quality image descriptions [153]. For one example, the images in Figure 1.1 come from a Buzzfeed tweet, accompanied by the text "19 of the most disrespectful sandwiches ever made." The computer-generated alt text

[1] This dissertation was conceived and almost entirely written before the popularization of ChatGPT. Thus, when discussing automatic approaches to alt text generation in this document, I am not referring to methods using ChatGPT or similar generative AI.

is only 52% sure it is accurate in describing it as “a close up of a sandwich.” On the other hand, a human-generated version of the alt text reads: “The image is two photos of open-faced sandwiches. The first is an English muffin with a tiny piece of ham on one half and a tiny fried egg on the other half. The second is two slices of bread and one tiny piece of cheese on one half” [153]. The second description unquestionably captures better the reason why these images were likely chosen to accompany this tweet, even though the automated alt text was (at least in this case) technically correct. Large companies can afford to pay employees to make sure their content is both accessible and of a higher quality, but the enormity of the backlog of undescribed images online makes accessibility difficult to pursue unless there is a large amount of resources and attention devoted to the problem.

Figure 1.1: Examples of photos given low-quality alt text

Crowdsourcing has the potential to be a powerful tool to combat image inaccessibility [200] but the balance between the valuable accessibility gains and the labor

issues associated with crowd platforms [119] have yet to be studied. Platforms like Amazon Mechanical Turk (AMT) rely on underpaying workers [240], while volunteer-crowdsourcing projects like Wikipedia rely entirely on the free labor of already-burdened workers. Additionally, crowd workers are disproportionately disabled [245]. People with disabilities are historically tasked with more of the invisible labor necessary to make accessibility legible and easy for non-disabled people around them [37].

An additional quality issue that comes up with automatically generated alt text, along with some human-generated methods, is the description of human subjects in the images. Even if computer vision technology can identify a human face, it takes another level of consideration to assess or describe social identities such as race, gender, and disability status. Automated means for categorizing image subjects' identities have been criticized by a range of scholars [103,105,125,210] with some particularly highlighting the importance of avoiding assumptions about identity based on detected visual appearance in image descriptions [105]. Describing social identities is much more complicated than an automatic label can capture. In fact, automated and human generated alt text has to consider how to inclusively write alt text about unknown subjects. For example, do you assume gender or pronouns of a subject? What about race? What are the possible social consequences of making the wrong guess or using the wrong term, both for the subject, the screen reader user, and the platform hosting the image?

No method has proven perfect for addressing the problem of images lacking alt text. However, considering the research as a whole, there is an additional oversight in what topics it covers that guides my own investigations. Overwhelmingly, research on alt text focuses on interventions: creating, proposing, or testing new potential solutions to the

issue of image accessibility. Or, alternately, they focus on understanding or producing rules and guidelines for what constitutes high quality alt text, based primarily on feedback from screen reader users themselves (e.g.,[18,153,182,185,227]). What is only beginning to be researched (e.g., [255]) are the current methods and processes used to create alt text. This flips the ideal research process on its head; understanding the current ways users and content creators on the web create alt text should, logically, come before the interventions, so that the interventions themselves can be tailored to how the current practices work and how to improve them. This dissertation addresses that oversight by presenting three studies alt text production in practice and contributing guidance to academics, practitioners, and activists on how to face image accessibility as an ongoing modern problem. Taken together, the three studies making up this research are intended to serve a dual purpose. The overarching motivation is the need for research that explores different methods by which alt text is currently made. But the studies also seek to explore how, if at all, inclusive image descriptions can be made using those methods.

1.2 Overview of Studies

As stated, this document is based on three separate but interrelated research studies. I have formatted the dissertation document so that Chapter 2 covers Study 1, Chapter 3 covers Study 3, and Chapter 4 covers Study 3. In each chapter I describe the background and related work, methods used, findings, and takeaways from the studies. Then, Chapter 5 discusses potential new avenues of work suggested by this research. In this introduction, however, I want to briefly introduce the three studies that make up

Chapters 2-4, so that you can have a basic understanding of the direction for the dissertation as a whole when reading the early chapters of the document.

First off, Study 1, originally described in a conference paper for ASSETS 2021 [72], considers a set of inclusive images produced by Google. Through a collaboration with Google, Study 1 sought to assess the quality of the representations, including representations of disability, with actual disabled users. However, to do so, the question arose of how to properly create alt text for the images that reflected the inclusive nature of the visual depictions. Thus, Study 1 became an exploration of how to co-design inclusive alt text with users through focus groups and interviews, resulting in general recommendations for how to write alt text for similar images in an inclusive way. Study 1 importantly found that users understood different aspects of the image description to be attributable to different roles in the creation process. The initial roles identified were image artists, alt text writers, and alt text readers, all of which contributed in some way to the final inclusive understanding of the image.

The second study, published as a journal article in TACCESS [71], continued the collaboration with Google. In Study 2 I investigated different alt text and inclusive imagery creation processes used at Google. Through interviews with artists and accessibility practitioners who worked with Google, I identified and described four processes used to create imagery at or with Google, including the co-design process I had used in Study 1 as a Google collaborator. The study also revealed that alt text creation was a burden often shouldered by a few isolated stakeholders, rather than integrated into the organizational processes and supported at all levels of the product development cycle. Additionally, as one of the Study 2 participants brought up, despite being situated in a strong position in the

industry, the company only ultimately had control or purview for creating alt text for their own content, not for the vast amount of images that passed through their platform but were created or uploaded by users.

This led me to Study 3, which documents a digital ethnography of a volunteer-based online community dedicated to collaboratively crowdsourcing image descriptions for the purpose of improving accessibility on a well-known social media platform. The community, operating on Tumblr but largely organized on Discord, provides a case study of how alt text might be made collectively by and for a community. Based on observations, interviews with community members, and quantitative analysis of trends in interactions, I describe the way crowdsourced image accessibility works in this community, the motivations, advantages, and drawbacks of participating, and the ways different members of the community interact and encourage accessibility to be a shared responsibility.

1.3 Contributions

This dissertation concludes with a discussion of the implications and value of studying image accessibility in practice. In particular, I contribute two implications for future design interfaces, two key findings for accessibility researchers going forward, and one cross-cutting implication for future work. First, I argue these studies suggest more adaptable and sophisticated alt text interfaces must be designed for social media sites and other image-hosting platforms. Alongside interfaces for viewing alt text, my work makes clear that there need to be more tools for content creators designed to help with creating and testing alt text with users. Improving accessibility and inclusion as it relates to images improves the overall user experience of the website. For researchers, I argue more work

must be done to study and base future interventions on the many current processes and contexts in which alt text is being created. Secondly, there must be more research into alt text as a theoretical construct and as a type of digital content that can be read and analyzed like any artifact. Lastly, I suggest there should be more interfaces, digital protests, and artistic provocations designed to make inaccessibility visible to the non-disabled public. This directed visibility, which I term obtrusive design, could meaningfully raise awareness and prompt culture shift so that the labor of making spaces accessible would no longer fall disproportionately on users with disabilities.

In addition to presenting the findings and implications for researchers, in Appendix A I include three pragmatic contributions designed to be of interest to different types of audiences that my work is relevant to. Based on the data collected in Study 1 and Study 3, I provide guidelines for inclusive alt text creation based on type of content, subject, and amount of background information provided. I present general guidelines for alt text creation, before inviting readers to follow my guide in deciding what course of action to take or type of information to include in the alt text for a given image. This first contribution helps accessibility-based peer production communities and accessibility advocates within companies streamline their production processes and eliminate unnecessary cognitive burden. The second included contribution is useful to both accessibility researchers and corporate audiences. Based on findings from Study 2 and Study 3, I provide a list of six different alt text production processes, the roles involved in each, and their relative pros and cons. Lastly, for accessibility researchers and activists, I include an "owner's manual" describing how to set up and run a functioning volunteer-based accessibility crowdsourcing community. This final contribution draws from Study 3

and provides grounded advice for potential future leaders of crowdsourcing efforts based on the case study of the accessibility community I observed. Taken together, I hope these contributions mean that the work I have done will reach beyond the limitations of academic relevance and help induce the creation of more, and more inclusive, alt text in society as a whole.

1.4 A Note on Theoretical Paradigm and Methodology

While I discuss several methodologies used in accessible computing literature in the following chapter and describe several methods that I put into practice throughout Chapters 2-4, I want to briefly note here the overall paradigm I utilized in this project. I conducted these studies from a constructivist approach, in which I did not take for granted that there was an objective truth which was available to study, even in a limited sense. I take the three studies of this dissertation as three qualitative case studies, from which understanding is derived from exploring the topics within a limited scope and context, rather than trying to generalize or claim objectivity in my own observations [13]. I approach this research with the assumption that what I am studying is socially constructed, where the case being studied is itself influenced and produced both by the actors and participants in it, by my observation of it, and by the overall social context of both. My findings are intended to be read always as contingent on and shaped by the context of their individual case. I draw conclusions at the end of this dissertation from all three studies with the assumption that there are disjoints and gaps, that after understanding processes in one context the phenomena can be described but not precisely replicated to another.

I balance my assumptions around contingent and constructed knowledge against an earnest belief in pragmatism. That is why I focus on paradigm here but do not linger on it in

other areas of the dissertation. Computing researchers, disability activists, and technology companies are driven, in my experience, more by pragmatism and findings than they are by theoretical navel-gazing. The constructedness of all reality does not render unimportant the experienced outcomes of ableism, labor constraints, or social and ethical pressures to treat one another fairly. Thus, I focus my writing primarily on these outcomes and on the contributions described above but acknowledge here that the understanding I have of these case studies are unavoidably, intentionally limited. If we refuse to base action on incomplete understandings, then I cannot imagine we will ever be able to act at all.

CHAPTER 2: Study 1 – User Perspectives on Elements of Inclusive Alt Text

This chapter will describe a qualitative interview study I conducted with 25 users with disabilities in which we discussed their perspectives on important elements for inclusive image descriptions. The findings were originally published and presented at ASSETS 2021 [72] and some of the following passages borrow from that original publication.

2.1 Background: The Avatar Project

The Avatar Project, which formed the basis for my first study, began as an imagery project spearheaded by Emily Blank, Art Director for Google's Material Design team. Blank, along with Google Staff Researcher, Michael Gilbert,[2] had held a series of workshops at Google prior to our collaboration starting early 2020. These workshops, involving researchers and designers at Google, had been concerned with how to properly represent a more diverse set of users, including users with disabilities, in various design and marketing assets. The search for inclusive imagery to represent users had begun with using photographs of Googlers who had given permission for their image to be used in such documents. However, concerns arose around branding and ethical implications of having photographs of real people used in contexts beyond their direct control. Thus, when I began working with Blank and Gilbert in 2020, they had already conceived of and commissioned an initial set of three artists to create the Avatar Project images.

The intention of the Avatar Project images was to function as reusable brand-compliant components that could be used by all designers (including blind designers and

[2] Both Blank and Gilbert were co-authors and collaborators on the initial publications from Study 1 and Study 2.

developers who often lack accessible tools [232]) to represent users in company contexts. They were to be included in the company's design system so they could be used by designers and developers for internal documents such as mockups and personas. They would also be viewed by users in user testing environments and external advertisements. The intention was that by commissioning artists to create fictionalized user avatars, there could be a diverse set of potential users represented, prompting designers and engineers to "think more critically about inclusivity, diversity and representation" [149] without requiring real people to give away control of their own likeness. Thus, the three original artists, plus a fourth artist who worked on the project after our collaboration had begun, were commissioned to create a diverse set of illustrations of fictional users with the size, style, and content of the different artists' contributed sets varying.

My collaboration with Blank and Gilbert began because they were interested in understanding how real users responded to the Avatar Project illustrations. Since the Avatar Project images included people with disabilities, and my lab specializes in accessibility, we were consulted to see if we could find out how people with disabilities responded to the representations of disabilities in the Avatar Project set. Study 1 involved a subset of images from one of the Avatar Project artists (see **Error! Not a valid bookmark self-reference.**) that depicted visible disabilities as well as assistive aids. That artist's collection as a whole included representations of fictionalized people of many ages, races, genders, ethnicities, professions, sexualities, disabilities, and other characteristics. Together the images, as well as the process undertaken by the artists, Googlers, and collaborating researchers, to make them formed the basis for my study of *inclusive imagery*, a key term for the following chapters.

Table 2.1: Images Included in Study 1 Feedback Sessions

A. 	E. 	I. 
B. 	F. 	J. 
C. 	G. 	
D. 	H. 	

My working definition of *inclusive imagery* applies firstly to digital images that are accessible (via alt text or image descriptions), but which, secondly, depict marginalized

identity, and, thirdly, do so in a way that is non-stigmatizing or -stereotyped. The potentially surprising detail here, if you are thinking of inclusive imagery from a media studies or representation standpoint, is that under my definition it is not inclusive imagery if it is not accessible to blind or low vision (BLV) users who could benefit from a written description of visual details. I consider accessibility of the images as a defining characteristic of whether they count as inclusive imagery because inherently inclusivity, the process of including everyone and not leaving any groups of people out of the decision-making or feedback process, means that people with disabilities should be included as an audience of the images.

2.2 Related Work

Study 1, and indeed all the studies making up this dissertation, fit within the disciplinary umbrella of accessible computing (AC) under the wider field of Human-Computer Interaction (HCI). Methodologically, accessible computing is broad, pulling from both design approaches that are not specific to the AC space (such as user-centered design [96], participatory design [211], and value sensitive design [87]) as well as several specific approaches within and arising from studies of accessibility. The latter category includes: universal usability [219] which seeks to promote designs that work for the broadest range of people, ideally all people; user sensitive inclusive design [172] which reacts to universal design by arguing for attention to be paid to specific individuals with complex and shifting needs; ability-based design [258] which seeks to undermine the stigmatizing or othering framework that puts disability (i.e. the lack of normative ability) at the center of AC and instead reframes technology design as being about abilities in all their many forms. Further approaches include design for social accessibility [216], which argues that designers must

not just make the most usable technology but also the technology that fits disabled users' desire for social acceptability and avoiding stigma, and design for user empowerment [136], which is based on the premise that research and design should actively include and foster users with disabilities as co-creators of technology.

One way of organizing existing AC literature might be to categorize based on type of intervention or method. Common types of research methods include interviews (e.g., [243]), surveys and questionnaires (e.g., [155]), focus groups (e.g., [5]), lab studies (e.g., [176]), observational studies (e.g., [218]), user tests and evaluations (e.g., [148]), diary studies (e.g., [140]), and design probes (e.g., [226]). These studies produce types of research products that include gestural based innovations (e.g., [145]), more accessible touch screens (e.g., [123]), systems with sound/voice input or output (e.g., [9]), haptic feedback and/or portable devices (e.g., [114,224]), games and playful interventions (e.g., [163,194]), mobile applications (e.g., [33]), and web accessibility insights (e.g., [34,178]), to name a few. Key themes and design motivations that appear in many studies of disability within the computing sphere include: independence and empowerment (e.g., [117,124]), acceptance and empathy (e.g., [19,186,256]), rehabilitation and mobility/navigation (e.g., [51,257]), social integration and engagement (e.g., [14,107]), learning and memory (e.g., [139,142,191]), and safety (e.g., [2,35]).

In addition to having a variety of methodologies, accessible computing as a field also covers a wide range of topics and populations. Generally speaking, AC concerns the needs and desires of people with disabilities as users or stakeholders in technological tools or systems. Key populations of interest include people who are blind or have low vision, deaf or hard of hearing (DHH), autistic or neurodivergent, or those have an intellectual or

developmental disability, or motor or physical impairment. Lastly, older adults, who may or may not have a disability but are often subject to age-related changes and limitations, are considered to fall under the accessibility umbrella as far as AC is concerned.

Within the user communities AC studies, there is not an equal distribution of research and attention, according to surveys of the accessibility research conducted over the past 25 years [152]. Instead, the needs of BLV users were addressed in nearly 45% of the accessibility papers surveyed, compared to 14.2% for the next most commonly studied community, people with motor or physical disabilities [152]. This proportional overrepresentation of visual accessibility concerns in accessible computing means, among other things, that alt text has been an extensive topic of research among accessibility researchers.

2.2.1 Image Accessibility: Alt Text and Image Descriptions

Image accessibility broadly refers to the rendering of digital images as accessible to screen reader users. It is an important aspect of digital accessibility broadly and is listed as the very first requirement for basic website accessibility, according to the W3C [250]. There are several formats for accessible imagery [166]—including alt text, image descriptions, long descriptions, and captions.

Alt text is an HTML image attribute that provides a brief textual explanation of image contents [247]. Alt text is only accessible via the intermediary of a screen reader in most cases, although it also appears when images fail to load. Alt text is indexed by search engines as an element in search engine optimization [100]. Long descriptions were an accessibility feature also based in the HTML and included in prior versions of the WCAG guidelines [265]. Embedded in the “longdesc” HTML tag, long descriptions provided links

to separate pages with in-depth descriptions for complex images such as graphs, maps, infographics, and data visualizations. However, long descriptions as a HTML feature are now depreciated and were "never well supported in screen readers" [265]. In comparison, image captions can be used for image accessibility and are very common. However, captions traditionally do not directly describe the visual elements of an image. Instead, they can provide supplementary information, for example artistic credit for the image [56], or context that may help a user understand the image's relationship to the text. Captions are not, strictly speaking, an accessibility feature but I include them here partially because several studies have used the term "captions" to refer to descriptions of images, particularly those generated via machine learning (e.g., [141,153,215]. An image description, unlike alt text or long description HTML attributes, is usually included in the body of the webpage or social media post so that screen reader users, people who use magnification, and sighted users can all make use of it. The content included in alt text, image descriptions, captions, and long descriptions can, but will not necessarily, overlap. In this dissertation I generally use the terms alt text and image descriptions, and I use them more or less interchangeably. Exceptions include a few points in Chapter 4 where my participants themselves argue there is an important difference between those two different image accessibility solutions.

2.2.2 Studies of Alt Text Guidelines and Requirements

Accessible computing literature has studied the different requirements and guidelines for alt text in different contexts [227] for nearly 20 years. In 2005, researchers conducted interviews with five blind and low vision participants and found that context was relevant for determining proper content for alt text [182]. Since then, numerous

studies have looked to define and test guidelines for alt text creation based on feedback from screen reader users [18,153,185,227]. For example, Salisbury, Kamar, and Morris generated guidelines for improving the quality of alt text created by crowd workers [201]. These guidelines were then formalized into a rubric and used by actual crowd workers in Gleason et al. [95]. Academic alt text guidelines are often based on user research studies, but rarely have they been put to the test by actual novice alt text writers, making the Gleason et al. [95] example fairly unique.

Beyond academic guidelines drawn from user studies, there are also non-academic guidelines from a variety of sources, including museums [56], non-profit accessibility initiatives [63], and individuals [49,144]. Importantly for my discussion of The Avatar Project, some of these guidelines deal with issues around describing either drawings and works of art [63], or photographs of people and their identities [49], but none discuss artistic depictions of marginalized identity in depth [56]. Within the HCI field there are general inclusivity guidelines for researchers interested in writing about disability [10] or gender [209]. But they do not discuss proper language to use in alt text specifically.

In recent years, building on prior work [182], accessibility researchers have focused in on investigating the complexities of writing image descriptions for BLV users in different contexts [227] and with different marginalized identities [18]. Stangl et al. [227] investigated the different informational preferences for alt text based on different sources and surrounding context. Having looked at a total of seven sources (news websites, social networking sites, eCommerce websites, employment websites, online dating websites, productivity apps, and e-publications), Stangl et al. found that social networking sites, dating sites, and news sites, were sources where the most information was sought [227].

They noted that for social networking sites, the image poster's purpose was relevant to BLV users. They found that images requiring a "description of a person's physical appearance" brought "subjective or sensitive" topics such as race, gender, and disability to the forefront in ways that required further study. The study implied but did not discuss in depth the competing needs of alt text readers (i.e., BLV users, screen reader users) versus image subjects (i.e., people in the images, when people were present).

Bennett et al. [18] focused on the problem of describing marginalized identities in photographs through their study of screen reader users who also had additional marginalized identities (Black, Indigenous, Person of Color, or trans or non-binary, or some combination of those labels). Along with their inquiry into participants' non-visual methods for understanding identities of people online, they found four areas where participants felt information about identity was particularly important: 1) when identity was the topic of conversation, 2) when assessing the identities being represented in media, 3) when assessing the audience or potential shared identities of people around them, and 4) when specifically seeking out the perspectives of marginalized people. Bennett et al. also discussed the limitations and potential consequences of AI-powered image descriptions, specifically the risk of misidentification, and concluded there should be more studies targeting professional content creators, rather than everyday users [18].

The body of literature on identity in alt text is still growing, with few guidelines or studies focused on the specific challenge of representing social identities, particularly marginalized identities, in alt text. And yet already there is research supporting the fact that describing identities can be important for BLV users in certain contexts and guidelines suggesting the intricacies of describing identity inclusively must be studied.

2.2.3 Representations of Disability

Beyond including relevant and inclusive descriptions of social identity in alt text, there is another way in which disabled people have a stake in imagery such as The Avatar Project: media representations of disabilities. People with disabilities are underrepresented across many forms of mainstream media. For example, a survey of top-grossing movies found that out of over 4,000 named characters or characters with speaking roles, only 1.6% had a disability [223]. Within media studies and disability studies, there are discussions of how people with disabilities are written about in print news and depicted in advertisements [102,193] and how fictional characters with disabilities are portrayed in film, television, and literature [61,89]. The intricacies of inclusive language for disabilities have resulted in resources for both general (e.g., [6,254]) and HCI-specific audiences (e.g., [10,106]).

Existing studies of disability representation in media largely agree on two points. First, representations of disability have historically been limited [61]. Secondly, what representations of disability do exist in media are often one-dimensional, stereotypical, patronizing, or otherwise offensive [38,77,115]. Structural ableism not only causes the marginalization of disabled people in the workplace and educational sphere, but also leads to ableist depictions of disability in media. There are however increasing examples of disability activists expanding the visibility of or adding depth to conversations around depictions of disability, for example in the documentary *Crip Camp* [173], videos by Haben Girma [93], or in podcasts like *Reid My Mind* [189]. Making sure depictions of accessibility are themselves accessible has been a particular area of discussion, with social media becoming a site of outcry when there is representation of blindness, for example, on TV, but

no attempt to make the show accessible to actual blind people [76]. The accessibility of media has been studied by HCI researchers, investigating both professionally authored subtitles or audio descriptions for movies and TV [67], and multimedia approaches to making user generated content accessible (e.g., [126,205]).

2.2.4 Images of Disability in Design

Depictions of disability within HCI or design contexts are generally understudied. Some work has discussed the implicit representations of older users as formulated in the design process [170]. A larger body of literature looks at personas and scenarios, as common design tools, and has studied to topic of representing marginalized users through these tools(e.g., [40,42,109]). Personas depicting diverse user groups have been discussed as a way to engage and educate designers on underrepresented groups, working to foster empathy and combat stereotypes [109,156,244]. Even the challenges and tensions with adapting personas to represent disabilities and other complex overlapping identities have also been discussed [73,157]. However, the visual aspect of personas, and user representations in design more broadly, are rarely discussed. Two studies [202,203] have looked into how using multiple persona photos affect the perception of and engagement with the personas. They found the inclusion of multiple photos for a single persona increased the perceived informativeness of the document but led to confusion for designers. However, personas are only one of many ways to visually represent users in the design process. Research on user representations in mockups, placeholder images, or other images in product development is scarce.

A growing body of work (e.g., [103,169]) makes the argument that designers and technologists must consider disability and other identities at every stage of their design

process in order to create more accurate, fair, and equitable systems. Studies of algorithmic fairness have addressed systems that perpetuate and automate racial injustice [15,174], limitations and harms done to transgender and non-binary users of technology [26,125], and effects of algorithmic decision making on people with disabilities [241]. Biased or unfair technologies, such as algorithms that automatically tag pictures of black people as "gorillas" [212], highlight the need for designers to see diverse representations and consider diverse potential user groups for their technology.

Thus, this first study of my dissertation focused on the important topic of representations of disability in design images and how to make those images accessible by creating image description that describe social identities inclusively.

2.3 Methods: Co-Design through Interviews and Focus Groups

The first study was primarily a qualitative interview study, with interviews and focus groups providing the means for co-designing alt text as well as generating data for thematic analysis.

2.3.1 Research Question

The research question for Study 1 was:

1. From the perspective of users, specifically users with disabilities, what are important elements in defining inclusive image descriptions in fictionalized depictions of disabilities?

The study focused on fictional depictions of disabilities in image descriptions for several reasons. For one, as discussed above, prior research had focused on depictions of real people in photographs, concluding that the most inclusive option in these

circumstances was to defer identity descriptions in photos to the preferences of the subject themselves [18]. Thus, when images were of fictionalized subjects, the question of how to describe subjects became an interesting question. Who was the authority for defining proper terminology when describing fictional subjects, who could not actually be consulted themselves? The research question's focus on users and particularly users with disabilities was a result of trying to test the received representation in the images with members of the same marginalized social identity as the fictionalized subjects (refer to The Avatar Project, which formed the basis for my first study, began as an imagery project spearheaded by Emily Blank, Art Director for Google's Material Design team. Blank, along with Google Staff Researcher, Michael Gilbert, had held a series of workshops at Google prior to our collaboration starting early 2020. These workshops, involving researchers and designers at Google, had been concerned with how to properly represent a more diverse set of users, including users with disabilities, in various design and marketing assets. The search for inclusive imagery to represent users had begun with using photographs of Googlers who had given permission for their image to be used in such documents. However, concerns arose around branding and ethical implications of having photographs of real people used in contexts beyond their direct control. Thus, when I began working with Blank and Gilbert in 2020, they had already conceived of and commissioned an initial set of three artists to create the Avatar Project images.

The intention of the Avatar Project images was to function as reusable brand-compliant components that could be used by all designers (including blind designers and developers who often lack accessible tools [232]) to represent users in company contexts. They were to be included in the company's design system so they could be used by

designers and developers for internal documents such as mockups and personas. They would also be viewed by users in user testing environments and external advertisements. The intention was that by commissioning artists to create fictionalized user avatars, there could be a diverse set of potential users represented, prompting designers and engineers to "think more critically about inclusivity, diversity and representation" [149] without requiring real people to give away control of their own likeness. Thus, the three original artists, plus a fourth artist who worked on the project after our collaboration had begun, were commissioned to create a diverse set of illustrations of fictional users with the size, style, and content of the different artists' contributed sets varying.

My collaboration with Blank and Gilbert began because they were interested in understanding how real users responded to the Avatar Project illustrations. Since the Avatar Project images included people with disabilities, and my lab specializes in accessibility, we were consulted to see if we could find out how people with disabilities responded to the representations of disabilities in the Avatar Project set. Study 1 involved a subset of images from one of the Avatar Project artists (see **Error! Not a valid bookmark self-reference.**) that depicted visible disabilities as well as assistive aids. That artist's collection as a whole included representations of fictionalized people of many ages, races, genders, ethnicities, professions, sexualities, disabilities, and other characteristics. Together the images, as well as the process undertaken by the artists, Googlers, and collaborating researchers, to make them formed the basis for my study of *inclusive imagery*, a key term for the following chapters.

Table 2.1 for the specific Avatar Images that were assessed).

The additional reason for focusing on fictionalized depictions was that the Avatar Project images were already designed as fictionalized depictions for the specific purpose of being more inclusive. Google had not used depictions of real people to avoid the ethical difficulties related to assuring subjects retained power over their image. Other companies and content creators may very well be expected to take the same approach, making the question of inclusive descriptions of fictionalized depictions of disability to be particularly important. In the context of growing numbers of content creators seeking to create inclusive depictions of disability, the question of fictional depictions of disability is only likely to become more pragmatically relevant.

However, I do my best to report in the findings section of this chapter the important elements of image descriptions that participants discussed not just for fictionalized design images but that applied to image descriptions more broadly. While the context of the Avatar Project and design images is important to situate my findings, participants used various other touchpoints in conversation to explain their opinions on image descriptions. In particular, discussions of social media and other situations where participants encountered images, both fictionalized and real, were common. The goal of the study was to identify what users with disabilities felt were important when creating inclusive image descriptions in both the specific case of the Avatar Project, as well as important elements that apply to image descriptions broadly.

2.3.2 Recruitment

Study 1 was reviewed and approved by the UCI Institutional Review Board (IRB). For the study, the research team and I recruited 25 people with disabilities (see Table 2.2)

Table 2.2: End User (U) IDs and Demographics

ID#	Session	Age	Gender	Race	Disability
U1	S1+	18-29	Non-binary (ve/ver/vis)	Black or African American	MIID
U2	S1	18-29	Non-binary (they/them/theirs)	Black or African American	MIID
U3	S1+	30-39	Woman	White	P, MIID, BLV (sr)
U4	S1+	18-29	Man	Asian	DHH
U5	S2+	18-29	Woman	White	P, MIID (multi)
U6	S2+	18-29	Non-binary (they/them/theirs)	Asian	MIID
U7	S2	18-29	Non-binary (they/them/theirs)	Asian	MIID
U8	S3	18-29	Woman	White	MIID (multi)
U9	S3+	30-39	Woman	White	P, MIID
U10	S4+	30-39	Woman	White	BLV (sr)
U11	S4+	18-29	Woman	White	MIID (multi), P
U12	S5	40-49	Woman	White	P
U13	S5+	30-39	Woman	White	BLV (mag)
U14	I6+	18-29	Woman	White	BLV (mag)
U15	I7+	18-29	Woman	Hispanic or Latinx	BLV (sr)
U16	S8+	18-29	Woman	Black or African American	P, MIID
U17	S8+	30-39	Man	White	P, BLV (sr)
U18	S8	18-29	Woman	Asian, White	MIID
U19	I9	30-39	Woman	Black or African American, White	DHH, BLV (sr)
U20	S10	30-39	Man	White	BLV (sr)
U21	S10	30-39	Man	White	BLV (sr)
U22	S11+	40-49	Man	Hispanic or Latinx	P
U23	S11+	18-29	Non-binary (they/them/theirs)	South Asian	MIID (multi)
U24	S12+	30-39	Woman	Hispanic or Latinx, White	BLV (sr)
U25	S12+	18-29	Man	White	BLV (sr)

Key: + is participated in a follow-up interview, DHH is D/deaf or Hard of Hearing, P is physical disability, MIID is Mental Illness or other Invisible Disability, (multi) is multiple disabilities in this category, (sr) is uses screen reader some or all of the time, (mag) is magnification user

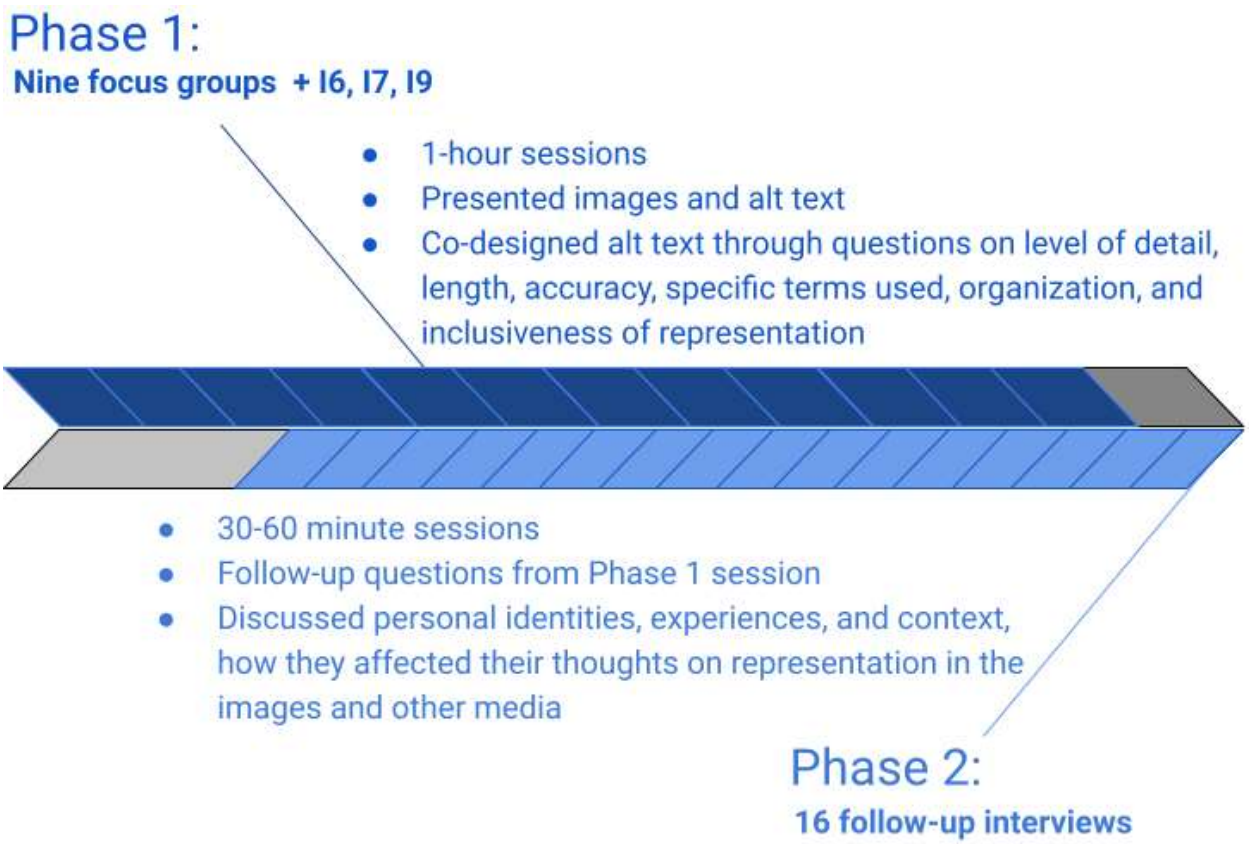

Figure 2.1 Diagram and summary of Phase 1 and Phase 2 of Study 1

from a combination of previous participant pools, our existing social networks, and through regional disability groups. We also used snowball sampling so that early participants could connect us with new participants for later data collection sessions. We included participants based on their self-disclosure of having a disability in a pre-screening survey. I intentionally prioritized the recruitment of people who had additional marginalized racial or gender identities and those who had multiple disabilities. This was so that results could potentially speak to inclusive alt text practices along dimensions other than disability. I also prioritized, in the final few sessions of the study, recruiting screen reader users, as I wanted to assure that the study reflected the opinions of PWD who were experienced with alt text in their daily life.

2.3.3 Data Collection

Study 1 consisted of nine focus groups and nineteen interviews, totaling nearly 28 hours of dialogue. All focus groups and interviews were conducted remotely via phone, Zoom, or Skype, and participants were compensated at a rate of $20 per hour. I attended all sessions, usually with one or more additional researchers present, and all sessions were video- and audio-recorded after obtaining written or verbal consent from participants.

2.3.3.1 Phase 1

For nine focus groups and three interview sessions (I6, I7 and I9), we used the same protocol—lasting about 1 hour each—to elicit participants' opinions and reactions to a subset of images from the Avatar Project (see The Avatar Project, which formed the basis for my first study, began as an imagery project spearheaded by Emily Blank, Art Director for Google's Material Design team. Blank, along with Google Staff Researcher, Michael Gilbert, had held a series of workshops at Google prior to our collaboration starting early 2020. These workshops, involving researchers and designers at Google, had been concerned with how to properly represent a more diverse set of users, including users with disabilities, in various design and marketing assets. The search for inclusive imagery to represent users had begun with using photographs of Googlers who had given permission for their image to be used in such documents. However, concerns arose around branding and ethical implications of having photographs of real people used in contexts beyond their direct control. Thus, when I began working with Blank and Gilbert in 2020, they had already conceived of and commissioned an initial set of three artists to create the Avatar Project images.

The intention of the Avatar Project images was to function as reusable brand-compliant components that could be used by all designers (including blind designers and developers who often lack accessible tools [232]) to represent users in company contexts. They were to be included in the company's design system so they could be used by designers and developers for internal documents such as mockups and personas. They would also be viewed by users in user testing environments and external advertisements. The intention was that by commissioning artists to create fictionalized user avatars, there could be a diverse set of potential users represented, prompting designers and engineers to "think more critically about inclusivity, diversity and representation" [149] without requiring real people to give away control of their own likeness. Thus, the three original artists, plus a fourth artist who worked on the project after our collaboration had begun, were commissioned to create a diverse set of illustrations of fictional users with the size, style, and content of the different artists' contributed sets varying.

My collaboration with Blank and Gilbert began because they were interested in understanding how real users responded to the Avatar Project illustrations. Since the Avatar Project images included people with disabilities, and my lab specializes in accessibility, we were consulted to see if we could find out how people with disabilities responded to the representations of disabilities in the Avatar Project set. Study 1 involved a subset of images from one of the Avatar Project artists (see **Error! Not a valid bookmark self-reference.**) that depicted visible disabilities as well as assistive aids. That artist's collection as a whole included representations of fictionalized people of many ages, races, genders, ethnicities, professions, sexualities, disabilities, and other characteristics. Together the images, as well as the process undertaken by the artists, Googlers, and

collaborating researchers, to make them formed the basis for my study of *inclusive imagery*, a key term for the following chapters.

Table 2.1). We formed focus groups based on participants' availability as indicated in the initial screening survey, scheduling a focus group when at least two participants were available. The three interviews included in this section were intended to be focus groups but due to cancellations only one participant attended.

Before the start of each of these initial sessions, to prompt thinking about representation in images and descriptions, we asked participants to provide either a sketch or one-sentence textual description of themselves from the shoulders up. At the start of each session, participants and researchers verbally shared their self-representation, as researchers displayed the illustrations and/or descriptions via the teleconferencing system's screen sharing and chat features. After introductions, we presented an example interface mockup that used three avatars, to give participants context for the intended audience and use cases for the images and alt text we were discussing (i.e., they would be used by industry designers in mockups). Next, we displayed four to six slides, each with one avatar's illustration and its associated image description. For the sake of time, we could only present a subset of the nine avatars in each session; our selection was such that participants would be shown avatars most closely related to their disability identity. For our participants who used screen readers, we pasted the image descriptions into the text chat as we introduced the image and gave participants time to listen to the description.

For each image description, we asked participants what they thought about the level of detail, length, accuracy, specific terms used, organization, and whether it fairly

represented the identities portrayed. We also asked how they felt about the image and description being used by a potentially non-disabled technology designer. At the conclusion of the session, we asked more generally about how social identifiers such as age, gender, race, and disability should be described in image descriptions; what should be considered when writing image descriptions and assessing quality; and whether invisible identities or disabilities should be included in the images or descriptions. As a form of theoretical sampling [47], we asked participants in later sessions to reflect on the participant opinions and our interpretations from earlier in the study.

2.3.3.2 Phase 2

Of the 25 participants who took part in the initial focus groups and interviews, 16 participants were able to participate in a follow-up interview—lasting between 30 and 60 minutes—to provide us an opportunity to probe further the comments they made in their session in a setting where discussing personal identities would potentially be more comfortable. These semi-structured interviews covered what social and personal identities our participants deemed important, whether their importance changed based on certain contexts, and whether these identities informed their reaction and responses to the illustrations and image descriptions shown in the previous session. We discussed participants' experiences with representations of their disability or disability in general in text or images[3].

[3] In Session 11, no image descriptions were shown, and instead only visual representation of disability was discussed. That session and the associated follow up interviews with those participants were included in the dataset, despite not discussing image descriptions, because they provided general feelings about disability representation in media.

2.3.4 Data Analysis

Focus group and interview recordings were transcribed through institutionally authorized automatic transcription services and edited for accuracy by the researchers. I, along with the other primary researchers, analyzed all transcripts through line-by-line inductive open coding. After creating an initial codebook, we developed an affinity diagram using MURAL, an online collaborative whiteboarding tool, to form axial codes and subsequent themes. Researchers met weekly during analysis to iterate on codes and themes.

While analyzing the comments from participants, we noticed participants seemed to understand and limit their comments to the original context (fictional representations of disability used by designers) to differing degrees. Many participants discussed image descriptions more generally, including their experiences with them on social media. These context switches were taken into account during analysis and the fact that participants used a variety of different contexts to discuss their expectations of important alt text elements supports the generalizability of findings beyond design contexts.

2.3.4.1 Iterative Interviews and Focus Groups as Co-Design

Study 1 used interviews and focus groups to collect general information about alt text elements and actors that were important from the perspective of users, but the interviews and focus groups also served as a medium for conducting alt text co-design with users. Co-design, and the more broad term co-creation, is "any act of collective creativity" particularly that which relies on "designers and people not trained in design working together in the design development process" [204]. Co-design exists on a spectrum, with the depth and breadth of user involvement and leadership in the design process differing

from project to project. We did not explicitly discuss the focus groups as co-design workshops at the time. However, the act of iteratively changing the product (alt text) based on the feedback from participants was always a primary part of the method used for Study 1.

The first version of the image descriptions was the original two- to three-word descriptions provided by Google when they gave us access to the Avatar Project images. Based on those descriptions, and with attention paid to existing guidelines on image description and identity in particular, a first 'prototype' version of the alt text was created and shown to the first focus group participants. Following that, after every session one of the primary researchers returned to the set of images and descriptions we were using for each session and changed wording or added elements based on the feedback we had received from the participant(s) in that session. Thus, the final version of the alt text produced in Study 1 was our attempt to synthesize all the feedback into a single image description that included all the elements that users mentioned were important in the context the images were intended to be used in.

2.4 Findings

I have included here all findings relevant to the overall study of inclusive image descriptions and how they are made. This includes two subsections. The first is focused on important elements of image descriptions, as described by disabled users in the study. The second describes the important actors and their roles in image descriptions, as we discovered through conversations with users that actors and their interpretational power were key elements of how alt text was made and received.

2.4.1 Important Elements of Image Descriptions

"Elements" is an intentionally vague term. I use it to capture all pertinent aspects of image descriptions that users discussed, from word choice to image composition to representation of marginalized identity. I include in this section at least a brief mention of all the elements that users discussed with us in their evaluations of the Avatar Project images and alt text, arranged roughly in order from those elements that were particularly important to users or were discussed by the greatest number of participants, to the least important elements. For details on the participants U1-U25 as they are quoted, please refer to Table 2.2. The only demographic detail that I mark the user IDs with is an asterisk (*) after their user ID number if they use a screen reader all or part of the time, a caret (^) if they use screen magnification, and no modifier if they are not a BLV participant.

2.4.1.1 Bare Minimum

One important element of alt text that might get assumed but which our participants were sure to point out was having alt text at all. For many of the participants, encountering alt text in the wild was rare (or at least far rarer than it should have been) and several participants made sure to point out that fundamentally having any alt text present at all for most images would be a step up from the current practices. For example, U17* explained that in his experience, "Very, very, very few people bother doing alt text of their profile pictures…[for] the vast majority of people, [alt text]'s just not something they think is relevant or important or interesting… the kinds of alt text that you guys have done [in this study] with the cartoon images is a lot closer to the kind of headshot descriptions that I would like to see, that I don't hardly see any of. And the people who do them tend to be like disability activists types and nobody else bothers. So, I think it's really cool that you guys

are even trying to do this." U3* mentioned that many sighted people may not even know about alt text: "I still think it's so new that people don't understand it or necessarily how to do it." This is backed up by the unfamiliarity that some of the sighted users in our study had with alt text, despite having disabilities themselves. For example, U11 stated that "before we really did this [focus group], I hadn't thought too much about image descriptions." Ultimately, before the users in our study could even talk about elements that alt text should include, there needed to be an increase in the basic availability and understanding of alt text in online spaces.

2.4.1.2 Purpose

Beyond simply having alt text at all, the next most important element users discussed was that proper alt text should be attuned to the purpose of the image. Six participants mentioned explicitly that the purpose of an image was important in deciding what to include or focus on in image descriptions. For example, U11 felt that image descriptions should "essentially, prioritize the information based on... what is the photo trying to convey." The intent of the image creator or poster, the reason why they posted it, was important, from U11's perspective, in determining what to include in the alt text. I'll return more to the issue of intent in the next section of this chapter. When discussing the decision to use gender neutral language for an image of a young skateboarder with a prosthetic hand (see Figure 2.2), U17* commented that "representing people as ambiguous is good in general, especially when it's not necessarily relevant to what you're trying to portray... The point of the image is that you're showing that amputees do things that everybody else does... It doesn't matter what gender they are." U17*, a cisgender man, felt the artist's purpose was to represent disability in a normalizing way, and thus disability

should be referenced more explicitly than gender in the image description. U19*, a cisgender woman, similarly agreed with the interviewer when asked if the general purpose of the image was more important than specific details that describers could choose to include.

Six participants discussed the role explicitly describing disability in image descriptions could play in the Avatar Project images, since they were intended to serve an educational purpose, benefitting non-disabled professionals in design contexts. U14^ argued that "because you're using [the avatars] for mockups, it's very important for developers to take [disability] into consideration. So, I think maybe [disability] should be front and center." U1 agreed that describing disability was valuable, "to remind [designers] that these [disabled] users are going to be using [their] product."

Figure 2.2 Example of normalizing imagery of disability while de-emphasizing other identities

A colorful Claymation-style illustration of a youth wearing a helmet and holding a skateboard with a prosthetic hand. Their other hand is tucked into their red shorts as they look off to the side, smiling.

However, participants were careful to delineate the difference between the immediate purpose of having good representation and the much broader purpose of being inclusive generally. U24* stated, "Good representation... [is] something more focused on taking away the stereotypes, whereas inclusive means 'let's not put a flat screen ATM machine out there.'" In another example, U3* believed the image descriptions would help implicitly educate designers: "companies... kind of put it on the user to [create image

descriptions] ... [because] I just don't think they have the education around it." But she believed that this education could lead to more inclusive technologies to support alt text. U2–similarly to U24*–felt that avatar images, and image descriptions by extension, "are absolutely not enough, but avatars create... awareness in society about people's disabilities." In other words, for the overall purpose of inclusivity to be effectively served, inclusive image descriptions were themselves only one element of what should be larger corporate support for accessible technologies.

2.4.1.3 Context

One important element that should be considered when writing image descriptions, according to eleven participants, was the context the image was posted in. For example, discussing the image and description in Figure 2.3, U10* felt that if the image was posted in the context of "talking about this person going on a trip to pursue whatever hobby, it wouldn't necessarily, I feel… make sense to expound so much on the vitiligo… If it's maybe an article about vitiligo or about this person's experiences with vitiligo, it makes sense for it to be in the first line of the description… if it's not, I feel like that part could move a little farther down." Context was important when deciding how much focus to put on certain elements of the image in the description. According to U19*, some details, such as information about body size, can be left out entirely if they are not contextually relevant: "If it's just a general image then maybe it doesn't need to be mentioned whether or not they're skinny but if… for instance, I was in a Facebook group that had fitness goals and if you posted a before picture and after picture, obviously it would need to be described whether or not the person is heavier or lighter."

Figure 2.3 Description of avatar's vitiligo in inclusive or educational context

An avatar with dark-colored eyes and vitiligo so that parts of their medium-dark skin tone has lost pigment and is a very light brown tone.

In some cases, there is contextual information about the image subject that has to be included in order for an image to make sense. U17* mentioned that he was in a Facebook group that described memes and "you can't just say, 'a person,' you have to say 'it's this character from such and such'... or else the whole meme doesn't make any sense. And if you don't know that character, you could know to skip it, because… you know that you won't understand what the joke is… There are details–non-apparent details–that are important to know in cases like that, where it's not just a person but it's a person in a particular universe." There can be context around where the image originally came from that users need to know in order to fully understand that image when it is reused or placed in another context, as with a meme.

The context also influences the intended audience, which then dictates what information viewers are likely to need. For example, U16 talked about an image description that included "curly hair" and the details that are most important to her as a Black woman. She argued, "It just depends on what type of information I need to know at that moment... if it was in a hair magazine, like a Black hair magazine... [The reader] might want their hair like that." Similarly, U15*, gave social media as an example where use case may affect

description content, likening the level of depth needed to "the difference between glancing at a picture and looking at one." She continued: "if you're scrolling through Facebook, you're not going to look real intensely at your friend's Facebook profile. You know what they look like." Thus, the context, both in terms of the surrounding information and goal of the image consumer, was relevant to what alt text should include.

Disability in particular was an area where participants felt that information should only be included if contextually relevant. Ten participants mentioned that disability was important to their identity and how they describe themselves, but also emphasized that it's not always relevant. For example, U9 described how she sometimes cropped her disability out of profile images. She explained, "Sometimes I struggle. I don't typically display [my disability], but I guess If I was using an app that relates to it, then I would... I think it would be really situational." For this participant, their hesitance to include disability in some representations related directly to the platform or community she was engaging with. This insight applied to representations of both visible and invisible or non-apparent disability. In another example, U11, who has multiple non-apparent disabilities, felt that when deciding how she wants to be represented, "there are nuances that are dependent on what type of site is being used... Is it Tinder? Is this a professional site? Am I using this on my Canvas when I'm replying to my professors?" Participants generally agreed that, regardless of whether disability is visible in the image, elements of context like the audience may make it irrelevant or inappropriate to include disability in the description.

2.4.1.4 Accuracy & Completeness

Another element that multiple participants mentioned was important when considering the quality of alt text was accuracy and completeness. U20* remarked that one

of the issues with automatic alt text on platforms such as Facebook was that they had to sacrifice the completeness of the description for the accuracy of it.

> I'd like Facebook to even just get to the point where I know… 'a person in a suit with a background of grass.' … The problem is, Facebook is being fairly cautious because of the risk of misidentifying details… they're basically giving very limited amounts of information for the exact reason of not wanting to give misinformation or run the risk of giving misinformation.

On the other hand, according to U17*, more concise language can sometimes paint a more accurate picture of exactly what is in the image. For example, referring to the image and description in Figure 2.4, U17* said "in regard to 'a service dog' versus 'a dog wearing a harness,' a dog wearing a harness doesn't convey anything because lots of dogs wear harnesses that aren't service dogs. … [Saying someone has a physical disability], that's even more broad than 'a dog wearing a harness' because it could literally be anything." Accurate and complete alt text was therefore a balance between precise language, when possible, versus avoiding including details that were based on incorrect assumptions.

Figure 2.4 Language accurately describing a service dog

An avatar of a service dog with a green collar and silver harness, held by a pale hand.

Thirteen participants mentioned the importance of disability, in particular being portrayed accurately in image descriptions. U22 described how representations of disability "ha[ve] to be accurate. It is very important, because if not, you are sugarcoating it ... you don't want to sugarcoat, because that is not reality." He explained how representations should not be unrealistically optimistic to the point of changing the meaning a viewer would draw from that depiction of disability. U11 discussed representations of different bodies by saying, "Literally, I want to be accurate, like my body shape … sometimes [descriptions] are like, 'here is someone with an hourglass figure,' and that does not match how I perceive myself." Writing about disability in an idealized way, "sugarcoating it," was seen as a disservice. It functioned to make representations more palatable to non-disabled audiences, rather than accurately capturing and making visible disabled bodies.

2.4.1.5 Objectivity

Another important element of image descriptions was objectivity, or at least the attempt to achieve it. U19* wanted alt text writers to remember: "Don't be like, 'this picture, it is a very attractive woman.' I hate when people do that… be as objective as you possibly can." U15* argued that image descriptions needed to "use sort of this general objective language" if they wanted to be trusted as representations of disability, as well as explaining that in image descriptions written about her, she would "rather it be neutral and objective" in tone.

One way that participants felt objectivity could be supported was by focusing on strictly visible elements. U19* suggested instead of saying "attractive woman", describe "what is she wearing? What is she looking at? What expression do you think is on her face?

Stay away from as many adjectives as you possibly can." U8 argued that writing an image description should follow a "protocol of just writing down what you can see and not trying to infer anything more." U3* said that image description authors should avoid "guessing age or race or things like gender, [instead] keeping things just as descriptive as possible." By avoiding assumptions about things like age or relative attractiveness, details that were objectively visible in the image could be focused on.

And yet, perfect objectivity may well be impossible when writing alt text. U11 argued that image descriptions are just "another form of media... like any story, they're going to be subjective to the point of view of the person who tells it." Even in U19*'s earlier quote, she specified the goal should be "to be as objective as you possibly can." She further explained that "three people can look at the same picture and draw different conclusions." U15* felt that the best way to reach the "neutral and objective" alt text she wanted would be through "a collaborative image [description], almost a compilation of a bunch of people looking at an image of me and if some people find me super serious, or stuck up or whatever it is, and some people just think that I'm calm [it would even out]." Thus, some level of subjective interpretation may be unavoidable, but attempts to remain as objective as possible are still valuable.

Figure 2.5 An early iteration with assumption-based description of age

An avatar of a 13-year-old athlete with a prosthetic hand.

2.4.1.6 Personal Relationship to Social Identity

Describing social identities in image descriptions is made particularly difficult by the fraught nature of marginalized identities in many spaces. Thus, an important element that participants identified when discussing how to describe identities such as disability and race was the personal relationship the subject had to their identity. For example, fifteen participants in our study commented on having specific preferences for the language used to describe their own disability. U9 mentioned not feeling represented by medical language, as it did not accurately describe her experience with her disability, "I always have to look up [my condition's name], it is all medical. And then, with the condition, everyone looks completely different. There is not a 'normal' way to describe it." U10* and U13^ mentioned feeling more connected to the "blind" community, rather than the "disability" community, more broadly. Neither very specific medical terminology nor the very broad language of "disability" felt right for these participants.

The personal relationships participants had to their disabilities often came up as factors that made them hesitant to prescribe one correct way of describing disability. All our participants disclosed that they had a disability in conversation, but in their role as image subjects, only three of the sixteen who provided image descriptions of themselves

chose to mention disability in those descriptions (See Table 2.3 for descriptions). This may be due, in part, to the stigma associated with disability. U19* said that she typically does not describe herself as blind because it has "a negative kind of connotation," and that "many people get hung up on the fact that you're blind that they don't notice anything else." Similarly, U15* said: "I want to be seen, I think, as somebody, a person, before I want to be seen as a blind person." She was one of six participants in the study who explicitly stated that they did not consider or want to be portrayed as if their disability was the most important thing about them. U4 mentioned a hesitation to describe his disabilities in writing, because of how it tied him to a community he didn't feel he could fully claim, "I don't feel that sense of pride over being deaf to the same degree as the other [non-mainstreamed] side of the deaf community."

In contrast, U3* explicitly included both her visible and invisible disabilities in her image description introduction at the beginning of the focus group. She explained in a follow-up interview that she does this with images of herself because "this was one of my ways of saying, 'okay, well these are things you might not see in the picture, but this is something I want you to know about myself in this description because it's important to my identity.'" U15* mentioned that, in order to accurately describe her visual disability, her disability has to be "something that [the image describer and I] talked about versus making an assumption about who I am ... because everybody's eyesight is very different." These examples show that participants have unique and complex relationships with their own disability identity, and that this makes describing disability in image descriptions a highly personal process.

Table 2.3: Three Differing Styles of Self-Description Featuring Disability Disclosure

U3's upper body is shown. She has a brace on her right hand, to stabilize her permanent injury. She wears heavy makeup, covering an eyelid mutation called BPES, and has lived her entire life visually impaired. She also lives with the invisible illness Complex PTSD.	I am an African American woman in my late twenties who has crooked wrists and club feet due to being born with a physical disability. (U16)	45 year old Latino cisgender male. Short dark brown hair, brown eyes, clean-shaven. Wearing a long-sleeved shirt, conservatively dressed. Sitting in a power wheelchair. (U22)

While almost all participants were comfortable discussing aspects of their own identity related to age, race, gender, and disability, they particularly worried about or avoided definitively labeling these identities for fictional characters or images with unknown subjects. For example, U3* suggested that, normally when generating image descriptions about disability, you should simply consult the subject for "permission." Similarly, U5 asked the interviewers if the avatars depicted real people, so there was a means to confirm their gender, to avoid misgendering them. U11 argued that, when she reads image descriptions, she prefers that they are constructed by the subject: "It does become 'who is telling your story?'... because if I read it from the source of the image, I am going to feel like I have a better sense of what they want me to understand about them, [compared to] somebody who is looking at it as an observer and imposing their own biases onto the image." These examples emphasize how challenging it may be to create image descriptions in some cases, such as when there is no access to a subject who could validate a description or give permission.

2.4.1.7 Explicit or Implied Identity

U10* felt one element to consider when describing identities in images was how explicit or emphasized those identities were. On the one hand, explicitly stating identity can be important. For example, U10* felt that explicitly labeling an image subject's disability could be valuable for specific contexts such as "some kind of disability pride event... or a whole bunch of pictures of people with different disabilities." Rather than explicitly labeling a fictional subject's disabilities, a third of participants discussed including assistive aids in alt text instead (see Figure 2.6 for an example). U10* explained that naming assistive devices was more informative than identity labels because "it could be someone with a [service] dog or someone with a [white] cane and either way, it would be 'a blind adult.'" U12 and U13^ both agreed that describing assistive devices such as wheelchairs and white canes was preferable to using identity labels that the subject may not or cannot identify with. But even when mentioning relevant assistive devices in the image, U13^ clarified that one should "include the tools, but maybe not have the tools that those people use be the central focus." Even given the educational opportunity the avatars represented in the design context (see 4.1.1.10 below), participants felt that disability identity labels were so personal and complex (see 4.1.1.6 above) that they hesitated to explicitly assign one to these fictional representations.

Figure 2.6 Example implying disability by describing assistive aids

An avatar wearing sunglasses, holding a white cane in one hand and a steaming coffee in the other.

This alternative method, implying disability status by describing visual signifiers of it such as assistive aids, avoided an issue that U21* identified where if you choose to explicitly label an identity that is not strictly visible (such as disability identity, in this case blindness), "how would you know? Because you're then kind of insinuating they have a look about them that's blind, which a lot of visually impaired people find insulting. 'You look [blind].' Well, what does that mean?" Thus, focusing on strictly visible elements such as assistive aids and letting those markers imply a disability identity avoids the issue of assumptions the alt text writer might be making about the subject.

2.4.1.8 Hedging Language

One particular element of describing identities that might not be strictly visible, was a type of 'hedging' language that some participants used or recommended using. For example, two participants specifically used gender-neutral language such as "person" and "they" or "them" pronouns to discuss the image subjects. Instead of assuming that the people in the image (see Figure 2.7) are women, U7 suggested describing them as "appear[ing] to be aesthetically feminine." U8 stated that describing other traits like one's clothing or accessories can allude to a person's gender as well, but also admitted this approach is based on assumptions she has about gender. U19* argued that "if you don't

know the [subject's] gender or if they're non-binary, you could always use the word like 'they appear to be a woman' or a man or whatever, so that way you're not saying they are or they're not."

Figure 2.7 Early hedged iteration describing potential gender of couple holding a baby

An avatar of an interracial couple, who appear to both be women, holding and smiling at a baby.

Concerning race, participants largely preferred describing skin tone only, to alleviate the potential for mis-categorization. U6 also believed that since race is "more of a construct," describing the subject's skin tone made the image description more "indisputably true." U20* pushed back slightly on describing the skin tone. He agreed that skin tone did not always line up with race, but he felt that describing skin tone, "doesn't give me a heck of a lot of information. Whereas, at least if you say someone appears to be African American, that is much more evocative than... medium dark skin tone or whatever." Because the avatars used in this study all depicted fictional characters, virtually all participants emphasized a need to describe visible characteristics as opposed to assuming identity labels. One element of providing evocative descriptions based only on visual information may be the use of hedging language to emphasize that these identities are being guessed at but have not been confirmed by the subjects.

2.4.1.9 Level of Detail

Figure 2.8 Example of too-detailed information about groceries for wheelchair using avatar

A colorful, 3D cartoon-like illustration of an adult with a light skin tone and a blond beard and ponytail using an electric wheelchair to carry some groceries. On their lap is a brown paper bag holding a loaf of bread, a large carrot, some cheese or other orange package, and celery. A blue milk carton peeks out of the top of the reusable bag hanging from the back of their chair.

As might be expected, many participants mentioned that conciseness in image descriptions is important. In our study, eight participants (seven BLV, one sighted) noted that more details, especially more "irrelevant" details, can make image descriptions difficult to consume. For example, U14^ noted, "You don't want to have too much information in the alt text, because then you can get overwhelmed and then it's hard to ... understand what's going on in this picture." U16 also preferred clear and concise descriptions because "not everyone wants to listen or read all of that." When discussing the avatar holding a white cane and a coffee, U17* said, "I've seen ridiculously long image descriptions … so I'm personally perfectly fine with 'white cane'... I like economy over minute detail unless there's some kind of reason for including minute detail." As we will see below, the educational purpose in design contexts may constitute a reason for more detail. U19* felt that personal preferences come into play "because there's those of us who like that detail, and there's some people who don't." Ultimately, participants couldn't easily articulate what the line was between too much or not enough detail but said that it was something they could form their own opinion about in the moment of encountering specific alt text. A

potential example of “too much” detail, based on a comment from U13^, can be seen in Figure 2.8.

2.4.1.10 Educational Value

Five participants mentioned how details included in image descriptions can be educational for people not familiar with that particular identity. Demonstrating the exact educational opportunity that might necessitate “minute detail,” U8 was confused why the term “white cane” was used when “part of the cane is green.”. Rather than recognizing “white cane” as the formal term of an assistive device commonly used by BLV people to navigate, she attributed the term to the color of the cane itself. The inclusion of the term “white cane” was enough to prompt questioning, but further detail in the description could have done more to educate. Similarly, when reading the description of the person with vitiligo, U19* mentioned that she was unfamiliar with the term “vitiligo.” She believed that most people would not be familiar with the term and suggested that it was more important to describe what vitiligo looks like visually rather than use the medical term. While viewing the avatar signing “Thank You,” (see Figure 2.9) U21* noted, “The person is clearly hard of hearing, owing to the hearing aid. You describe the motion [in the text]. So, for somebody who's reading this image, for a blind person, now they know how to say ‘thank you’ in sign language.” We see how including more detail in image descriptions can provide a learning opportunity for readers about the disability being represented.

Six participants suggested that explicitly describing disability in image descriptions could serve an educational purpose, benefitting non-disabled professionals in design contexts. U14^ argued that “because you're using [the avatars] for mockups, it's very important for developers to take [disability] into consideration. So, I think maybe

[disability] should be front and center." U1 agreed that describing disability was valuable, "to remind [designers] that these [disabled] users are going to be using [their] product." U3* believed the image descriptions would help implicitly educate designers, encouraging them to create features in their software to support image description, adding that: "companies... kind of put it on the user to [create image descriptions] ... I don't necessarily think people want to be excluding [disabled] people, you know, I just don't think they have the education around it." However, participants were careful to delineate the difference between having good representation and having good accessibility. For example, U24* stated, "Good representation... [is] something more focused on taking away the stereotypes, whereas inclusive means 'let's not put a flat screen ATM machine out there.'" U2 similarly felt that "avatars are absolutely not enough, but avatars create... awareness in society about people's disabilities." In other words, to be considered effective, good representations of disability in image descriptions had to be backed up by corporate action, in the form of accessible mainstream technologies.

Figure 2.9 Educational description of avatar signing "Thank You"

The avatar has their fingers together and hand positioned upward, resting against their chin. They motion their hand forward, smiling slightly. With this motion, they are saying thank you in sign language.

2.4.1.11 Customizable Interfaces

Related to the Level of Detail section above, multiple participants mentioned or reacted positively to the idea of having customizable "levels" to image descriptions. This would allow them to potentially have a short initial version of the alt text and then be able to decide if they wanted to go on to hear the longer description. U20* felt "a long-term solution" that many platforms needed to be thinking of ways to implement "will allow for customization of how much detail they want." He even went on to describe having each type of detail be coded as "individual items that can be called" which would allow for people to get just the details they were interested in, for example "maybe someone wants the first half of this [description]... some people don't care about what the person's wearing." This level of customizability is somewhat in line with the verbosity settings that many screen readers already allow but would have to be enabled and accommodated for by the image describer as well as the platform hosting the image. U19* said that having multiple versions of alt text with differing levels of detail "would be a great option" to deal with the differing personal preferences people had on how much detail they wanted. When we prompted U17* to think about this idea, brought up already by other participants, he had a slightly more complicated reaction to the idea of multi-level image descriptions. He said, "I can't think of that many areas where you would want multiple levels of image descriptions" although he specified that video games or "anything dealing with animation basically" were areas where having interactive image description interfaces would be valuable. He didn't himself play video games however so he "didn't know how technically feasible it would be."

2.4.1.12 Normalizing vs. Othering Language for Out-Group Legibility

Another element participants thought should be considered specifically when describing marginalized identities was how and to what degree the description was legible to out-group members. U10* argued that "there's that balance of not making it a bigger deal when it doesn't need to be, but also it doesn't have to be ignored either." Essentially, she felt the balance had to be between including disability information, so as to normalize disabilities for those unfamiliar with them, versus overemphasizing them to the point of leading non-disabled users to exclude or other people with disabilities. U19* also pointed out that it's important to use language that is understandable to outgroup members because, for example, "white cane" as a detail that signals blindness to others hinges on the question of "is that [term] widely known." She explained "I'm not non-disabled. So, I assume everybody knows that if you see a white cane, you're blind, or if you see a hearing aid, you're hard of hearing or Deaf. Maybe not everybody knows that." Similarly, U20* and U21* discussed how "light skin" as a description could be ambiguous depending on if the person reading the description happens to know that there are people who are "light skinned but they're not white." In these cases, considering how an out-group member might interpret the image description is an element of crafting an unambiguous and/or positive description of that identity.

2.4.1.13 Descriptions Beyond Visible

One element of image descriptions that blind participants particularly brought up, was that not all visual details were necessary, or that other types of details could be valuable as well. For example, U10* felt "the thing I care most about in the image" was the inclusion of detail that "conveys this person's personality, how they come across."

Personally, potentially because she is blind, U10* said that "I'm not going to care as much about how they look… when I encounter a person in real life, I'm going to notice [and] care more about how they come across and whether they're laughing, smiling, that kind of thing." This was something that made her sad that "it's all photos basically on social media… Everything's photos," when what she wanted was to be "more inundated with audio clips." U19* similarly felt that for sighted people "visually there's probably so much information that you guys receive that… you can get into this thing where it's too much detail." Some blind participants also noted that identity labels could be less evocative if people didn't have visual frames of reference. U21* argued that taking "somebody who's visually impaired or who has never seen before" and giving them a term such as Sikh "really tells me nothing about that man because I don't know what Sikh looks like. I don't know what the skin tone is, I don't know [anything about him] unless I go and look it up." The overemphasis on visual details led blind participants in particular to feel that sighted people did not always prioritize the right information when writing image descriptions. This also aligned with several participants' decision to include non-visual information such as non-apparent disabilities in their descriptions of themselves.

2.4.1.14 Consistency

Consistency was not a common topic for our participants, but it was mentioned as an element of alt text by some. U19* talked about it in a variety of ways. She said, for example, that images in a set with a similar style did not need to have the style described exactly the same each time. Instead "if it's at the beginning it should be fine" rather than "keep repeating it." Specifically in terms of describing identities, U19* remarked that she had "noticed in books, if they describe someone, and they're describing people of color,

they will always describe the [skin tone of] the person of color, but they will never describe the white man or woman as being white" so that in alt text "if you're going to describe the skin tone, that needs to be across the board, it needs to be pale or medium skin tone. It has to be everybody. You can't [just] do it with one particular race." U19* as a person of color herself felt that race was potentially overemphasized and it wasn't necessarily relevant to include it in image descriptions, however, if race was mentioned at all it had to be described consistently across subjects.

2.4.2 Important Actors and their Roles in Image Description Creation

In addition to the many elements listed above that participants felt were important to consider when writing image descriptions, there were also key elements that were tied to specific actors and their roles in the image description creation process. The three sets of actors that came up in discussion with users were the artists or image creators, writers or image describers, and users or description readers. Each of these sets of actors had different roles they played or elements they contributed to the final image description. For a diagram of the roles and layers of interpretation, see Figure 2.10.

2.4.2.1 Artists or Image Creators

2.4.2.1.1 Intent

One topic that has appeared in the elements discussed above but which is particularly relevant to discuss when talking about the role of artists is their intent. As U17* puts it "if the artist has gone to pains… to make them not white, you should have some way of reflecting that in the alt text." The fact that the author put in the effort and intended for the image to depict a person of color makes it more important to describe that detail in the alt text. This was tied up closely with the idea of the image's purpose (discussed above),

to the point that 8 participants mentioned the artist or intent behind the image was part of their final interpretation of the image description.

2.4.2.1.2 Translation

Another role the artist has, in addition to being able to convey an intended message or identity through their description, was the element of translating existing ideas and identities into visuals. U8 considered what would happen if the artist was given just the alt text and told to recreate the original image. “If you were to use race as a way to describe the images, and you’re using the text to create the image… that might change how in the end the image looks, depending on the artist.” This was, as we’ll see in Chapter 3, actually a very accurate way of thinking about how artists who are commissioned for images like the Avatar Project actually work. Artists often have to base their art on at least an agreed-upon set of ideas, characters, or identities, and then have leeway within that to put their own stylistic details into the image. But at the most basic level, the artist is translating the basic written or verbal description of the images agreed to with the project funder into the actual visual representations of subjects seen in the images.

2.4.2.2 Writers or Image Describers

2.4.2.2.1 Interpretation or Bias

Various examples of how the unchecked biases of image description authors can impact alt text surfaced in our conversations. U19* shares: “a lot of stereotypes and racism and homophobia get still passed down in the visually impaired community because you rely so... heavily on what is described to you. So, if [my father says], ‘oh my god, these two women, they have this baby’, you're going to have a negative connotation... because of his tone.” U11 distinguished between overt prejudice and more subtle bias or subjectiveness,

stating, “we need to make image descriptions unbiased, in a sense that they should be free from prejudices. But I don't know that they need to be unbiased, in the sense that they're free from a sense of who the [author or subject] is or their personality”. U10* felt that when creating alt text there are two levels: “there’s some objective details one can describe. And then in some ways it is subjective because it’s dependent on what the person looking at the picture is seeing.” There may be benign forms of subjectivity that appear in image descriptions that can be acceptable or even just unavoidable, but overt ableism, racism, homophobia, and other prejudices should be avoided.

2.4.2.2.2 Familiarity with Subject

One strategy that participants identified to reduce bias was to seek an image description author who was familiar with the subject of the image. For example, U11 explained how, “[image descriptions] speak a lot to the person who makes the image description – what they prioritize. That is why I think it’s descriptions [I see on social media] done by the people who are in the pictures are oftentimes better.” This implies that image descriptions produced by the subjects of images may be of better quality. But for BLV participants, they did not always have the option to independently author their own image descriptions. For example, U15* relies on assistance from a sighted partner. She explains: “That [description] is basically the way that whoever you are trusting [to help create the description] sees you.” Thus, the familiarity of the describer with the subject, or the degree to which the describer’s interpretation of them matches the subject’s own, can be an important element of making sure an image describer is playing their role correctly.

2.4.2.2.3 Personality or Tone

Another element of alt text that an image describer can affect is the tone. Particularly, U10* mentioned that "for better or worse" she can really get a sense of the personality of some image describers, if she's familiar enough with them as people. She explains: "I have friends on Facebook who some of them just write these delightful image descriptions. And they're just fun to read because [of] the way they use language, but also, it's because I know who that person is so when they're describing the picture… I just hear 'Oh, that's my friend Gary!' Just like, 'oh, that's so him, the way they were describing [the image].' And I love it. It's super great because it just feels like this is the person's personality." Alongside benign forms of subjectivity in image descriptions, in certain situations, particularly when the readers are familiar with the describers, a personalized or playful tone to image descriptions may be warranted.

2.4.2.3 Users or Description Readers

2.4.2.3.1 Interpretation

Beyond the biases that can appear in the writing of image descriptions, seven participants referenced the fact that description readers played their own role in the final interpretation of the image, based on their prior knowledge or experience. Upon reading an image description of an interracial couple holding a baby, U17* said he assumed they were a same-sex couple even though it wasn't explicitly stated because he lives "in a certain time and place and [has] had certain experiences... but someone living... halfway across the world might not... it depends." U11 believed that "a lot of image descriptions rely on you knowing things" and that this was intrinsic rather than something to be avoided. U10* acknowledged that "how [she interpreted] the picture might not be how [she] would

interpret it if [she] could see it," but she did not imply that this was a sign of a poorly written description, necessarily, just a factor in reading image descriptions in general.

2.4.2.3.2 Misunderstanding

Another role that readers could potentially play in the process of creating the final meaning from an image description, is the possibility of misinterpreting or misunderstanding the image. More than a subtle interpretation of ambiguity described above, misunderstanding is the potential for image descriptions to fail, linguistically, to achieve the meaning intended by the artist. For example, U20* and U21* discussed how there could be misunderstandings on multiple levels. U21* argued the goal was to "define [race] in a way that someone who has never seen before could understand… I have low vision, I know what an African American man looks like, I know what a Latina woman looks like. I know what that skin tone is. When somebody says caramel skin, I kind of know what they're talking about. But if you go up to somebody… who's never seen before [they're] not gonna know what I mean." Thus, depending on the experiences of the reader, even seemingly objective visible details (such as "light skin" discussed in Normalizing vs. Othering Language above) can be misinterpreted. Also, demonstrated in fact in the focus group with U20* and U21*, there can be opportunities for wording to explicitly be misinterpreted. U20*, listening to one of the image descriptions said, "discussing the way the baby's being held… jointly they're holding [it]...I kind of read that as they were holding hands" and upon clarification from me as the moderator he said, "I thought I was being nitpicky about my wording, not realizing that I was totally misunderstanding the entire concept." Partially a product of unclear wording (and thus related to the role of image

describer) it is ultimately on the user's shoulders to navigate the wording and either correctly or incorrectly interpret what the image is depicting.

2.4.2.4 Other Actors

2.4.2.4.1 Platform

There were also small mentions in the interviews and focus groups of other actors who played a role in how image descriptions ultimately worked. For example, as discussed in the Accuracy and Completeness section above, U20* argued that the spare automatic descriptions he encountered on Facebook was a result of Facebook "being fairly cautious." In turn, "the way Facebook does it with their algorithm" means that image descriptions cannot be updated or completely rewritten manually by users. Thus, the platform itself, as a supplier of automatic alt text and as a technical interface to allow people to enter their own alt text, plays a role in what the final user experience of the images on that platform is.

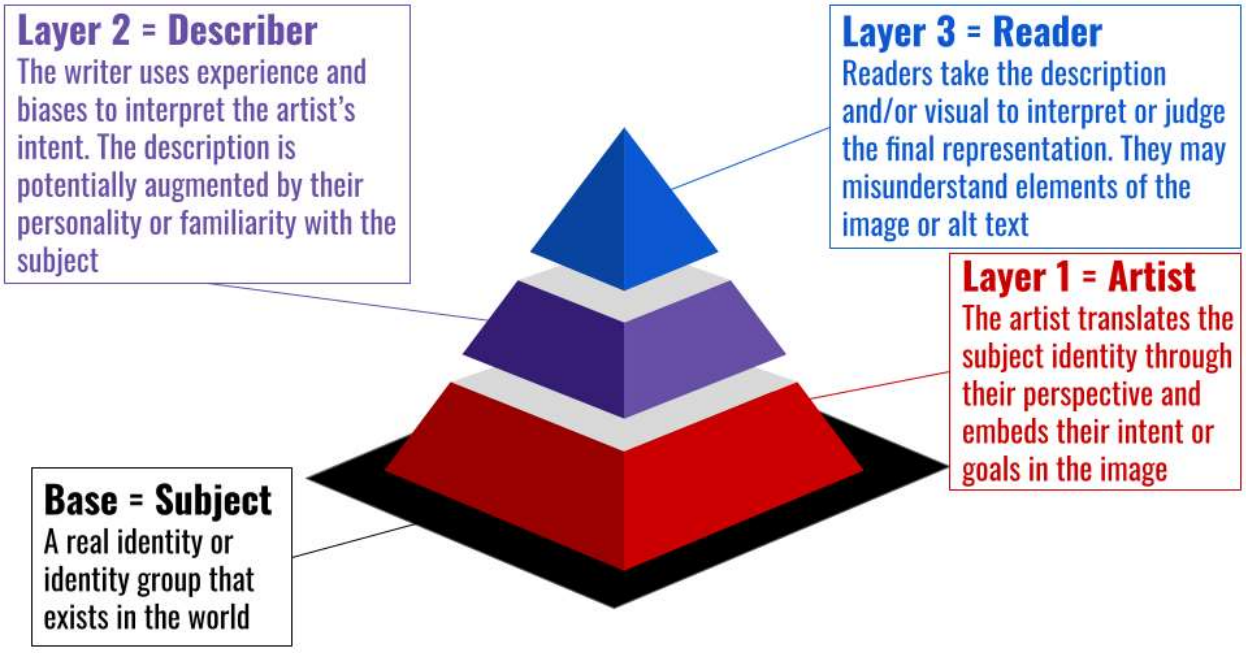

Figure 2.10 Pyramid showing successive layers of interpretation in creation and reception of accessible imagery

2.4.2.4.2 Technical Interface

The technical interface of the screen reader also played a role in how image descriptions were understood. Such as when U20* and U21* both mentioned that their screen reader pronounced "Sikh" like "sick." U21* said "I'm going to have to go into my JAWS dictionary and change that." While the context of the alt text in this case made it clear that the word was referring to the religious group rather than the physical health status, both participants were well aware that their screen readers could mispronounce words and that this could lead them to not understand the correct meaning of the alt text.

2.5 Takeaways

I want to take a moment to synthesize and summarize the many elements of image descriptions I just discussed. I also want to highlight key takeaways from the study, both in terms of discussion-level conclusions from all the data, and in terms of a short reflection on the strengths and limitations of this Study. I will also have a brief section explaining the future work that this study sets up, both the future studies of this dissertation, and the larger research paths that Study 1 opened up for anyone to follow.

2.5.1 Framework for Alt Text Creation

Prior work [18], published while this study was underway, discussed the representation of self in photographs and their alt text posted to social media, and the complex ways people with multiple marginalized identities had to negotiate identity disclosures in online spaces. Bennet et al.'s overall guidance was that the subject of the photograph should be the final authority on how to describe their identities. Users in our studies seemed to agree. They felt the personal relationships subjects had to their identity

impacted how you should describe those identities. However, in the case of the Avatar Project images, the subjects were not real people who could be consulted on their preferred language. Users in our studies reacted to that in several ways. On one hand, several blind users in particular advised describers to use hedging language such as "appears to be" to avoid passing along assumptions uncritically and instead signal that visual information such as skin tone was being used to guess at a non-verified identity such as the subject's race. In some other cases, such as describing disability in particular, participants felt that image descriptions could best balance the difficulty between explicit and implicit identity terms by using visual signifiers such as assistive aids to imply an identity. This led to some difficulty when the readers were not familiar with the assistive aids, potentially because they were an out-group member or because there were not enough educational details included in the image description to match the intended audience / context of the image.

However, one particularly interesting response to there not being real subjects to consult was that participants discussed the role of the artist in shaping the image and thus the description. Many participants did not feel comfortable projecting—or suggesting image describers should project—identities onto fictional characters that they did not believe were visually verifiable. Media studies has engaged deeply with the concept of reader response theory—the idea that "meaning does not reside in the text, but in the mind of the reader" [268]. However, for our participants, in the situation where description authors sit between the media and the reader, the "readers of the image" were not seen as a final authority for decoding the subject's identity. Instead, participants located the ultimate

authority for correct alt text language (in the case where subjects themselves could not be asked) in the hands of the image creator, i.e., the artist.

Even in this description of how different elements of alt text work together, you see I repeatedly refer to a few different actors involved in the final understanding of the image and description: image subjects, image creators, image describers, and description readers. Ultimately, examining the key elements of inclusive image descriptions is incomplete unless we consider these important actors and their roles in constructing the final interpretation of the image description. The framework (see Figure 2.10), as we interpreted from users' discussions of image descriptions and how they should be made inclusive, revealed each of these three sets of actors played a variety of roles.

Artists (i.e., image creators) were the source of the intended identity being depicted (if the characters were being invented wholesale), or in some cases the translation of subjects' identities into the given artistic medium (if based on existing subjects or agreed-upon character types). Writers of image descriptions had the potential to play the role of semi-objective or actively biased interpreter of visual information. Beneath this larger role, their familiarity with the subject and their personality had the ability to modify their interpretive role, for better or worse. Lastly, the users consuming the images and descriptions played their own interpretive role, using their prior knowledge and assumptions to fill in gaps in the language of the description.

They also, potentially, could have a role in misinterpreting or misunderstanding the image description, based on the wording, their familiarity with the terms used, and the attention they may or may not pay to particular details. There were other actors that were

implied in one or two places to have some role in the final image description. The organizational and technical actor of the hosting platform played a role in understanding the alt text. And the technical interface of the screen reader also acted as a mediator that could affect the final understanding of the image description.

2.5.2 Reflections on Strengths and Limitations of Study 1

It's difficult now, after several years of working with and reinterpreting the data from Study 1, to look back on it with new eyes. There are many small things that I would have done differently. Everything from recruiting more participants (and particularly scheduling more participants per focus group session) to keeping more careful track of each and every suggested or implemented change to the co-designed image descriptions. One of the limitations of Study 1, particularly as it was originally published, was that it was co-design "via committee" [28] and I don't feel we as researchers took enough of a stance on editing and streamlining the alt text for peak readability.

I also feel that Study 1 was limited in its specificity. While all good research has a relatively small scope, so as to be able to gather enough data to say something meaningful, the important elements identified in this chapter, as well as the guidelines from the original publication of this study, are quite specific to the context of the Avatar Project. The images were a set of illustrations, created by the same artist, for use by developers and designers. All of these details limit the generalizability of the alt text creation framework described above, despite logical arguments that could be made to its general applicability.

I fear, on the other hand, that the guidelines from the original publication were not clearly enough explained as applying only to the specific context of the images we studied.

This potentially has led readers of the research to take a "more is better" approach to image descriptions, when careful analysis of the data clearly shows that this is only true if the detail in question is warranted based on the purpose of the image. Additionally, users did not always agree in this study. While I tried my best to represent those dissenting opinions in the findings, I fear that the judgment calls that we as researchers ultimately had to make may be taken as the definitive "correct" way to deal with tricky issues such as describing identity, when in fact the user data is conflicting, and no simple answer may be applied to all cases.

Another limitation of the study is the limited pool of identities represented by participants. While we had some diversity in terms of gender and race, our sample included more white people than is an accurate proportion of the general population. We also did not have any older adults who participated in the study, which meant we had a dearth of data related to age and its intersection with disability, even though older adults form a large percentage of the disability population at large [238]. We engaged with participants who were fluent in English and were based in or had emigrated from the contiguous United States.

Despite these drawbacks or limitations of the study, both as it was originally published and as I have re-analyzed and presented here, there are many strengths to the study design and implementation. For one thing, while the research team debated it at the time, looking back I think it was the right choice to include both sighted and BLV participants in the study, as long as they had some sort of disability. I believe this choice allowed us to understand much more broadly how image descriptions were understood

and approached by users as a whole, rather than only screen reader users. This ended up being highly valuable because it parallels the image description experience available to users on Tumblr (see Chapter 4). Additionally, the broad range of users we spoke to allowed us to get feedback from people with the same identities as those depicted in the Avatar images (with the exception of older adults, which was a critical user group we failed to recruit). Feedback on how identities should be described by people with those identities is I think the most valuable of all, because they are the most aware of existing stereotypes and issues to avoid. Indeed, if I could have found an easy way to do so, I would have represented opinions here (such as those on race) with a simple visual indicator if the feedback was coming from someone with know-how on the issue (such as a person of color). Finally, I believe this study was valuable because of the depth of information that was gathered on such a small subset of images. I can't speak to all types of images, or all possible contexts they may be posted in, but for the particular topic of inclusively depicting marginalized identities in artwork, we were able to get very specific answers.

2.5.3 Future Work

This study suggests several avenues for future work. For one thing, it suggests that there are many more identities and identity intersections than those we covered and that they all deserve careful, detailed attention on how to represent them inclusively in image descriptions. For another, the relative importance or relationship between competing elements identified by users could be investigated in more detail. How does one balance educational opportunities in alt text with not othering marginalized subjects to out-group members, for example? If a describer has to decide between hedging language or implying identity, when is one more appropriate than the other? How fundamentally different could

users practically imagine alt text for the same image being given different enough audiences, context, and purpose? All of these questions deserve further study to understand the nuances of what proper alt text looks like.

However, none of those directions describe the actual work that followed up on Study 1. The more interesting direction stemming from this research, for me, was not the specifics of image descriptions being written in for this purpose under these conditions. Instead, the larger structure of how image descriptions are made and received seemed valuable. No one, to my knowledge, had laid out a framework to describe the actors and their roles in alt text creation and reception. I believe there are valuable insights and implications to be derived from the theoretical examination of alt text creation as a process involving multiple groups of people interacting indirectly through the medium of the image and image description. However, considering the framework as I described it here was only based on the perspective of users, the next step for this research path was to investigate if this framework held up when looking at alt text creation in practice. Additionally, follow-up work had to balance the users' perspectives represented in Study 1 with the perspectives of image describers and image creators who could provide different insights into the alt text creation process. To address these gaps, Study 2 was designed and executed.

CHAPTER 3: Study 2 – Perspectives on Alt Text Creation Processes from Industry

This chapter supplements the existing work on image accessibility and alt text interventions by looking specifically at the processes through which practitioners make alt text in situ. The study revolves around the central research questions of 1) how accessible and inclusive images are created in the context of a technology company and 2) how various actors negotiate the elements of alt text during the creation process. I approach the rarely discussed topic of imagery and alt text creation in practice by presenting findings from interviews I conducted with three artists and three accessibility practitioners working with Google. These interviews are supplemented slightly with the data from 25 end users (those discussed in Chapter 2). I first describe four processes of alt text creation used by Google—The User-Evaluation Process, The Lone Writer Process, The Team Write-A-Thon Process, and The Artist-Writer Process—and unpack their potential strengths and weaknesses as they relate to access and inclusive imagery. I also discuss further themes that complicate the prior version of the alt text creation framework or that imply new directions for image accessibility research and practice. I conclude with a discussion of how alt text elements are negotiated by differing actors, an updated version of the framework for alt text creation, and a list of takeaways that alt text researchers and industry professionals can learn from considering alt text in industrial practice.

3.1 Related Work

3.1.1 The Lack of Alt Text and How to Fix It

Lack of alt text has been an important accessibility issue for over 15 years [182] and, despite improvements in alt text coverage over time [100], it remains a large problem

facing the BLV community. Image description prevalence on popular websites increased from less than 40% in 2006 to 72% in 2018 [23,100]. Although some estimates put the number of images lacking alt text as of 2022 at closer to 55% [249], with many more having low-quality alt text such as file names or the word “image” [100]. And in some cases, prevalence is much lower, with sites primarily comprised of user generated content such as Twitter found to have alt text on only 0.1% of images [95]. Lack of alt text in an increasingly visual online world can lead to lower participation of BLV users on those sites [167]. It also leads to a large gap between the experiences and information available to sighted users versus screen reader users [3], with many screen reader users lacking important visual information on a variety of topics, while sighted internet users often are not even aware image accessibility is necessary [214].

To address pervasive inaccessibility of digital imagery, the accessible computing field has explored several potential solutions, ranging from fully automated to highly social methods for creating and propagating image descriptions. Several studies have leveraged computer vision to automatically label images [44,101,259]. While these automatic approaches are valuable for their scalability, research shows they cannot promise the same quality as human-generated descriptions [18,95,153]. Researchers have also explored automatic methods for finding alt text or captions presentin other instances of images and providing them to users seeking image descriptions [100]. In the realm of human-generated image descriptions, crowdsourcing [22,32,200], friendsourcing [30,159], and social microvolunteering [31] have all been explored as potential methods for providing image descriptions. There are two primary considerations when determining the appropriateness of human-generated methods. The first is how to provide high quality

information (i.e., how to teach friends, crowd workers, or volunteers to write image descriptions that follow accepted guidelines and address the particular information needs of BLV users) with as little burden on the describers or oversight from experts. The other consideration is how to create systems that allow humans to generate alt text on-command while reducing lag time between request and response. Some researchers have attempted to use multiple methods in combination to achieve high quality results [95].

Currently, research continues in multiple directions, with attempts to improve quality of automatic methods, research into motivating and teaching crowd workers, and further research (as discussed in Chapter 2) into what methods for generating alt text might be preferred based on the information needs of BLV users in differing contexts. No one method will likely triumph, but finding which image description problems are best solved with AI versus with human-powered labor is still an ongoing challenge. What research does suggest is that there are two distinct issues to address if image accessibility as a whole is to be significantly improved. First, the issue of scale, particularly for user generated content, is significant. Billions of images are uploaded to the internet every day by individual users [80], and yet the vast majority of individuals online are not even aware of image accessibility as an issue[214]. Secondly, there is the issue of quality and professionally generated content. While user-generated content is by far the larger portion of images posted online, Bennett et al. [18] suggest professional content creators are an important part of the struggle to increase the quantity and quality of alt text for important or high-profile images. In studying the inclusive imagery creation processes at a large technology company, Study 2's findings particularly address the second of these two concerns.

3.1.2 The (Lack of) Research on Production Processes for Alt Text

In general, studies of alt text primarily focus on one or more of the following: assessing the current state of image accessibility, including what proportion of images are accessible (e.g., [167,182]); offering an automatic tool, human-in-the-loop process, or other intervention to populate more images with alt text (e.g., [95,100,259]); or discussing alt text quality, including tools for assessing alt text quality or studies of what makes good alt text (e.g., [18,197]). Studies assessing the current state of image accessibility generally focus on either the prevalence of alt text or the experience of screen reader users encountering alt text in naturalistic settings. Even when the topic of alt text production does come up, as in studies of how alt text should be made better quality, there is not a thorough baseline for how users currently make alt text against which these ideal methods can be compared. There are a few exceptions to this, such as [150,221,255] which have, as part of the larger studies, sections of their procedure dedicated to examining how alt text authors currently create alt text. However, all three of these publications were published in the 2020s and even most of these papers [150,221] don't focus heavily on the process of creating alt text, only on the alt text quality, issues, or understanding authors demonstrated when creating alt text.

Thus, the topic of how alt text is made in practice appears to just be starting to gain the attention of researchers and remains an underexplored area. This is an oversight, as having a thorough baseline understanding of how alt text is currently produced in different contexts and under different processes can allow future researchers to better understand why alt text is currently lacking or poor quality, and how to design more effective future interventions.

3.1.3 Defining Communities of Practice (COP)

Beyond addressing the topic of alt text, Study 2 is also a study of what is called a "community of practice." The definition of communities of practice is fairly straightforward: "groups of people who share a concern or a passion for something they do and learn how to do it better as they interact regularly" [78]. To further explain this short definition, communities of practice have three defining characteristics: 1) an identity defined by a shared domain of interest within which there are levels of expertise that can be achieved (not necessarily formally), 2) joint activities and discussion within a community space (be it virtual or physical), and 3) a shared repertoire of stories, tools, and ways of addressing problems, that helps them practice their shared activity [78].

One of the many bodies of research on communities of practice focuses on organizations and workplaces as sites of communal practice. Communities of practice have been identified as sometimes taking the form of people within an organization who work together [147]. However, just belonging to the same organization does not make people a community of practice, they must work and collaborate together. Workplace communities of practice have been shown to display Wenger's 14 indicators of a community of practice including:

1. Sustained mutual relationships
2. Shared ways of engaging in activities together
3. The rapid flow of information and propagation of innovations
4. Absence of conversational preambles, as though interactions are simply continuations of an ongoing process
5. Quick setup of a problem to be discussed

6. Substantial overlap in participants' understanding of who belongs
7. Knowing what others know, can do, or how they can contribute to a goal
8. Mutually defining identities
9. Ability to assess the appropriateness of actions
10. Specific shared tools, representations, or other artifacts
11. Local lore, shared stories, inside jokes, knowing laughter
12. Jargon and shortcuts to communication
13. Certain styles recognized as displaying membership
14. Shared discourse reflecting a certain perspective on the world [50]

In HCI research, it is common to talk about "HCI practitioners" i.e., designers, developers, project managers, and other people involved in the technology industry. Whether all references to "HCI practitioners" or the more specific sub-group of "accessibility practitioners" align one-to-one with the definition or indicators of a community of practice is difficult to say. However, [50] have shown that a learning community focused on technology and accessibility at the very least can meet all 14 of the Wenger indicators.

3.1.4 Example Community of Practice: Accessibility Practitioners

HCI practitioners, and the research surrounding them, often focuses on the gap between research communities and industry practices. This issue has been studied from a number of different angles (e.g., [41,54,59,98]). Within the larger body of literature on HCI practitioners, there are an increasing number of studies looking specifically at accessibility in practice. Some of these studies involve evaluating existing guidelines and accessibility evaluations for the challenges they pose to designers and developers who need to put them

into practice (e.g., [55,137,230]). Evaluations of existing guidelines often conclude that considering designers and developers are the “users'' of the guidelines—that is, the people who need to understand and put into practice the information contained in the guidelines—their pragmatic needs and perspectives are not properly considered in the design of accessibility guidelines. This has led to several studies focused on developing tools to bridge the gap between accessibility guidelines or evaluations and the practitioners who have to instrumentalize them (e.g., [88,213]). Other studies focus on pedagogical considerations for students and early career HCI professionals and how they learn or understand accessibility (e.g., [43,97]).

There are many studies that directly consult tech industry professionals to understand their approach to or understanding of accessibility, however almost all of them are survey-based [69,86,118,184,187,242,260]. Interview studies or qualitative investigations of accessibility in practice are few and far between (exceptions include [8,90,246]). How industry communities address image accessibility in particular has rarely been investigated. Giannoumis and Nordli briefly describe the basic knowledge of image accessibility their participants demonstrated and how it was requested and delegated in the context of their case study of a Norwegian broadcasting company [90]. Williams et al. [255] describe the workflows that computing researchers used when creating alt text for their publications, noting issues with complicated figures, tight deadlines, and technical tools used to insert alt text. However, prior to the publication of Study 2, there were–to my knowledge–no studies directly looking at the overlap between professional tech industry communities and how they practice image accessibility.

3.2 Method: Email and Video Interviews

Building on Study 1's investigation of important alt text elements according to users, Study 2 focused on how alt text was understood or created differently in industry contexts and by industry professionals.

3.2.1 Research Questions

Study 2 sought to answer the following research questions:

1. How are accessible and inclusive images created in the context of a technology company?
2. How do various actors, including image creators, alt text creators and end users negotiate the attributes of accessible and inclusive images for their platform?

The second research question was summative and sought to synthesize the findings from the first study and the second study into a larger understanding of how the important actors identified in Study 1 articulated and negotiated the elements of image descriptions that were important in their given context.

3.2.2 Recruitment

Study 2 was approved under the exempt category by the UCI IRB. To recruit participants for Study 2, I first reached out to all artists involved in the Avatar Project illustrations. My Google collaborators shared their contact information, and I reached out to them via email to discuss their participation in the study. I ended up conducting interviews with three of the four artists who had worked on the project. Two of the artists, due to work schedules, time zones, and other constraints, were not available for video

interviews but were willing to participate in email interviews which took place over the course of multiple months. The third artist agreed to participate in a video interview.

In addition to the Avatar Project artists, our collaborators in the Google Material Design team put me in contact with co-workers that they knew had experience with accessibility or writing alt text. I also used snowball sampling as the Material Design interviewees discussed other co-workers they knew from previous positions or teams and I had the early interviewees share my contact info with these colleagues, thereby broadening my participant pool.

Table 3.1: Details about Artists (A) / Image Creators

ID#	Image Set Details	Background	Additional Notes
A1	Created largest set of Avatars; Included explicit depictions of disability; Colorful but realistically shaded style	Of Asian descent but didn't explore that in the Avatar Project; Had family members with disabilities	Worked with a character designer on their images; had finished their illustrations over a year before interview
A2	Included explicit depictions of disability; Combination of lineless and thick-lined shapes	Had experience as a teacher working with children with disabilities	Finished their illustrations well before interview; Based in Italy but spoke in English for interview
A3	More abstract style, focused on depicting/using everyday objects; Included images inspired by disability communities/concepts	Trained in Industrial design; Had created images for various US editorial pieces	Was in the midst of drafting and redrafting their set when interview was held

Table 3.2: Details about Alt Text Writers (W) / Accessibility Practitioners

ID#	Years in Job Title	Background / Job Focus	Experience With Accessibility
W1	UX Writer for 3 years	In a technical writing program; In charge of crafting the words in the UI; Oversaw contractors on copy editing and put data into the internal content management system (CMS)	Created image descriptions for a new launch of an existing product website including blog posts; Had previously organized a collaborative team activity to create alt text for figures that were missing it.

ID#	Years in Job Title	Background / Job Focus	Experience With Accessibility
W2	Program Manager for 15 years at various companies	Background in computer science; Focuses particularly on engineering and technical problems; Also handles general project management tasks.	Is fully blind; Worked in accessibility outreach and consultancy; As a program manager, oversees various attempts to organize efforts to produce alt text as well as having personal experience as an alt text consumer.
W3	UX Writer for 5 years	Before Google, had worked in human resources at a large institution working with internal tech teams; Has always been passionate about language; Sees job as understanding and guiding a user through their goal.	Had had been introduced to alt text as a temp but didn't have extensive experience with it until she took responsibility for generating alt text for a large set of images that were going to become available on Google products.

In total I recruited 3 artists and 3 accessibility practitioners; their relevant information is summarized above (see Table 3.1 and Table 3.2).

3.2.3 Data Collection

In the artist interviews, I discussed the artists' process for creating the images, particularly any images representing disability, and their[4] artistic approach to the project. I asked questions regarding how, if at all, their own personal identities or experiences influenced their process or artistic products and how or if they tried to explicitly depict marginalized identities or if the race, gender, or age were intended to be ambiguous. The interviews finished with asking the artists their feelings on receiving feedback from consumers of the images.

[4] In an effort to assure anonymity for the artists, giving one artist was named in one of the initial publications, all artists are referred to with the gender-neutral "they" pronoun throughout.

In the accessibility practitioner interviews, I asked the industry professionals about their general job title and duties, about how accessibility was done at the company from their perspective, and how or to what degree accessibility was introduced as part of their job responsibilities. I also discussed with them their thoughts or experiences with alt text including processes they used to write alt text or used to determine or affirm what should be included in alt text.

All the accessibility practitioner interviews and the one artist interview over video were conducted over video conferencing, either Zoom or Google Meet, and recorded. Each video interview lasted between 30 and 45 minutes.

3.2.4 Data Analysis and Google Review

All interviews were automatically transcribed and then edited for clarity. I then performed line-by-line inductive coding and sorted data into overarching themes. After analysis, data was written up into an initial draft of the study, which was reviewed by Google. We went through several rounds of revision to make sure the discussion of Google's accessibility practices was not misleading and ended up cutting much of the general description of accessibility practices in industry from the version of Study 2 that was submitted to TACCESS.

3.3 Findings

In Chapter 2, I described the methods used to create inclusive and accessible imagery in collaboration with Google (as part of Study 1). However, as I spoke to research collaborators and began to embark on the interviews with industry stakeholders that make up this study, I found that outside of research collaborations such as Study 1, in-depth iteration on alt text with feedback from users was rarely if ever used by practitioners.

While our own collaboration demonstrates that companies sometimes are involved in full studies focused on evaluating alt text for users (what will be called The User-Evaluation Process below), there were three more types of alt text creation described by practitioners at Google: The Lone Writer Process, The Team Write-A-Thon Process, and The Artist-Writer Process. For all four processes I will here describe the key steps involved in the process, how relevant actors (users, alt text writers, artists, or designers, etc.) interact or inform the final alt text, and best features or largest drawbacks of the creation processes compared to one another. Following that, I briefly describe three additional themes that point toward areas for improvement in future alt text production processes in industry

3.3.1 Image and Alt Text Creation in Industry

Before I present the four different models of alt text creation, I will begin by describing the general context in which the models our interviewees discussed were situated. The accessibility practitioners I spoke to had varying degrees of familiarity with organizational approaches to accessibility outside the specific projects and teams they had worked on (see Table 3.2 for background and contextual details of the accessibility practitioners, referred to here as W1, W2, and W3). According to W2, Google's approach to accessibility "[works] on multiple levels… There's a bottom-up approach; there is a bit of a top-to-bottom approach." This top-down and bottom-up approach was demonstrated in the alt text creation processes as well. W1 described the organization-level requirements for alt text passed down to product teams: "In any product launch, you have the accessibility review and a legal review, amongst others like engineering code reviews… [so] you have a [Quality Assurance] tester—accessibility tester—then going through the content [to assure image descriptions] make sense in the context of the task that they're

trying to accomplish in order to test it." In addition to these top-down processes, more bottom-up resources were used. For example, W1 and W3 both discussed being aware of or using a variety of alt text guidelines existing in different forms and formats across teams. The manifold nature of the organization–as a large tech company with many teams and divisions–required a similarly manifold approach to accessibility and alt text generation in particular, demonstrated in the differing methods that accessibility practitioners described.

The role of individual interest and advocacy as a complement to existing organizational structures was also an ongoing theme in discussions of accessibility practices with the practitioners. W2 pointed out there were both formal and informal structures to support accessibility, with many individuals–particularly engineers–having their personal "curiosity" supported by provided training and networking opportunities: "Notwithstanding the fact that people can just connect with each other [through] bug bashes and workshops and events… there's plenty of opportunities to learn about accessibility." As we will see in the next section, personal passion and dedication among accessibility practitioners was particularly valuable in some alt text generation methods such as the Lone Writer process. This overarching context for the four alt text creation methods also highlights the importance of organizational support and formal policies in enabling robust alt text creation.

3.3.1.1 The User-Evaluation Process

I will begin with the alt text generation method I am calling The User-Evaluation process. Demonstrated in Study 1, I am calling this the User-Evaluation Process because it involves the direct evaluation and revision of image descriptions with users. In our enactment of the process in Study 1 we particularly focused on centering the perspectives of users with

disabilities through interviews and focus groups. However, theoretically the core feature of the User-Evaluation process could be seen as including users in direct evaluation and revision of image descriptions through almost any method (e.g., surveys, usability tests, A-B testing) depending on the resources and time available or the amount of qualitative feedback desired.

An important element of the User-Evaluation method is how knowledge and responsibility is distributed. Acting as the alt text writers in study 1, we based our original version of the alt text on the descriptions given by the artist, but otherwise had creative control over the alt text and how user feedback was integrated. User comments impacted the final alt text, but (at least in our case) did not filter back to the artists. A2 particularly said any feedback "would be a great enrichment and I would gladly change any part of my illustrations." A1 was similarly open to changing the images based on user feedback. Because the User-Evaluation process took months, coming back to the artists and asking them to redraft the images was logistically impossible. It is important that, unlike the Artist-Writer process (see below), it can be hard to iterate on both imagery and alt text using this process because, as Study 1 shows, the final received understanding of the alt text is impacted not just by the alt text but the imagery itself. Thus, a drawback of the User-Evaluation process may be the close collaboration and careful management of timelines needed to properly distribute the feedback from users to all actors in the imagery creation process.

The practical feasibility of the User-Evaluation process may be its largest drawback, as buy-in at least was very high for the method. User feedback was something all of the

artists (A1-A3) and accessibility practitioners who had experience with directly creating alt text or images for user consumption (W1 and W3) discussed as valuable. W3 explained that, “If possible, it would be great to get some actual feedback from users… based on these descriptions that I wrote, how confident do they feel in selecting an image as their profile picture? That it was an image that met their intention? Were they worried that the image wasn't quite what they wanted? … Was it helpful? Should I have been more detailed? … Or was it okay to be really spare?” The issue of user feedback may be particularly pertinent for W3 because, as we’ll discuss in the next section, she used the Lone Writer process and may have felt more fully responsible for the overall alt text. She said, “I think if I were going to do a big set again, it would be good to actually get some feedback.” The enthusiasm for user feedback among our participants suggests that this process may be ideal for situations where buy-in from different stakeholders is important, such as image descriptions with broad cross-product or cross-organization applications.

3.3.1.2 The Lone Writer Process

In contrast to the many potential stakeholders involved in the User-Evaluation process, the Lone Writer process is in some ways the most straightforward. The process is characterized by a body of images being passed to a single alt text writer who creates the alt text for all the images before they are passed down to users. The Lone Writer process, as it was described to us by W3, was: “Working on a collection of 800 illustrations where I did write the accessibility descriptions for all of them” Describing her process, W3 said: “I sat down and just started going through as many as I could… I think I did a calculation that it would take me about a minute to write each one so whatever 800 minutes is… yeah, I ended up working a couple of weekends to get this done.” The Lone Writer process is

defined by this single solitary workstream, where rather than collaboration within or across roles (users, writers, artists) the content flows unidirectionally through the layers of alt text creation.

One of the features of this process was that W3 was able, because she was the only writer of the alt text, to keep a certain level of consistency between images. “I dug up the guidelines that I already knew about from my first team, where I was a temp and read through them… Some of the guidelines for that, they were going to make things too long, because it was you know, like ‘10 words’ and I was like ‘Oh, I think, like six maybe for this’ just because I don't want [it] to take four hours for somebody to go through the whole library… one of my guidelines for myself was like, don’t refer to colors [because they could be customized by the user]… I tried to limit myself to four to six words generally and using present tense.” These adaptations of existing guidelines–and the resulting similarities in length and tense across the 800 images–were simple to decide on and implement as W3 was solely responsible for the alt text. Whether this degree of consistency is ultimately a positive in terms of user experience was not assessed. However, as discussed in Chapter 2, some types of consistency within and knowledge of the entire set of images was desirable to users (see 4.4.1.14).

However, one clear logistical drawback of the Lone Writer Process is the potential, as seen in this example, for a high burden to be placed on one individual. W3 at multiple points described the effort put into the 800 images over two weeks. When asked if this was a task she more or less assigned to herself, she agreed “it definitely was.” She explained that she worked on it over the weekends “just because it was really important to me to do it.”

She felt personally responsible for creating more accessible images, to the point that she put in time outside of working hours. And she described if a similar situation happened in the future, "we've got to have these descriptions and I will write them even if I have to do it over two weekends. I think overall it's important to do this work." The Lone Writer process by its nature means that a high degree of creative burden, both in terms of time and in terms of responsibility, falls on one person.

Comparing W3's experience with the Lone Writer process to our experience with the User-Evaluation Process for the Avatar Project also suggests some potential tradeoffs of this method. W3 characterized the images as including a wide range of subjects, including real life locations, objects arranged in a room, and humans or pets engaged in activities intended to be relatable to users. This is very different from the detailed portraits of individuals or concepts included in the Avatar Project, which were intended specifically to encourage reflection on inclusion and identity. This may partially explain the significantly shorter amount of time–approximately one minute per image–that W3 took to create the alt text compared to the several months long and more than half-dozen iterations the User-Evaluation method required. This suggests that for complex images the Lone Writer process may not scale, or conversely, that for large sets of images changes would have to be made to the User-Evaluation process as we conducted it.

3.3.1.3 The Team Write-A-Thon Process

W1 and W2 both had experience with the second type of alt creation model that we're here calling The Team "Write-A-Thon" Model. The essence of the Team Write-A-Thon process is a group of alt text writers collaboratively developing alt text. W1 spoke to this model when she described a "fix-it" day she had organized for her co-workers to

participate in alt text creation. She describes "fix-its" as "an engineering team concept… where you get everyone together and you have a list of tasks and you just run through them all because that's the best and easiest way to get the job done that has not been figured into anyone's priorities yet. No surprise, that's always alt text." She said for the specific "fix-it" she ran she "got anyone, not just designers but anyone from our team, like forty people maybe, to sit down and describe the images and … then [co-worker] and I [would] then review it and add it to the CMS." The fix-it approach was described as one way that multiple alt text writers can have a chance to sit down together and work on accessibility as a group.

The Team Write-A-Thon Process, as a collaborative approach, has the potential to reveal how varied people's opinions are about writing image descriptions. As W1 described, "It was also a good experience if you want to see just how different two different people can think the alt text should be… someone is going to say, 'this is a mobile UI with this, this, this.' Someone else is gonna say like, 'a news app and five buttons.' And you're like, 'I dunno? Would that be the same thing?'" The subjective nature of alt text production becomes more obvious when, as in W1's case, there are many writers describing images with "a lot of repeated constructs" but the alt text produced is very different. W1 felt that the fix-it day resulted in understanding "a few different standardized ways of describing things… [Because] there's like three different ways of saying what really should be the same thing. So actually, consistency is a big chore across it all." Thus, the Team Write-A-Thon model may also create a level of consistency similar to the Lone Writer model, but through consensus rather than a single person's decision. It may also have the advantage of placing less of a burden of work on just one person.

Although getting agreement from multiple people may result in more clear and consistent understanding across the group, it can be very time consuming and laborious trying to reach that level of agreement. W2 describes some of the drawbacks of a collaborative approach: "When we're redesigning the [product] website, it's got a lot of diagrams and things like that. So, the question did become pretty quickly… 'I'm going to be spending the next half a year labeling 1000 images, or is there a better way we can spend our time?'… Now you've got 20 designers sitting in the room trying to say like 'Well, this is the right way to say it. Do I give enough information? Do [I] not give enough information?' You see now that becomes what I alluded to earlier: [a] philosophical exercise. Like UX writing exercises. I think, ideally, we should be aiming at the best we can, but there also comes a bottleneck that you need to be balancing priorities through." Essentially, while a collaborative approach between many writers "[done] in a very standard fashion" may ideally lead to a distributed load that "[does not] feel like 15 designers are creating the images" in practice it can lead to more time and effort being spent coordinating the writers and getting consensus from everyone.

3.3.1.4 The Artist-Writer Process

This brings me to the last alt text production process I am going to describe: The Artist-Writer Process. For the prior processes, I did not discuss the role of the artist because the role of the artist was always minimal, with a clear hand-off from the artist well before and completely separate from the alt text creation process. The Artist-Writer Process, in contrast, occurs when there is either direct collaboration between image creators and alt text creators, or when the image creator and the alt text creator are the same person. A key finding here was that who counts as an "image creator" is actually more

complicated than assumed. Two of the three artists we interviewed explained that their role as an artist was not simply to create an illustration, but to collaborate with a person writing the words accompanying the image. A2 said very clearly: "There is always a lot of mediation and collaboration between the art director, the author, the illustrator and the reader." A3 described "the traditional model" as being a collaboration between "the copywriter and the art-director." When an artist becomes part of an Artist-Writer process, they become a more involved co-creator of the final meaning of the alt text.

The Artist-Writer process has two variations. The first is having the image creator and the alt text writer be the same person. W1 explains this is, in her experience, rare: "[my colleague is] not like any other person who contributes imagery in that she does write the alt text herself." The more common practice in W1's team before the Team Write-A-Thon method was attempted, was the second version of the Artist-Writer process. Under this version of the model, a single author wrote the text of the webpage, made decisions behind the images, and "[it] was not a complete submission until it has alt text, designers must contribute to alt text." This meant there was still a single source of creative control over imagery and alt text, but that it took the form of one person overseeing the artist and writing the text, rather than actually making the imagery themselves.

Having a single author in charge of content, imagery, and alt text had drawbacks, however. W1 explained: "The goal… is that the person who is responsible for the information about a given page or topic would also write the alt text for the image, [because they are] the subject matter expert on it. But in practice it never works." The reason it didn't work, according to W1, was that general writers did not have specialized

knowledge of or experience with alt text: “I feel like the amount of training versus the frequency someone would do it… the overhead was too high. Like someone might actually only end up writing alt text once a year [so] they are not going to get good at it.” In theory, W1 felt, “[it’s] important education-wise, but they are not going to be a good contributor.” W1 knew “I’m going to have to rewrite it.” Thus, the Artist-Writer process had tradeoffs to either quality or conservation of effort.

However, there are still valuable elements to the Artist-Writer approach. For one thing, the educational element of distributing alt text production labor across a large team should not be dismissed. There is an implication that if a minimum alt text quality could be assured, it would avoid the time or manpower bottlenecks seen in the Lone Writer or the Team Write-A-Thon processes. Lastly, the unified authorial intent that the Artist-Writer process facilitated was closest to the understanding of “intent” we and end users had as our baseline understanding in Study 1. The artist, for example, may be able to include non-visual elements of an image in the alt text because they know these details in a way a writer simply interpreting the image may not. Thus, artists have more detail and certainty about what they were intending to portray that can be captured in the Artist-Writer process but not in other processes. The fact that users understood the artist’s intent as it functioned in the Artist-Writer Process may make it a process that could be easily integrated with elements of the User-Evaluation process because, as we’ll discuss in the next section, clear delineations between the processes are not always necessary or entirely accurate to the in-situ experience of them.

3.3.1.5 Deconstructing Delineations Between Types of Alt Text Creation

Despite labeling and describing these four alt text creation processes as distinct, in practice the delineations between different types of alt text creation processes were not always so clear cut. W3 explained even in the Lone Writer process, that she didn't work entirely alone: "There were a few images where it was a lot harder [to write alt text] … so I did set aside some of them [to give to a colleague] as like 'I'm having trouble… What feedback can you give me?'" Even, or perhaps especially, writers who take full responsibility for alt text don't work in complete isolation, making the distinction in practice between the Lone Writer process and the Team Write-A-Thon process not always easy to determine. Similarly, the User-Evaluation process, as we practiced it in Study 1, involved writing the iterations of the alt text not just in collaboration with the users, but in collaboration with different researchers on the team, making it reminiscent of the Team Write-A-Thon approach.

Additionally, even in methods that did not directly elicit user feedback–as is standard in the User-Evaluation process–artists and accessibility practitioners discussed prior knowledge and experiences with disabilities informing their approach to accessible imagery in a way that complicates the simple forward march of ideas from creators to end users. A1 explained: "Since my sister is autistic and is non-verbal, I was exposed early on to neurodiversity and different disabilities…. My uncle has a prosthetic leg so that was just a very normal sight as a child to play hide and seek and see a leg hanging around the house… These things really opened my eyes about the concept of accessibility for all different kinds of people as well as the constant daily challenges that exist everywhere because our society was built for neurotypical, able-bodied people." W3 similarly said that her personal

relationship to assistive technology made her more passionate about making imagery accessible: “I don’t have to use a lot of assistive features… but there are things that are helpful for me even though I don’t need them, like automatic captioning of YouTube videos… if I want to watch a YouTube video in bed without waking up my spouse, it's lovely to be able to actually do that without having to get out of bed and go find earbuds. [Accessibility features] could be helpful to everyone.” People with disabilities, even in more direct models where accessible imagery is not being created with direct consultation with users, still impact the way accessibility is done, because all actors in the imagery and alt text creation process are filtering information through their past experiences, including–for many–personal experiences with people with disabilities.

3.3.2 Further Themes from Inclusive and Accessible Imagery in Practice

In this section, I will highlight three additional themes from the interviews with image and alt text creators that ultimately shape the overall process of inclusive imagery creation in the studied company. First, I will describe the way artists and accessibility practitioners fulfilled the expectations users in Study 1 had by considering the intended use case when creating imagery and alt text. Next, I will describe how the alt text creation processes discussed above were impacted by the macro pragmatic and micro personal considerations of different actors, including the organization. Lastly, I describe how the inclusive imagery as it was done at Google had clear champions for the words and the images but no set role to champion the alt text, leaving a potentially fillable gap in future alt text creation processes.

3.3.2.1 Intended Use Cases and Identity Depictions in Inclusive Imagery

Just as users in Study 1 believed it ought to be, I found the intended use case for the images impacted the entire inclusive imagery and alt text production process. For example, W3 thought about the use case when deciding how or if to explicitly mention race in alt text: "We have an image of a person riding a motorcycle… [and the alt text just says] 'Person riding motorcycle.' You're just trying to keep it pretty generic…so that these images could appeal to anyone." Whether aracial images or image descriptions actually do appeal to anyone is a much larger question, but what is relevant is that W3 felt that because these images were to be user icons, she needed to keep the descriptions of people ambiguous.

The artists working on the Avatar Project likewise said the imagery's intended use as an inclusive tool was an important factor when deciding how to depict identity. A2 summarized the project as "Google asked me to create a series of avatars that represented diversity… the goal was simply to represent the widest possible spectrum of human beings" including an image of disability that emphasized "an extremely shareable and ordinary moment" where "the lack of limb [was] visible but not overwhelming." A1 similarly felt that they "didn't want our depictions to tokenize or degrade any person or specific disability, or flatten them to just that… it was the goal for each and every avatar to show a multi-dimensional story of the character… not centering a disability as a single identity." Despite the inclusive purpose, artists discussed wanting to keep a degree of intentional ambiguity. A1 explains that "we had specific ages/genders/races in mind, but the final product keeps a lot of them pretty ambiguous… Just because I assigned a specific gender to a character in my mind, doesn't mean that people will/should see it the same way." A3 felt that "any type of successful art of any kind has these layers… where it doesn't hit you over the head, but it

kind of suggests something and leaves a little bit of room [for interpretation]." Inclusive imagery, in the minds of artists, did not always require explicit representation of identities.

What was interesting was that, despite the role of artists in depicting identities, there was not a clearly defined path of intent from artist to artwork. A1, talking about themselves and the character designer they worked with, said: "I feel our art was adapting to the purpose/subject matter… This art was for a client, and it had a specific purpose to represent people inclusively." A2 similarly argued that "the job of an illustrator (at least in my case) is to represent an idea, a concept, a state of mind, that is usually some else's, the writer of the text." The pre-figured intent coming from the project sponsor therefore informed both the alt text and the artwork (see Figure 3.1).

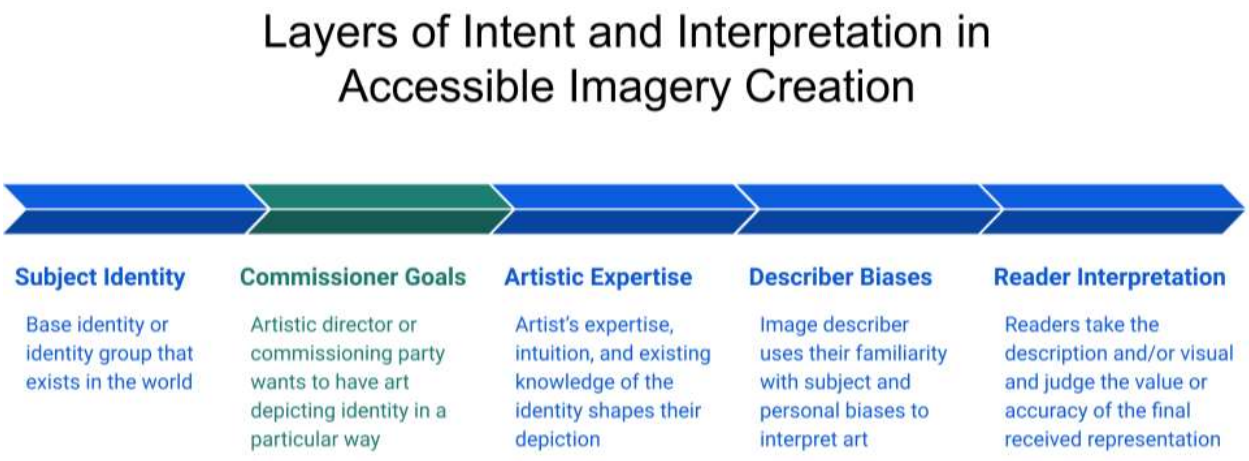

Figure 3.1 Steps in imagery creation process now including commissioner goals

3.3.2.2 Balancing Pragmatic Considerations Versus Personal Intuition and Expertise

Another tension I saw at play when artists and accessibility practitioners were creating accessible imagery was the need to balance pragmatic concerns of the organization and users, with the individual agency and expertise of the contributing image and alt text creators.

There were a host of pragmatic considerations that accessibility practitioners considered when deciding how to create alt text. W2 valued human generated alt text, but explained he would never be the person to say, "'Okay let's leave everything else and just work on the alt text,' because we've got 10 thousand pages full of images and we're going to be spending the next year labeling those images." W1 emphasized that pragmatically when getting new people up to speed on writing alt text, she tells them to focus not on descriptions themselves but on what about the image is "meaningful" or "most salient" which "[is] tricky because [what] you're asking… someone to do is say 'What is this image's role in the [webpage]? What is it there for?'" Thus, both the scale of a large company, and the complexity of producing useful alt text when writers are inexperienced affect the methods and approaches used to create alt text.

There are also personal levels of expertise that inform how individuals create inclusive and accessible imagery within the context of these larger constraints. For example, A3 explained that when they decide to move from researching a topic to creating the images themselves there is a certain amount of intuition that they have as an artist. "It's just a balance…I have to read about and at least sort of understand what's going on…, but I think it's just sort of intuitively knowing when… I've done enough of that." W1 has a similar level of expertise in her role as a UX writer. She explains that "[she] take[s] responsibility, happily, for the way language and image work in our product" and that in that role of responsibility, she understands that "oftentimes the secret is that someone made the image without a clear intent or without a purpose and that's when alt text becomes really hard." From both artists and writers their personal experience and expertise in their field allows them insight into how accessible and inclusive imagery needs to be produced and this

balances with the tangible time constraints and pragmatic scale considerations to create the final processes and deliverable user experiences.

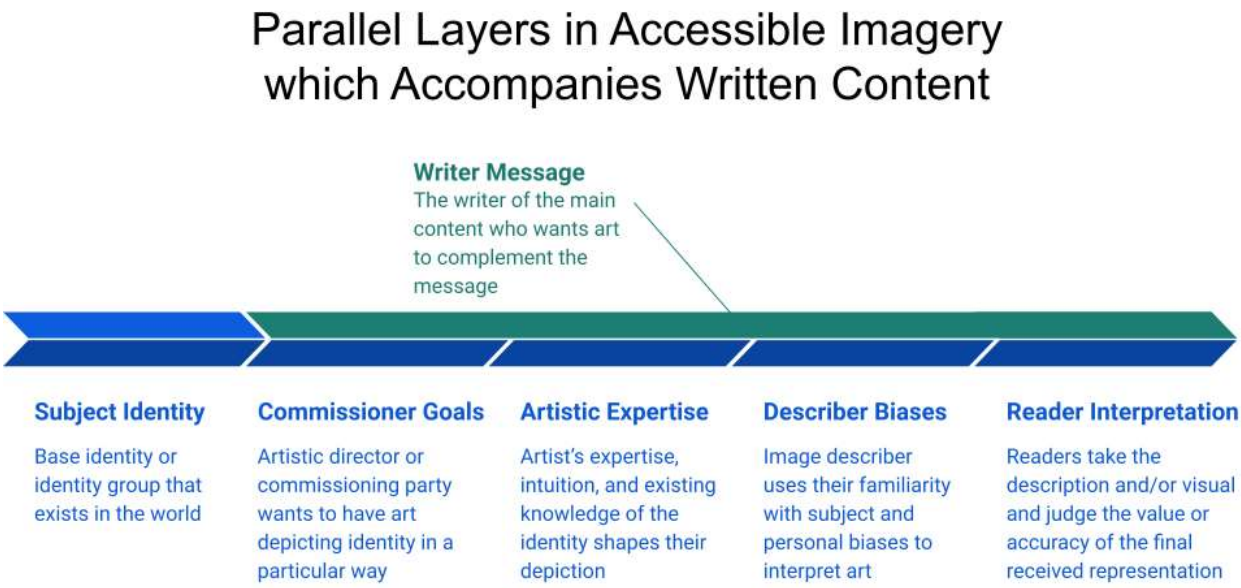

Figure 3.2 The steps of accessible imagery creation, now including the content writer's message

3.3.2.3 Understanding the Relative Importance of Words, Images, and Alt Text

Another element that influenced the approaches image creators and alt text creators took to creating accessible imagery was how each of the different players in the process valued words, images, and alt text. W1 explained that asking someone "to only do the alt text is like clearly missing a big part. Because like what is the surrounding? The adjacent text is a big part of that experience too." Describing the way images, surrounding text, and alt text work together, W1 argued that: "A visual example, even if it just reinforces the concept is probably still useful…[but] the adjacent text is probably already saying exactly what I would want [in] the alt text…[it is] this in-between thing where… this is not so essential to meaning, [but] it's not decorative." The parallel relationship between the written words, imagery, and alt text is shown in Figure 3.2. Although, as W1 shows, the written text, the imagery, and the alt text, are not—in practice—always clearly defined in

relation to each other. W2 even gives the example of describing one's appearance at a conference to add an additional layer to the puzzle: "Some people will say… 'oh you should describe yourself to other people'… I feel like if people hear my voice, I want them to experience just my voice without knowing how I look… There's a beauty in not knowing too much about [a] person." Thus, image description processes, in practice, are about balancing the overlapping information sources that are available to different users and considering how to best combine words, images, and image descriptions to achieve the desired result.

Perhaps unsurprisingly, in cases where images, words, and alt text are all working together, the perspective from which you look at the problem determines what is seen as the most important. For example, A1 explains "it took a lot of thought to figure out what disabilities would translate visually and the most clearly." W1 explained one area where she put considerable time and effort was perfecting the words and images as part of the overall experience of the product she was working on:

> I will say that something I'm proud of with this site, even with the compressed time… We tried to create identities… even though alt text treats the imagery as though it were arbitrary, It's not. We were very deliberate in choosing a few different people, different ages, different identity, and then also building out a little mini world for them too because once you see someone's phone screen it implies like a whole social circle…this is a more complex representation of identity than the previous generation of our products… there's an intentional little story there… But now I'm wondering… the alt text sphere has no idea we made all these choices, right?... Sometimes the same thing is true for... for all users, which is that I wish

> there were ways of scaling its complexity, [making versions which are] more sophisticated… because right now I can only treat it as this thing where I have to say as little as possible, as efficiently as possible.

Alt text was ultimately thought about last and least because there was not an expert on alt text or screen readers' experiences to weigh in. It never "figured into anyone's priorities" as W1 said earlier. In the overall production process, images and words both have champions, the artists and the UX writers, but the alt text is lacking the same kind of organizational champion role.

3.4 Takeaways

Following this overview of the findings from Study 2, I want to cover some of the larger takeaways from this data. First, I will explain how the study of industry alt text creation processes led to a more complex version of the alt text creation and reception framework originally presented at the end of Chapter 2. Secondly, I will particularly discuss what Study 1 and 2 together tell us about Study 2's second research question, regarding how elements are negotiated by different actors, including users. Then, I describe the implications these findings have for image accessibility research and practice, calling on researchers and practitioners to consider alt text creation as a part of the overall UX design process. Lastly, I'll cover reflections on and limitations of Study 2, and outline some potential future work.

3.4.1 Complicating the Framework for Alt Text Creation

The findings from Study 1 complicated the guidance suggested by Bennett et al.'s study, which focused on self-presentation in photography and argued that the subject of a photograph should be the final authority on how to describe their identities [18]. I found

that users saw fictionalized representations as defined by the intent of the image author who, alongside the users themselves, determined what the final received interpretation of disability identity was. My study of industry practices and inclusive imagery production processes complicated this further by exploring the more elaborate relationship that artists played in conveying intent through images.

My findings originally suggested at least three layers at work in the final reception of an image description. First, an artist renders a fictional interpretation of disability or any other social identity. Then, the image describer interprets the fictional subject into text. Finally, the reader of the description interprets the text. Our user interviews revealed that the identities of the "interpreters" at each stage of the translation layered on their own biases, assumptions, or perceived meanings. However, artists are not actually alone in creating a representation of disability or other social identities. In fact, at least in the case of the Avatar Project, the artists were collaborators attempting to stay faithful to the inclusive intent of the imagery project as defined by their client, Google. Character identities originated from a combination of that foundational intent and the artists' own personal experiences and inspirations. Even artists themselves did not always intend a specific identity to be depicted, but sometimes intentionally depicted figures ambiguously. And this imagery, no matter how well-conveyed or accurate, is often running in parallel to the main written content of the webpage, product, or app. Thus, the final reader of the alt text is basing their understanding of the representation on a gestalt of main written text and the layers of interpretation and intent at play in creating the accessible imagery (Figure 3.3, below).

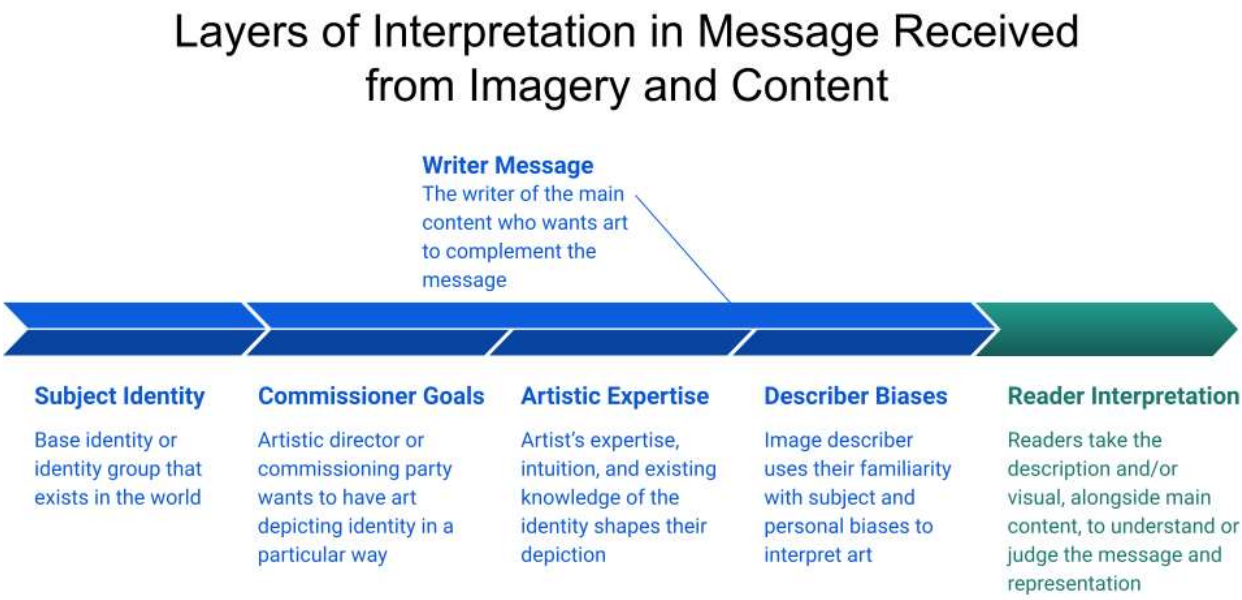

Figure 3.3 Final accessible imagery and content creation process, including reader's synthesis

3.4.2 Negotiations and Differing Perspectives Between Different Actors

Understanding the more complex framework for alt text creation I just discussed, makes it more clear how final meanings, at least in the Avatar Project images, are ultimately a negotiated process between many actors. The original commissioner or art director first thinks of a general idea or purpose for which images could be useful. Then there is a negotiation over the exact details of what images are to be created and which identities or characters to depict. This can involve the artist, character designers, art director, and potentially the writer of the content the images will accompany. The artist brings to these discussions their own prior knowledge about disability and other social identities, their particular points of inspiration or research they've done into the topic, and their stylistic preferences. Filtering this through the guiding purpose given by the commissioning party, the artist will end up first with a set of concepts that they agree on with the art director. In the Avatar Project image set discussed in Study 1, this took the

form of one-to-three-word descriptions of the characters that were passed onto us as researchers in the form of image file names. For example, the avatar of the fancily dressed woman with vitiligo, was originally saved with “Model” in the filename. The images are then iterated on with feedback from the commissioner until they are all finalized.

The image describer, coming after the artist, may–like we did–have access to some part of the image creator team, either the commissioner or the art director, the artist or simply text explaining the purpose of the image. From here, the writers have to consider how to best convey the many visual elements (as U19* mentioned in Study 1, for sighted people “visually there’s probably so much information that you guys receive”) in a brief piece of alt text. This consideration may, depending on the process being used, involve direct feedback and prioritization of that feedback from users. It may also involve negotiating with co-workers over particular wording, the degree to which the purpose has been achieved, and context the images are likely to be seen in. Through all of this, an image description is produced that, hopefully, in some way conveys the same information, experience, or emotion as the image.

However, the final understanding is not produced until a user encounters it. At which point the user will also be considering how the image description’s meaning exists within the context of their own experience, the surrounding text, the caption, and any visual information they may have access to (if they are sighted or have low vision). This leaves users to interpret all the information—as carefully or briefly as their personal preferences and interest dictate—and decide on their final received meaning of the alt text.

This includes to what degree they feel the identities of the image subjects are depicted inclusively.

Ultimately, the takeaway here is that in practice there can be many stakeholders, some not clearly delineated by the artist, writer, or user roles. The final, received impression of the alt text can never be guaranteed until it finds its home in the mind of a user. Yet, concretely, every day, many people across the digital world create alt text that functions to some degree, even if it is not received perfectly. Just as with many processes of the world–such as, for example, the human body–an almost bafflingly complex system nonetheless continues to function. And like the human body, the degree of function ranges widely, with many options existing between perfect transmission of electrical signals and complete stoppage of the process at all.

3.4.3 Implications for Image Accessibility Research and Practice

Beyond strengthening the model of alt text creation that I proposed based only on Study 1, the analysis of alt text production at Google points towards some other interesting implications for image accessibility as a field of study and practice. The four creation processes we described highlight different areas where improvement in image accessibility is warranted.

3.4.3.1 Platforms Needed for Direct User Feedback on Alt Text

The wide support among accessibility practitioners and artists for the User-Evaluation method demonstrates that organizational stakeholders care about and want to prioritize user feedback. Alt text writers we spoke to understood that the text should respond to the intended use case of the images but didn't necessarily have the data to assure them that their intended message was being received. This is because, despite buy-

in, it's not always possible to choose an alt text creation method based only on if it creates the best quality results. Instead, issues of scale, time needed to get new alt text writers up to speed, and incompatible deadlines all end up constraining what methods are practically available for any given project. Thus, one potential route for either a research or practical contribution to image accessibility would be the creation of a central user research platform or other screen reader user feedback system. If there was a way to practically get feedback from screen reader users on the scale and with the short turn-around time that large companies like Google need, the quality of the alt text created by these large companies would no doubt leap forward.

3.4.3.2 Alt Text Creation is Part of UX Design, and Needs Similar Support

The Lone Writer Process demonstrates that alt text writing in industry can sometimes require high labor burdens be placed on individual workers, particularly those that personally care or feel responsible for accessibility being included in the final products. While the dedication of these workers is admirable, and potentially has some advantages in terms of consistency and turn-around time, this practice shows there is more work to be done on an organizational level. The Lone Writer process, as well as the dedicated advocates for images and user text that we discussed in 5.2.3, shows that alt text production is actually part of UX design and needs to be supported as an area of UX design. This could mean dedicated job duties or even job titles focused on alt text or other accessibility design. Accessibility champions have a role in industry, both as the experts who new practitioners can turn to when they are tasked with creating alt text despite not having experience, or as the people putting in the effort themselves to give alt text the fine-grained attention that imagery gets from art directors and UI text gets from UX Writers.

3.4.3.3 As Part of UX Design, Alt Text Creation Follows UX Principles

The Team Write-A-Thon process, as I described it here, points toward the fact that image accessibility research does not yet consider consistency to the degree it should. While consistency of alt text within the same page or website is discussed by WCAG [1], large organizations face the issue of assuring not simply that individual images and pages are described accurately or consistently. Instead, there's the additional challenge of making sure all team members or contributors understand and follow the same guidelines across all product teams, interconnected websites, and brand experiences. Writing alt text is unavoidably subjective, and leaning into the advice that there is no single "right" way to write alt text can be valuable for individuals. But when an organization or team is trying to create a single cohesive user experience across their entire brand, a wholly different level of consensus and clarity of guidelines is needed to assure that repeated elements or figures are described with the same wording regardless of how many writers are involved in creating the alt text.

Consistency is part of the established canon of UX principles for designers to follow [146]. Thus, considering what other UX principles can apply helpfully to alt text creation is an important future avenue for researchers and practitioners. Current alt text guidelines usually attempt to be generalizable, functioning as writing guides, but the Team Write-A-Thon process demonstrates that ultra-targeted guidelines for a single organization, team, or collective can be almost if not more valuable in practice. Or, alternately, considering how alt text guidelines could be rephrased or restructured to be followable in the same way UX guidelines are, may be valuable.

3.4.3.4 Just like UX Design, Iterating Cross-functionally May Support Alt Text Creation

The Artist-Writer Process, as it was described at Google, also demonstrates that at least in some contexts there can be additional stakeholders, layers of intent, or overarching frameworks that are invisible to users, but which nonetheless impact the finished product. Users, as demonstrated in Study 1, found it easy to locate intent in the goals of the artists. This works well in one version of the Artist-Writer process, where the two roles are actually the same person, but in organizational practice a project goal may come well before any individual artists are involved. This potentially means that image accessibility can likewise be built into imagery projects from their earliest points of conception. There is an opportunity for final alt text to be created through an iterative back-and-forth between art director, writers, and artists, to the same degree and in parallel to the creation of the art itself.

3.4.3.5 Guidelines and Automated Tools Needed to Streamline "Alt Text Experience Design"

Ultimately, our data only shows a tiny piece of the puzzle. All our participants came from the same organization, and even then, the different methods we identified were not always clearly delineated from each other. The complexity we saw just in this small sample of alt text creation methods in practice, demonstrates the way image accessibility can be anything from a negotiation between many actors, to a sole worker blazing through 800 images over the weekend. The way guidelines exist now, as sets of vague tips and rules directed towards writers rather than designers, cannot properly serve the needs of organizations that need to make many images accessible. The complexity of alt text creation, as a form of UX design that needs to be reactive to small differences in goals, context, constraints, team makeup, timelines, etc. leads to a new question. How could we

design guidelines, or even potentially create automated or semi-automated tools to truly standardize the experience of creating alt text in the manifold ways a large tech company requires? How can we as practitioners and researchers take the staggering complexity of tech industry image accessibility practice, and turn it into an asset for better user experience for alt text consumers?

3.4.4 Reflections on Strengths and Limitations of Study 2

There are several clear limitations to this study. First and foremost, it does not claim and is not intended to be comprehensive. I spoke to a limited selection of artists and accessibility practitioners, with only White accessibility practitioners and Asian-American or White artists represented. I spoke to a total of six professionals involved in accessible imagery production. There is no way, from that, to draw conclusions about all types of accessible imagery processes at this large company, much less at other companies. However, I believe the only way to begin charting the complexity of industry accessibility practices is to start somewhere and build on it.

The data for Study 2 also focused largely on one project and only used direct reported practice as the basis for my understanding. While reported practice is an accepted approach to understanding accessibility practitioners [8], it remains to be seen if a full embedded ethnographic study of accessibility practices in industry would reveal more nuanced understandings of processes and where actual processes differ from reported ones. Focusing on one project was important for depth of information and familiarity, but academic research cycles are deeply out of sync with industrial practices and as user interviews, publications, and other scheduling conflicts took more and more time, it became increasingly difficult to collect data from previously highly engaged stakeholders. I

want to emphasize, therefore, that tech companies' research teams should and likely are doing similar or better work examining industry practices in a way that keeps pace with companies' rhythms. I believe academics could stand to learn more from industrial research practice that is, unfortunately, more often than not hidden behind closed doors due to liability or business concerns.

In reflecting on these limitations, I feel some of them were functionally unavoidable. Recruiting more–and more diverse–participants from inside Google would have been hugely beneficial, but getting access to Googlers was entirely dependent on word-of-mouth between our direct collaborators and people on their team. There were the additional constraints of finding time to speak to Googlers within their busy schedules and without pressuring them to discuss details they weren't comfortable with. I don't fundamentally know how the number or depth of these interviews could have been improved without substantially more buy-in from multiple teams or Google leadership. Even with the interviews I conducted, there was information that I had hoped to include in the publication that ended up being removed as part of the editing and Google review process, which is disappointing to me.

However, despite that, I find the strengths of this study to lie in its uniqueness and practical challenges. This study is, to my knowledge, the very first study focusing on different processes of alt text writing, particularly in the daily practice of accessibility practitioners working in industry. As a research area the study of alt text production practices is only just starting to ramp up (e.g., [255]). I hope that studies such as mine are picked up by the community for the valuable insights into both industry accessibility and

alt text that they provide. If they are, perhaps future studies can ask similar research questions and start finding out the wide variety of possible answers they will receive.

3.4.5 Future Work

This study suggests several avenues for future work. The most obvious of these is more work on industrial accessibility practices (particularly image accessibility) as this research only began to scratch the surface of image accessibility complexity in practice. While this research took the form of focusing on practitioners at a single company (so as to begin understanding how the inclusive and accessible imagery process worked from start to finish) it may be valuable to approach the problem from another way. For example, surveys have been successfully used in the past to gather data from a wide sample of practitioners on their accessibility knowledge and experience (e.g., [69,118,260]). Potentially, distributing an anonymous survey to UX Writers asking about their role in and understanding of image accessibility practices at work would avoid some of the hesitance around disclosure that might have been present when connecting to participants through their employing company.

However, another less obvious but equally important direction for research coming out of this study is to begin to examine how non-industry practice of writing image descriptions work. While there is beginning to be work on the process computing professionals follow to write alt text [255], there are so many other groups of people who write alt text. For example, as users in Study 1 pointed out, most often the people championing image descriptions on social media are disabled activists. Similarly, work shows that direct familiarity with disabilities (such as having a friend, partner, or family member who is blind, in this case) leads non-disabled users to consider accessibility more

[214]. Logically it follows that there is probably a fairly significant group of users on most social platforms that create alt text for their posts because they know someone who needs that accessibility accommodation. Studying how these different groups write alt text–including their process, their means of determining quality, and any barriers they face in their alt text writing practices–can likely provide alternate versions of the alt text creation framework as I present it here.

In fact, the research I chose to pursue following this study follows that general path. During the process of collecting data, writing up submissions, and publishing Studies 1 and 2, I noticed increasing activity and discussion around image descriptions on Tumblr. To supplement the understanding of industry practices discussed in this chapter, I wanted to see how everyday users of a digital platform went about writing alt text. Upon identifying that there was in fact a growing community on Tumblr and Discord that worked together to generate and share accessible content, I saw that the deeper embedded ethnography that I was not able to do at Google was possible in this online accessibility community. I will discuss the findings from this digital ethnography of an online accessibility community in the next chapter.

CHAPTER 4: Study 3 – Digital Ethnography of an Online Accessibility Community

This chapter will focus on the third and final study of my dissertation. This study, building on the prior study's focus on alt text creation methods used in an industry context, looked at the question: How are image descriptions created in the context of an online accessibility community? In this case the online accessibility community was the People's Descriptions Discord Server (PDDS) which I will introduce below. The digital ethnography of PDDS included two other research questions:

1. How do participants in an online accessibility community describe the benefits, motivations, and challenges of creating community-generated alt text?
2. What kinds of relationships exist between the roles of image artist, alt text writer, and alt text reader in the context of image descriptions created by an online accessibility community?

In this chapter, I will describe the relevant literature and methods used for this study, and then report my findings on each of the three research questions above. In reporting these results, I will interweave data collected across the four phases of the study, including a mixture of quantitative results, themes, and findings from the interviews with PDDS members, and examples or patterns from my observations and field notes on Tumblr and PDDS specifically. I will save the discussion of takeaways from this capstone study for the next chapter.

4.1 Background

Before introducing Discord as a platform and the PDDS in particular, I would like to introduce Tumblr. Describing Tumblr as a platform provides necessary context for

understanding the PDDS and is directly relevant to the image descriptions made by PDDS members.

4.1.1 Introduction to Tumblr

Tumblr is a microblogging platform which combines aspects of other popular social media platforms such as Twitter and Facebook with more traditional "blogging" features seen on sites such as WordPress. Active since 2007, Tumblr as a platform was considered culturally and technologically relevant roughly from 2010 to 2015 [239]. It was known as a platform for younger users, particularly those interested in social justice or media [160]. In 2013, when Tumblr's popularity was at its peak, Yahoo purchased the platform for $1.1 billion dollars [239]. This marked the beginning of the end, by some users' reckonings, as mismanagement and a misunderstanding of Tumblr's unique strengths led to a ban on adult content in December 2018 [48]. Before this, Tumblr had been one of the few places that still welcomed adult content, but after the Tumblr app was suddenly and unexpectedly banned from Apple's App Store, a highly criticized automated system was quickly implemented to catch porn and nudity [158]. The announcement, including its highly clumsy aping of inclusive language that explained the ban affected "female-presenting nipples," caused huge pushback from Tumblr's community, including a mass exodus of many users to other platforms like Twitter [70]. Visits to Tumblr's desktop site and mobile apps dropped over 40% from October 2018 to October 2021 [239]. Bought by Automattic (the company behind WordPress) in 2019 for only $3 million, Tumblr has been struggling to break even, much less be profitable, for years [179].

Yet, Tumblr remains. It apparently grew its revenue by 55% from mid-2021 to early 2022 [48] and downloads of Tumblr's app from Apple's App Store went up by 62% the

week after Elon Musk officially took hold of Twitter [116]. As of writing this in April 2023, Tumblr reports having over 573 million blogs, including 57,660 which were "made yesterday alone" along with over 9 million posts made that same day [225]. And beyond what is captured in numbers, Tumblr has retained a thriving culture–the kind that made headlines in November 2022 for generating an elaborate shared fantasy around a fake 1973 Martin Scorsese film over the course of a weekend [53,128,233].

In the next section I'm going to first introduce some terminology that is necessary for understanding how Tumblr functions, then demonstrate the terminology in use by breaking down the elements of a Tumblr post, before discussing the content of the anonymized but real post that appears in the labeled screenshot (Figure 4.2).

4.1.1.1 Terminology and Platform Structure

First, when a user signs up for a Tumblr account they choose a username and an icon, and they obtain a free blog page with its own unique URL based on that username (i.e., https://username.tumblr.com) [70]. All posts a user makes appear in reverse chronological order (most recent first) on their blog, for others to browse either sequentially or filtered by tag. Users can also follow other Tumblr blogs and then see all the posts from blogs they follow together in reverse chronological order on the main Tumblr feed, the user's dashboard. These posts can take many forms, with text, photo, video, link, and audio posts all being common, along with many posts having a combination of these elements.

The primary function Tumblr users perform, besides posting, is reblogging, which is similar to Twitter's 'retweet' feature, where an attributed version of the post (a reblog) from another user's Tumblr will appear on the users' blog and will show up on their followers'

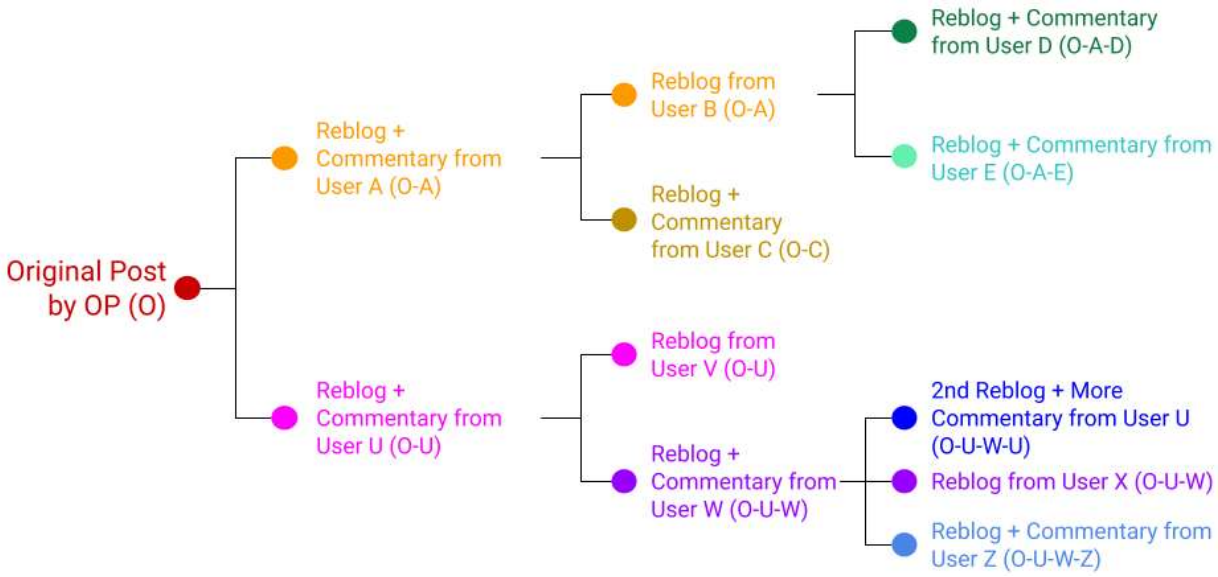

Figure 4.1 A visualization of Tumblr's branching post and reblog structure

dashboards. However, unlike retweeting, a reblog shows not only the original poster (OP), but the intermediary blog that the user reblogged from, creating a chain of reblogs that can theoretically be traced back through every step in the reblog chain until you return to the original post on the OP's blog. Additionally, similar to 'quote tweets' on Twitter, a reblog can contain additional content added on by that user in that particular instance of the post.

This results in every post on Tumblr taking on a branching tree structure (see Figure 4.1), where every reblog that includes commentary becomes its own measurably different version of the post. Users will only see one form of the post when they reblog (unless, for example, if User V follows User U and User A, in which case they may see two different versions of the post on their dashboards). If users reblog without adding anything to the body of the reblog, then it does not retain any evidence of their reblogging (outside of the notes, which track who reblogs, likes, and replies to a post) and is functionally the same post. Additionally, any commentary in a reblog is only added to that specific version

of the post, meaning you can reblog a post with commentary from User U, User W, and User Z, but you could not find a version of the post with commentary from both User U, User W and User C. This results in, among other things, many versions of a post proliferating (in Figure 4.1, including the original post (O), there are 10 different versions of the post circulating and rebloggable at the same time).

To give you a basic grounding on what Tumblr looks like to users and an example of accessibility in action, Figure 4.2 shows a simplified version of my own dashboard when logged into my @researcher-emory Tumblr account (more details in Section 4.3). If you examine this screenshot, you see the menu bar at the top, including the search bar on the top left and various icons in the top right, including the inbox and messages buttons, which I will describe in more detail at the end of this section. Moving down the screenshot, you see my icon displayed next to the interface used to create new posts, which is divided into the types of posts a user can make. Next are four tabs to see different versions of your dashboard. This is a relatively new feature and does not come up at all in Study 3.

Finally, you see the first post on my dashboard, a reblog from one of the community members I follow with a comic character in their icon (username colored in purple). This person I follow reblogged the post from another user called @accessible-art. All this information is displayed in the first line of the post. Moving down, there is the original post, a piece of art uploaded by a user with an Asian woman in her icon. In this case, the original

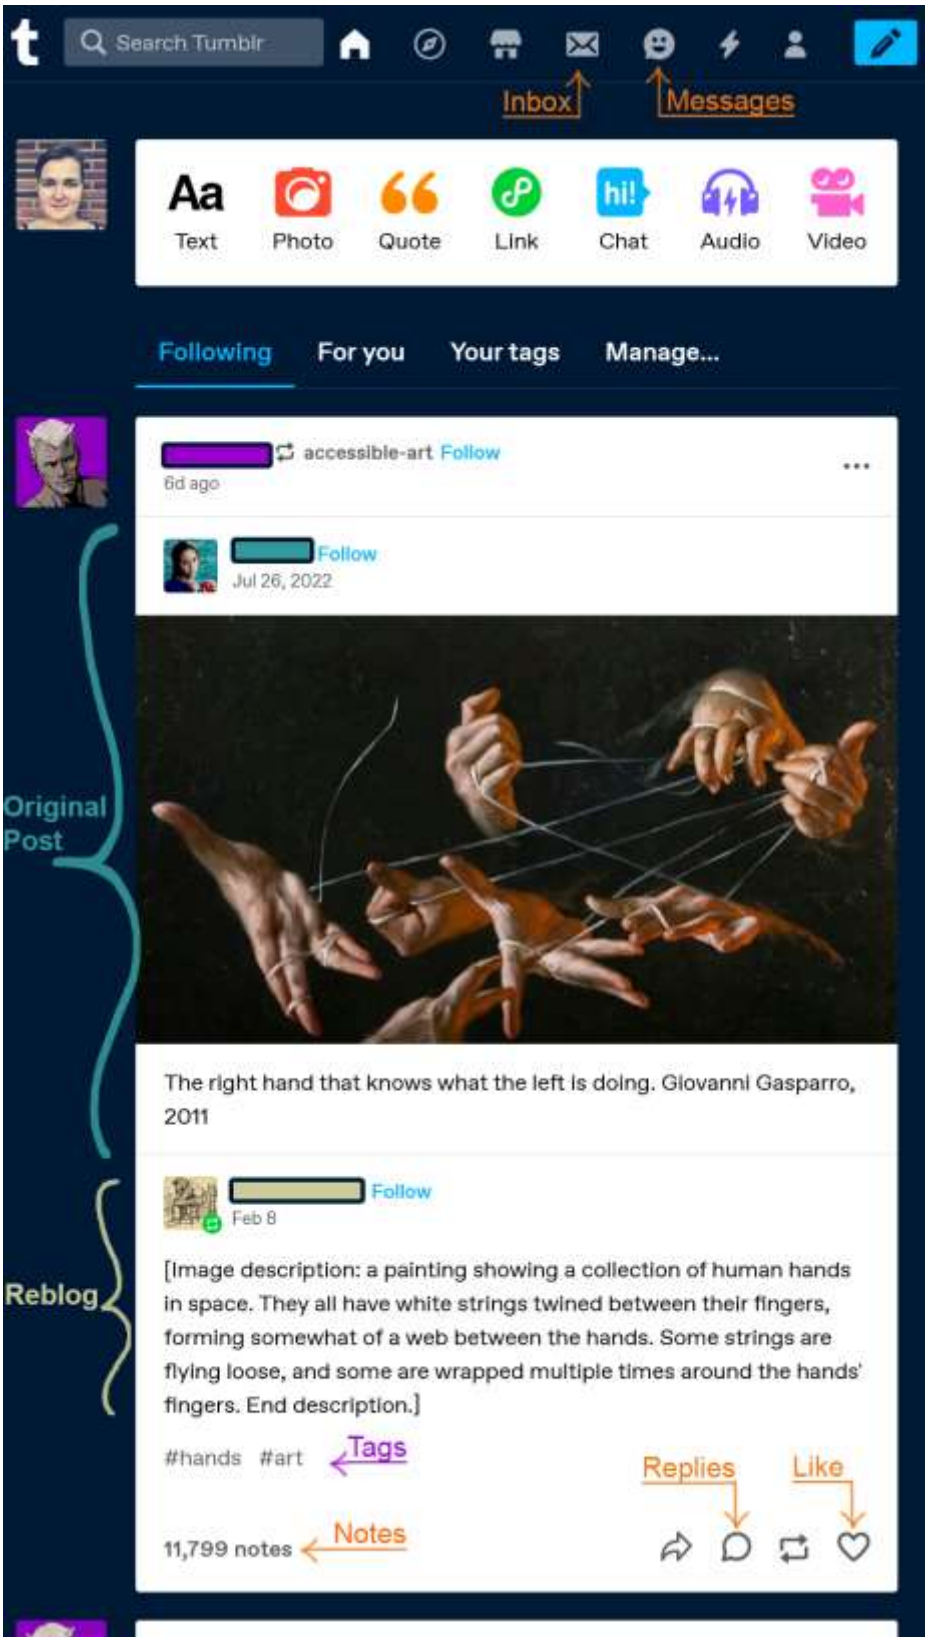

Figure 4.2 An anonymized version of Tumblr's Dashboard, with labels for key components

post has a small amount of text following the image, functioning as a caption and giving attribution for the art. There is then a line which separates the original post from a reblog by a user with a sketched icon who took the time to add an image description. Presumably, this version of the post was chosen to be reblogged by the @accessible-art account, for this reason. However, @accessible-art did not add any commentary so their

username does not appear anywhere in the body of the post. Below the reblog are tags from the user I follow (with the comic character in their icon).

Tags, although not demonstrated in their full glory here, are a step beyond Twitter's hashtag feature. While you can search and filter by them in the same way, Tumblr tags allow for spaces and any punctuation other than a comma (which is automatically interpreted as a tag delimiter) and a pair of quotes (where the content inside the quotes will be put into its own tag at the beginning of the list). This element of Tumblr's tagging system, along with the fact that tags do not get retained in downstream reblogs, has allowed it to become the de facto space for sharing one's thoughts and comments if they are directed solely to one's own followers. While recent updates to Tumblr have made tags viewable by anyone browsing the notes (i.e., a list of any sort of interaction with a post, including replying, reblogging, or liking) of a post, historically tags were considered a space for "asides" to the user's followers [29].

Following the most recent user's tags are several standard portions of any post, including a count of the notes, which can be clicked for more details (see Figure 4.3) and various ways to interact with the post on the bottom right. You can also see the beginning of the next post in the feed, which allows users to scroll down and see increasingly old posts until they either stop scrolling or reach the first post they had seen at the start of their last scrolling session on Tumblr. I will now describe in a little more detail the features I've labeled in the menu bar and at the bottom of the post. These buttons function as other ways Tumblr users can interact with each other, beyond reblogging and tagging or commenting in a reblog. Replies to posts are short messages that are included in the notes

of the post (across all versions). They traditionally result in a notification being sent to the OP and the person who reblogged the version of the comment that received the reply. To like a post, users click the heart symbol next to the reply button. Through the menu, users can send messages to other people, which are similar to a DM on Twitter. Lastly, users can send asks to users' inbox. Asks are short messages that can be answered by the recipient in the form of a post.

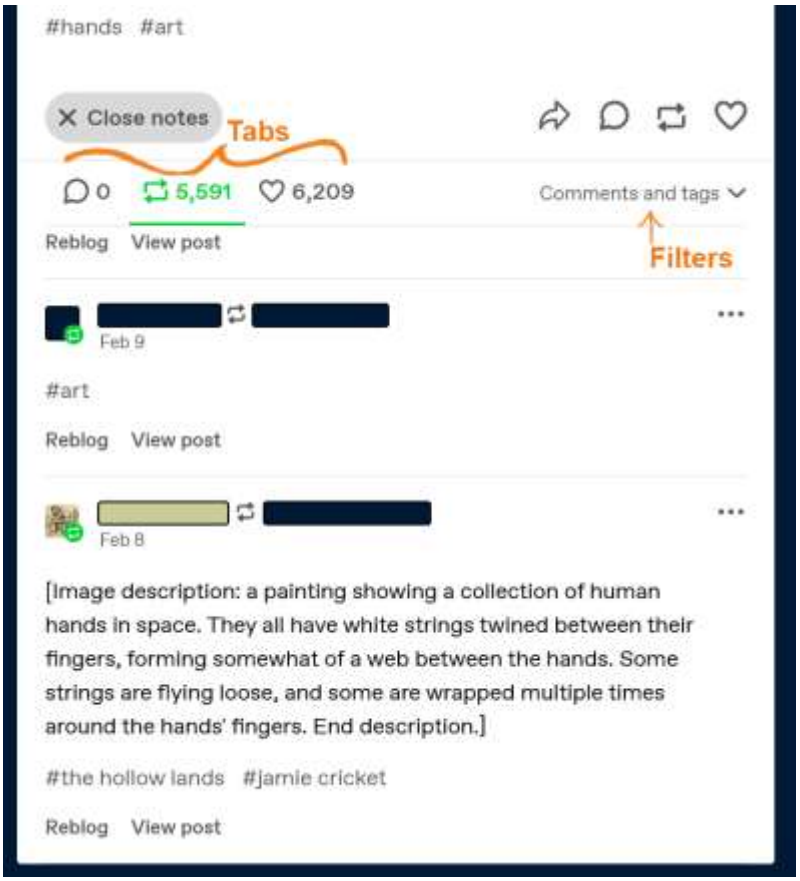

Figure 4.3 Zoomed-in view of Notes interface, with elements labelled

Additionally, when making or reblogging a post, users also are given the option to publish the post at that moment, save the post to drafts (a list of all posts deemed not yet ready to be published), to queue a post (send it to a list of posts set to be published at user-defined intervals until the queue is empty), or schedule the post to publish at a specific later date or time (which also sends it to user's queue but only to be published at the defined time). When making, reblogging, drafting, or queueing a post, Tumblr users also have an option to complete this action via their main blog, or any of their sideblogs (additional blogs run by the same user, who still has a single dashboard, but additional separate blog pages for each of their additional URLs). For the last bit of terminology, if two

users mutually follow each other, they may refer to each other (and collectively everyone they are mutually following) affectionately as mutuals [70].

Before I go on, I want to give slightly more detail about notes and the interface to view them, as they currently exist on Tumblr. Currently, if you click on the button on the bottom left hand side of a post which shows the number of notes (see Figure 4.2), additional content will drop down (displayed in Figure 4.3), allowing users to click one of the three tabs to either see all the replies to the post (in the case of Figure 4.3, none), or scroll through all the reblogs (this is the currently active tab in 4.3), or see a list of the usernames of everyone who liked the post. Within the reblog tab, reblogs can be filtered to only show those with comments, to show those with comments or tags (as in Figure 4.3) or "other reblogs" (i.e., those without tags or comments). One can then scroll down to see all the notes within that category, although depending on how many notes there are, the list may have to load multiple times as you scroll.

4.1.2 Introduction to Discord and the People's Descriptions Discord Server

Discord is "an online, persistent group-chatting application" where users can communicate directly with one another via text, voice, or video [108]. Or, more commonly, users create and organize themselves around servers, i.e., multimodal virtual chatrooms that can be anything from invite-only groups of close friends to publicly shared communities of thousands or even millions of strangers coming together to share information on topics such as AI image generators [269]. Originally known primarily as a home for gaming communities [190], Discord has 150 million monthly active users as of 2023 [270] and is an increasingly popular platform for communities of all kinds to organize. Anyone can make a Discord server, with individual servers being divided into

voice/video or text channels, i.e., separate "rooms" for ongoing conversations around a specific topic (or any other organizing principle the server owner desires). In addition to defining the channels, the server owner has the ability to deputize moderators and create other customized roles that determine what types of users can do what on the server.

The People's Descriptions Discord Server (PDDS) is a Discord server created to collaboratively describe or generally make accessible any social media posts (primarily from Tumblr) shared by the users. The PDDS creator has publicly posted several links to join the Discord on their Tumblr blog to advertise the community and let Tumblr users know it is a resource for people trying to learn about or practice accessibility. The PDDS is made up exclusively of text channels, which cover topics such as rules, resources, questions, off-topic chatting, and different types of accessibility requests that users can make (image descriptions, video descriptions, audio transcripts, or translation). My study focused primarily on the image describing activities occurring on the PDDS, which generally take the form of a member of the server posting an image or link to a Tumblr post and asking if anyone can help write a description for the images. Image description requests can also take the form of posting a draft description or general query and asking for other members to weigh in on word choice, accuracy, etc..

The specifics of how the PDDS functions and how it came to be in the context of Tumblr's larger ecosystem is the topic of my third study. For now, I will say that all PDDS members are volunteers, who donate their time and energy to collaboratively make Tumblr a more accessible place. The community therefore falls under the volunteer-based crowdsourcing category of communities discussed earlier. Additionally, because members

are actively producing image descriptions that are valuable to outsiders, they can be considered a peer production community. Members are often active in fan communities on Tumblr and fandom content is commonly included in the posts being made accessible. The PDDS is an accessibility community that also functions as a community of practice and exists at the interesting locus of the communities discussed in this chapter. The PDDS thereby brings to the forefront labor issues around crowdsourcing, questions about distribution and integration of disabled users within social media communities, and new potential alt text production processes that can address the issue of quality versus quantity with human-generated methods. Embedded in the context of Tumblr as a platform and functioning through the medium of Discord, my work considers how the complicated social circumstances of the PDDS enabled the community to develop and allowed members to effect new types of accessibility outcomes than previous collaborative approaches to accessibility.

4.2 Related Work

4.2.1 Communities of Practice and Social Media

In this section, I will return to the topic of communities of practice (discussed in Chapter 3) but will specifically cover online fan communities and communities on social media. Some of the types of communities I will discuss in this section are communities of practice, fandoms, peer production communities, and volunteer and paid crowdsourcing communities. These terms are not mutually exclusive, but instead have significant overlap. For example, the PDDS introduced above may be interpreted through multiple lenses: as a peer production community, as volunteer-based crowdsourcing, and as a community of practice.

4.2.1.1 More Communities of Practice: Fandom

In addition to being grouped around organizations or disciplinary boundaries (such as accessibility practitioners described in Chapter 3), communities of practice can be social groups (e.g., [91,112,192]) as long as they otherwise meet the definition outlined by Wenger-Trayner. These social groups can be based in physical or virtual space, as long as they have the three qualities of a community of practice listed above: a domain of interest, a community space, and shared practices. Some social groups that have been identified as meeting the definition of a community of practice are those within fan spaces [82,112].

Fan communities, also referred to as "fandoms," have been defined as groups made up of people (i.e., fans) who "shared common interests and confronted common problems" [121] in order to "develop networks and institutions and create a common culture" [68] based around "the regular emotionally involved consumption of a given popular narrative

or text" [57]. Embedded in this composite definition of fandom is a domain (the popular narrative of text), a community (fans form networks and institutions), and a practice (confronting problems and creating culture around shared consumption) that marks it as a community of practice. While not universally accepted, some scholars argue fandoms can be "a site of ideological and cultural resistance to the heteronormative and patriarchal values often shaping mass media" [122]. Fan activism is a common topic of research [21,39,133], and learning and sharing new ways to practice their political beliefs is a feature of many fan communities. Although not extensively studied, there has been research on the overlap particularly between fandom and disability, particularly how some fan spaces can practice an "activism of care" to "destigmatize mental illness and celebrate neurodivergent participants" [143].

Another common, although not required, feature of fan communities is their interwoven nature with social media. Fandoms existed before the internet [121] but since the creation of the web, as well as social media specifically, fandoms have flourished on a variety of platforms from Usenet to Tumblr [81]. Tumblr in particular has been a documented home for many fans and fannish communities for nearly a decade [110,111]. I already covered Tumblr in the Background section of this chapter, but I will additionally say that as of writing this, Tumblr's official iOS mobile app is listed on the App Store as "Tumblr – Fandom, Art, Chaos" and describes itself as "a place to connect with others over shared interests" [266]. Particularly given changes over the past year resulting from Elon Musk's purchase of Twitter, numerous popular press articles suggest that Tumblr is undergoing a resurgence and that in 2022 and beyond it may prove to be the site most welcoming of digital society's fringe groups [48,83,116,235,239]. Young activists, fanatical

media consumers, and disability communities may find Tumblr to be their preferred site of communal practice compared to many other social media platforms.

4.2.1.2 Peer Production Communities

In addition to work that has looked at specific fandoms as communities of practice [82,112,188], Wikipedia–or "wikis" generally as a type of knowledge production community–is another type of community of practice. Wiki communities, as producers of shared knowledge, also fall under another category: peer production communities [130]. Peer production communities are fairly similar to communities of practice, in that they are defined by coordinated work done by many individuals either toward a common goal or on a common platform [199]. However, as the name implies, a key element to decide if something fits under the peer production label is whether the members produce something such as "information, knowledge, or cultural goods" [16]. While communities of practice definitionally produce "a shared repertoire of resources: experiences, stories, tools, ways of addressing recurring problems" [78], these shared understandings may be of little to no value to those outside the community.

Such may be the case with some fandom artifacts for example: does a fandom-specific meme constitute a cultural good or piece of knowledge if it never breaches the boundaries of the fandom itself? Wikis, on the other hand, definitionally produce a resource for others outside of the community to refer to and use. The existence of fandom wikis, along with fan communities that create widely consumable content such as fanfiction [79], shows that there is overlap in many of these community types, which I have diagrammed (see Figure 4.5). Some fan communities are most accurately described as peer production communities, but not all of them are. But all fandoms, I would argue, count as communities

of practice because, as described earlier, fandoms by definition include the three requisite elements: domain, community, and shared practice.

Peer production can take the form of volunteer-based labor, where large numbers of people collaborate to create or accomplish something as a force of free crowdsourced labor [132]. It can also take the form of paid crowd work where a group of individuals are paid to complete many small tasks that in combination generate valuable data [4]. One fundamental consideration surfaced by researchers is the degree to which peer production participants fundamentally control or shape the platform they are contributing to [199]. For example, Amazon's Mechanical Turk is a crowd labor platform where all of the peer production is directed towards and controlled by a corporate sponsor [119]. This leads to ethical and labor issues related to crowd work [119], which are an ongoing area of research in HCI [221,240].

Even in volunteer-based crowdsourcing, labor issues are not altogether ameliorated. As studies of academic labor show, beyond the bounds of traditional paid responsibilities, there are many forms of invisible or assumed labor that are technically "volunteer", but which nonetheless require time and energy to maintain [138,248]. Many companies or

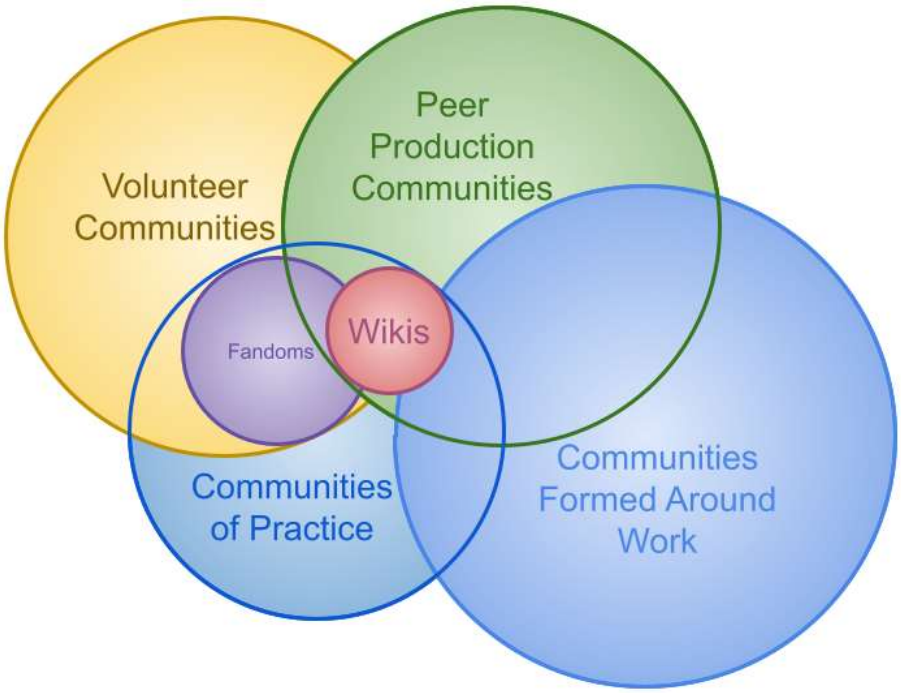

Figure 4.4 A Venn Diagram depicting overlaps between different types of communities

institutions directly or indirectly profit from the invisible or uncompensated labor of workers [65], meaning even volunteer-based crowdsourcing and peer production systems may present an unfair labor burden on participants. Additionally, the burden of crowd work and some types of invisible labor are disproportionately placed on people with disabilities.

4.2.2 Disability Studies and Communities

In this section, I will cover disability theory and communities, particularly online communities. Studies of disability communities provide the basis for my definition of the PDDS as an online accessibility community.

4.2.2.1 Disability Studies and AC

Beyond the design and research approaches already outlined in Chapter 2, Accessible Computing locates another key source for theoretical or methodological lenses in disability studies. Disability studies, later referred to as Critical Disability Studies by some scholars, is theorizing and research that supports "the struggle for social justice and diversity" on many "plane[s] of development" including "social economic and political, but

also psychological, cultural, discursive, and carnal" [62]. The primary thing that sets disability studies apart from more traditional rehabilitation or special education disciplines is that it prioritizes the "social constructionist model" of disability as well as "a more complex conceptual understanding of disability oppression… that nevertheless still employs key ideas about disability that saw the light of day with the ascendance of the social model" [62].

Since at least 2010 [154], disability studies has been used as a critical or theoretical framework within accessible computing. Recent work has only pushed this integration further with the use of concepts discussed in disability studies such as critical realism [84], interdependence [17], care work [20], reflexivity [113], and intersectionality [73] among others. In addition to the theoretical approaches, accessible computing also borrows from disability studies an emphasis on the lived experience of disability [113]. Speaking to actual people with disabilities is considered a key method of generating insights on accessible technologies. Included in studies of the lived experience of disabled people is work that looks at disability communities where they live (e.g., [194,195]).

4.2.2.2 Disability Communities on Social Media

Outside traditional HCI venues, studies of disabled communities, particularly those existing on or organized through social media, are common (e.g., [52,74,75,162]). Social media is a progressively more common venue for disability community building [234] as it does not have the inherent physical access issues that in-person community building has. From instant messaging and blogs, to Twitter, Facebook, and YouTube, every social media platform or medium has people with disabilities present on it. Tumblr, as discussed above, is commonly seen as welcoming users who do not fit in other places. It's not surprising,

therefore, that Tumblr has been a site of research studying disabled users and the disabled community (e.g., [45,99,129]). When Tumblr introduced a "poll" function in 2023, a widely circulated post asking if people were autistic received almost 29,000 users voting that they were either autistic, allistic but otherwise neurodivergent, or maybe autistic but not sure. While a Tumblr poll, like any post circulated on Tumblr, obviously does not receive attention from a random or unbiased sample of Tumblr users, it's still interesting that out of almost 33,000 respondents, only 11.5% responded "no" to the question [85].

In addition to specifically disabled communities, the presence of disabled people embedded in an un-networked manner within other communities is also a topic of study. Research shows that crowd work, for example, is a not-uncommon solution for people with disabilities seeking work [263], with over 60% of surveyed crowd workers reporting having one or more disabilities [245]. When people with disabilities have to work with non-disabled people, it's been documented that those with disabilities must expend extra energy to bridge the inaccessibility of technical and social systems in order to do their work [36,37,60]. This is the "invisible labor of access" that falls disproportionately on people with disabilities in working situations [248] and is part of a larger group of studies on the ways accessibility is mediated and constructed by both disabled people and their non-disabled friends, family members, and colleagues [11,151,159,180].

4.2.2.3 Collaborative Labor to Produce Accessibility

In addition to the work disabled people do to help with accessibility, as crowd workers or as collaborators in daily life, I also want to briefly touch on an example of non-disabled people working collaboratively to create accessibility for the disabled community: The Social Accessibility Project. Described in a set of studies published from 2008 to 2010

[206,207,236,237], The Social Accessibility Project set up a framework for web users to volunteer their time to collaboratively author metadata that would make websites more accessible to screen reader users. This is an example of what I'm terming an accessibility community, a group of volunteers collaborating online for the purpose of promoting accessibility generally or improving the accessibility of particular digital content. In The Social Accessibility Project, the volunteers or "supporters" came originally from the internal intranet the project was hosted on. They were incentivized to participate through a system where their volunteer activities would earn them points and they would be visibly ranked on the project page if they were among the top contributors. This project was, ultimately, more of an experiment in technical means for collaboratively authoring metadata than it was an example of a naturally developing, spontaneous accessibility community.

Another study looked at more naturalistic examples, this time of user generated content that was communally made accessible on two online platforms: Wikipedia and an online health forum for Lyme disease [135]. This study discussed the community culture of these accessibility communities but focused primarily on moderation strategies and moderators themselves as key players in the development of accessible content in these spaces. One thing that is worth noting is the Wikipedia example from this study highlights that an accessibility community is not always a standalone community but can be part of a much larger community united by shared goals, platform, governance, etc., just as the accessibility advocates existed within the context of Wikipedia at large. I'd also like to note that in both cases the content being made accessible was primarily text, rather than focusing on image accessibility.

4.3 Method: Digital Ethnography

The simplest definition I can articulate for digital ethnography as I practiced it in Study 3, is ethnographic methods applied to the digital sphere. While there are slight discrepancies between individual definitions of netnography, virtual ethnography, and digital ethnography, the terms have enough overlap that I choose to use them interchangeably. For example, I appreciate the way netnography was defined by Kosinets as "the culturally-oriented study of that technologically-mediated social interaction that occurs through the Internet and related information and communication technologies" [134]. The culturally-oriented aspect of this definition is the key connection to the older field of ethnography. The key features of digital ethnography as I understand them are a) focusing on observing social interactions and culture from an embedded or emic (insider) perspective [27], and b) being concerned with computer mediated communication, digital interactions and culture, or online communities [134].

The first key method of ethnography then, which I applied to the digital sphere in Study 3, is participant observation. Participant observation is considered "the cornerstone of ethnography" without which a methodological approach "may be legitimate and effective for exploiting any number of topics, but it is not ethnographic" [27]. Participant observation, put simply, is the embedding of the researcher in the social context in which they study. It is the "emplacement of the researching self in a fieldsite… [because] becoming directly involved in the activities of daily life provides an intimate view of their substance and meaning" [27]. Assumed in this "emplacement" is both longitudinal study of a specific community, and a certain level of reflexivity as a researching subject. The exact length necessary for proper participant observation, or the correct degree of reflexivity, is

not clearly defined. Participant observation, in particular as part of an ethnography of a peer production community, has been documented as taking as little as three months [222] or up to three years [181] or more to produce findings. The assumption is that you should engage to the degree necessary to understand the practices of the community as well as how that understanding is shaped by your own biases.

Interviews are another key component of ethnographic methods that are relevant to Study 3. Boellstorff et al. describe interviews as "so central to effective ethnographic research that we cannot imagine a project that did not include them" [27]. While participant observation allows a researcher to observe and analyze behavior and processes that participants may not be actively aware of (i.e., tacit knowledge), interviews capture the oftentimes deep knowledge that participants have of themselves and their own active thought processes. In addition to interviews, quantitative data collection, particularly descriptive statistics, is an accepted part of digital ethnographic practice. As Boellstorff et al. put it, "quantitative data can tell us about widespread trends" while qualitative data can be used in "explaining what such trends might mean" [27].

4.3.1 Research Questions and Overview of Phases of Research

The third and final study making up this dissertation sought to answer the following research questions:

1. How are image descriptions created in the context of an online accessibility community?
2. How do participants in an online accessibility community describe the benefits, motivations, and challenges of creating community-generated alt text?

3. What kinds of relationships exist between the roles of image artist, alt text writer, and alt text reader in the context of image descriptions created by an online accessibility community?

These questions were answered with a combination of observational, elicited, and quantitative data collection methods. This was so that I could triangulate between an empirical understanding of the process community members followed to create image descriptions, the self-reported processes and motivations of community members, and my own interpretation of the connections and social interactions that made up community-generated image descriptions.

4.3.2 Data Collection

To answer these three research questions, I broke Study 3 into four phases (see Figure 4.6). The first two phases of Study 3 consisted of different forms of participant observation to gain a basic understanding of how the community functioned and its context. Phase 1 took the form of continued but now direct participation in the Tumblr accessibility community, and Tumblr's platform culture in general. Phase 2 was a period of intensive participant observation of the PDDS. Particularly to answer RQ2, I dedicated Phase 3 to speaking to members of the PDDS community directly and seeing how they articulated their experiences as community members and image description writers.

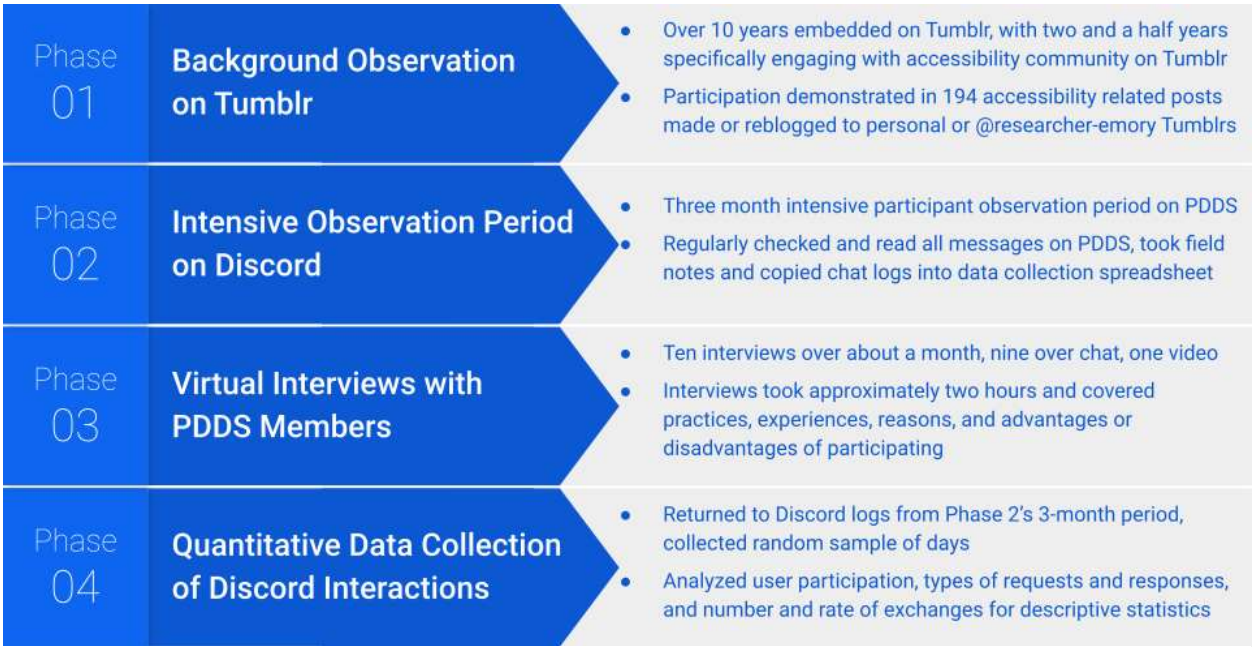

Figure 4.5 High level description of the four phases of data collection in Study 3

Focusing on establishing a direct answer to RQ1, I used Phase 4 to collect quantitative, descriptive data to give greater insight into what an "average" image description process on the PDDS looked like. This was intended to supplement the particularly interesting or novel image description examples that I had already collected during my participant observation phases. This phase, along with phase 3, also proved important, once analysis was complete, for answering the cumulative RQ3. In some cases, these phases overlapped, but generally they are listed in order of when the phases started relative to each other.

4.3.2.1 Phase 1: Participant Observation on Tumblr

I have already used ethnographic approaches to study Tumblr in the past [70] and have been using the platform as a user for over 10 years. During that time, my personal Tumblr account has been a site of personal and intellectual growth as an academic as well as functioning as an archive of posts which support and prompt my scholarly reflection. As

early as June 2020, my Tumblr record shows that posts about image descriptions were circulating on the platform and informing to some degree my approach to Studies 1 and 2.

While I created a dedicated Tumblr blog for the purposes of storing posts related to my participant observation of the PDDS, I found that I engaged just as much, if not more, with posts and conversations on accessibility through my personal blog during the course of this dissertation. Thus, I describe my participant observation as occurring on Tumblr over the two and a half years between June 2020 and January 2023 through both accounts. During that period, across both my personal Tumblr account and the "researcher-emory" blog dedicated to Study 3, I wrote or reblogged 42 posts discussing accessibility, I personally made 33 image or video posts accessible, and reblogged 119 accessible image or video posts. These numbers are intended to illustrate that I had ongoing engagement with the accessibility community on Tumblr during this time, with the original posts and tags on reblogs functioning similarly to field notes to keep track of my developing thoughts on the community practices around image accessibility.

4.3.2.2 Phase 2: Intensive Participant Observation on Discord

In addition to the ongoing observation, I engaged in as an active Tumblr user from mid-2020 to the beginning of 2023, I conducted an intensive 3-month period of observation of the PDDS from August through October 2021. As of August 2021, I had already casually been a member of the PDDS for at least 6 months, but the intensive participant observation began after receiving IRB approval for the study from the UCI IRB, at which point I officially announced my presence as a researcher on the server and used the username "researcher-emory" for the entire rest of my involvement in the server. During the 3-month period, I regularly checked for and read any new messages on the

PDDS, checking somewhere between several times a week to multiple times per day. After reading all the new messages, I would note any interesting exchanges, respond to requests for input or help with descriptions. Any important exchanges were noted down in a log I kept with plain text, anonymized copies of chat logs occurring on the PDDS channels, where I afterwards wrote in-depth field notes reflecting on the data that had been collected.

While I did not formally collect any demographic information about the observation participants, I did track how many unique users appeared in my data collection log, along with a few other details when they were available in the Discord user's profile or display name. To summarize this information, there were a total of 65 observation participants (i.e., users who appeared in at least one exchange recorded in my notes over the course of the 3 months of data collection. As far as available demographic information, 44 of those 65 participants identified through the Discord interface that they either a) used they/them pronouns, b) used multiple pronouns, c) used neo-pronouns, or d) used some combination of the a, b, and c.

I collected pronoun data pragmatically, so I could know how to refer to users in the text of my dissertation if or when their exchanges were described. However, I note it here for a purpose. Pronouns are not equal to gender identity and–as Tumblr users themselves will point out–anyone can use any pronouns regardless of if they are cis or trans. For example, it is completely possible that some of the PDDS members used they/them pronouns on the server not because they were non-binary but because they had privacy concerns that were assuaged by not revealing their gender. Additionally, the pronoun features I just described do not reveal any binary trans people who may have been solely

using he/him or she/her pronouns. Thus criteria a) through d) are not a 1:1 measurement of the number of trans or non-binary PDDS members. However, the pronoun numbers imply that there were a large percentage of PDDS participants who, at least based on the information available, seem to be trans or non-binary. Taking 44 of out 65 as a fuzzy estimate of PDDS members who were active during my observation period and were trans or non-binary, this means that the PDDS community as a whole is made up primarily of users with at least one marginalized identity. This informs my overall interpretation of the community and its values, as I will discuss in more detail in the Findings section below.

The only other demographic information recorded about observation participants (when the information was available) was if they were over or under 18. This was again recorded pragmatically, as only users over 18 were eligible to participate in the interview portion of the study and I wanted to mark those users who were or were not eligible. Thus, I can report that of the 65 observation participants, 26 identified themselves in some way as older than 18, 11 identified themselves as younger than 18, and the remaining 29 users did not mention or indicate their age in any way.

4.3.2.3 Phase 3: Virtual Interviews

In Phase 3, I ended up interviewing ten members of the PDDS community (see Table 4.1). The first interviewee I spoke to I reached out to directly, as e created the server and had a particularly important role in the community. For the other nine participants, I posted a call for participants on the channel of the PD Discord server dedicated to questions. I asked participants to react to the post with a particular emoji if they were okay with me DMing them the details of the study and the interview portion particularly. I received 12 responses to the post, of which 9 were actually able to and interested in

participating in the time frame I was conducting interviews. All of the participants were sent the study information sheet formulated for the interview portion of the study and they were given an option to conduct the interview over a video conferencing software or through Discord itself in the form of a chat interview (for more details, see next section). Of the 10 participants, only one opted to participate via video, all other participants preferred the chat interview option.

The participants agreed to have the interview recorded (in the form of copying the chatlog, in the case of chat interviews) and answered some basic demographic questions at the beginning of the agreed upon time (see Table 4.1 for demographics). The final two demographic questions (asking how long they had been on Tumblr and the PDDS) transitioned into asking about how they first heard about image descriptions or decided to get involved in making them.

I then asked participants why they wrote image descriptions or what kept them coming back to the PDDS. Then I moved on to asking about their experiences and practices around accessibility, including what typical interactions around accessibility on Tumblr or the PDDS looked like, how often they create or reblog accessible posts, how they decide if they have the energy to make a post accessible, their typical process for writing image descriptions, and how they decide if an image description is complete or correct. I also asked about the perceived frequency of accessible posts on Tumblr or social media more broadly.

Table 4.1: Demographics for People's Description Discord Server Interviewees (PDI)

ID#	Age	Gender	Pronouns	Race	Disability	Other Identities	Education	Income	Time on Tumblr	Time on PDDS
PDI1	31	Nonbinary	They/them or e/em/eir	White	Chronic Pain	Queer	BA/BS	{$50-99k	8 y^	3 y
PDI2	18	Nonbinary	They/them	Vietnamese	None	None	Some college	$0-24k	6-7 y	4 m
PDI3	24	None	They/them	White	None	None	Some grad	$0-24k	7 y	<1 y
PDI4	19	Nonbinary	Any or they/them	Arab	None	ESL (not taught formally)	Some college	$0-24k	4 y	6 m
PDI5	22	Nonbinary	They/them	White	Does not identify as such	Autistic, ADHD	Some college	$0-24k	8 y	<1 y
PDI6	20	Man	He/Him	White	ADHD, Autistic, Dyspraxia	Trans, Gay, Asexual	Some college	$0-24k	7 y	3 y
PDI7	N/A	Woman	She/her	Indian-American	None	N/A	Some college	N/A	3 y	5 m
PDI8	23	Nonbinary	Ze/zir	White	Anxiety, Depression	None	Some college	N/A	5 y	4 m
PDI9	20	Agender	They/them ey/em or it/its	White	Autistic	Home schooled	Some college	N/A	7 y^	~1 y
PDI10	18	Bigender	It/its	Indian & White	EDS, Autistic, Others	None	Some high school	{$50-99k	5 y	<1 y

Key: ^ is not including a period of at least six months of inactivity, { indicates household income is reported

This transitioned into a section discussing the interviewee's perception of the purpose or aim of image descriptions or the PDDS, how they saw their contributions as playing into the accessibility community on PDDS or Tumblr as a whole, if they feel their experiences around making posts accessible was overall a positive or negative experience (along with an example of a particularly positive or negative experience they had had), and

how they felt their own perspective or identities influence their creation of image descriptions.

The last question in that section asked interviewees to describe the best part of their experiences contributing to accessibility. This then transitioned into the final full section discussing challenges such as the hardest part of writing image descriptions or barriers they felt stopped people from joining the PDDS or writing descriptions. We also discussed changes they would make to the Tumblr or Discord interface and to the Tumblr or Discord community / culture to help with accessibility. I then wrapped up the interviews and thanked them for participating.

4.3.2.3.1 Chat vs. Video vs. Email Interviews

Over the course of these three studies, I used three different types of interviews: video interviews, email interviews, and chat interviews. I want to briefly discuss the methodological reasons for conducting interviews in these different ways and some of the advantages and disadvantages.

First, all of these interview methods were selected in contrast to in-person interviews. I made this decision because these studies took place during the initial Covid-19 pandemic outbreak and the years following it, making in-person interviews dangerous for both me and my participants in many cases. Additionally, because remote studies, remote work, and remote communities were the norm in this period, I felt remote interview methods were more understandable and acceptable to participants than they might have been in years prior. Particularly in Study 1, video interviews were chosen because they allowed disabled people to participate regardless of their physical location, physical access needs, and

comfort with traveling. The drawback of using video interviews, particularly during such a tumultuous time, was that there were several issues with scheduling participants which led to some sessions that were intended to be focus groups having only one participant and having to be changed to an interview at the last minute.

In the case of Study 2, email interviews were given as an option for several reasons. First, they are an accepted qualitative research method [58,120]. Secondly, they may even be the preferred method when, as was the case with the artists, respondents have differing scheduling and fluency needs [120,177]. Lastly, they allowed the artists the benefit of the time to think and reflect before responding, which considering multiple of the artists had worked on the Avatar project months prior, giving them time to consider, recall, and refer back to their experiences at the time allowed for more full answers. The drawback was that follow-up questions, while included in the email exchanges, were not able to fully tease out in-the-moment subtleties of how participants worded their responses and why.

In the case of Study 3, chat interviews were given as an option for two primary reasons. First, because the PDDS community was and is primarily text-based and community members were likely to feel more at ease discussing Discord text interactions over Discord chat. This was particularly important for any participants who may have had privacy concerns. Secondly, from my participant observation period, I had already noted that many community members mentioned having disabilities or being "low on spoons" (referring to Spoon Theory [164]) which meant I expected some or even most of my interview participants to prefer the reduced effort and social fatigue from text-based interviews.

There were several drawbacks to the chat interviews, however. For one thing, they took longer, with an interview protocol that would have taken less than an hour to run through verbally often taking two hours or more to cover via text, because of the increased time needed for participants to type their responses. Additionally, the chat interviews were more structured than I generally prefer. While there were follow-up questions for participants, the wording and order of the questions was less responsive to participants because I was worried about increasing the pause between responses. Also, because respondents were typing their responses, they may have been self-censoring more than they would have been if communicating verbally. This is in some ways a benefit, as it allowed reflection and reference to records, as with the email interviews, but also meant that the responses may not have been as revealing. It was useful, then, that the chat interviews were supplemented by observational data.

The last interview, I believe, proves that the chat format was overall worthwhile. The participant had missed two of the scheduled times we'd been set to conduct the interview, and said it was struggling with its health issues at the moment. I suggested sending the questions asynchronously, so that the community member could respond at its convenience, without having to sustain attention and effort for the entire two hours at a single time. It was grateful for this accommodation and was, in the end, able to complete the interview questions over the course of several delayed exchanges, where they probably would not have been able to do so otherwise. This demonstrates, to me, that the drawbacks of the method were balanced against increasing the accessibility of participation in the research, which as an accessibility-focused researcher, is of paramount importance to me.

4.3.2.4 Phase 4: Trace / Quantitative Data Collection

In my quantitative data collection phase, I sought a baseline understanding of "average" image description behavior by returning to the same 3-month period of the PDDS when participant observation took place. To collect a randomized selection of the image exchanges taking place during the 3-month data collection period, I took the 92 days of that period (31 days in August, 30 days in September, and 31 days in October) and selected a random set of 30 days for my sample. For each day in my sample, I re-read every message in the image description requests channel and collected every exchange related to an image description request (i.e., side talk, the few times it appeared, was omitted). For each exchange, I recorded the timestamp, wording, and any relevant image or link from the original request, and each response to the initial request, including if the original requester replied to other members of the community with thanks or clarifications. From this data, I was able to identify descriptive data such as the average number of messages involved in an image description creation exchange, and the average length of time between the first and last message in an ID exchange.

I also collected information about the community members involved in each exchange, namely the hashed display name, how many messages the community member had sent in the total exchange, and how long they had been a member of the PDDS server (this information is displayed automatically on Discord for any server member who was still a member at the time of data collection in late 2022). This allowed me to describe the average number of members involved in an exchange and what amount or length of activity in the PDDS they demonstrated.

Lastly, I collected what roles, if any, the users had chosen as part of their use of the PDDS server and which role they were playing in that exchange. In many Discord servers, "roles" are labels that server members can choose or indicate they would like to be assigned by the mods, that say something about their status in the server. In the case of the PD Discord server, the channel that incoming users were invited to read first allowed them to choose between two status roles: requester (i.e., someone who would be primarily asking others for help or feedback with accessibility attempts) or volunteer (i.e., someone who would be primarily responding to other people's requests for help or feedback). Thus, for each community member involved in an exchange in the image description channel during one of the sampled days, I recorded if they had chosen for themselves the "Volunteer" role, the "Requester" role, both or neither of them. Or, if the user was no longer a member of the server when this data collection was taking place (in late 2022), then the role information was no longer available, and I marked it as "N/A." I also marked, for each exchange, which user was the requester versus who was replying, and if there was evidence that the requester was also the image artist, post creator, or neither (or if it was unclear). This allowed me to make connections between the etic (outsider) concepts I was bringing from Studies 1 and 2 (i.e., image artist, image describer, user) and the emic ways community members identified themselves as requesters (i.e., consumers or users) of IDs, volunteer responders to ID questions (i.e., image describers / alt text writers) and image or post creators (i.e., image artists). The frequency and types of overlaps between these roles will be discussed as they relate to RQ3, along with the qualitative description of how these role overlaps appeared in the participant observation and interview data.

4.3.3 Data Analysis

After data collection was complete, I reviewed all data logs, chat logs, transcripts, and image-ID pairs that were captured. I inductively coded the observations and field notes, and then the interview logs. After data was given first-level codes, I read through all the codes and combined / refined codes into a second-round codebook which I then wrote up into memos and arranged into the structure of the findings. I also, at that time, calculated the descriptive statistics from the quantitative dataset and investigated how the statistics supported or conflicted with the understanding I had from the qualitative data. Finally, all findings were organized into which research question they related to, and I chose which insights to include in the summative takeaways section discussing all three studies (4.7).

4.3.3.1 Researcher Positionality

My analysis, and even data collection, across all studies was impacted deeply by my positionality as a researcher. Reflexivity is part of the key features of digital ethnography in particular [183] so I feel my position as a researcher needs to be acknowledged as part of the methodological approach to this work. As a white, trans person with multiple invisible disabilities, and as a researcher working with corporations and intended to continue into corporate research positions, I found myself torn between my pragmatic and idealist selves.

Particularly in Study 1 and Study 3, I felt compelled to fully convey the opinions of the people with disabilities I spoke to or had experience with. This led, in the case of Study 1, to potentially including too many details and specifics from different participants. I knew how I would feel if my access needs or desires were dismissed as secondary or unworthy of inclusion and did not want to make my participants feel that same way. In Study 3, I feel I

learned from my mistakes in Study 1, and held a more critical lens to the beliefs and practices of the online accessibility community I was participating in. While I still wanted to elevate the voices of people with disabilities in the community, I was aware that some of the actions of Tumblr and PDDS community members conflicted with the needs I'd heard screen reader users articulate in Study 1. As a long-time member of the Tumblr community, I was and am both aligned with many of their values, and simultaneously critical of their execution. As I've seen in many non-accessibility related discussions on Tumblr, the best of intentions can still lead Tumblr users to false conclusions, logical fallacies, and biased or inconsiderate actions.

This led me, in my data analysis, to balance criticism against good faith interpretations of potential flaws or limitations I observed in my participants' words and actions. I believe all ethnographic researchers, in their attempt to approach that emic understanding of a community or culture, risk losing sight of problems or biases that members of the community themselves do not see. And yet, having talked to and observed and interacted with these community members for months at a time, I knew that I shared many identities, beliefs, and history with them. To pass judgment on their practices was also to judge myself and everyone who attempts (however imperfectly) to increase accessibility and inclusion in the world.

These were all ways in which my tendency towards idealism fought with my pragmatic side. This was particularly important in my analysis of Study 2, where rather than being compelled towards generosity by the disabilities of the participants, I felt instead like I was compelled to some degree by my corporate connections and career interests to diminish

the types and severity of accessibility issues discussed or demonstrated in the discussions with image creators and accessibility practitioners. While I believe, as with the other studies, that all actors had the best of intentions and tried their best to increase accessibility, the overarching corporate structure seemed to stifle the powerful effects of the pure passion I saw in, for example, the volunteer image describers of the PDDS.

In the accounting of all three studies as I present them in this dissertation, I try to walk the line between the lofty principles of disability justice and inclusion for all, versus the hard practicalities of how accessibility is negotiated between the many people, platforms, organizations, and structures we operate beside and within.

4.4 Research Question 1: State of, and Process For, Creating Accessible Images on Tumblr

The starting point for understanding PDDS and the implications of their accessibility methods more broadly is to first explain the specifics of their work practices. To explain the process used by PDDS members to create image descriptions, I'm going to start by describing the general phenomenon as it exists on Tumblr. This includes the general state of accessibility on Tumblr, according to the reports of PDDS members. I'll also describe how this state was achieved and is maintained via the investment of time and energy by image describers such as those in the PDDS. I'll then describe the specifics of the process describers follow using both quantitative measures and reported practices of PDDS members. Finally, I'll explain the community aspects of crowdsourced accessibility, including the community norms I observed during my ethnography, and how describers talked about the overall ecosystem of the PDDS in the context of Tumblr and other social media sites.

4.4.1 What's Happening Here? Describing the Phenomenon

To begin, reviewing the state of accessibility on Tumblr is difficult. I cannot give a general report or systematic quantitative analysis of, for example, how many posts on Tumblr include image descriptions, or how many Tumblr users have read or written an image description. Due to the size and heterogeneity of Tumblr as a platform and community, these types of measures would require a very different type of approach than my ethnographic study of accessibility focusing on the PDDS. Instead of attempting to report ground truth about image descriptions on Tumblr, I instead sought to get PDDS members themselves to describe their own understanding of the prevalence of image descriptions on Tumblr and if or how it differed from other social media platforms they were familiar with.

I want to start with PDI1, who created the PDDS and is perhaps one of the individuals best situated to report on the state of accessibility before image descriptions became well known. E explained that: "When I first got involved in [a popular podcast] fandom, it was almost entirely inaccessible. Image descriptions were not a common practice. I remember searching the tag 'accessible fanart' and there was nothing there." Eir experience of fandom on Tumblr prior to creating the PDDS is important particularly because PDI1 admits that "it's somewhat true that a lot of this [accessibility community on Tumblr] was started and spread from me. [The podcast] as a fandom was straight up inaccessible before my blog, and these accessibility efforts have spread to many, many other fandoms on tumblr, as a direct result of the [podcast] fandom. I don't take credit for most of it but I do think the popularity of my blog sort of started a ripple effect."

In contrast, many of the other interview participants joined the PDDS server much more recently so they could better report on the state of accessibility more recently. For example, PDI2, a Vietnamese nonbinary 18-year-old, had been a member of PDDS for 4 months at the time of our interview. They explained that "i've noticed a lot more people making IDs [image descriptions] recently, especially when the images are something that'd be easy to write an ID for, like single tags or short tweets. i want to say 60 or 70ish percent of those easily ID'd posts will have image descriptions in the notes, but that percentage radically decreases as the difficulty of writing an image description increases." The state of accessibility on Tumblr does, as PDI2 said, seem to be changing quickly. PDI4, a 19-year-old nonbinary Arab, had been on PDDS for 6 months and Tumblr for 4 years when they explained that "prior to following someone who writes them [IDs] mid 2020 i had no idea they existed/never encountered them on posts." For context, the PDDS itself was only created in January 2019. PDI6, a disabled White trans man currently studying accessible pedagogy, did not feel that descriptions were "super widespread, unfortunately" but added that they were "getting way more common though!"

One important element of image accessibility on Tumblr that many participants mentioned was that it is not equally distributed across the platform. Instead "they're mostly concentrated to very specific fandoms or circles of people" (PDI6). PDI5, a 22-year-old White autistic person, stated: "The amount of people describing theirs and other's (sic) posts have definitely increased a lot but while a significant amount of posts I see are described or have a description in the notes I think that's definitely a product of the people I follow, so I don't think it's very widespread at all if you look at tumblr as a whole rather than my little corner." PDI9–who mentioned being homeschooled was a key part of the way

they interacted with the world, in addition to being agender and autistic –likewise said that "IDs are widespread in some spaces and some individual fandoms… but not in general" which ey believed was because "people have only started doing them recently." Beyond the "60 or 70ish percent" estimate that PDI2 gave for easily described images, the only interviewee who made a direct guess at the prevalence of image descriptions was PDI7, an Indian American woman who chose to use the pseudonym, Kay. She argued that "there's a growing awareness of them… [so] it's not hard to find descriptions on the *majority* of posts I see" (emphasis added).

Based on my own observations on Tumblr, these statements fairly well describe the experience of using Tumblr, if you are at least somewhat connected to one or more users who create image descriptions. PDI2's assessment that easier to describe content (short, screenshotted text or unambiguous images) has a higher rate of description reflects my own experiences, with the added dimension that content geared towards a general Tumblr audience, rather than a specific fandom or niche interest, are far more likely to be described. There will occasionally be posts with memes or reposted TikTok's that have clearly spread (evidenced by having upwards of 10,000 notes) and yet lack any accessibility description, but this is relatively rare. On the other side of the spectrum, posts that are older, have to do with a fandom where image descriptions have not caught on, or would require highly complex image descriptions, may very well have 40,000 to 100,000 notes without any image description present. I cannot, based on my own observations, estimate the experience of people completely disconnected to accessibility on Tumblr. I do not doubt that a fair number of the 135 million active monthly users on Tumblr have never seen an image description or do not understand what an ID is.

The experience of accessibility on Tumblr is at this point largely something users have to choose to opt into or seek out. Going back to the introduction to Tumblr given in above, the way Tumblr's post structure works means that even for posts which have image descriptions somewhere in their reblog chain, and are thus technically accessible, the majority of users interacting with those posts are likely not viewing, liking, or reblogging the version of the post with the image description attached. However, if someone needs that accessibility accommodation, or once someone does seek out those pockets of highly accessible blogs or fandoms, the experience of accessibility may be quite extensive.

Users I spoke to from the PDDS are likely some of the most invested in image accessibility on Tumblr. Even given that, it is remarkable to see the level of effort participants put into making their blogs accessible. PDI1 decided early on that e would make eir blog 100% accessible, meaning: "from the start of my blog in 2018, I included image descriptions and didn't (and still don't) reblog anything that isn't described." For context, PDI1, at the time of writing this chapter, has over 48,000 posts on eir blog, which indeed dates back to 2018 and has image descriptions dating back to those earliest posts. It is more than just the founder of the community who puts this much effort into making their blog accessible. PDI4 stated that "around 90% [of the posts on their blog were accessible] in total… 100% if they're okay to be reblogged not like. personal" with the additional explanation that they don't make every personal post accessible because they don't anticipate those posts are going to be spread anyway, so the likelihood that someone will come across them and need an image description is low. PDI6 likewise said that his "entire blog is described (other than stuff from like 5 years ago)." Kay explained that her high level of engagement with Tumblr means "I definitely describe around ten posts a day."

While not every participant I spoke to made sure their blog was entirely described, even those who were not able to make every post accessible mentioned the amount of effort they put into describing. PDI5 explained: “I'd preferably write an ID for every undescribed post I come across but often don't have the time or energy for it. If it's a simple image I can often describe it even when I'm low on time/energy but if it's a complexer one that's often not an option, sadly.” PDI2 estimated that “around 60 to 80 percent” of the images on their blog had descriptions because “it fluctuates based on how much of my brain is online at the time to write IDs.” As I’ll discuss in more depth in the Challenges section of this chapter, writing image descriptions can call for a lot of mental energy and labor which users may or may not have to spare. PDI10 was a bigender person with multiple disabilities who used it/its pronouns and identified as both White and Indian. PDI10 explained that it balanced the issue of energy versus accessibility by deciding to “just reblog things less frequently if they’re undescribed… being an entirely accessible blog has definitely changed the way I interact with tumblr.” In the next section I’ll get further into the details of what creating an image description consists of for PDDS members.

4.4.2 How Is It Being Done? The Process of Writing “IDs”

As Kay referenced in the prior section, half the experience of writing IDs for Tumblr is finding them. This is because, according to the PDDS members I spoke to, one of the first steps in the process of making a post accessible is “check[ing] the notes for an existing ID” (PDI2). Currently, Tumblr has implemented a way to filter notes by whether they contain tags, comments, or neither. However, at the time that I was speaking to PDDS members, this was not the case. Instead, as PDI6 described, he would “scroll around in the notes of a post for existing descriptions which can be a Process depending on how many random

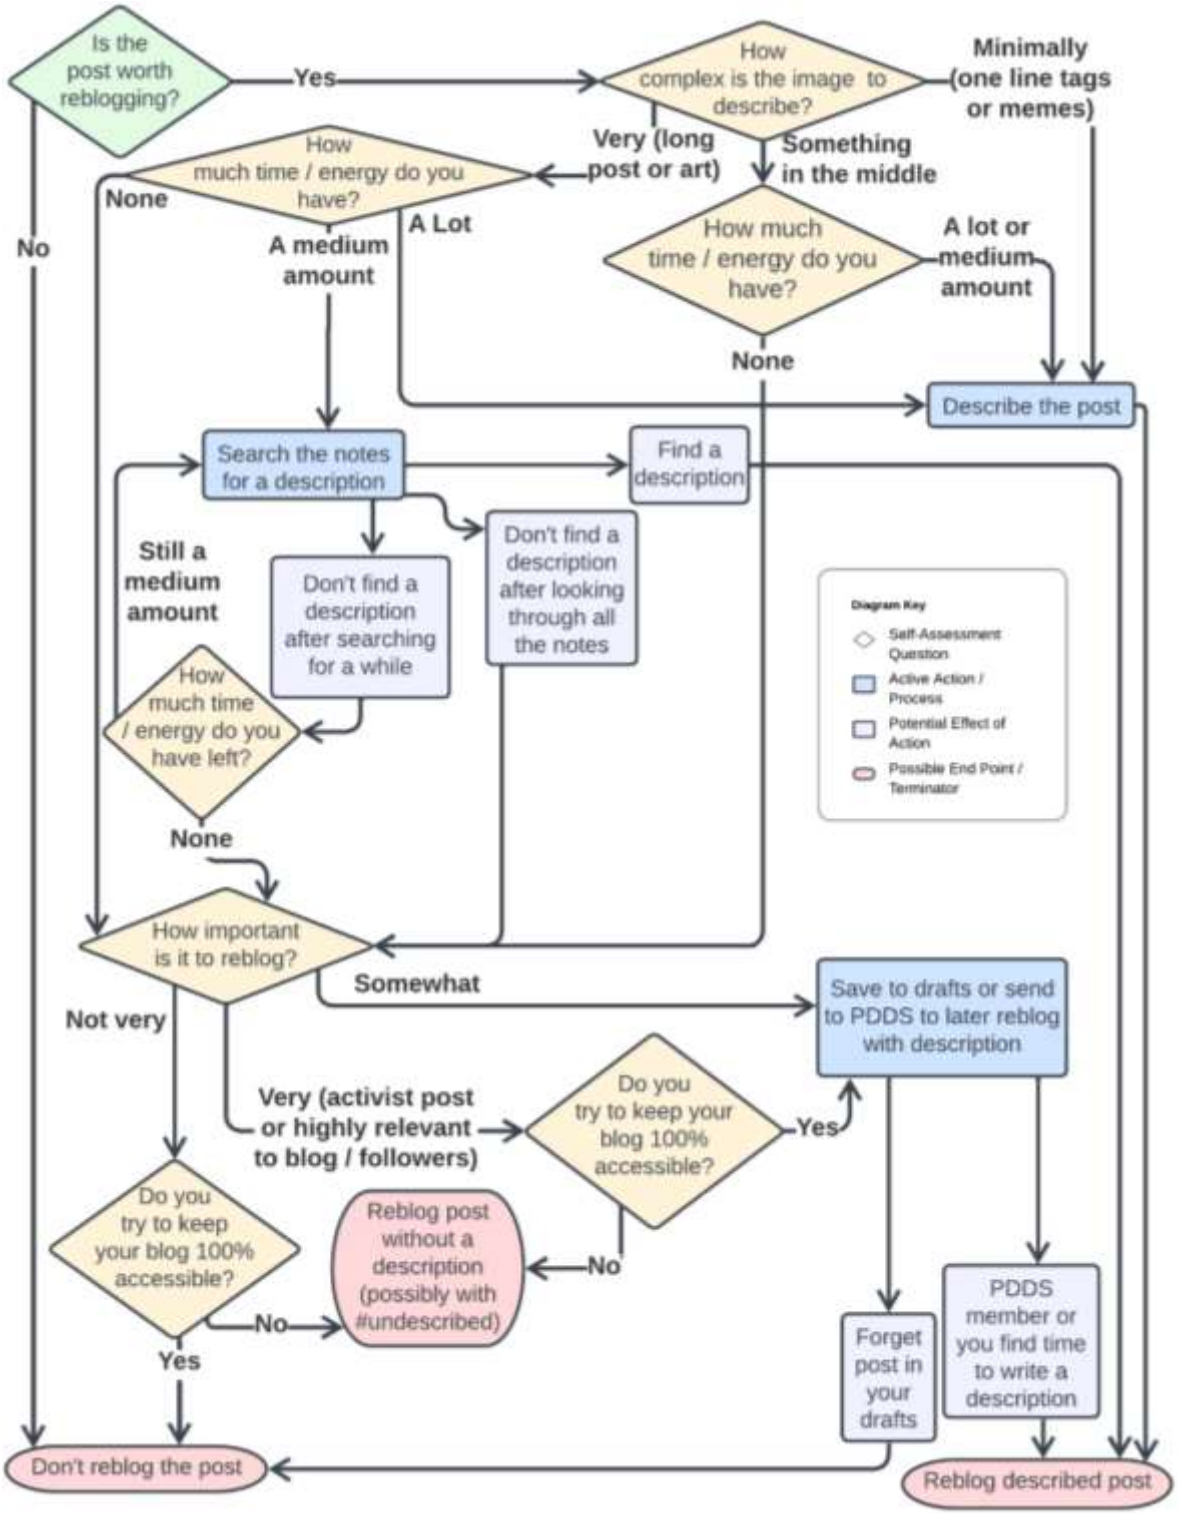

Figure 4.6 A flow chart depicting the process PDDS members described going through when working on an ID

weird comments people have left on a post." PDI3, the only participant who chose to be

interviewed over video chat, similarly said that when looking for IDs in the notes "sometimes they're easy to find, sometimes [it's] like 'no, I'm not finding this.'" The number of notes on a post, particularly the number of notes that include tags (prior to recent updates) or comments (with the current system as of writing), determines how tedious of a step looking for a prior description may be.

Based on my observations of the PDDS interactions, it may also be fairly common for duplicate descriptions to be made for this reason. Multiple times members requested help with a description of a post, only for someone to reply that they had actually already found an ID in the notes. In one particular case I recorded, a question was asked regarding what details to include in a description, and one PDDS user responded with two different reblogs he had found in the notes that described the post with different levels of detail so that the requester could decide which they wanted to use. The redundancy has some advantages, as it may mean new rebloggers have fewer notes to look through to find the image description, but it may also mean that effort is wasted to write a new ID when an existing one just wasn't looked for or found easily enough.

However tedious, finding existing IDs may be the most straightforward part of the process of creating an accessible post, as my interviewees highlighted. Many described following a flowchart-like decision-making process (see Figure 4.6 above) which took a great deal of cognitive effort, in addition to the literal time expended on crafting the ID. Users were often checking in with themselves to determine if they have enough energy or time at that moment to write the ID, or if they should save the writing process for later. PDI3 explained that a post "sits in my drafts for a long time. And then whenever I have free

time… I normally do it in big chunks so like i'll describe 50 or so images [at a time]." Some interviewees described the ID writing itself as fairly simple: "the process really is just like I see the image I look at it, I try to give like a brief one or two sentence overview of it like this is what the thing is, and then like i'll go into more details as needed" (PDI3). While other participants, such as PDI4, discussed that depending on the image, writing a description "gets tricky cause there's a lot of details you never know if you're missing something or if the artist wanted to highlight something and you're missing [it]."

Participants discussed having knowledge of general image description style guidelines, as well as demonstrating that they followed guidelines principles. For example, PDI9, who wanted to be referred to as Sil, argued that "there's a higher emphasis on detail than purpose on tumblr that in general style guidelines, which I feel comes out of the uncertainty of the audience" adding that ey were "most familiar with" the Diagram Center guidelines for image descriptions [63]. Despite Sil's concerns that Tumblr describers may not always clearly convey or focus on purpose when writing image descriptions, another participant, PDI5, explicitly mentioned that they "try to keep the function of the image in mind" when deciding what to include in the ID. Ultimately, the process participants discussed (represented in a composite flowchart in Figure 4.6) was complex, with some unique elements, but it also shared many features with the image description writing processes seen in Study 2.

The depth of interviewees' responses allowed for a great deal of detail into how these processes function in practice. For example, in addition to the Diagram Center guidelines Sil mentioned, there were other tools and resources that members discussed

using as a community or individuals. Kay and a friend created one of these resources, a Google Doc listing many common memes and descriptions of the key elements so that users could use them as templates and fill in any changed details in the particular iteration of the meme they were describing. PDI8 discussed using this exact document as a way ze personally approached IDs: "I also try to remember if I have a template for that specific image (if it's a meme there's a compilation of templates I have access to, if I've described another piece that's related to it…). If I have it I copy the template or the other description and I start modifying it." PDI8 also mentioned that ze would "use an OC [optical character] Reader for the text in the picture" when there is too much text to easily write it by sight. Sil described the PDDS itself as "pretty firmly in the tool category for me" functionally the PDDS was a labor resource when ey didn't have the "spoons" for writing a full description. (I'll discuss spoons more deeply in the Challenges section.)

Another detail of the image describing process for Tumblr posts that several participants mentioned was that it would occasionally require switching between or working with multiple devices. PDI3 said that "sometimes i'll have like the image itself up on my phone or like on a separate screen to look at the image and then be typing on another one." PDI6 explained that "I usually write my descriptions on my laptop… [but] I mostly use the tumblr app so i'll like a post i want to describe and then run to my computer to describe it before I forget and like too many more posts." From the perspective of effort extended, this definitely seems more involved than simply browsing or scrolling a feed, as most users of both Tumblr and other social media sites are probably used to.

Lastly, the interviews revealed some individual preferences that users had in their own ID writing practices. Kay for example told me that she was more active on a sideblog rather than her main blog as "I prefer to keep my blog mostly pg-13, and I like keeping it lighter, so almost no real-world stuff goes there" which makes her "feel strange about describing posts for issues I find heavy…sometimes I'm capable of describing a post but don't want to have it on my blog." She personally solved this through having different sideblogs with different types of accessible content. Sometimes these individual practices are shaped by the identities of the describers themselves. PDI8 explained: "I'm nonbinary so I am set to [sic] not assign pronouns/ gender to strangers, even if they are fictional… I also try to never use [sic] femenine, androgynous or masculine because they are so subjective and culturally tied that if it's used to talk about someone you are assuming that either the person of the author is okay with being labelled as either one of them." In a slightly different direction, PDI6 explained that "being autistic means that i like to get Every Detail Exactly Correct, which means i put a lot of care into figuring out exactly *what* i'm describing. i also pay more attention to describing emotions in pictures than other people might because that's something i struggle with parsing anyways." PDI4, as a nonbinary person of color themselves, felt that "it does make me pay more attention to race and skin tone and how to describe that in particular" and it explains their preference for "using pretty much gender neutral language unless like, you know how x person ids [identifies]."

All these details, as well as Figure 4.7, show the process of writing IDs from the perspective of individuals reporting their practices. However, beyond the internal thought processes and decision making, there is also the question of objectively what does the process of writing IDs with the help of the PDDS look like? To answer this, I'll review some

of the codes and metadata extracted from my quantitative sample of user exchanges on Discord. The most common type of request, occurring in over half (52.5%) of the exchanges I analyzed, was a request for or a response of a full ID written by a volunteer. This usually took the form of someone linking an image and asking if anyone could ID it for them. The large percentage of full ID interactions, compared to general questions, which only occurred in 45.5% of exchanges, may be partially biased by the fact that the sample was pulled from the "image descriptions" channel rather than from the "questions" channel where users might have posted questions relating to image descriptions more often, but they would also be mixed in with interactions not related to image descriptions at all. The second most common category of interactions (occurring in a little over 26% of the interactions) were asking for help related to the subject of the image. For example, asking PDDS members how to describe a person or character's pose, clothing, expression, or simply asking if anyone knew who the subject was (i.e., what piece of media a character came from).

Some of the requests I observed were for incredibly specific terms that could not easily be found through other means unless you happened to be familiar with the terminology involved. For an example, see the interaction in Table 4.2 where four total users struggle to name a specific element in an image. None of the suggestions are actually the exact term, and even I, analyzing the exchange after the fact, couldn't find the correct term as it is almost impossible to Google. This is despite the fact that multiple people recognized the decoration and that it would likely be a construct known by either a fashion expert, a party planner, or perhaps a cake decorator. The interactions captured in the quantitative sample therefore ran the gamut from very broad (what I labeled "General

Quality Check") questions to ultra specific. Note that the percentages I've reported add up to greater than 100% because many interactions had multiple codes, for cases such as if someone asked for help on describing a pose and received a full ID in return.

Table 4.2: Observed Interaction with PDDS Members Seeking Specific Terminology

	[Requester:] the word is on the tip of my tongue but what do you call the swooping ribbons?
	[Volunteer 1:] frills seems more like a gathered thing to me? i guess i could say like 'swoops of ribbon connecting the roses' but it seems wordy?
	[Volunteer 1:] Ohh i misunderstood your first message
[Volunteer 1:] These look like bits of rope-ish material i think?	
[Volunteer 2:] oh that is so frustrating im pretty sure theres a name for it as well, but i cant for the life of me remember and google isnt helping. even if its a tad wordy, i think the swoops of ribbon/rope thing works well at least!	
[Volunteer 3:] (40 minutes later) bunting?	

The time it took to receive responses also varied greatly. While the minimum time to receive a response was less than a minute, the longest time from first quest to last exchange was over two and a half days. Indeed, while fairly rare, there were nine requests out of the 80-interaction sample that didn't receive any response from the community. These requests simply slipped through the cracks in the overlapping attention spans of the PDDS members. Some requests that might have otherwise fallen into this category were brought back up later (the original requester asking again or replying to their own message to ask if anyone had seen it) and then responded to. Indeed, beside the outliers with long or no responses, 50% of the interactions lasted between six and a half minutes and two hours, 25 minutes, with a median of 27 and a half minutes. These numbers are also somewhat

skewed as they include all messages in the exchange, even though 53 out of the 80 interactions (66.3%) included at least one message of additional pleasantries such as thanks or apologies for the wait. Thus, the actual lag time between requests and the substantive responses may be less. Overall, PDDS members did not deliver a uniform experience, as might be expected for an Amazon Mechanical Turk request for example, but for a completely volunteer body of image describers, delivered results on a variety of topics, including niche questions about characters and terminology that could not easily be googled, with a fairly quick turnaround time, on average.

4.4.3 What Does Crowdsourced Accessibility Look Like? Community Norms and Practices

Despite the personal quirks and individual ways describers' identities shaped their process, the interviews also surfaced some more common ideas and norms in the community. The community norms I received from participant responses are by no means an exhaustive list. However, I highlight a few here to demonstrate what makes the PDDS (and Tumblr's accessibility community more broadly) different from an internal company's or cultivated volunteer organization's approach to image accessibility.

One of the simplest commonalities among many IDs on Tumblr is a particular formulation of the opening to a description. PDI5 described it as "'[character] from [media] doing Y'" (presented exactly as typed). PDI2 said a "basic description" would cover "stuff like 'a tweet by username' or 'fanart of this character from media.'" It's hard to say whether this is more a stylistic commonality resulting from Tumblr describers learning most of their description tips from one another, or if it's a logical starting block for most IDs considering the types of images usually on Tumblr. I observed PDDS members explaining

they had reached an agreement that mentioning who the character is and what media they are from at the beginning of the description is beneficial because "if its his canon appearance you might not need to go too into detail because its kind of assumed that thats how he looks." Essentially referencing the character and media immediately conveys a shared frame of reference to the description reader, allowing the actual description of the image to be shortened. The literal formatting of image descriptions on Tumblr is another thing that participants largely agreed on. Kay explained that "the majority of describers I've seen know to avoid special fonts in their descriptions, and despite some variation, the formatting of most IDs seems consistent!" Sil argued that "the two main places I see differences in the written format are that some people mean ID as image description (which is what I use) and some mean it as identification, so they'll either open [the image description with] ID or Image ID, and then some people put a slash after the description, like /end ID]." One other norm for many people in the community is, particularly if they try to have a high percentage of their blog accessible, to actually tag those posts that aren't accessible. Sil, PDI8, and PDI2 stated that if they reblog something without an ID they have a specific tag such as "#undescribed." The logic to this practice is that people may filter out posts natively through Tumblr based on tags (e.g., if they don't wish to see a certain fandom, or topic that bothers them) so if someone is trying to create a completely accessible Tumblr experience theoretically they could follow only blogs with a high degree of accessible posts and then filter out these few undescribed posts based on the tag someone uses for them.

A few participants also discussed some of the common advice that I have seen directed towards describers on Tumblr. For example, PDI8 felt that "people end up learning

as they practise [that] IDs aren't only for blind/ visually impaired people (so you should take that into account, like spacing the description better for people with dyslexia) or that they shouldn't be used to make jokes (like people inserting their own commentary or changing things to make it 'funnier')." PDI6 similarly felt that of the agreed-upon describing rules "a big one is 'keep descriptions neutral'. i see people occasionally saying things like 'an adorable puppy!!' or 'this art is gorgeous!!' in their descriptions, but most people who are regular describers stick to the basics of what's in the image." Another somewhat unique element of image descriptions as advocated for on Tumblr is to include some sort of ID more or less always, even if it's very simple and/or redundant with surrounding text. It is commonly explained in Tumblr posts advocating for writing descriptions that it is just as important for a blind user to know what an image is not as to know what it is (i.e., if there is a single image in a thread of written replies, the image could potentially be a meme or sarcastic aside that renders the entire rest of the exchange in a different context, without a description at least saying that it's a gif of someone nodding, for example, blind users won't know if they agree with the post as written). This is in direct contrast to official guidelines for alt text on most websites, which state that if a graphic is decorative or does not add information to the written words then its alt text should be left blank.

Beyond these stylistic similarities and accepted norms between accessible blogs and posts on Tumblr, when describing things that were common or agreed upon in the community, participants overwhelmingly talked about the spirit of support and growth that the PDDS members fostered. For example, PDI10 said the best part of the community was "the willingness everyone has to help out" and that it's agreed on "that there's always space to improve your own accessibility." Kay stated that in her experience "questions are

always answered quickly and compassionately." PDI5 said that they had never "encountered a describer who didn't offer to help others if they were struggling with writing IDs." PDI8 recalled "starting and doing mistakes and doing things that aren't as convenient for screen readers and now I have more experience and I do things a bit differently. That's because other people in the community shared their knowledge and tips and people are welcome to grow… the 'no one size fits all' mindset [is one] that helps people make mistakes but also do things on a different pace and grow." As we'll discuss more in the Motivations and Benefits sections, participants in the PDDS felt that they were united in and benefitted from a strong dedication to accessibility.

Overall, the research question on how image descriptions were made in this online community has a complex answer. First and foremost, it is clear that part of the process of making image descriptions is actually a process of advocacy and education, where community leaders and participants take on the role of modeling and explaining the necessity of accessibility work. This outreach element of how the accessibility community functions leads to a noticeable shift in how widely available or understood image descriptions are on the Tumblr platform, which in turn fosters more engagement. This positive feedback loop does much to overcome the inherent complexity and cognitive burden of actually writing and sharing image descriptions. Through a series of check-ins with themselves, estimations of time available and required, and collaboration with other members of the community, participants made accessible a large and continuously growing number of posts on Tumblr. Part of the way image descriptions are made through the PDDS also relies on certain agreed-upon standards, norms, and shared practices. These norms

lead to the image descriptions that are created bearing some unique characteristics and identifiers that separate them from alt text in other contexts.

4.5 Research Question 2: Reported Motivations, Benefits, and Drawbacks of Community-Generated Accessibility

In this section, I am going to focus on the second research question defining this study, namely: How do participants in this community describe the benefits, motivations, and challenges of creating community-generated alt text? We have already touched on some of these topics in the prior section, but here I'll dedicate each subsection to one of the topics (Motivation, Benefits, and Challenges). Because this research question focuses on how participants describe the benefits and challenges, I will mostly draw from my interviews with PDDS members. However particularly in the Challenges section I'll use some examples from my observations of PDDS interactions to explain why I asked interviewees if they had certain emotional drawbacks to engagement with IDs.

4.5.1 Why Do They Do It? Participant Motivations, no "Clout" Detected

From discussions with PDDS members, there were broadly speaking two different types of users with two slightly different angles to their accessibility approach. The first type I'm going to talk about are people with outside knowledge of image accessibility. Either people who personally knew someone who was blind or visually impaired, or who engaged in offline activism that led them to become informed about and interested in participating in describing images online. The second type of participants were those who found out about and engaged with description primarily through Tumblr and those social circles, interested in contributing to the growing cultural tide. Generally, whichever

personal motivations people had for engaging in PDDS, they seemed generally aware of the same overarching motivations of the community.

The founder of PDDS is, perhaps logically, a perfect example of those types of users who brought their outside interest in accessibility into the online community space. PDI1 explained, before starting eir fandom blog in 2018: “I've been involved in (mostly) offline activism since 2016, and that's where my interest in disability justice started. I'm a member of a disability justice org that's local to me and this is where I learned about accessibility in a more meaningful way, and also made friends who are blind. So I had already pretty much gone through the process of making my personal spaces online accessible in ways they weren't before.” PDI6 put it similarly: “one of my super super close friends is blind and also a huge irl disability advocate (she does a ton of lobbying and work with national organizations), and we talk all the time about irl [in real life] accessibility, so i was already tuned in to being aware of accessibility issues irl and it was super natural for me to fall into online accessibility work.” There were other users who sat somewhere close to but not exactly in this category. PDI3 noted that they “knew some people with visual disabilities… but [because] they weren’t really like on tumblr or anything” those personal connections were not a strong motivation for online accessibility work. In an odd twist, Sil explained that ey “had the personal experience of what it felt like to not be able to access almost all content on my dash” but not because of a vision disability. Instead, it was because their internet connection was slow enough that “images wouldn’t load on my phone most of the time.”

On the other side of the spectrum, PDI2 described how, “i saw some of the people i follow adding image descriptions to their own blogs; that + they were reblogging/making posts about why others should also do that” and that that was as simple as the equation needed to be for them to decide to get involved themself. PDI5 explained that they “saw people writing image descriptions on tumblr and saw a couple of posts detailing why it was important, which were all reasons I agreed with, so I gradually started doing image descriptions as well.” Kay could very clearly remember why they started writing image descriptions, approximately a year before: “I realized I'd been seeing a lot of IDs on my dash, and I decided that even if I wasn't going to write them myself, the least I could do was look in the notes to reblog ones others had written. Some weeks after that I started writing my own, which made my confidence in my ability to do so grow, and soon I was writing them regularly and connecting with others in the community!” PDI10 could even remember which specific blog had first educated them about image descriptions. This user “was very vocal about image descriptions. I really just started because it seemed the logical thing to do, in my mind.” In the case of these “converts” to the practice of writing image descriptions, it’s hard to say if users on different platforms would have the same likelihood of joining in with writing descriptions. While it was not discussed by any of my participants, based on my extensive experience with Tumblr I think it’s fair to say that the general orientation towards social justice and digital activism present on the platform might have led to more widespread and fast adoption of image descriptions than you would expect on other social media sites.

One interesting thing that many participants referenced, although none talked about it explicitly, was that Tumblr users might choose to start making descriptions because it

was popular in the fandom they were in. As mentioned earlier in “Describing the Phenomenon,” participants reported that image descriptions definitely clustered in certain fandoms. PDI6 explained it as “in some chinese drama show circles where a handful of us were writing descriptions and then suddenly a bunch of people noticed that and now there are significantly more folks who write descriptions!” It’s unclear whether this is more because the accessibility message, like most messages on Tumblr, traveled effectively through the networks of fandom blogs and mutuals until it caught on with more people, or because there was a certain desire to get in on a potentially new and interesting thing to add to your posts. Even if accessibility is partially gaining momentum as a trend, this does not necessarily mean that the effects won’t be long-lasting. An even less cynical explanation for why fandom bloggers may be more likely to join in on writing IDs is presented perfectly by PDI1’s experiences. E said: “when I was hyperfixating on [this podcast], all the work I did to make the fandom accessible didn't really feel like work lol. There was a period of time where I was adding image descriptions to 40+ posts a day.” Essentially the same passion that fans are already extending to their media of choice may be easily harnessed, when presented by fellow bloggers from their own fandom, to serve both fannish and accessibility goals.

Indeed, several participants explicitly renounced anything but moral obligation or personal belief in accessibility’s importance as motivations for writing IDs. PDI3 stated that they “don’t have a huge following of people that reblog my stuff. So I don’t really do [accessibility work] for clout or anything like that.” PDI4 explained that their reasons for describing are simple: “it's only fair that other people be able to have access to the same content/infomation/art/etc (sic) as i do. it's also more of a "why not?", if im able to then

why not make one extra post accessible to others… its definetly (sic) not a 'adding an id guarantees my post gets more notes' kinda thing." One thing I find interesting here is that it's actually quite unlikely that any one user knows how many people see or benefit from their image descriptions. Tumblr's notification system shows reblogs directly from your blog. But unless you created the original post, it doesn't show any downstream reblogs (i.e., people who reblog something from someone who reblogged it from you), making it almost impossible to know if a post that is circulating and receiving a lot of notes happens to be one with your addition (be it image, witty rejoinder, or ID).

By far the largest motivation participants discussed or demonstrated was simply that they made image descriptions because they empathized with the experiences of blind users or other people with accessibility issues. PDI3 explained: "I'm not really visually impaired to where I would need like a screen reader or something. But I find that I see that image descriptions for me are really helpful just because if a gifset doesn't load, then I can just like know what the images [are], but also like I feel like I get more out of it by reading the description" and that this alone is motivation enough to put time and effort into describing posts.

PDI6 stated that "i know a few folks on PA who are visually impaired or blind… it's important to me to have people who actually use image descriptions to be active in the community because they deserve to have a say in what their accessibility needs are and how they get those needs, esp because disabled folks' voices are erased by abled people so frequently." He went on to add, that people who don't write image descriptions don't feel it is their responsibility to create accessible content on Tumblr, "they don't go on tumblr to

advocate for disability rights or whatever, so why do they have to care? which is awful and lame cause disabled people also go on tumblr and then they can't access half the images on their dash." As I'll discuss more later, it's fairly common for consideration of accessibility to lead to frustrations with those who don't take the time to learn.

However, more common still is that caring about people with disabilities returns its own kinds of rewards. Kay said that "it feels validating to know that the descriptions aren't going to waste and are helping people. I think it's also a good exercise in mindfulness for me, because I can't know who needs IDs and who just cares about accessibility when an ID I wrote gets reblogged, which reminds me of why it's important to make my blog as accessible a place as possible!" Sil explicitly stated "the emotion I draw from [is] empathy, but I'm also motivated by being able to be useful to other people and improve their experiences."

4.5.2 So What Do They Gain? Benefits and Desired Outcomes of Collaboration

In addition to motivations PDDS members largely agreed on the benefits and ideal outcomes of their work. The reasons people choose to participate (motivations) and the anticipated or actual benefits they understood from participation had a high degree of overlap. In some ways many of the benefits recounted in this section could be rephrased as motivations, and vise versa for the motivations from the previous section. However, I have tried to delineate them as best I can so that benefits are those commonly understood as effects. Paired with them are other outcomes that PDDS members articulated as overarching goals of the community's accessibility work, rather than individual motivations for their participation.

PDI10 was one participant who couldn't actually explain why it wrote IDs "other than a sense of 'this is what i've done and will continue to do'." Describing-as-habit is one type of creation model that is particularly interesting to discuss how to cultivate, as it could get around the feelings of emotional and energy burden that I'll discuss more later. It is exciting to consider the possibility of expanding from a community of practice into a communal habit, where shared activities are done as a matter of course, without a clear delineation between practitioners, experts, or dabblers. While this was not yet realized for many of the PDDS members, given the degree of complexity and effort required for the task (as described in the first section of this chapter's findings), PDI10 hints that this may be an achievable goal.

The largest reported benefit of engaging in digital accessibility was the positive emotional benefit or sense of feeling useful. Kay explained that their desire to be useful, surprisingly, had them request descriptions less often "since I tend to feel bad about making them unless I've also recently volunteered a description myself. I like being able to answer questions, though, and I'm less hesitant about asking them in the questions channel too." The behavior Kay describes, where a member will respond to a previous request before posting their own, was one that I observed quite frequently over the three months of participant observation on PDDS. PDI4 emphasized that while "the goal is simply for accessibility no other thing… i feel like writing ids is an enjoyable activity for many." PDI3 stated that "I almost kind of find them relaxing [to do]. And also like I feel like I pay attention better to like a piece of art when I have the ability to describe it" adding that they also enjoy the fact that "the people are really friendly…[and] we're doing a good thing… to help other people enjoy content." Giving an interesting overview of the benefits of writing

IDs, PDI2 said that they continue to write them for three reasons: “i want to make my blog more accessible” (i.e., intrinsic benefit), “it’s nice having people reblog IDs i write and feeling like i’ve been helpful” (i.e., emotional benefit), and “there are people who i’m pretty sure stated following me for the IDs and knowing this, i’d feel way too guilty if i ever stopped” (i.e., emotional obligation). This slightly complicates the straightforward narrative of emotional benefits and sense of fulfillment.

Similarly complex, Kay described the process of making IDs with other PDDS members as emotionally “a positive feedback loop” where “I’ll take cues from someone else on, say, transcribing special-font captions, and then someone else will tell me that they started doing that because I did, and I can image the ripple effect spreads from there! So it’s REALLY nice to think about the concrete good you’re doing with IDs in that way.” What’s interesting about this is the example Kay used is actually one of the accessibility practices used on Tumblr that I have not seen guidelines or accessibility resources otherwise mention. Describers will rewrite text if it is in a font or size that is not accessible. For example, if someone uses Tumblr’s small size text for an ID, other users may copy and paste the ID in a reblog so that it appears twice, once in smaller font (the original post) and once in whatever default font size the user has set (in their reblog). It remains unclear to me if this practice is altogether more beneficial than it is disruptive–is the unpleasant redundancy of repeated text for screen reader users outweighed by the potential benefits to low vision users who don’t use screen readers? The positive feedback loop Kay mentioned therefore produces a feeling of “concrete good you’re doing” but whether those practices are in reality most beneficial is not always verified.

As mentioned before by PDI10, the community commonly demonstrated a growth mindset wherein one of the benefits, cyclically, of engaging in accessibility was becoming better at doing so in the future. PDI4 mentioned that they "joined the server cause i thought it would be good to be able to ask for help or help where it could be needed. i often struggle with aspects of a description and not the whole and had no one to ask for help." They explained that learning the ropes of a new practice like describing can be "a very hard barrier to get over, but then like once you can do it, it's easier so like the first couple posts that you do, are going to be really hard. Obviously, because you've never done it before" but that getting help with those early forays into IDs can build confidence. PDI5 explained that "there's no such thing as a perfect description" and "I sometimes think 'oh shit I could have just done that' after seeing a description for an image I struggled with and then discarded." Essentially, both the individual IDs and the process of breaking down an image into a coherent description are opportunities for learning and growth. PDI6 mentioned how he "used to get super details with every description i wrote, but now i usually only write extremely detailed descriptions for big art pieces." This is a pattern I observed on Tumblr, that as a community, describers learned and adapted to critiques regarding inappropriately long and complex descriptions. PDI6 explicitly said that "finding how to describe things the best way possible is a learning process" both backing up the assumption that the more you describe the better you will become, while also subtly reinforcing that there is no single "best" image description, only attempts to reach "the best way possible."

Another widely agreed upon benefit of image descriptions, on a wider scale than personal feelings of accomplishment, was the potential to help a variety of people through accessible posts. PDI2 argued that when discussing those who benefit from image

descriptions “people who use screen readers, people who don't but maybe have bad eyesight, and people with bad internet connections for whom images take forever to load are the groups i see mentioned the most, but also kind of everybody, since IDs can also be useful when you're just not sure what you're looking at.” PDI3 said: “I don't think it's just one demographic that it's aimed towards, I think it’s aimed at multiple. People with visual disabilities, or processing disorders, things like that.” Discussing both image descriptions and video transcripts (another type of accessible post that the PDDS members create), PDI5 stated that posts have the potential to help people who are “blind, visually impaired, deaf/Deaf and HoH, those with bad internet, those with concentration issues (often neurodiverse people), people with specific triggers (as for example spiders can be filtered out when there's an image description, but not if it's just an untagged image of a spider), and those in a controlling or abusive environment (by being able to read a transcript rather than play audio out loud or being able to scroll past an image and just read the description).” PDI8 had even more additions and clarifications for the list of potential beneficiaries of image descriptions: “autistic people (since a lot of times the description points out many things, like that the picture is a meme, or taht (sic) the tone of the character is meant to be sarcastic), dyslexic people (so they can read it with their own fonts), people who are learning/ don't have that much experience with the language (so they can search the words easily or even copy pasting the text to translate it), people with low connection or scripts disabled (because images won't load or will take too long)and people that want avoid certain content (because if it's in the description it's easier to filter it out, many times things aren't easy to tag or warn, so if they are described, it's easier for the filter to catch any relevant words because they will be written in the post).”

While some of these interviewees clarified that serving these additional groups are "more of a side benefit" (PDI2), it was sometimes unclear if participants saw these as examples of the curb cut effect [104]–where additional people could benefit from an accessibility feature intended to serve a specific group–or if they equally valued the needs of all these potential users when considering potential conflicting access needs. PDI1, as the founder of PDDS, mentioned that e "don't know how much [of the community's understanding of who IDs are for] has come from me" because "I am very upfront with the fact that image descriptions are primarily for blind people and I will talk about that until I'm blue." Although e mentioned that e "have found it useful to include ALL the ways image descriptions are useful as a tool to reach more people… a lot of people are surprised to learn that THEY benefit from image descriptions too. A meme gets explained in a description, or tumblr won't load images but you still get to read the descriptions, and suddenly you (general you) start to understand all the ways this tool can benefit you. and that makes people care more." E argued that "I don't think [explaining the broad usefulness of IDs]'s a bad thing! … it's about community. It's not about being a savior for poor disabled people, you know? Our freedom is tied up together… Image descriptions aren't something you include because you feel sorry for poor blind people. Image descriptions and accessibility in general is for you, too."

PDI1 discussed another common site-wide benefit that PDDs members felt IDs had: more exposure bred more awareness and more accessibility. PDI1 explained that "a few people writing descriptions in a huge fandom isn't going to have much of an effect most of the time. [This podcast] was perfect because my blog really took off and the fandom wasn't HUGE, so my posts got a lot of traction and exposure right away." In eir case, once image

descriptions started getting exposure many "jumped on board with honestly very little prodding." Once there was a leader or role model for users to look to for resources and motivation regarding accessibility practices, things spread easily from there. Accessibility-as-social-contagion if you will.

Speaking to other PDDS members, this was borne out. PDI6 personally remembered he found out about image descriptions "in junior year of high school" because he was into a podcast and "in the process of hyperfixating and following a fuck ton of people on tumblr i came across [PDI1]'s blog" and that was all he needed to join the server. PDI4 observed that if "bloggers with a larger following make their blogs more accessible [then] others follow suit. like my friend says, image descriptions [are] trickle down economics lol." PDI2 stated: "i'm hopeful that if i just leave enough IDs on popular posts and reblog enough PSAs about why accessibility is important that some of those people [who don't support accessibility] will turn to the light." Returning to PDI6, he was enthusiastic about writing IDs as advocacy, "i know i've gotten people to care about descriptions just by doing it and reblogging resource posts, so spreading the word and teaching people is super super important to getting more folks to do it… in the fandoms i've been in where it really caught on, most of the work was a handful of people yelling about how important it was and bothering big blogs and everything, and then people started to tell other people about it and finally more people started to care." Kay mentioned that "IDs are a really visible (ha), eye-catching way to provide an accessibility aid, so you can educate people just by being prolific with them."

While big blogs modeling accessible behavior was overall seen as a benefit by most interviewees, before we move fully on to the challenges section of this chapter, I want to

mention that PDI1, as an important figure in the community, felt that the impulse to follow a more notable poster's advice was a double-edged sword for em. "I do feel uncomfortable," e said, "being treated as an authority. I get that I'm often seen that way, and I can't control how people perceive me. But it's definitely weird. Especially because literally every single thing I know has been taught to me by someone wiser and smarter than me in the disability community. I don't ever want to be seen as the absolute authority on anything." Ultimately, the cause as a whole benefited from the way accessibility caught on via Tumblr, but individual bloggers might have noticed how challenging modeling accessible behavior could be.

4.5.3 What's Stopping Them? Technical Challenges and Labor Costs

When considering the challenges of practicing accessibility on Tumblr as a community, they fall into broadly three categories. First, processual challenges: things that are difficult about writing IDs in general practice or are particularly challenging on Tumblr due to technical limitations. Second, there were challenges related to labor or energy expenditures that describing requires. And lastly, there are personal considerations that can lead to particular challenges for some individuals.

Starting off with the processual challenges, picking out which details to describe or how to describe something specific or complicated were very common issues people had (demonstrated in both the interviews and the requests I logged). PDI2 stated that describing immediately becomes harder "if the art is really complex and/or i can't tell what's happening in the image." PDI3 likewise said "things that are really, really long or detailed are really hard for me Just because I'm trying to not make the post like 10,000 times longer but also still trying to get the point across of what it is." PDI4 felt the largest

challenge for them was “describing and phrasing the id” so it is not “written with convoluted and overly complicated sentences that are confusing to hear” while including all relevant information because “maybe something to you is relevant but someone else might gloss over that.” For PDI8, the challenge is often what specific term to use to describe something: “I usually have to look up things through colour archives, names of specific styles, types of clothes, tools, furniture... I like to give precise terms that are well known. There's no point in saying that the chair is a glass tile green if no one is going to know what type of green that is. It's more useful to say it's a pale grey green.” PDI6 mentioned that he is one of many “who struggle with understanding facial expressions” which makes IDs useful to him but also means it’s a challenge when writing his own. “Other than facial expressions,” he said the hardest thing was “fully conveying what the artist intended with their art through words. sometimes there are things you just can’t really write in words the best way.” Sil mentioned: “I’ve been doing [IDs] long enough now that I’ve figured out how to get around stuff that I used to have trouble with.” However, ey still have difficulties and “generally, if I ask for help with something, it’s how to describe a facial expression/emotion.”

Participants also gave examples of how writing image descriptions were particularly difficult at times or in certain contexts because of Tumblr’s technical features. As already discussed, “tumblr is notorious for being shit at loading images” (PDI6) which means that, Kay discussed, “it can be difficult describing multiple images on mobile because the viewing window gets smaller in the reblog screen and sometimes images suddenly stop loading and being displayed.” Sharing tags from a previous version of the post is very common practice on Tumblr, however there is no easy way to do this accessibly (rather than through a

screenshot of the tags that then needs to be described) because "it requires going to the blog that made the tags, then copying the link for that post into [a mobile] browser and copying the tags from there" (Kay). As discussed, finding image descriptions in the notes is a key part of the accessibility practices of PDDS members, however PDI5 points out with "the notes section of Tumblr, it's very frustrating that it won't always show all of the notes, making it very hard to find existing descriptions." This means that very often users cannot use their browser's search feature to highlight any notes that include "ID", for example. PDI3 explained the user experience this way: "when i'm trying to find an image description scrolling up and then like having to wait for it to load more like that's really hard to deal with." Or for PDI4, who uses both desktop and mobile versions of Tumblr, "on mobile most notes dont show up so i either save it to drafts to look later… or if I have the patience for it, i open the website on the phone [browser] and have a look there."

There are also technical limitations that describers might have difficulty with when translating specific images into text on Tumblr. For example, PDI2 explained that they have experience "replacing highlighted/underlined text with bolded text, encasing it in [highlight] [/end highlight] brackets, or noting it elsewhere in the ID, since tumblr is missing the functionality to highlight text or underline it." Or how "Information inside a table has to be rewritten in sentences because tumblr doesn't let you create marked up HTML tables" (PDI2). PDI2 mentioned that they would like "to have CSS styling on posts so i could use text-transform: uppercase to write screenreader-friendly all caps, but i don't think it's likely that staff would ever let us edit the CSS of individual posts." The limitations of Tumblr as a fundamentally web 2.0 experience meant that describers could not follow

WCAG guidelines as they were created for website designers/developers rather than users of preexisting platforms.

In fact, users had a wide range of ideas for how Tumblr could be improved to address accessibility challenges, some of which have since been implemented into Tumblr's design and some which have not. PDI1 felt that not just Tumblr but "every website" should have "a wikipedia-style method of volunteers adding and editing image descriptions." A similarly optimistic approach from PDI6 was the idea that Tumblr should implement "a requirement for people who post images to have a really solid alt text/description to go with their image." He acknowledged that "you can't really do that on such big platforms, but i can dream!" Discussing more small-scale improvements to Tumblr, particularly its mobile app, PDI4 mentioned a desire for "the ability to increase font+ change contrast. Also maybe a 'show alt text' [button]... an automatic caption for videos or the ability to add one… maybe some flashing warning or settings [for] photo sensitivity." Some of these suggestions have been partially addressed (Tumblr mobile has a range of "themes" similar to dark mode which can help with contrast, and Tumblr also now has an "ALT" button on the bottom left corner of images that users can click to see a pop-up with any alt text that the original poster added to the image), while others remain under the purview of other technical systems (There are automatic captioning for videos built into many phones, and third party extensions regularly circulate Tumblr to help those with photosensitive epilepsy at least on desktop and some mobile browsers). Kay similarly wanted "the opportunity for alt text to be made obvious with a screenreader so people who need IDs but don't user (sic) screenreaders could take advantage of descriptions in alt text" which Tumblr has since integrated as a feature. Two participants (PDI8 and Sil) mentioned "the option to separate

the reblogs into reblogs with comments on them (where the IDs are), reblogs with only hashtags and reblogs with nothing" or even more specifically pull out IDs from other comments. The former feature was already implemented and PDI8 called it "a good change". Some of the technical challenges or missed opportunities for accessibility on Tumblr may therefore be ameliorated in future updates, although the large-scale suggestions may be far off.

The challenges for describers that may not be as easy to fix fall under the other categories, particularly the overarching challenge of labor burden and burnout. Describing requires a certain amount of free time, such that I saw a noticeable dip in the frequency of PDDS interactions during the tail end of my observation period (October) when most schools were back in session for the year. One finding that was demonstrated in those who chose to speak with me, as well as being actively discussed by participants, was that the task of advancing accessibility disproportionately fell on people with disabilities themselves. PDI1 remarked that "most people who get involved in accessibility efforts are themselves disabled. That's a pattern I've noticed both on- and offline." This trend actively applied to em as well; they mentioned their disability directly in the demographic questions, along with 5 other participants from the PDDS interviews (plus a 6th who stated they were autistic and had ADHD but didn't identify as having a disability). Beyond just those I spoke to, PDI5 agreed that "[their] impression" was that a fair number of PDDS members are neurodivergent. PDI6 likewise said "most of my friends who do image descriptions are also neurodivergent or disabled and have personal stakes in describing."

The disabilities members had actively affected their engagement with IDs and the online accessibility community. For example, PDI5 stated that “I don't feel I have the spoons or knowledge required to help others as often as I thought.” Four interviews and five separate exchanges over the three-month period actively mentioned “spoons” [164] or other disability barriers that contextualized requests or responses from volunteers. There are blind and visually impaired members of the community to consider as well. PDI1 pointed out “one of the mods of the [PDDS] is blind” and PDI10 stated “a lot of people I follow for image descriptions are blind/VI.” As PDI6 pointed out “some folks genuinely can’t really write ids because the ids are for them.” For those users, there is not the option of running an inaccessible blog: either they only reblog described posts, or they reblog posts with only context clues to tell them what the images contain. PDI1 felt that it was important to “lift up the voices of blind and visually disabled folks often. But also, it’s work [to create accessibility]. And it’s worth that definitely shouldn’t solely or even mostly be theirs [because] blind people have to fight for access in their personal lives constantly.” Even non-blind people may have difficulty functionally writing IDs, as Sil pointed out people often discuss having “a lack of spoons (since most people I follow are disabled), where they might want to, but they can't read faces or they can't order their thoughts enough to do more than reblog (these are both challenges I have/do face).”

However, even for people who did not have a disability, they frequently mentioned how tiring and effortful describing was. PDI2 mentioned that they consider “how tired my brain is” and “whether I have the time to write an ID” when deciding if they can write a description. They even broadly hypothesized that “writing IDs is something that can take a lot of time and energy, and while I think most humans like feeling helpful they’re a lot more

likely to do so when… it doesn't take a lot of effort, which is not the case here." Kay noted that "IDs feel draining if I'm stressed, if they're long, or if I'm unsure how to describe an element" so she considers if she's going to answer a request on the PDDS "based on my energy level at the time, how long the post is, and if I'm familiar with the material in the post." PDI3 mentioned that they don't always have time to write descriptions because they're busy with work, leading them to store up a batch of posts and then describe them in one go. They explained: "I sometimes do it when i'm avoiding my actual work so it's like, ' Yes, this is work, this is productive work. Great, i'm doing a thing.'" Many others mentioned the labor involved in writing descriptions.

Description is, as PDI1 put it "a lot of (unpaid) work." E explained: ""I am an organizer at heart, even though it's VERY easy for me to get burned out." E try to do projects that are "a group effort so if I don't have the energy, the project can continue for awhile without me." For PDI2, deciding if they're going to write an ID now or later "depends on how much effort writing an ID will involve and how much energy i have at the moment" which had led them to have "149 drafts right now and some of them are old as hell… posts end up kind of just languishing in my drafts forever" potentially because they never have the time or energy to go back to those more complicated posts. PDI8 stated explicitly that: "I am not always putting the same time or effort on all descriptions because otherwise I would burn out because they are time consuming and also tiring like any other type of work." In fact, for zem, the hardest part of IDs is "finding the energy and balance to write them, because there's a lot of undescribed content." For PDI6, "it doesn't take me a ton of energy to write descriptions but sometimes i just don't have any energy at all for things." Potentially, as PDI1 suggested, even the feeling of IDs being low energy is because

Tumblr bloggers are already used to intense passion for their interests so “directing some of that fan energy into an accessibility project wasn't that big of a leap.”

One thing that came up, quite unexpectedly, in the interviews and observations, was that image accessibility and the effort it required could lead to frustration with others. PDI5 remarked that: “it’s especially frustrating when I've described multiple of that person's posts and in some cases have even offered to describe posts for them and they still post without a description.” They noticed they will often respond to this by “try[ing] even harder to be helpful, in a perhaps almost passive aggressive way, by describing the image as soon as I see it pop up, trying to get it done before the post takes off so that hopefully the original poster can add the description to the post and people will reblog the described version.” But this sometimes led to individual bloggers becoming “a drain on my resources” where they “unfollowed them to spare myself the energy and frustration.” PDI8 explained that ze needed to “vent with people that do descriptions because they also know the annoyance that I have” when other people don’t adopt accessible practices. This venting was something I saw repeatedly in the PDDS exchanges. One user sparked a conversation by asking on the PDDS, “Is it just me or does anyone else get annoyed when artists complain that people describe their posts wrong Like. You could have added one yourself, you know.” Someone responded with “I'm definitely of the opinion that once you create inaccessible art, you've p much forfeited your right to complain about others trying to fix the problem You created” which received six reactions expressing approval with this belief. The same responder added that “I'm trying to get into the habit of thanking artists or other describers in the tags for their work. I know it's difficult, so rather than getting frustrated at those who don't, I want to use positive reinforcement on those who do Plus it leaves me

feeling less annoyed." These emotional reactions fall into the final category of personal challenges describers face.

Other personal challenges faced by describers are shaped by their individual identities. For example, PDI4 stated that: "i often struggle with a somewhat limited vocabulary in areas considering english is not my native language so sometimes [an] id feels maybe too simple." PDI8 must weigh "for complex art pieces…How much can I describe without being incredibly tiring? How Can I not get too tired myself?" which ze address with personal adaptations such as looking to see "if I have a template for that specific image (if it's a meme there's a compilation of templates I have access to, if I've described another piece that's related to it...). If I have it I copy the template or the other description and I start modifying it." Ultimately, the way personal needs butt up against accessibility practices is best typified in PDI10's comment that "sometimes I have to choose between describing things and my own personal sensory and/or emotional triggers." This element of describer's personal needs clashing with the broad accessibility dictates of the community was discussed on the PDDS itself. Halfway through October, the final month of data collection, one PDDS member brought up "is it wrong to block someone from your personal blog when you have a lot of image descriptions there?" to which one user replied "[in my opinion] not at all. Your comfort over anything else" with four people reacting in agreement and another person chiming in that "Your blog is your blog." Thus, one of the final challenges individuals faced was just the fact that their personal experiences, comfort, triggers, or beliefs did not always align with what would make their blog (or Tumblr as a whole) most accessible. The problem, essentially, is that people are people, and not robots

whose only purpose is to describe posts, as I will discuss more when discussing implications for AI generated alt text.

Addressing the guiding research question for this section, participants of the PDDS described their main motivation for participating as simply care for accessibility. They preempted and dismissed claims that image accessibility was about getting more engagement on their posts, and instead explained that they felt personally motivated or even obligated to contribute to what they felt was a good cause. Some PDDS members were already engaged in offline accessibility activities, while others converted to the practice through their exposure to it on Tumblr. In either case, they wanted to practice accessibility in their online lives because they felt personally, socially, or sympathetically connected to people who benefited from image descriptions.

On the topic of benefits, PDDS members articulated the benefits of community-generated image descriptions as much broader than alt text advocates generally argue for. Not only did members report personal benefits, positive emotional fulfillment, and a growing sense of expertise from participating in the online accessibility community, they also argued that by particularly focusing their efforts on image descriptions (visible to all users) their accessibility efforts benefited far more than just screen reader users. The high visibility nature of image descriptions meant they could serve user groups who needed written descriptions of things like facial expressions or triggering content. Additionally, because they were so high profile, every example of an image description was an opportunity to educate Tumblr users on what image accessibility is and why it's needed.

Lastly, participants in this online accessibility community described a variety of challenges to their participation. From the inherently complex process involved in creating image descriptions, to the technical limitations of Tumblr's interface that forced users to implement workarounds, the smaller drawbacks of the system ended up feeding into the overall time and effort requirements of participating. Members explained that there was a commitment involved in describing images, particularly for many users who already had disabilities. The endless optimism and dedication that participants displayed in the motivations and benefits to community-based image accessibility sometimes turned sour in the face of their energy going unappreciated. As I'll discuss more in the next section, for many members any Tumblr user could be an image describer, so having users stubbornly remain unengaged or uneducated led to the final challenge in participation: frustration with other Tumblr users.

4.6 Research Question 3: Understanding the Overlapping Roles of Artists, Writers, and Users

In this section, I will discuss the final research question which guided my study of PDDS as an online accessibility community: what kinds of relationships exist between the roles of image artist, alt text writer, and alt text reader in the context of this community? For that, I'm going to cover three main ways that I saw artists, writers, and users' roles overlapping. First, there was the community support for and demonstrations of the Artist-Writer process where the artist themself was encouraged to write a description. Secondly, the community approached and thought of every Tumblr user as a potential, not-yet-but-soon-to-be writer of image descriptions. And finally, I cover the quantitative results of the sampled request exchanges to show that, in practice, requesters (i.e., users) and

responders / volunteers (i.e., writers) were not mutually exclusive groups but that most members filled both roles.

4.6.1 Do You Know What I Mean(t)? Community Support for An Artist-Writer Process

Just as I observed in the case of the Avatar Project images, understanding the intent of the author can be hugely valuable when it comes to determining what to include in an image description. Many PDDS members also supported and advocated for artists or post creators to write their own image descriptions. For example, PDI1 argued that the most important thing to change in Tumblr's community is to make sure "more people saw them as necessary to include with their original posts. the original post part of that is key. What I and so many people do on tumblr by including descriptions in reblogs is honestly subpar. It's not ideal. The ideal thing would be that everyone includes descriptions with their original content. That would also cut down the workload for everybody." Image descriptions, as something included in a reblog, are always something that could be removed or missed depending on which version of the post a user sees. However, if more posts included alt text or image descriptions from the get-go (i.e., added by the original poster when creating the post) there would be no need for downstream accessibility efforts. PDI3 stated that things would improve easily "if more people could be open to making their content more accessible." Kay found herself frustrated "very frequently" because she has "a lot of artist mutuals who are prolific and never describe their art, and sometimes it's really annoying to describe a lot of someone's art and never get an ID copied into the original post no matter how many times I ask nicely in the tags." PDDS members support the creation of accessible posts at the source, i.e., by the post creator or artist.

Figure 4.7 Art of an original character created by PDDS member

And at least for some artists, this was seen in practice. In my observations of the PDDS, there were a few examples of artists going to the server for help describing their own art. One artist posted a picture they'd made (see Figure 4.7) and asked, "how should i describe the purple garment she has on? its meant to be like a coat but im at a loss for words on how to accurately describe its appearance" and then after a suggestion was made by a frequent PDDS poster, the artist returned with: "does this ID sound alright? [Image ID: A digital drawing of a human character in an angular style. All of the lineart is done in bright pink and purple. She is skinny, has pure white skin, pink spiky hair that is pulled into a ponytail, and purple lipstick. He is standing, facing the viewer with her legs together and toes pointed out. One hand is raised holding a champagne flute with a yellow liquid inside. His other hand is lifted and pointed at her own face as she smiles with his eyes closed in a pleasant expression. She is wearing a long purple coat with sharp shoulders, short puffy sleeves, a jagged hem, and is belted at the waist. The front of the coat is a deep V shape that exposes most of his bare

chest. He is also wearing tight yellow pants, pink stilettos, yellow bracelets, and a few rings on each hand. There are yellow sparkles around her, and the background is a flat dark blue. /end ID]." Someone else chimed in that "i don't feel great abt the 'pure white skin' the association of white with pure is very uncomfortable and i wld always rec another descriptor before that one." Someone else agreed and suggested saying "something like 'paper white' if you want to be more descriptive" which the artist agreed to do. Essentially, this is a mixture of an artist-writer, team write-a-thon, and even user evaluation process (considering the responses were from users who I, contextually based on their other interactions in the server, believe were people of color) in action through just one four-person exchange on the PDDS.

4.6.2 Why Not Try It? Every User is a Potential New Writer

It's important to note that fundamentally, in the moment of writing a description, the average PDDS member goes from acting as a Tumblr user (i.e., consumer of images posted by others) into a description writer (i.e., someone adding content to the platform via their description). This instantaneous bit of magic happens every day for hundreds if not more Tumblr users who write descriptions in reblogs. Thus, it may be unsurprising that a fundamental element of the PDDS philosophy is that every Tumblr user is someone not yet convinced to undergo the transformation into describer.

This belief is visible in two ways in my interviews with PDDS members. First, is the basic belief that inaccessibility stems from ignorance, that anyone who is sufficiently informed would pursue accessible practices. PDI1 stated that:

> “I have a theory that there are four kinds of people when it comes to making their content accessible. 1. People who need to know or personally connect to accessibility efforts in order to start. 2. People who need to repeatedly be exposed to image descriptions and educated on why they're important, over and over, before they start. 3. People who are ableist assholes who will never do it. 4. People who are disabled themselves and can't do it. I think most people fall into 1, 2, and 4. In my opinion, exposure and normalization is 100% the best tool in getting people to join in accessibility efforts. Spreading awareness is key. (This is partly why I don't like alt text descriptions. I think descriptions NEED to be visible to sighted people in order to spread awareness.)”

PDI2 similarly felt as far as barriers to accessibility went “ignorance probably is the biggest one: not knowing why IDs are important, or knowing that but not knowing how to write one or how to learn how.” This applies not just to IDs broadly, but to the PDDS specifically as PDI3 notes “for people not joining the discord I think it may just be that, like not a lot of people know about it.” PDI6 felt that “the best way to increase accessibility is to make people realize that their actions re: making content accessible affect so many people and can benefit more people then they might realize!” Kay noted that if she could change anything about Tumblr it would be to “increase awareness of what IDs are and why they're necessary!! In both fan servers I'm in and general tumblr, I get the feeling that people don't think about adding IDs to things because they just aren't aware there's a need. If more people appreciated the necessity of IDs, I believe that they'd be ready to seek them out and make them themselves.”

Underlying a lot of the assumptions that education will fundamentally change most people's approach to accessibility is the fact that the PDDS members themselves were once in those ignorant shoes before becoming educated. PDI6 summarized it quite well, saying that for barriers "a big one is feeling overwhelmed by the idea of having to describe images all the time. if you've never done it before, describing seems like a HUGE effort (and it is! but people don't always understand exactly how to describe so they might get more overwhelmed)." The addition of "and it is!" highlights that PDI6 empathizes with the difficulty and stress of first starting to write IDs, even as he encourages people to try anyway. PDI4 similarly said "in regard to starting in writing ids i think it's a daunting thing. i think ppl are scared about writing it incorrectly. also maybe a lack of education on them or their importance." The compassion users had for their fellow Tumblr users, especially given how many of the PDDS members themselves started writing descriptions because of their exposure through Tumblr (see 6.2.1 above) highlights that the distinction between describers and not-yet-describers is very thin from their perspective.

4.6.3 Is This What You Call Mutual Aid? A Requester/Writer Dichotomy is Not Welcome Here

Indeed, the very elastic distinction between writers and users is demonstrated not just in the individual moments of describing but can also be tracked and quantified in the PDDS dataset. First, some background information to contextualize the results. When I collected the 80-request sample of exchanges from the Discord I also noted down, for each user: their username (hashed, so as to be unique but non-identifiable), how long they had been a member of PDDS (which is available in the user profiles of all users who were still currently members or had left and rejoined the server), and two other pieces of

information. I recorded if, in practice, the user was requesting or responding to a request (and if they were requesting, I used details such as the form or content of the request to estimate if the requester was also the artist of the image, the creator of the post, simply sharing a link from another's blog, or unknown). I also recorded which "role" they had chosen when they first joined the community. They could select to have their profile on PDDS (which lasted as long as they were still a member of the server and hadn't left) mark them as a "requester" or a "volunteer" or they could select both at once (they could also select neither, although few people did this).

Discussing the results of this analysis, I'll first say, out of the 80 requests, 12 of them (15%) were from artists or original post creators. An additional 15% might have been artists or post creators but I could not definitively identify it (i.e., I labeled the requesters' status "unknown"). Additionally, removing the requesters who had left the server or who never chose a role, out of the 66 remaining requests, 9 requests (15%) were actually made by those with a volunteer role designated. Similarly, out of the 86 times someone responded to a request, 9.3% of the time the responders had designated themselves as a requester. This highlights that, among other things, the roles users chose (requester or volunteer) did not always align with how they acted in practice (as a requester or responder, i.e., a user or a writer).

In addition to this, I divided the participants in the exchanges up by ID until I had identified 67 unique participants who had interacted in some way within the 80-request sample. From there I identified, for each participant, if in their interactions they solely acted, in practice, as a user (i.e., requester) or writer (i.e., responder), or if they took on

multiple roles. What I found was that 26 participants (38.8%) only acted as writers, 14 participants (20.9%) only acted as users, and the largest percentage, 40.3% (27 users) acted in multiple roles. Because some users engaged more than others, the total number of engagements for each user group also heavily swayed towards participants playing multiple roles (PPMR) (see Table 4.3).

Table 4.3: Unique Users and Their Engagements Groups by Roles Played

Type of User in Practice	User	Writer	PPMR
# of Unique Participants	14	26	27
% of Unique Participants	20.9%	38.8%	40.3%
# of Engagements	52	98	194
% of Engagements	15.1%	28.5%	56.4%

Looking at the distribution of engagements for the Users, Writers, and Participants Playing Multiple Roles (PPMR), there are some clear outliers to the data, including one user who requested 14 times (the next largest is 9), three writers who responded more than 10 times (next largest is 8), and two PPMR who interacted more than 18 times (the next largest is 12). After removing these outliers, the new numbers and percentage of engagements, along with the mean and median number of engagements is shown below in Table 4.4.

Table 4.4: Engagement Counts, Percentage, Means, and Medians for Participant Groups (Excluding Outliers)

Type of User in Practice	User	Writer	PPMR
# of Engagements	29	60	153
Average # of Engagements Per User	2.42	2.61	6.12
Median # of Engagements Per User	2	2	5
% of Total (Non-Outlier) Engagements	12%	24.8%	63.2%

What is clear from the data, once outliers are excluded, is that participants playing multiple roles are by every measure between 2 to 3 times more active than the participants who only act as requesters or writers. The community, therefore, is demonstrably made up

primarily of people playing the roles of both users and describers, not one or the other. In terms of alt text creation, that means that beyond the belief that with enough education users may always be turned into describers (see 6.3.2.), in practice PDDS members wear multiple hats. The PDDS members in the sample moved back and forth frequently between offering advice and feedback as co-authors of image descriptions, and requesting help from others to write descriptions they as users want but cannot create by themselves.

The third and final research question of Study 3 revealed that the roles of artist, describer, and user overlapped and complimented each other in the context of the PDDS accessibility community. The community pushed for and featured examples of artists also describing their own work, sometimes with the help of the PDDS volunteers. The greater insight the artist had into the artwork was seen as an asset to and element of the description that could be best served when the artist also described their artwork. There were, in my observations, examples of artists—used to a visual medium—struggling to articulate in words what their art was meant to depict. This was another advantage of the back-and-forth nature of PDDS interactions, where artists could weigh in or edit suggested alt text but didn't have to produce it wholesale on their own. Similarly, it allowed volunteers, acting as readers of the alt text, to troubleshoot possible misunderstandings or mis-phrasings in the image description so they could be corrected before the author posted the original image. This was, as the PDDS members understood it, the best method (given Tumblr's branching reblog structure) to assure that downstream instances of the artwork would remain accessible.

Additionally, PDDS members both advocated for and provided themselves as examples of the fluid nature of the relationship between describers and users. Returning once again to the visibility and advocacy element of image descriptions, participants I spoke to explained that exposure to image descriptions as a consumer was what primed and eventually prompted them to begin writing their own descriptions. They learned by example and through practice what image descriptions were and how to create them, turning users and writers from separate roles into something else: an on-ramp, a chance for users to transition as seamlessly as possible into writing descriptions themselves.

Lastly, observing and quantitatively charting the relationship between the roles PDDS members officially claimed on the Discord server, versus the roles they played in practice as they interacted with other members, revealed that PDDS members largely played multiple roles. They worked as both volunteers or co-writers of image descriptions and as requesters or users of image descriptions. By flipping between these roles, they demonstrated that user feedback from anyone, even those who are not screen reader users, has some value for evaluating and improving image accessibility. Essentially, they operationalized a cooperative describing process similar to the Team-Write-A-Thon method from Study 2 on a much larger scale. The difference being that in the context of a crowdsourced approach to image accessibility, the strict consistency requirements seen at Google were not at play.

4.7 Takeaways: Making Connections Between Studies

From the findings of the three studies of this dissertation, there are a few points of resonance or connection to draw out. First, Study 1 discussed normalizing versus othering language as it relates to either in-group signaling or out-group education. Considering this

in the context of the alt text writers I spoke to in Study 2, it becomes clear that the further separated—socially, culturally, literally—the alt text writer is from the subject and from the reader, the more potential there is for misunderstanding or othering. Reflecting on W3's comment that she chose to make the image descriptions aracial because the intent was for the images to be "relatable," it seems corporate image describers are, as members of techno-capitalist system founded on racism [175], likely to reinforce color-blind racism and other examples of "tech liberalism" that make blind people of color feel unwelcome or on-edge in digital spaces [18]. The perspectives and practices that (largely White) accessibility practitioners in the tech industry follow are still, no matter how vocally focused on inclusivity, likely to further marginalize users who are already at the intersection of marginalized identities [73].

The desire to consider the points of view of marginalized users of all kinds, including but not limited to disability, was present but ultimately unrealized in Study 1 as originally conducted. By requiring that all participants had a disability, but not seeking out participants with non-disabled identities that were represented in the Avatar Project set, I privileged the experiences of White disabled users over non-disabled users of color. However, after exploring the role of alt text writers more, in Study 2, it became clear that there is more to evaluate when considering the inclusivity of imagery than simply the concerns of screen reader users. This implies that there is a difference between analyzing a set of images for maximum accessibility and an evaluation to understand the inclusion of the images in relation to any and all identities.

While inclusive imagery is defined by the claim that accessibility is part of inclusion, accessibility and consultation with disabled users is not the summit of inclusion. Instead, as Study 3 showed, all users have a unique combination of perspectives, experiences, and knowledge which gives them insight into the appropriate way to depict or describe identities in inclusive imagery. While the organization of PDDS is not necessarily optimal for getting responses from the most qualified users for each particular image, the three studies together suggested a new potential imagery creation process that intentionally sought out users who could speak to all marginalized identities depicted in the images.

Additionally, evaluating the accessibility of imagery should ideally consider more the user experience for more than just screen reader users. As surveys of accessible computing show, studies of blind and low vision users represent a disproportionate amount of research on accessibility [152,171]. And yet, as prior researchers [171] and the participants from Study 3 suggest, imagery has accessibility implications for more than just blind and low vision users. Thus, the studies point toward a gap in work on image accessibility wherein disabilities other than just blindness may be important to consider.

I also wanted to return to one particular element considered in Study 1 and connect it to Studies 2 and 3: depicting non-visual identities. What Study 2 reinforced was that the Artist-Writer process, later included in the broader Formal Iteration Method, has a clear advantage not discussed elsewhere. When artists work with or as describers, the identities they intended to depict, including invisible identities such as “model” or non-apparent disabilities like chronic pain, can be described accurately and completely, without guesswork, additional research or cognitive load falling on the describer. Study 3 and the

PDDS online accessibility community provided examples of how non-apparent identities could be found out in other cases. If the image being described depicted a character with a canonical non-apparent disability, then the describers became responsible for knowing that additional information and deciding when or if mentioning a non-apparent disability was relevant to the image description.

Study 1 and Study 3 also resonated in interesting ways. The fact that U15* from Study 1 literally described the ideal way she would want an image of her to be described was through crowdsourcing directly compliments the work of the PDDS. U15* explained that having multiple people collaborating to describe an image of her would lead to their opinions averaging out. This exact process, a sort of "wisdom of the crowd" [131], was observable in how PDDS members course-corrected each other away from incorrect or biased descriptions. Similarly, the advice U17* gave based on his experiences with memes on Facebook (namely, that people should describe the source of the work first so that users can know if the meme is interesting to them) exactly echoed the advice PDDS members gave for the first sentence of image descriptions depicting fictional characters. These examples of users in Study 1 preemptively justifying and supporting the methods for describing images observed in Study 3 are a wonderful case of completely independent verification of the PDDS system.

The online accessibility community studied in this chapter made explicit a nascent distinction seen in Study 1 regarding elements of the image accessibility issue. As I discussed in Chapter 2, the first and greatest accessibility problem, by most participants reckoning, was simply a lack of alt text. Research supports the conclusion that high quality

alt text could not possibly be delivered by most users because for the most part non-blind or non-disabled users have no base knowledge of alt text or image accessibility broadly [214]. The PDDS members I spoke to understood this clearly. For the PDDS the image accessibility problem was actually two-fold. There is the obvious pragmatic problem of producing and distributing alt text or image descriptions. But there is also the problem of education or outreach, whereby everyday users (particularly on a site rich in user generated content such as Tumblr) have to be made aware of what image accessibility is, why they should care, and what they can do as individuals to support it.

In short, Study 3 detailed not simply the processes used to create accessible imagery, but it also explored the process by which image accessibility could be supported as a shared cultural practice. The promise of the PDDS approach lies not simply in how they create image descriptions, but how they move image description creation from a niche area of effortful expertise within a community of practice, into what Bourdieu calls "habitus": the unselfconscious norms that are "internalized as second nature" within a certain social field [46]. Habitus refers to "ways of seeing, moving, talking and so on" which is highly relevant here. Essentially, when PDDS members begin to internalize the need for image descriptions, to complete them through habit rather than conscious effort, they are reorienting themselves to a blind way of seeing despite not using a screen reader themselves. They are shifting the shared "common sense" within their community (and eventually, hopefully, Tumblr at large) towards a disabled point of view. They are asserting that how blind people interact with a website should be a required and self-evident consideration when sharing or uploading content on a digital platform, in the same way

screen reader users themselves cannot interact with any digital interface without considering visual accessibility.

And yet, this potentially influential change is occurring in a community already composed of many disabled users. Whether it is traditional paid crowdsourced labor, through Amazon Mechanical Turk for instance [119], or if researchers try to replicate the community of practice I documented in Study 3, it's important to remember that crowdsourcing can never be the perfect solution as long as the labor is distributed inequitably. That is, as long as the unpaid labor of accessibility [37,60,248] or paid crowdsourcing labor [245,263] falls disproportionately on the shoulders of already marginalized (particularly disabled) users, it will not be an equitable solution to accessibility remediation. And yet, disabled users as a whole have the most expertise when it comes to accessibility. This paradox is worth struggling with.

Disabled users should not be the sole or even primary drivers of accessibility remediation, given they are already burdened by inaccessibility and ableism in their personal lives. Yet simultaneously, accessibility practice should be informed as much as possible by people with lived experience of disability and not simply imagined or simulated [19] approaches to what the experience of inaccessibility is like. This was the exact same problem that PDI1 presented when asked how or if people with visual disabilities should be part of the PDDS community. Essentially eir answer was, yes, they should be able to, but no, they shouldn't be required to. This is currently the best answer for the question of integrating people with disabilities into accessibility practices at companies as well.

Because, as Study 2 showed, the greatest need for the expertise of disabled people is where there are currently the fewest pathways to get it.

Large companies such as Google can't or don't take the steps to find and interact with disabled users to the degree they should, or even to the degree that many of the Artists and Writers I spoke to desired. User feedback is wanted and needed by practitioners in professional content creation contexts, and yet the exact users who have the expertise in image descriptions and accessibility are creating the accessibility themselves rather than connecting with companies as consultants. The disconnect between the participants in Study 3, who were so passionate and well-practiced but lacked the resources, time, or support to fully invest in image descriptions, versus the organization supporting Study 2, with its vast resources simply not funneled into supporting inclusive imagery creation, is frustrating and disheartening. Even the writers within Google were not as supported as they could have been, organizationally, in their drive to create accessible imagery. There should be platforms or technologies created to make sure larger content creation engines are tapped into and learning from people with expertise in accessibility, including the volunteer crowd workers that make up Tumblr's accessibility community. Additionally, to offset the unpaid labor burden taken on largely by people with disabilities, these collaborations between companies and users should and must present value (ideally monetary value, along with other kinds of value) to the users participating in them.

Study 3 also suggested that there is a place for cultivating knowledge and empathy in non-disabled accessibility advocates. PDDS community members–as fellow Tumblr users, as fellow converts to the topic of accessibility, and often but not always as fellow

people with disabilities–excelled at empathy. They demonstrated empathy both towards screen reader users, and towards people uninformed about accessibility. Community-based approaches to accessibility, where people with disabilities are not simply consultants or users but fellow members of a shared space, have promise for addressing both the shortage of knowledgeable accessibility practitioners and the traditional problems with empathy [19]. The important element of the paradox then is not that people with disabilities be given the option to disengage, but rather that those who do choose to engage are met fully and reciprocally by the organizations seeking to ameliorate inaccessibility. In that sense the solution is not a platform or even a process to support user feedback, it's a genuine community-based engagement between technology workers and users as equals. How best to enable such an engagement, or how to assure they remain non-exploitative, is a question for future researchers.

The final connection to be drawn between image accessibility in practice by technology companies and by online accessibility communities is simple but important: practice is needed. From the content writers who W1 said "only end up writing alt text once a year [so] they are not going to get good at it", to the PDDS members first getting acquainted with how image descriptions work, image accessibility is made through practice. Thus, even without community-based engagement with disabled users, improvements may be made in technology companies' image accessibility simply by fostering communities of practice among their employees. Communities of practice are highly valuable for learning and sharing the everyday expertise of how to complete a complex job [78]. And if corporations such as Google want advice for how to support and develop communities of practice, there are resources and research on how organizations

can cultivate them [147,253]. For image accessibility in particular, a community of practice makes sense as a starting point because it facilitates collaborative describing processes as well as assuring that members get exposure to and practice with the specific guidelines and nuances of writing alt text.

Ultimately, the three studies of this dissertation dovetail into a cohesive look at how image accessibility works as a process and practice in different contexts. From a co-design project with disabled users, to the multiple modes of working demonstrated at Google, to the volunteer crowdsourcing of PDDS online accessibility community, all three studies demonstrate that collaboration is essential for inclusive and accessible imagery to be produced.

In the next chapter, I will discuss implications and future work arising from this dissertation.

CHAPTER 5: Conclusion

This dissertation has provided three studies of alt text creation processes and the complexities therein. Moving from the first study onward, my work allowed for an increasingly complex view of how individuals and communities of practice create alt text in different contexts. In this chapter, I discuss the implications this dissertation as a whole and its findings present for future work by image accessibility researchers and social media sites.

5.1 Implications for Future Work

5.1.1 Implications for Design: Pushing Social Media and Other Image-Hosting Sites Towards Sophisticated Accessibility Features

Given the pragmatic aims of this project and the importance of quality image descriptions to BLV Internet users, this section provides a set of feature recommendations arising from these three studies and their implications for companies involved in hosting or displaying alt text.

First, interfaces for alt text viewing and alt text authoring must become more sophisticated. The example U15* gave of "glancing" at a photo as you're looking at a friend's Facebook profile leads to a feature suggestion for different versions of alt text to be delivered to users based on context. This is supported by research such as Stangl et al. [227,229], which argues context should be considered more fully when displaying alt text. Stangl et al. [227] focused on the idea that images within the context of different websites have different informational requirements, but even on the same website—as U15*'s example shows us— different pages might imply different levels of engagement with images on that page. Having a friend's profile picture accompanied by lengthy alt text every

time it appears next to a post is redundant at best and annoying or even disruptive at worse [92]. Instead, having images that appear multiple times on a page only deliver the full alt text the first time a screen reader encounters them may be more useful. Or even having a short version of the alt text be the only kind available for profile pictures on a main feed, saving the more detailed version for a profile page. These profile pages may then also have more complex alt text features (e.g., blind users could interact with voice recordings of users narrating their own alt text [168]) that would satisfy users like U10* who wanted to be "inundated with audio clips" instead of images. There are examples on the other end of the spectrum, such as the scenario W1 described, where artists and image creators are "very deliberate in choosing" what to include in imagery in a way that might be meaningful to screen reader users if not for the fact that the "alt text treats the imagery as though it were arbitrary." Thus, it may be valuable to produce alt text interfaces where a user can access different kinds of alt text. Writers would also thereby be afforded opportunities to create concise to more sophisticated versions. Such advanced interfaces for alt text have been discussed by alt text researchers [166], but to date there are few if any examples to point to. The practical elements of how to institute these interfaces or provide distinct versions of alt text for the same image are still to be explored.

Once you open the opportunity for different versions of the alt text for the same image, more feature options are suddenly possible. For example, there could be value in alt text "localization" from companies with a large enough reach. Research shows that other common website features such as emoji [127] in fact have different interpretations based on culture. So it may make a meaningful difference if, for instance, UK users were able to access alt text that used the more locally popular identity-first terminology [154] while

users in the US could consume the more generally accepted person-first language in their alt text. The alt text and image descriptions studied in this dissertation were exclusively written in English. And yet, for public-facing images such as those produced by Google or for posts on an international social media platform such as Tumblr, there is no reason not to consider if and how alt text in other languages or for other cultural contexts could be produced and incorporated into image accessibility practice. On a purely technical level, an image's HTML "alt" tag, where alt text is stored for most websites, has only one field and contains only one string. There is no way, to my knowledge, for users to set language or location preferences on a screen reader and then be presented with only the alt text that is in their desired language. Even if such a feature is possible on screen readers today, considering getting broad coverage of alt text in even one language is still so problematic [25,167], the likelihood that non-English-speaking screen reader users can currently access many popular images through image descriptions in their native language is slim to none.

There's an additional consideration for companies that produce alt text. Alt text is part of the overall user experience of an interface. And generally, user experience researchers take for granted that users must be involved in assessing user experience. Yet, alt text or accessibility generally is not necessarily tested as part of general user testing procedures. In a study discussing ways that user researchers include accessibility in their work, only 18 of the almost 60 interviews with user experience researchers included mention of conducting actual usability testing with people with disabilities [1]. The number of those usability tests that specifically assessed alt text as an element of user experience is likely much lower. Referenced in the takeaways from Study 2 and pulling heavily from the finding in Study 1 that users primarily knew the right length for alt text only when

encountering it, my research suggests users can't necessarily say in advance how much detail is the correct amount. Instead, it's more an opinion formed in the moment of consuming it. Thus in-situ user research methods and prototyping tools should be created so that screen reader users, can give feedback on alt text in the context of a product.

Included in and complementing tools such as these, there's a need for developer- or designer-facing features for "previewing" alt text with various screen readers. Better website accessibility is better user experience in general [3], so there is a natural incentive for companies to improve alt text. Yet actually testing alt text with screen readers, as members of PDDS found and as organizations such as WebAIM suggest, is very complicated for the average, non-screen-reader-using, technologist to accomplish [251]. In fact, a 2021 survey of web accessibility practitioners found that only 13.2% of respondents without a disability reported being very proficient using a screen reader [252]. Operating a screen reader even successfully enough to test your alt text can be difficult. If there were more features that allowed designers, UX writers, developers, and content creators to simply test the alt text element of their page with different types of screen readers, then many issues could be caught before they make it to the ears of screen reader users themselves.

Along with testing alt text, one PDDS member (PDI1) suggested a novel alt text design tool feature worth discussing. E proposed that having an interactive interface for collaboratively creating alt text on a social media site such as Tumblr could be useful for alt text consumers and writers alike. Just since I conducted my interviews, much of the "low-hanging fruit" that PDDS members described needing for image accessibility, have already been addressed by Tumblr staff. The major adjustment made is the addition of a visual

indicator on the bottom left-hand corner of images if they have alt text. This button can also be pressed to see the alt text, allowing alt text to have some of the same benefits that PDDS members argued were specific to image descriptions. However, no matter how improved the features are for displaying and signaling the presence of alt text, the PDDS community functionally could not rely on Tumblr alone for their collaborative alt text authoring because it lacks the right kind of general interactive features. That is why Discord, as a supplementary platform, proved necessary for PDDS volunteers to coordinate and collaborate through. Having ways to crowdsource image accessibility, ask questions about images, or up-vote and down-vote image descriptions within Tumblr itself could make for a unique combination of user engagement and accessibility feature. There is a need, therefore, for a new type of interactivity that combines the collaborative, in-depth accessibility improvements that can be headed by cohesive groups [207] with the direct embedded nature of PDDS's engagement with and through Tumblr.

Customizable interfaces for both alt text viewing and alt text authoring are deeply needed. For example, screen reader users should be able to set preferences, either universally or per application or website, for more details or less detailed alt text to improve the user experience. W1's explanation that sometimes images are neither completely informative nor completely decorative, but something in between, implies that there may be need for a setting (at the time of alt text creation and consumption) to flag such images and handle them differently—a type or style of alt text purely for those semi-decorative, or skimmable images. Or even, as many describers on Tumblr argue, it may be useful for alt text to support an indicator of what an image *is not*. Screen reader users are impacted not just by the information they do not have access to, but also by not knowing

what type of information they are missing in the first place [24]. And yet, official guidelines for alt text say that decorative images should simply have an empty alt tag, so users who encounter these images cannot necessarily know if an image lacking alt text is decorative or if alt text was mistakenly omitted for that image. An image on Tumblr can be an incredibly valuable interpretive tool, or it can be near random and meaningless. There should be a way of letting screen reader users know at a "glance" what images are or are not relevant, with the option to dig into even those relatively irrelevant images if they are interested.

5.1.2 Implications for Image Accessibility: Future Topics and Considerations for Researchers

While this project as a whole is guided by pragmatic interest in the way alt text is made in practice, I find the research implications to be highly important to the outcome of this project. Researchers must consider processes and roles with a focus not only on the end product, but also on how every player in the system is being treated, burdened, or replaced. Together with research on collaborative creation of accessible content [150,159], this dissertation marks a new area of study for accessible computing researchers focusing on how alt text is created in different contexts, by different actors, and within different limitations. Researching how alt text is made in high stress situations versus carefully considered content creation processes, by paid crowd workers versus disability activists, or within corporate structures versus by unconstrained individuals, all allows attention to be paid to both the human actors in the system and the ethics or inclusivity of that larger system. Some may say that ultimately it is the outcome, the accessibility of images, that matters. But accessible computing, and HCI more broadly, has argued that innovation

without care [20,198] and labor without ethics [119,262] can be worse than useless. If we continue as a field to study only the newest technical approaches and the fastest available methods for creating alt text, then we run the risk of creating systems that are not just poor quality, but actively hurt marginalized users, disabled users, crowd workers, content creators, and anyone involved in digital imagery creation or sharing.

To return to a point touched on in Chapter 4, an algorithm may make less (or more likely simply *different*) errors than human volunteers such as those participating with the PDDS. Algorithms for creating alt text will and do make errors of mis-categorization and misrecognition [153,210]. Humans, at least based on the examples PDDS participants gave in Study 3, are more likely to make errors of incompleteness or hesitance. An AI cannot have an emotional response triggered by an image or piece of artwork. But if a human is trying to create image descriptions, we have to contend with the fact that they might, and indeed should, avoid certain content that is disturbing to them. Or they may limit the spread of their accessible images by blocking someone who they feel personal discomfort around. These are, in a way, flaws in a perfectly accessible system that are introduced only when real people are tasked with creating it. And yet, human error is part and parcel of human care. While this dissertation did not speak on the topic, as this document is being finalized ChatGPT and other generative AI promise a new era of automatic alt text, where both the accuracy and the availability could be significantly improved. However, even setting aside issues of making sure high quality and unbiased training data is available or that the AI could be relied upon not to "hallucinate" details of an image [7], the most technically accurate approach to creating alt text will never be the most inclusive. Subjects' real identities cannot be categorized by any computer vision algorithm, no matter how

advanced. Because the messiness of people's identities defies distinction into separate, visually determinable boxes [210]. When a person cannot or will not function as their ideal cog in the perfect accessible system—for example, by blocking another user from accessing their content on personal grounds even if it hinders the free-flow of accessible content—it is proof that existing within a network of human beings (society itself) will always be more complicated (but also more worthwhile) than creating a network of perfectly coded machines.

The inherent messiness of human data is relevant to image accessibility research in two ways. First, as research has already described, it hopelessly complicates any attempts to generate neutral or objectively accurate datasets [208] from which to build automatic alt text generation systems. Secondly, it directs our attention back to the workers who make up the backbone of manual generation of alt text [95,261]. The issue of triggering content was dealt with compassionately by the PDDS, which is unsurprising when considering Tumblr's history of social justice [161]. But how might the same issues be considered, or not, when contracting traditional micro-workers to create image descriptions? Considering the abundance of private [228], violent, or otherwise disturbing data that content moderators across the world are already exposed to [196,231], there is little doubt that workers engaged to create image descriptions would be called upon to describe potentially similarly challenging content. Unlike in the PDDS where the overarching motto was that "it's your blog" and therefore your own mental health and well-being comes first, how will crowd workers, particularly in research interventions such as the long-term VizWiz project [22,32,221] be implicitly or explicitly pressured into viewing content that disturbs them? If

a system delivers accessibility at the cost of the well-being of the people creating the accessibility, then I don't believe it is a good solution.

Despite the drawbacks of algorithmic methods and the datasets they are based on, my research on alt text creation in industry contexts does suggest the need for more work on the specific circumstances and contexts where automated or semi-automated tools may be helpful. For example, creating and testing tools to help standardize language in contexts where repeated constructs, portions of interfaces, or other similar elements appear again and again might be valuable in industry contexts. These tools could decrease the time and effort required to reach consensus and guarantee consistency with collaborative image describing processes such as the Team Write-A-Thon Process. The details of how such a system could work and how to assure technical accuracy is an entire research project in and of itself. However, more important than researchers developing tools and measuring success in terms of the number of image descriptions generated, is developing an understanding of *when* and *how* automated or semi-automated methods can actually benefit alt text quality and facilitate human connection. Research has already found that presenting automated suggestions for alt text and then having a human edit it leads to lower quality results than simply having the human write it from scratch [150]. Thus, along with developing the tools, there is also research to be done into how these tools can be designed, presented, and integrated to work with existing creation practices in complementary rather than distracting ways. Alt text creation processes such as the initial four described in Chapter 3 do, to varying degrees, work as they are. Researchers can now begin to investigate how to support and streamline rather than degrade or replace these existing alt text creation practices and roles.

Returning a moment to the importance of roles in alt text creation processes, one of the largest problems with existing research (e.g., [95,100,259]) was not just that it focused on interventions, but that it took a very simplistic view of how alt text functions theoretically. Bennett et al. [18] were among the small minority of image accessibility researchers who even directly considered the role of the subject in creating alt text. The vast majority of alt text interventions focused only on the screen reader users consuming the alt text and the writers creating it. But as the three studies of this dissertation demonstrate, the reality of who is involved in delivering accessible imagery to end users is much more complicated. Being aware of these additional layers of interpretation or potential players implicit in image accessibility opens up a wide range of new research questions. For example, from the standpoint of algorithmic fairness, how can the intentions of the artist be systematically considered when creating datasets that include non-photographic imagery? How can imagery consumers (i.e., screen reader users) with different orientations towards the image (in-group, out-group, sharing an identity with the subject, sharing a marginalized identity in particular) be properly considered or accommodated for when creating or teaching crowd workers to make alt text? Researchers should consider not simply the website alt text is embedded in [227] or even the scenario of the user's information needs [229], but how the context of posting plays a role in perception of imagery. Is it user generated content or was it part of a professional content creation process? Was the describer under time or social pressures? If yes, how does that change readers' expectations for quality level or imply certain societal bases for bias in the resulting alt text? And finally, research should consider the context of the user's life

experience, social identities, mood, and/or physical surroundings at the time of consuming the alt text and how that might impact the way alt text is interpreted or understood.

In the takeaways from Study 1, reader response theory was mentioned. Studies of alt text are begging for more attention to be paid to the interpretive complexities of accessible images as artifacts. There should be research that studies alt text on Twitter not just to see if it is there or determine its basic quality level [94], but to study it as a source of meaning. What topics are "accessible Twitter" talking about? How are they talking about them? What sorts of language, formality levels, perspectives, and norms do accessible images on Twitter or any other platform demonstrate or reflect? Taking alt text as writing, as poetry [271], as meaningful content able to be and worthy of being analyzed and studied in as much detail as any novel or any pop cultural artifact, is something that image accessibility researchers have not yet begun to do.

Additionally, inspired by the discussion of "habitus" in Chapter 4, there is an opening for new research that takes the argument of the PDDS seriously: there in no accessibility without visibility, no win condition that does not include cultural shift. AI-generated and on-demand crowdsourced alt text (e.g., [95]) focuses on improving the in-the-moment experience of screen reader users. This is, without a doubt, an important goal—and indeed considering how AI can be used to scaffold alt text creation may be an interesting topic to study—but if we focus on "just in time" interventions, we will always be lagging two steps behind the content creation itself. We will be creating accessibility only after the content is posted, sometimes even later, when the image is encountered by a screen reader user. Embedded communities, such as accessibility advocates on Tumblr, can

do better. They can still do their best to catch up with new content, but often have the chance to do so the moment they see it, rather than when directly requested. And online accessibility communities are doing more than addressing specific instances of inaccessibility, they are changing the user consciousness around accessibility on their target platforms. Current AI approaches can without a doubt be helpful for addressing the enormous backlog of images posted online, as W2 pointed out. But even platform-level technical features such as Facebook's automatic alt text will never fundamentally change the culture that produces inaccessibility. An accessibility community like PDDS wants enough people to practice accessibility that further outreach, education, and dedicated labor to produce accessibility is unnecessary. They want to shift the platform's community norms and even the wider societal norms such that accessibility is passively practiced as a matter of course. This is in stark contrast to traditional alt text interventions, which are designed to be unobtrusive, automated, and streamlined.

5.1.3 More Accessible Futures Through Obtrusive Design

This dissertation, and particularly the PDDS's emphasis on image descriptions over alt text, suggests one additional orientation towards accessibility that might be valuable for both designers and researchers. I'm calling this approach *obtrusive design*. I define obtrusive design as new features or creative use of existing features within technical systems to make (in)accessibility unignorable for the average (non-disabled) user. The PDDS members use the existing reblog feature (rather than alt text) for image descriptions because doing so makes the labor of accessibility, and the consequences of inaccessibility, starkly visible to anyone who consumes the post, rather than just those using screen readers. Tumblr's "improved" approach of having a small button to indicate there is alt text,

and which you can click to display it, is effective if the user knows it is there. It is integrated enough into the system that it functions as intended when people use it. But the intentionally unobtrusive nature of this and other common accessibility features means that users have to go out of their way to find it or find out about it. This fundamentally does not force a culture shift in the same way the visible and unignorable presence of image descriptions in the body of Tumblr posts has. While there is research that importantly points out the importance of social acceptability in assistive technology design [217], I'd like to suggest we at least consider an alternate approach. Obtrusive design, in the form of hyper-visible alt text displays or sophisticated features for co-authoring image descriptions, would not just improve accessibility on specific social media platforms, but establishes a design approach that may be a valuable potential avenue for future work by image accessibility researchers.

Given the value of cultural shift that my research of the PDDS revealed, perhaps non-disabled people *need* to see accessibility in action. Perhaps representing disabilities in imagery isn't enough. How, then, could future image accessibility researchers develop interventions that did not fade into the background? What if the very interfaces you used every day did not allow you to forget that disabled people used them? If not just designers, but every technology user, had to think about disabilities as they went through the world every day, how long do you think it would take before they became attuned–as disabled users are already attuned–to the inaccessibility all around them? How long after that would they be willing to change their own actions in order to address it? In the style of critical or adversarial design [64], how could interventions intentionally make visible the labor of access?

There are ways to engage attention in the short term, as a form of disruption or protest harkening back to the 504 sit-ins of disability activists before us [220]. For example, one Tumblr post I encountered in my data collection for Study 3 was a series of five large images. In large, black, capitalized text against a grey background each one read "Alt Text Unavailable For This Image". Following the images (which due to their size took a decent amount of time to scroll past) was a brief image description, a note, "reminder that this is the typical online experience for blind folks," a genuine plea for people to add image descriptions to "every image that you share. here and on every other social media platform. every time," and some resources including a link to join the PDDS.

Obtrusive design is not only newly designed features, but also provocative usage of existing features with the goal of increasing the awareness of accessibility. For example, imagine a massive protest on Twitter[5] or Instagram or a similar platform, where all accessibility researchers, advocates, and allies for a period of time posted purely placeholder images with alt text describing the original content. Leveraging only the existing image uploading and alt text writing features, combined with social connections and interpersonal investment in introducing and explaining the motivation for the protest, could in this way very well spread the message of image accessibility to far more people and in a far more unambiguous fashion than past attempts at education. Or imagine having a platform, even for a few days, refuse to allow users to upload images unless they had some kind of alt text. I can almost imagine the frustrated videos people would post of the error message repeatedly popping up informing them that they could not upload the image

[5] Officially re-named "X" at present.

without alt text. While this particular idea runs into the issue PDI1 raised that there are people who cannot write alt text because they are the people for whom alt text is *for*, the potential for massive disruption and displeasure among sighted users is intriguing to me. These users would, just for a short period of time, be forced to face the inaccessibility that they are otherwise allowed to passively contribute to.

Obtrusive design, as a form of protest, has the potential to be very divisive, but the potential rewards are also great. Critical design and more recent off-shoots of it have already convincingly argued that there is value in using technical systems to call attention to [66] and prompt reflection on [12] our current societal issues. And yet, to my knowledge, there is no "critical accessibility design" or critical design focused on raising awareness of and support for accessibility. I urge image accessibility researchers to consider how art pieces, provocations, and interventions, could increase the awareness of everyday social media users, in both the short term and long term, of their role in producing digital (in)accessibility.

Beyond the provocative actions imagined under "obtrusive design," researchers should consider if there are other ways to follow in the lead of the PDDS. Are there other potential solutions to image inaccessibility that are highly visible and seamlessly integrate non-disabled and non-marginalized users into co-creating a more equitable and accessible future? If future image accessibility efforts and interventions can elicit the same empathy and engagement from general users that the PDDS has begun to achieve on Tumblr, then accessibility has the chance to move from charity (non-marginalized users helping marginalized users out of pity), past PDDS's mutual aid (marginalized users helping each

other reciprocally) and into actual a new frontier of social change. This next frontier of activism and research may entail efforts to shift accessibility into habitus: expected, common sense, taking no more additional effort than it takes to exist in social spaces generally.

5.2 Conclusion

This dissertation has presented three case studies investigating different processes by which accessible imagery is made in practice. Study 1 covered a study wherein researchers, collaborating with a technology company, co-designed image descriptions with 25 users with disabilities. Study 2 documented processes for creating alt text used among a small sample of Google artists and accessibility practitioners. Study 3 investigated an online accessibility community that collaboratively produced accessible posts for a social media site through volunteer crowdsourcing. Finally, this chapter suggested directions for future innovation and research around image accessibility. I emphasized designing more sophisticated alt text interfaces and more common alt text authoring or testing tools. I suggested studying different alt text creation contexts and alt text's theoretical and textual value. And lastly, I argued it was valuable to explore "obtrusive design" to spread awareness of image accessibility to broader audiences and prompt cultural shifts not possible as long as accessibility is hidden from most internet users.

REFERENCES

[1] Accessibility Guidelines Working Group. 2018. Success Criterion 3.2.4 - Consistent Identification. *Web Content Accessibility Guidelines (WCAG) 2.1*. Retrieved May 5, 2023 from https://www.w3.org/TR/WCAG21/#consistent-identification

[2] Dragan Ahmetovic, Roberto Manduchi, James M. Coughlan, and Sergio Mascetti. 2015. Zebra Crossing Spotter: Automatic Population of Spatial Databases for Increased Safety of Blind Travelers. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '15), ACM, New York, NY, USA, 251–258. DOI:https://doi.org/10.1145/2700648.2809847

[3] Amaia Aizpurua, Simon Harper, and Markel Vigo. 2016. Exploring the relationship between web accessibility and user experience. *International Journal of Human-Computer Studies* 91, (July 2016), 13–23. DOI:https://doi.org/10.1016/j.ijhcs.2016.03.008

[4] Ayad Al-Ani and Stefan Stumpp. 2016. Rebalancing interests and power structures on crowdworking platforms. *Internet Policy Review* 5, 2 (June 2016). DOI:https://doi.org/10.14763/2016.2.415

[5] Aloha Hufana Ambe, Margot Brereton, Alessandro Soro, Laurie Buys, and Paul Roe. 2019. The Adventures of Older Authors: Exploring Futures through Co-Design Fictions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, ACM Press, Glasgow, Scotland Uk, 1–16. DOI:https://doi.org/10.1145/3290605.3300588

[6] American Psychological Association (APA). 2022. Disability. *https://apastyle.apa.org*. Retrieved February 5, 2023 from https://apastyle.apa.org/style-grammar-guidelines/bias-free-language/disability

[7] Razvan Azamfirei, Sapna R. Kudchadkar, and James Fackler. 2023. Large language models and the perils of their hallucinations. *Crit Care* 27, 1 (December 2023), 1–2. DOI:https://doi.org/10.1186/s13054-023-04393-x

[8] Shiri Azenkot, Margot J. Hanley, and Catherine M. Baker. 2021. How Accessibility Practitioners Promote the Creation of Accessible Products in Large Companies. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1 (April 2021), 1–27. DOI:https://doi.org/10.1145/3449222

[9] Shiri Azenkot and Nicole B. Lee. 2013. Exploring the Use of Speech Input by Blind People on Mobile Devices. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '13), ACM, New York, NY, USA, 11:1-11:8. DOI:https://doi.org/10.1145/2513383.2513440

[10] Tim Babinszki, Anna Cavender, Michael Gower, Jeffery Hochi, Darcy Lima, Erich Manser, and Shari Trewin. 2019. Inclusive Writing. In *Web Accessibility: A Foundation for Research*, Yeliz Yesilada and Simon Harper (eds.). Springer London, London. DOI:https://doi.org/10.1007/978-1-4471-7440-0

[11] Mark S. Baldwin, Sen H. Hirano, Jennifer Mankoff, and Gillian R. Hayes. 2019. Design in the Public Square: Supporting Assistive Technology Design Through Public Mixed-Ability Cooperation. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (November 2019), 155:1-155:22. DOI:https://doi.org/10.1145/3359257

[12] Eric P.S. Baumer. 2015. Reflective Informatics: Conceptual Dimensions for Designing Technologies of Reflection. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), Association for Computing Machinery, New York, NY, USA, 585–594. DOI:https://doi.org/10.1145/2702123.2702234

[13] Pamela Baxter and Susan Jack. 2008. Qualitative case study methodology: Study design and implementation for novice researchers. *The qualitative report* 13, 4 (2008), 544–559.

[14] Ron Beacker, Kate Sellen, Sarah Crosskey, Veronique Boscart, and Barbara Barbosa Neves. 2014. Technology to reduce social isolation and loneliness. In *Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility - ASSETS '14*, ACM Press, Rochester, New York, USA, 27–34. DOI:https://doi.org/10.1145/2661334.2661375

[15] Ruha Benjamin. 2019. *Captivating Technology: Race, Carceral Technoscience, and Liberatory Imagination in Everyday Life*. Duke University Press.

[16] Yochai Benkler and Helen Nissenbaum. 2006. Commons-based Peer Production and Virtue. *J Political Philosophy* 14, 4 (December 2006), 394–419. DOI:https://doi.org/10.1111/j.1467-9760.2006.00235.x

[17] Cynthia L. Bennett, Erin Brady, and Stacy M. Branham. 2018. Interdependence as a Frame for Assistive Technology Research and Design. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '18), Association for Computing Machinery, Galway, Ireland, 161–173. DOI:https://doi.org/10.1145/3234695.3236348

[18] 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*, ACM, Yokohama Japan, 1–19. DOI:https://doi.org/10.1145/3411764.3445498

[19] Cynthia L. Bennett and Daniela K. Rosner. 2019. The Promise of Empathy: Design, Disability, and Knowing the "Other." In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (CHI '19), ACM, New York, NY, USA, 298:1-298:13. DOI:https://doi.org/10.1145/3290605.3300528

[20] Cynthia L. Bennett, Daniela K. Rosner, and Alex S. Taylor. 2020. The Care Work of Access. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, ACM, Honolulu HI USA, 1–15. DOI:https://doi.org/10.1145/3313831.3376568

[21] Lucy Bennett. 2012. Fan activism for social mobilization: A critical review of the literature. *Transformative Works and Cultures* 10, (2012).

[22] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, and Samual White. 2010. VizWiz: nearly real-time answers to visual questions. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, 333–342.

[23] Jeffrey P. Bigham, Ryan S. Kaminsky, Richard E. Ladner, Oscar M. Danielsson, and Gordon L. Hempton. 2006. WebInSight:: making web images accessible. In *Proceedings of the 8th international ACM SIGACCESS conference on Computers and accessibility - Assets '06*, ACM Press, Portland, Oregon, USA, 181. DOI:https://doi.org/10.1145/1168987.1169018

[24] Jeffrey P. Bigham, Irene Lin, and Saiph Savage. 2017. The Effects of "Not Knowing What You Don't Know" on Web Accessibility for Blind Web Users. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Baltimore Maryland USA, 101–109. DOI:https://doi.org/10.1145/3132525.3132533

[25] Syed Masum Billah, Vikas Ashok, Donald E. Porter, and I.V. Ramakrishnan. 2017. Ubiquitous Accessibility for People with Visual Impairments: Are We There Yet? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ACM, Denver Colorado USA, 5862–5868. DOI:https://doi.org/10.1145/3025453.3025731

[26] Rena Bivens and Oliver L. Haimson. 2016. Baking Gender Into Social Media Design: How Platforms Shape Categories for Users and Advertisers. *Social Media + Society* 2, 4 (November 2016), 205630511667248. DOI:https://doi.org/10.1177/2056305116672486

[27] Tom Boellstorff, Bonnie Nardi, Celia Pearce, and T. L. Taylor. 2012. *Ethnography and virtual worlds: a handbook of method*. Princeton University Press, Princeton.

[28] Bert Bos. Design by committee (An essay on W3C's design principles). Retrieved April 18, 2023 from https://www.w3.org/People/Bos/DesignGuide/committee

[29] Elli E. Bourlai. 2018. 'Comments in Tags, Please!': Tagging practices on Tumblr. *Discourse, Context & Media* 22, (April 2018), 46–56. DOI:https://doi.org/10.1016/j.dcm.2017.08.003

[30] Erin L. Brady, Yu Zhong, Meredith Ringel Morris, and Jeffrey P. Bigham. 2013. Investigating the appropriateness of social network question asking as a resource for blind users. In *Proceedings of the 2013 conference on Computer supported cooperative work* (CSCW '13), Association for Computing Machinery, New York, NY, USA, 1225–1236. DOI:https://doi.org/10.1145/2441776.2441915

[31] Erin Brady, Meredith Ringel Morris, and Jeffrey P. Bigham. 2015. Gauging receptiveness to social microvolunteering. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 1055–1064.

[32] Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P. Bigham. 2013. Visual challenges in the everyday lives of blind people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13), Association for Computing Machinery, New York, NY, USA, 2117–2126. DOI:https://doi.org/10.1145/2470654.2481291

[33] Danielle Bragg, Nicholas Huynh, and Richard E. Ladner. 2016. A Personalizable Mobile Sound Detector App Design for Deaf and Hard-of-Hearing Users. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '16), ACM, New York, NY, USA, 3–13. DOI:https://doi.org/10.1145/2982142.2982171

[34] Giorgio Brajnik. 2008. A comparative test of web accessibility evaluation methods. In *Proceedings of the 10th international ACM SIGACCESS conference on Computers and accessibility - Assets '08*, ACM Press, Halifax, Nova Scotia, Canada, 113. DOI:https://doi.org/10.1145/1414471.1414494

[35] Stacy M. Branham, Ali Abdolrahmani, William Easley, Morgan Scheuerman, Erick Ronquillo, and Amy Hurst. 2017. "Is Someone There? Do They Have a Gun": How Visual Information About Others Can Improve Personal Safety Management for Blind Individuals. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '17), ACM, New York, NY, USA, 260–269. DOI:https://doi.org/10.1145/3132525.3132534

[36] Stacy M. Branham and Shaun K. Kane. 2015. Collaborative Accessibility: How Blind and Sighted Companions Co-Create Accessible Home Spaces. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), ACM, New York, NY, USA, 2373–2382. DOI:https://doi.org/10.1145/2702123.2702511

[37] Stacy M. Branham and Shaun K. Kane. 2015. The Invisible Work of Accessibility: How Blind Employees Manage Accessibility in Mixed-Ability Workplaces. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '15), Association for Computing Machinery, New York, NY, USA, 163–171. DOI:https://doi.org/10.1145/2700648.2809864

[38] Emma Briant, Nick Watson, and Gregory Philo. 2013. Reporting disability in the age of austerity: the changing face of media representation of disability and disabled people in the United Kingdom and the creation of new 'folk devils.' *Disability & Society* 28, 6 (September 2013), 874–889. DOI:https://doi.org/10.1080/09687599.2013.813837

[39] Melissa M. Brough and Sangita Shresthova. 2012. Fandom meets activism: Rethinking civic and political participation. *Transformative works and cultures* 10, (2012).

[40] E. Brulé and C. Jouffrais. 2016. Representing Children Living with Visual Impairments in the Design Process: A Case Study with Personae. In *Designing Around People*, Pat Langdon, Jonathan Lazar, Ann Heylighen and Hua Dong (eds.). Springer International Publishing, Cham, 23–32. DOI:https://doi.org/10.1007/978-3-319-29498-8_3

[41] Elizabeth Buie, Susan Dray, Keith Instone, Jhilmil Jain, Gitte Lindgaard, and Arnie Lund. 2010. How to bring HCI research and practice closer together. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems - CHI EA '10*, ACM Press, Atlanta, Georgia, USA, 3181. DOI:https://doi.org/10.1145/1753846.1753951

[42] Daniel G. Cabrero. 2015. User-created persona: Namibian rural Otjiherero speakers. In *Proceedings of the 33rd Annual International Conference on the Design of Communication - SIGDOC '15*, ACM Press, Limerick, Ireland, 1–6. DOI:https://doi.org/10.1145/2775441.2775484

[43] Jim A. Carter and David W. Fourne. 2007. Techniques to assist in developing accessibility engineers. In *Proceedings of the 9th international ACM SIGACCESS conference on Computers and accessibility - Assets '07*, ACM Press, Tempe, Arizona, USA, 123. DOI:https://doi.org/10.1145/1296843.1296865

[44] Vítor Carvalho and Diamantino Freitas. 2015. Automatic Description of SVG Images for the Visually Impaired: A Gestaltic Approach. *Procedia Computer Science* 67, (2015), 2–11. DOI:https://doi.org/10.1016/j.procs.2015.09.243

[45] Sarah Cavar and Alexandre Baril. 2021. Blogging to Counter Epistemic Injustice: Trans disabled digital micro-resistance. *Disability Studies Quarterly* 41, 2 (2021).

[46] Daniel Chandler. 2020. *A dictionary of media and communication / Daniel Chandler and Rod Munday.* (Third edition. ed.). University Press, Oxford.

[47] Kathy Charmaz. 2014. *Constructing grounded theory*. sage.

[48] Kyle Chayka. 2022. How Tumblr Became Popular for Being Obsolete. *The New Yorker*. Retrieved April 5, 2023 from https://www.newyorker.com/culture/infinite-scroll/how-tumblr-became-popular-for-being-obsolete

[49] Alex Chen. 2020. How to write an image description. *Medium*. Retrieved March 14, 2021 from https://uxdesign.cc/how-to-write-an-image-description-2f30d3bf5546

[50] Jeeyoung Chun and Thomas O. Williams. 2021. A Community of Practice for Professional Development in Technology Integrations for Accessibility: A Case Study of a Faculty Inquiry Group. *College Teaching* 69, 3 (July 2021), 126–137. DOI:https://doi.org/10.1080/87567555.2020.1832435

[51] Gabriele Cimolino, Sussan Askari, and T.C. Nicholas Graham. 2021. Beyond Fun: Players' Experiences of Accessible Rehabilitation Gaming for Spinal Cord Injury. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Virtual Event USA, 1–13. DOI:https://doi.org/10.1145/3441852.3471227

[52] Coppélie Cocq and Karin Ljuslinder. 2020. Self-representations on social media. Reproducing and challenging discourses on disability. *Alter* 14, 2 (June 2020), 71–84. DOI:https://doi.org/10.1016/j.alter.2020.02.001

[53] Linda Codega. 2022. Martin Scorsese's Goncharov (1973) Is the Greatest Mafia Movie Never Made. *Yahoo News*. Retrieved April 7, 2023 from https://news.yahoo.com/martin-scorseses-goncharov-1973-greatest-183000223.html

[54] Lucas Colusso, Cynthia L. Bennett, Gary Hsieh, and Sean A. Munson. 2017. Translational Resources: Reducing the Gap Between Academic Research and HCI Practice. In *Proceedings of the 2017 Conference on Designing Interactive Systems*, ACM, Edinburgh United Kingdom, 957–968. DOI:https://doi.org/10.1145/3064663.3064667

[55] Chetz Colwell and Helen Petrie. 2001. Evaluation of Guidelines for Designing Accessible Web Content. *SIGCAPH Comput. Phys. Handicap.* 70 (July 2001), 11–13. DOI:https://doi.org/10.1145/501078.501082

[56] Cooper Hewitt Smithsonian Design Museum. 2019. Cooper Hewitt Guidelines for Image Description. *Cooper Hewitt Smithsonian Design Museum*. Retrieved August 8, 2020 from https://www.cooperhewitt.org/cooper-hewitt-guidelines-for-image-description/

[57] Francesca Coppa. 2014. Fuck yeah, Fandom is Beautiful. *The Journal of Fandom Studies* 2, 1 (April 2014), 73–82. DOI:https://doi.org/10.1386/jfs.2.1.73_1
[58] Emma Dahlin. 2021. Email Interviews: A Guide to Research Design and Implementation. *International Journal of Qualitative Methods* 20, (January 2021), 160940692110254. DOI:https://doi.org/10.1177/16094069211025453
[59] Peter Dalsgaard and Christian Dindler. 2014. Between theory and practice: bridging concepts in HCI research. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, ACM Press, Toronto, Ontario, Canada, 1635–1644. DOI:https://doi.org/10.1145/2556288.2557342
[60] Maitraye Das, Darren Gergle, and Anne Marie Piper. 2019. "It doesn't win you friends": Understanding Accessibility in Collaborative Writing for People with Vision Impairments. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (November 2019), 1–26. DOI:https://doi.org/10.1145/3359293
[61] Lennard J. Davis. 2013. *The disability studies reader* (4th ed ed.). Routledge, New York, NY.
[62] Lennard J. Davis (Ed.). 2017. *The disability studies reader* (Fifth edition ed.). Routledge, Taylor & Francis Group, New York.
[63] DIAGRAM Center. 2019. Image Description Guidelines. Retrieved March 14, 2021 from http://diagramcenter.org/table-of-contents-2.html/
[64] Carl DiSalvo. 2015. *Adversarial design*. Mit Press.
[65] Julie McDonough Dolmaya. 2011. The ethics of crowdsourcing. *Linguistica Antverpiensia, New Series – Themes in Translation Studies* 10, (2011). DOI:https://doi.org/10.52034/lanstts.v10i.279
[66] Lynn Dombrowski, Ellie Harmon, and Sarah Fox. 2016. Social Justice-Oriented Interaction Design: Outlining Key Design Strategies and Commitments. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems - DIS '16*, ACM Press, Brisbane, QLD, Australia, 656–671. DOI:https://doi.org/10.1145/2901790.2901861
[67] Carlos Duarte and Manuel J. Fonseca. 2019. Multimedia Accessibility. In *Web Accessibility: A Foundation for Research*, Yeliz Yesilada and Simon Harper (eds.). Springer, London, 461–475. DOI:https://doi.org/10.1007/978-1-4471-7440-0_25
[68] Stephen Duncombe. 2012. Imagining no-place. *Transformative works and cultures* 10, (2012).
[69] Pınar Onay Durdu and Zehra Yerlikaya. 2020. The Perception of Website Accessibility: A Survey of Turkish Software Professionals. *AJIT-e* 11, 41 (August 2020), 42–71. DOI:https://doi.org/10.5824/ajite.2020.02.003.x
[70] Emory James Edwards and Tom Boellstorff. Migration, non-use, and the 'Tumblrpocalypse': Towards a unified theory of digital exodus. *Media, Culture & Society* 0, 0 , 0163443720968461. DOI:https://doi.org/10.1177/0163443720968461
[71] Emory James Edwards, Michael Gilbert, Emily Blank, and Stacy M. Branham. In Review. How the Alt Text Gets Made: What Roles and Processes of Alt Text Creation Can Teach Us About Inclusive Imagery. *ACM Trans. Access. Comput.* (In Review).
[72] Emory James Edwards, Kyle Lewis Polster, Isabel Tuason, Emily Blank, Michael Gilbert, and Stacy Branham. 2021. "That's in the eye of the beholder": Layers of Interpretation in Image Descriptions for Fictional Representations of People with Disabilities. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility*. Association for Computing Machinery, New York, NY, USA, 1–14. Retrieved January 4, 2022 from https://doi.org/10.1145/3441852.3471222
[73] Emory James Edwards, Cella Monet Sum, and Stacy M. Branham. 2020. Three Tensions Between Personas and Complex Disability Identities. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (CHI EA '20), Association for Computing Machinery, New York, NY, USA, 1–9. DOI:https://doi.org/10.1145/3334480.3382931

[74] Katie Ellis and Gerard Goggin. 2013. Disability and social media. In *The social media handbook*. Routledge, 134–151.
[75] Katie Ellis and Mike Kent. 2011. *Disability and new media*. Routledge.
[76] Katie Ellis and Mike Kent. 2016. *Disability and social media: Global perspectives*. Taylor & Francis.
[77] Katie Ellis, Debbie Rodan, and Pia Lebeck. 2013. Disability, Obesity and Ageing. (2013), 177.
[78] Etienne Wenger-Trayner and Beverly Wenger-Trayner. 2015. Introduction to communities of practice. Retrieved April 5, 2023 from https://www.wenger-trayner.com/introduction-to-communities-of-practice/
[79] Sarah Evans, Katie Davis, Abigail Evans, Julie Ann Campbell, David P. Randall, Kodlee Yin, and Cecilia Aragon. 2017. More Than Peer Production: Fanfiction Communities as Sites of Distributed Mentoring. 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, 259–272. DOI:https://doi.org/10.1145/2998181.2998342
[80] Rose Eveleth. 2015. How Many Photographs of You Are Out There In the World? *The Atlantic*. Retrieved April 16, 2021 from https://www.theatlantic.com/technology/archive/2015/11/how-many-photographs-of-you-are-out-there-in-the-world/413389/
[81] Casey Fiesler and Brianna Dym. 2020. Moving Across Lands: Online Platform Migration in Fandom Communities. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1 (May 2020), 1–25. DOI:https://doi.org/10.1145/3392847
[82] Casey Fiesler, Shannon Morrison, R. Benjamin Shapiro, and Amy S. Bruckman. 2017. Growing Their Own: Legitimate Peripheral Participation for Computational Learning in an Online Fandom Community. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17*, ACM Press, Portland, Oregon, USA, 1375–1386. DOI:https://doi.org/10.1145/2998181.2998210
[83] Rebecca Fishbein. 2022. Welcome to Tumblr. Now Go Away. *The New York Times*. Retrieved April 5, 2023 from https://www.nytimes.com/2022/11/23/style/tumblr-twitter-elon-musk.html
[84] Christopher Frauenberger. 2015. Disability and Technology: A Critical Realist Perspective. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility - ASSETS '15*, ACM Press, Lisbon, Portugal, 89–96. DOI:https://doi.org/10.1145/2700648.2809851
[85] freddie / @fruitquake. 2023. trash can mushroom man. *Tumblr*. Retrieved April 6, 2023 from https://www.tumblr.com/fruitquake/708710547308298240/ive-been-noticing-tumblr-getting-dubbed-the
[86] Andre P. Freire, Cibele M. Russo, and Renata P. M. Fortes. 2008. A Survey on the Accessibility Awareness of People Involved in Web Development Projects in Brazil. In *Proceedings of the 2008 International Cross-disciplinary Conference on Web Accessibility (W4A)* (W4A '08), ACM, New York, NY, USA, 87–96. DOI:https://doi.org/10.1145/1368044.1368064
[87] Batya Friedman, Peter H. Kahn, and Alan Borning. 2008. Value sensitive design and information systems. *The handbook of information and computer ethics* (2008), 69–101.
[88] Ombretta Gaggi and Veronica Pederiva. 2021. WCAG4All, a tool for making web accessibility rules accessible. In *2021 IEEE 18th Annual Consumer Communications & Networking Conference (CCNC)*, IEEE, Las Vegas, NV, USA, 1–6. DOI:https://doi.org/10.1109/CCNC49032.2021.9369484
[89] Rosemarie Garland-Thomson. 2017. *Extraordinary bodies: figuring physical disability in American culture and literature* (Twentieth anniversary edition ed.). Columbia University Press, New York.
[90] G. Anthony Giannoumis and Lars Henrik Nordli. 2020. Institutionalizing Universal Design: How Organizational Practices Can Promote Web Accessibility. In *HCI International 2020 – Late Breaking Papers: Universal Access and Inclusive Design* (Lecture Notes in Computer Science), Springer International Publishing, Cham, 87–102. DOI:https://doi.org/10.1007/978-3-030-60149-2_8

[91] Sarah Gilbert. 2016. Learning in a Twitter-based community of practice: an exploration of knowledge exchange as a motivation for participation in #hcsmca. *Information, Communication & Society* 19, 9 (September 2016), 1214–1232. DOI:https://doi.org/10.1080/1369118X.2016.1186715

[92] Stéphanie Giraud, Pierre Thérouanne, and Dirk D. Steiner. 2018. Web accessibility: Filtering redundant and irrelevant information improves website usability for blind users. *International Journal of Human-Computer Studies* 111, (March 2018), 23–35. DOI:https://doi.org/10.1016/j.ijhcs.2017.10.011

[93] Haben Girma. Haben Girma - YouTube. Retrieved February 7, 2023 from https://www.youtube.com/@haben_girma

[94] Cole Gleason, Patrick Carrington, Cameron Cassidy, Meredith Ringel Morris, Kris M. Kitani, and Jeffrey 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 on - WWW '19*, ACM Press, San Francisco, CA, USA, 549–559. DOI:https://doi.org/10.1145/3308558.3313605

[95] 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. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, ACM, Honolulu HI USA, 1–12. DOI:https://doi.org/10.1145/3313831.3376728

[96] John D. Gould and Clayton Lewis. 1985. Designing for usability: key principles and what designers think. *Communications of the ACM* 28, 3 (1985), 300–311.

[97] Colin M. Gray. 2014. Evolution of design competence in UX practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, Toronto Ontario Canada, 1645–1654. DOI:https://doi.org/10.1145/2556288.2557264

[98] Colin M. Gray, Erik Stolterman, and Martin A. Siegel. 2014. Reprioritizing the relationship between HCI research and practice: bubble-up and trickle-down effects. In *Proceedings of the 2014 conference on Designing interactive systems - DIS '14*, ACM Press, Vancouver, BC, Canada, 725–734. DOI:https://doi.org/10.1145/2598510.2598595

[99] Frances J. Griffith and Catherine H. Stein. 2021. Behind the hashtag: Online disclosure of mental illness and community response on tumblr. *American Journal of Community Psychology* 67, 3–4 (2021), 419–432.

[100] Darren Guinness, Edward Cutrell, and Meredith Ringel Morris. 2018. Caption Crawler: Enabling Reusable Alternative Text Descriptions using Reverse Image Search. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Association for Computing Machinery, New York, NY, USA, 1–11. Retrieved January 31, 2022 from https://doi.org/10.1145/3173574.3174092

[101] Ankush Gupta and Prashanth Mannem. 2012. From Image Annotation to Image Description. In *Neural Information Processing*, Tingwen Huang, Zhigang Zeng, Chuandong Li and Chi Sing Leung (eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 196–204. DOI:https://doi.org/10.1007/978-3-642-34500-5_24

[102] Beth A. Haller. 2010. *Representing Disability in an Ableist World*.

[103] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M. Branham. 2018. Gender Recognition or Gender Reductionism?: The Social Implications of Embedded Gender Recognition Systems. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18), ACM, New York, NY, USA, 8:1-8:13. DOI:https://doi.org/10.1145/3173574.3173582

[104] Aimi Hamraie. 2017. *Building access: Universal design and the politics of disability*. U of Minnesota Press.

[105] Margot Hanley, Solon Barocas, Karen Levy, Shiri Azenkot, and Helen Nissenbaum. 2021. Computer Vision and Conflicting Values: Describing People with Automated Alt Text. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, ACM, Virtual Event USA, 543–554. DOI:https://doi.org/10.1145/3461702.3462620

[106] Vicki L. Hanson, Anna Cavender, and Shari Trewin. 2015. Writing about accessibility. *interactions* 22, 6 (October 2015), 62–65. DOI:https://doi.org/10.1145/2828432

[107] Hamilton A. Hernandez, Mallory Ketcheson, Adrian Schneider, Zi Ye, Darcy Fehlings, Lauren Switzer, Virginia Wright, Shelly K. Bursick, Chad Richards, and T.C. Nicholas Graham. 2014. Design and Evaluation of a Networked Game to Supportsocial Connection of Youth with Cerebral Palsy. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '14), ACM, New York, NY, USA, 161–168. DOI:https://doi.org/10.1145/2661334.2661370

[108] Daniel G Heslep and PS Berge. 2021. Mapping Discord's darkside: Distributed hate networks on Disboard. *New Media & Society* (December 2021), 14614448211062548. DOI:https://doi.org/10.1177/14614448211062548

[109] Charles G. Hill, Maren Haag, Alannah Oleson, Chris Mendez, Nicola Marsden, Anita Sarma, and Margaret Burnett. 2017. Gender-Inclusiveness Personas vs. Stereotyping: Can We Have it Both Ways? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*, ACM Press, Denver, Colorado, USA, 6658–6671. DOI:https://doi.org/10.1145/3025453.3025609

[110] Serena Hillman, Jason Procyk, and Carman Neustaedter. 2014. "alksjdf;Lksfd": tumblr and the fandom user experience. In *Proceedings of the 2014 conference on Designing interactive systems - DIS '14*, ACM Press, Vancouver, BC, Canada, 775–784. DOI:https://doi.org/10.1145/2598510.2600887

[111] Serena Hillman, Jason Procyk, and Carman Neustaedter. 2014. Tumblr fandoms, community & culture. In *Proceedings of the companion publication of the 17th ACM conference on Computer supported cooperative work & social computing*, ACM, Baltimore Maryland USA, 285–288. DOI:https://doi.org/10.1145/2556420.2557634

[112] Matt Hills. 2015. The expertise of digital fandom as a 'community of practice': Exploring the narrative universe of *Doctor Who*. *Convergence* 21, 3 (August 2015), 360–374. DOI:https://doi.org/10.1177/1354856515579844

[113] Megan Hofmann, Devva Kasnitz, Jennifer Mankoff, and Cynthia L Bennett. 2020. Living Disability Theory: Reflections on Access, Research, and Design. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '20), Association for Computing Machinery, New York, NY, USA, 1–13. DOI:https://doi.org/10.1145/3373625.3416996

[114] Jonggi Hong, Alisha Pradhan, Jon E. Froehlich, and Leah Findlater. 2017. Evaluating Wrist-Based Haptic Feedback for Non-Visual Target Finding and Path Tracing on a 2D Surface. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '17), ACM, New York, NY, USA, 210–219. DOI:https://doi.org/10.1145/3132525.3132538

[115] Ella Houston. 2019. 'Risky' representation: the portrayal of women with mobility impairment in twenty-first-century advertising. *Disability & Society* 34, 5 (May 2019), 704–725. DOI:https://doi.org/10.1080/09687599.2019.1576505

[116] Kalley Huang. 2022. Twitter's Rivals Try to Capitalize on Musk-Induced Chaos. *The New York Times*. Retrieved April 5, 2023 from https://www.nytimes.com/2022/12/07/technology/twitter-rivals-alternative-platforms.html

[117] Amy Hurst and Jasmine Tobias. 2011. Empowering Individuals with Do-it-yourself Assistive Technology. In *The Proceedings of the 13th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '11), ACM, New York, NY, USA, 11–18. DOI:https://doi.org/10.1145/2049536.2049541

[118] Yavuz Inal, Frode Guribye, Dorina Rajanen, Mikko Rajanen, and Mattias Rost. 2020. Perspectives and Practices of Digital Accessibility: A Survey of User Experience Professionals in Nordic Countries. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, ACM, Tallinn Estonia, 1–11. DOI:https://doi.org/10.1145/3419249.3420119

[119] Lilly C. Irani and M. Six Silberman. 2013. Turkopticon: interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13), Association for Computing Machinery, New York, NY, USA, 611–620. DOI:https://doi.org/10.1145/2470654.2470742

[120] Nalita James. 2016. Using email interviews in qualitative educational research: creating space to think and time to talk. *International Journal of Qualitative Studies in Education* 29, 2 (February 2016), 150–163. DOI:https://doi.org/10.1080/09518398.2015.1017848

[121] Henry Jenkins. 1992. *Textual Poachers: Television Fans & Participatory Culture*. Routledge, New York ; London.

[122] Henry Jenkins and Sangita Shresthova. 2012. Up, up, and away! The power and potential of fan activism. *Transformative Works and Cultures* 10, 1 (2012).

[123] Shaun K. Kane, Jeffrey P. Bigham, and Jacob O. Wobbrock. 2008. Slide Rule: Making Mobile Touch Screens Accessible to Blind People Using Multi-touch Interaction Techniques. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility* (Assets '08), ACM, New York, NY, USA, 73–80. DOI:https://doi.org/10.1145/1414471.1414487

[124] Shaun K. Kane, Chandrika Jayant, Jacob O. Wobbrock, and Richard E. Ladner. 2009. Freedom to roam: a study of mobile device adoption and accessibility for people with visual and motor disabilities. In *Proceeding of the eleventh international ACM SIGACCESS conference on Computers and accessibility - ASSETS '09*, ACM Press, Pittsburgh, Pennsylvania, USA, 115. DOI:https://doi.org/10.1145/1639642.1639663

[125] Os Keyes. 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW (November 2018), 1–22. DOI:https://doi.org/10.1145/3274357

[126] Lawrence H Kim, Abena Boadi-Agyemang, Alexa Fay Siu, and John Tang. 2020. When to Add Human Narration to Photo-Sharing Social Media. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Virtual Event Greece, 1–3. DOI:https://doi.org/10.1145/3373625.3418013

[127] Philippe Kimura-Thollander and Neha Kumar. 2019. Examining the "Global" Language of Emojis: Designing for Cultural Representation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ACM, Glasgow Scotland Uk, 1–14. DOI:https://doi.org/10.1145/3290605.3300725

[128] Madison Malone Kircher. 2022. The Fake Scorsese Film You Haven't Seen. Or Have You? *The New York Times*. Retrieved April 7, 2023 from https://www.nytimes.com/2022/11/22/style/goncharov-scorsese-tumblr.html

[129] Andrew Kirkconnell. 2019. "Where Everyone Waddles Like Me": An examination of the autistic community on Tumblr.com. Thesis. Retrieved April 6, 2023 from https://macsphere.mcmaster.ca/handle/11375/24989

[130] Aniket Kittur and Robert E. Kraut. 2010. Beyond Wikipedia: coordination and conflict in online production groups. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work - CSCW '10*, ACM Press, Savannah, Georgia, USA, 215. DOI:https://doi.org/10.1145/1718918.1718959

[131] Aniket Kittur and Robert E Kraut. Harnessing the Wisdom of Crowds in Wikipedia: Quality Through Coordination. 10.

[132] Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the 2013 conference on Computer supported cooperative work - CSCW '13*, ACM Press, San Antonio, Texas, USA, 1301. DOI:https://doi.org/10.1145/2441776.2441923

[133] Neta Kligler-Vilenchik, Joshua McVeigh-Schultz, Christine Weitbrecht, and Chris Tokuhama. 2012. Experiencing fan activism: Understanding the power of fan activist organizations through members' narratives. *Transformative Works and Cultures* 10, (2012).

[134] Robert V. Kozinets. 2010. *Netnography: ethnographic research in the age of the internet* (1st ed ed.). Sage Publications Ltd, Thousand Oaks, CA.

[135] Katie Kuksenok, Michael Brooks, and Jennifer Mankoff. 2013. Accessible online content creation by end users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, Paris France, 59–68. DOI:https://doi.org/10.1145/2470654.2470664

[136] Richard E Ladner. 2015. Design for User Empowerment. *Interactions* (2015), 6.

[137] Chris M Law, Ji Soo Yi, Young Sang Choi, and Julie A Jacko. 2006. Are disability-access guidelines designed for designers? Do they need to be? (November 2006), 4.

[138] Brandi Lawless. 2018. Documenting a labor of love: emotional labor as academic labor. *Review of Communication* 18, 2 (April 2018), 85–97. DOI:https://doi.org/10.1080/15358593.2018.1438644

[139] Amanda Lazar, Caroline Edasis, and Anne Marie Piper. 2017. A Critical Lens on Dementia and Design in HCI. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*, ACM Press, Denver, Colorado, USA, 2175–2188. DOI:https://doi.org/10.1145/3025453.3025522

[140] Jonathan Lazar, Aaron Allen, Jason Kleinman, and Chris Malarkey. 2007. What Frustrates Screen Reader Users on the Web: A Study of 100 Blind Users. *International Journal of Human–Computer Interaction* 22, 3 (May 2007), 247–269. DOI:https://doi.org/10.1080/10447310709336964

[141] Jaewook Lee, Jaylin Herskovitz, Yi-Hao Peng, and Anhong Guo. 2022. ImageExplorer: Multi-Layered Touch Exploration to Encourage Skepticism Towards Imperfect AI-Generated Image Captions. In *CHI Conference on Human Factors in Computing Systems*, ACM, New Orleans LA USA, 1–15. DOI:https://doi.org/10.1145/3491102.3501966

[142] Matthew L. Lee and Anind K. Dey. 2007. Providing Good Memory Cues for People with Episodic Memory Impairment. In *Proceedings of the 9th International ACM SIGACCESS Conference on Computers and Accessibility* (Assets '07), ACM, New York, NY, USA, 131–138. DOI:https://doi.org/10.1145/1296843.1296867

[143] Dean Barnes Leetal. 2019. Those Crazy Fangirls on the Internet: Activism of Care, Disability and Fan Fiction. *1* 8, 2 (April 2019), 45–72. DOI:https://doi.org/10.15353/cjds.v8i2.491

[144] Veronica Lewis. 2018. How To Write Alt Text And Image Descriptions For The Visually Impaired. *Veroniiiica*. Retrieved August 8, 2020 from https://veroniiiica.com/2018/01/31/how-to-write-alt-text-image-descriptions-visually-impaired/

[145] Mingzhe Li, Mingming Fan, and Khai N. Truong. 2017. BrailleSketch: A Gesture-based Text Input Method for People with Visual Impairments. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '17), ACM, New York, NY, USA, 12–21. DOI:https://doi.org/10.1145/3132525.3132528

[146] William Lidwell, Kritina Holden, and Jill Butler. 2010. *Universal principles of design, revised and updated: 125 ways to enhance usability, influence perception, increase appeal, make better design decisions, and teach through design*. Rockport Pub.

[147] G. Lintern, F.J. Diedrich, and D. Serfaty. 2002. Engineering the community of practice for maintenance of organizational knowledge. In *Proceedings of the IEEE 7th Conference on Human Factors and Power Plants*, 6–6. DOI:https://doi.org/10.1109/HFPP.2002.1042855

[148] Vera Louise Roberts and Deborah I. Fels. 2006. Methods for inclusion: Employing think aloud protocols in software usability studies with individuals who are deaf. *International Journal of Human-Computer Studies* 64, 6 (June 2006), 489–501. DOI:https://doi.org/10.1016/j.ijhcs.2005.11.001

[149] Janet Mac. 3D Illustration & Animation - Google Project. Retrieved April 13, 2021 from https://janet-mac.com/google-avatar-project

[150] Kelly Mack, Edward Cutrell, Bongshin Lee, and Meredith Ringel Morris. 2021. Designing Tools for High-Quality Alt Text Authoring. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Virtual Event USA, 1–14. DOI:https://doi.org/10.1145/3441852.3471207

[151] Kelly Mack, Maitraye Das, Dhruv Jain, Danielle Bragg, John Tang, Andrew Begel, Erin Beneteau, Josh Urban Davis, Abraham Glasser, Joon Sung Park, and Venkatesh Potluri. 2021. Mixed Abilities and Varied Experiences: a group autoethnography of a virtual summer internship. In *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '21), Association for Computing Machinery, New York, NY, USA, 1–13. DOI:https://doi.org/10.1145/3441852.3471199

[152] Kelly Mack, Emma McDonnell, Dhruv Jain, Lucy Lu Wang, Jon E. Froehlich, and Leah Findlater. 2021. What Do We Mean by "Accessibility Research"? A Literature Survey of Accessibility Papers in CHI and ASSETS from 1994 to 2019. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (CHI '21), Association for Computing Machinery, New York, NY, USA, 1–18. DOI:https://doi.org/10.1145/3411764.3445412

[153] 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*, ACM, Denver Colorado USA, 5988–5999. DOI:https://doi.org/10.1145/3025453.3025814

[154] Jennifer Mankoff, Gillian R. Hayes, and Devva Kasnitz. 2010. Disability studies as a source of critical inquiry for the field of assistive technology. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility*, ACM, 3–10.

[155] George Margetis, Stavroula Ntoa, and Constantine Stephanidis. 2008. Requirements of Users with Disabilities for E-government Services in Greece. In *Computers Helping People with Special Needs* (Lecture Notes in Computer Science), Springer, Berlin, Heidelberg, 438–445. DOI:https://doi.org/10.1007/978-3-540-70540-6_63

[156] Nicola Marsden and Maren Haag. 2016. Stereotypes and Politics: Reflections on Personas. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, ACM Press, Santa Clara, California, USA, 4017–4031. DOI:https://doi.org/10.1145/2858036.2858151

[157] Nicola Marsden and Monika Pröbster. 2019. Personas and Identity: Looking at Multiple Identities to Inform the Construction of Personas. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (CHI '19), ACM, New York, NY, USA, 335:1-335:14. DOI:https://doi.org/10.1145/3290605.3300565

[158] Paris Martineau. 2018. Tumblr's Porn Ban Reveals Who Controls What We See Online. *Wired*. Retrieved April 7, 2023 from https://www.wired.com/story/tumblrs-porn-ban-reveals-controls-we-see-online/

[159] Reeti Mathur and Erin Brady. 2018. Mixed-Ability Collaboration for Accessible Photo Sharing. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '18), Association for Computing Machinery, Galway, Ireland, 370–372. DOI:https://doi.org/10.1145/3234695.3240994

[160] Allison McCracken. 2017. Tumblr Youth Subcultures and Media Engagement. *Cinema Journal* 57, 1 (2017), 151-.

[161] Allison McCracken, Alexander Cho, Louisa Stein, and Indira Hoch. 2020. *a tumblr book: platform and cultures*. University of Michigan Press, Ann Arbor, MI. DOI:https://doi.org/10.3998/mpub.11537055

[162] Ryan A. Miller. 2017. "My Voice Is Definitely Strongest in Online Communities": Students Using Social Media for Queer and Disability Identity-Making. *Journal of College Student Development* 58, 4 (June 2017), 509–525. DOI:https://doi.org/10.1353/csd.2017.0040

[163] Lauren R. Milne, Cynthia L. Bennett, Richard E. Ladner, and Shiri Azenkot. 2014. BraillePlay: Educational Smartphone Games for Blind Children. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '14), ACM, New York, NY, USA, 137–144. DOI:https://doi.org/10.1145/2661334.2661377

[164] Christine Miserandino. 2013. The Spoon Theory. *But You Dont Look Sick? support for those with invisible illness or chronic illness*. Retrieved October 14, 2022 from https://butyoudontlooksick.com/articles/written-by-christine/the-spoon-theory/

[165] 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 (November 2015), 1–21. DOI:https://doi.org/10.1145/2764916

[166] Meredith Ringel Morris, Jazette Johnson, Cynthia L. Bennett, and Edward Cutrell. 2018. Rich Representations of Visual Content for Screen Reader Users. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, ACM Press, Montreal QC, Canada, 1–11. DOI:https://doi.org/10.1145/3173574.3173633

[167] Meredith Ringel Morris, Annuska Zolyomi, Catherine Yao, Sina Bahram, Jeffrey P. Bigham, and Shaun K. Kane. 2016. "With most of it being pictures now, I rarely use it": Understanding Twitter's Evolving Accessibility to Blind Users. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, San Jose California USA, 5506–5516. DOI:https://doi.org/10.1145/2858036.2858116

[168] Martez E Mott, John Tang, and Edward Cutrell. 2023. Accessibility of Profile Pictures: Alt Text and Beyond to Express Identity Online. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (CHI '23), Association for Computing Machinery, New York, NY, USA, 1–13. DOI:https://doi.org/10.1145/3544548.3580710

[169] Karen Nakamura. 2019. My Algorithms Have Determined You'Re Not Human: AI-ML, Reverse Turing-Tests, and the Disability Experience. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '19), ACM, New York, NY, USA, 1–2. DOI:https://doi.org/10.1145/3308561.3353812

[170] Louis Neven. 2011. Representations of the Old and Ageing in the Design of the New and Emerging: Assessing the design of Ambient Intelligence technologies for older people. University of Twente. Retrieved September 10, 2020 from https://research.utwente.nl/en/publications/representations-of-the-old-and-ageing-in-the-design-of-the-new-an

[171] Alexandre Nevsky, Timothy Neate, Elena Simperl, and Radu-Daniel Vatavu. 2023. Accessibility Research in Digital Audiovisual Media: What Has Been Achieved and What Should Be Done Next? In *ACM International Conference on Interactive Media Experiences*. ACM, 1–21. Retrieved May 19, 2023 from https://kclpure.kcl.ac.uk/portal/en/publications/accessibility-research-in-digital-audiovisual-media-what-has-been-achieved-and-what-should-be-done-next(acded962-ad2a-427f-8431-953e8f85a087).html

[172] Alan F. Newell, Peter Gregor, Maggie Morgan, Graham Pullin, and Catriona Macaulay. 2011. User-sensitive inclusive design. *Universal Access in the Information Society* 10, 3 (2011), 235–243.

[173] Nicole Newnham and James Lebrecht. 2020. Crip Camp | A Disability Revolution. Retrieved February 7, 2023 from https://cripcamp.com/

[174] Safiya Umoja Noble. 2018. *Algorithms of oppression: How search engines reinforce racism*. nyu Press.
[175] Ihudiya Finda Ogbonnaya-Ogburu, Angela D.R. Smith, Alexandra To, and Kentaro Toyama. 2020. Critical Race Theory for HCI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, ACM, Honolulu HI USA, 1–16. DOI:https://doi.org/10.1145/3313831.3376392
[176] Uran Oh, Shaun K. Kane, and Leah Findlater. 2013. Follow That Sound: Using Sonification and Corrective Verbal Feedback to Teach Touchscreen Gestures. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '13), ACM, New York, NY, USA, 13:1-13:8. DOI:https://doi.org/10.1145/2513383.2513455
[177] Shannon M. Oltmann. 2011. Telephone & email interviews: using the respondents' context to determine the best interview mode. In *Proceedings of the 2011 iConference*, ACM, Seattle Washington USA, 734–735. DOI:https://doi.org/10.1145/1940761.1940889
[178] Afra Pascual, Mireia Ribera, and Toni Granollers. 2015. Empathic Communication of Accessibility Barriers in Web 2.0 Editing. In *Proceedings of the 12th Web for All Conference* (W4A '15), ACM, New York, NY, USA, 23:1-23:8. DOI:https://doi.org/10.1145/2745555.2746642
[179] Nilay Patel. 2022. How to buy a social network, with Tumblr CEO Matt Mullenweg. *The Verge*. Retrieved April 7, 2023 from https://www.theverge.com/23506085/wordpress-twitter-tumblr-ceo-matt-mullenweg-elon-musk
[180] Yi-Hao Peng, Jason Wu, Jeffrey Bigham, and Amy Pavel. 2022. Diffscriber: Describing Visual Design Changes to Support Mixed-Ability Collaborative Presentation Authoring. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (UIST '22), Association for Computing Machinery, New York, NY, USA, 1–13. DOI:https://doi.org/10.1145/3526113.3545637
[181] Christian Pentzold. 2018. Grounding Peer Production in Practice: Editorial Routines and Everyday Engagement in the "Free Encyclopedia Anyone Can Edit." *Communication, Culture and Critique* 11, 3 (September 2018), 455–474. DOI:https://doi.org/10.1093/ccc/tcy010
[182] Helen Petrie, Chandra Harrison, and Sundeep Dev. 2005. Describing images on the web: a survey of current practice and prospects for the future. *Proceedings of Human Computer Interaction International (HCII)* 71, (2005), 2.
[183] Sarah Pink, Heather A. Horst, John Postill, Larissa Hjorth, Tania Lewis, and Jo Tacchi (Eds.). 2016. *Digital ethnography: principles and practice*. SAGE, Los Angeles.
[184] John R. Porter and Julie A. Kientz. 2013. An empirical study of issues and barriers to mainstream video game accessibility. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Bellevue Washington, 1–8. DOI:https://doi.org/10.1145/2513383.2513444
[185] Christopher Power. 2012. Guidelines are only half of the story: accessibility problems encountered by blind users on the web. (2012), 10.
[186] Cynthia Putnam, Maria Dahman, Emma Rose, Jinghui Cheng, and Glenn Bradford. 2015. Teaching Accessibility, Learning Empathy. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '15), ACM, New York, NY, USA, 333–334. DOI:https://doi.org/10.1145/2700648.2811365
[187] Cynthia Putnam, Kathryn Wozniak, Mary Jo Zefeldt, Jinghui Cheng, Morgan Caputo, and Carl Duffield. 2012. How do professionals who create computing technologies consider accessibility? In *Proceedings of the 14th international ACM SIGACCESS conference on Computers and accessibility - ASSETS '12*, ACM Press, Boulder, Colorado, USA, 87. DOI:https://doi.org/10.1145/2384916.2384932

[188] Claudia M Rebaza. 2009. THE MODERN COTERIE: FAN FICTION WRITERS AS A COMMUNITY OF PRACTICE.
[189] Thomas Reid. Reid My Mind | A podcast from a Black man adjusting to being blind. Retrieved February 7, 2023 from https://reidmymind.com/
[190] Jason G Reitman, Craig G Anderson, and Constance Steinkuehler. 2021. Tools, and Strategies. Discord Community Challenges, Tools, and Strategies. *Studies in Conflict & Terrorism* (January 2021), 1–22. DOI:https://doi.org/10.1080/1057610X.2020.1862850
[191] Luz Rello, Arturo Macias, Mariía Herrera, Camila de Ros, Enrique Romero, and Jeffrey P. Bigham. 2017. DytectiveU: A Game to Train the Difficulties and the Strengths of Children with Dyslexia. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '17), ACM, New York, NY, USA, 319–320. DOI:https://doi.org/10.1145/3132525.3134773
[192] Michelle C. Reynolds, Diana Caldwell, Natalie Boonchaisri, Katharine E. Ragon, and Susan B. Palmer. 2022. The Community of Practice for Supporting Families of Persons With Intellectual and Developmental Disabilities. *Intellectual & Developmental Disabilities* 60, 2 (April 2022), 85–100. DOI:https://doi.org/10.1352/1934-9556-60.2.85
[193] Charles A. Riley. 2005. *Disability and the media: prescriptions for change*. University Press of New England, Hanover, NH.
[194] Kathryn E. Ringland. 2019. A Place to Play: The (Dis)Abled Embodied Experience for Autistic Children in Online Spaces. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (CHI '19), ACM, New York, NY, USA, 288:1-288:14. DOI:https://doi.org/10.1145/3290605.3300518
[195] Kathryn E. Ringland, Christine T. Wolf, LouAnne E. Boyd, Mark S. Baldwin, and Gillian R. Hayes. 2016. Would You Be Mine: Appropriating Minecraft As an Assistive Technology for Youth with Autism. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '16), ACM, New York, NY, USA, 33–41. DOI:https://doi.org/10.1145/2982142.2982172
[196] Sarah T. Roberts. 2019. *Behind the screen*. Yale University Press.
[197] Silvia Rodríguez Vázquez. 2016. Measuring the impact of automated evaluation tools on alternative text quality: a web translation study. In *Proceedings of the 13th International Web for All Conference*, ACM, Montreal Canada, 1–10. DOI:https://doi.org/10.1145/2899475.2899484
[198] Chiara Rossitto, Henrik Korsgaard, Airi Lampinen, and Susanne Bødker. 2021. Efficiency and Care in Community-led Initiatives. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 1–27. DOI:https://doi.org/10.1145/3479611
[199] Agnieszka Rychwalska and Magdalena Roszczynska-Kurasinska. 2017. Value Sensitive Design for Peer Production Systems: Mediating Social Interactions. *IEEE Technol. Soc. Mag.* 36, 3 (September 2017), 48–55. DOI:https://doi.org/10.1109/MTS.2017.2728737
[200] 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. *Proceedings of the Fifth AAAI Conference on Human Computation and Crowdsourcing* (2017), 10.
[201] Elliot Salisbury, Ece Kamar, and Meredith Ringel Morris. 2018. Evaluating and Complementing Vision-to-Language Technology for People who are Blind with Conversational Crowdsourcing. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, International Joint Conferences on Artificial Intelligence Organization, Stockholm, Sweden, 5349–5353. DOI:https://doi.org/10.24963/ijcai.2018/751
[202] Joni Salminen, Soon-gyo Jung, Jisun An, Haewoon Kwak, Lene Nielsen, and Bernard J. Jansen. 2019. Confusion and information triggered by photos in persona profiles. *International Journal of Human-Computer Studies* 129, (September 2019), 1–14. DOI:https://doi.org/10.1016/j.ijhcs.2019.03.005

[203] Joni Salminen, Lene Nielsen, Soon-Gyo Jung, Jisun An, Haewoon Kwak, and Bernard J. Jansen. 2018. “Is More Better?”: Impact of Multiple Photos on Perception of Persona Profiles. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI ’18*, ACM Press, Montreal QC, Canada, 1–13. DOI:https://doi.org/10.1145/3173574.3173891

[204] Elizabeth B.-N. Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *CoDesign* 4, 1 (March 2008), 5–18. DOI:https://doi.org/10.1080/15710880701875068

[205] João Marcelo dos Santos Marques, Luiz Fernando Gopi Valente, Simone Bacellar Leal Ferreira, Claudia Cappelli, and Luciana Salgado. 2017. Audio Description on Instagram: Evaluating and Comparing Two Ways of Describing Images for Visually Impaired: In *Proceedings of the 19th International Conference on Enterprise Information Systems*, SCITEPRESS - Science and Technology Publications, Porto, Portugal, 29–40. DOI:https://doi.org/10.5220/0006282500290040

[206] Daisuke Sato, Masatomo Kobayashi, Hironobu Takagi, and Chieko Asakawa. 2010. Social accessibility: the challenge of improving web accessibility through collaboration. In *Proceedings of the 2010 International Cross Disciplinary Conference on Web Accessibility (W4A) - W4A ’10*, ACM Press, Raleigh, North Carolina, 1. DOI:https://doi.org/10.1145/1805986.1806024

[207] Daisuke Sato, Hironobu Takagi, Masatomo Kobayashi, Shinya Kawanaka, and Chieko Asakawa. 2010. Exploratory Analysis of Collaborative Web Accessibility Improvement. *ACM Trans. Access. Comput.* 3, 2 (November 2010), 1–30. DOI:https://doi.org/10.1145/1857920.1857922

[208] Morgan Klaus Scheuerman, Alex Hanna, and Emily Denton. 2021. Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 1–37. DOI:https://doi.org/10.1145/3476058

[209] Morgan Klaus Scheuerman, Katta Spiel, Oliver L. Haimson, Foad Hamidi, and Stacy M. Branham. 2020. HCI Guidelines for Gender Equity and Inclusivity. Retrieved from https://www.morgan-klaus.com/gender-guidelines.html

[210] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. 2020. How We’ve Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1 (May 2020), 1–35. DOI:https://doi.org/10.1145/3392866

[211] Douglas Schuler and Aki Namioka (Eds.). 1993. *Participatory design: principles and practices*. L. Erlbaum Associates, Hillsdale, N.J.

[212] Amanda Schupak. 2015. Google apologizes for mis-tagging photos of African Americans. *CBSNews*. Retrieved August 13, 2020 from https://www.cbsnews.com/news/google-photos-labeled-pics-of-african-americans-as-gorillas/

[213] Michael James Scott, Fotios Spyridonis, and Gheorghita Ghinea. 2015. Designing for designers: Towards the development of accessible ICT products and services using the VERITAS framework. *Computer Standards & Interfaces* 42, (November 2015), 113–124. DOI:https://doi.org/10.1016/j.csi.2015.05.004

[214] Letícia Seixas Pereira, José Coelho, André Rodrigues, João Guerreiro, Tiago Guerreiro, and Carlos Duarte. 2022. Authoring accessible media content on social networks. In *The 24th International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Athens Greece, 1–11. DOI:https://doi.org/10.1145/3517428.3544882

[215] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Melbourne, Australia, 2556–2565. DOI:https://doi.org/10.18653/v1/P18-1238

[216] Kristen Shinohara. 2017. Design for Social Accessibility: Incorporating Social Factors in the Design of Accessible Technologies. University of Washington.
[217] Kristen Shinohara, Cynthia L. Bennett, Wanda Pratt, and Jacob O. Wobbrock. 2018. Tenets for Social Accessibility: Towards Humanizing Disabled People in Design. *ACM Trans. Access. Comput.* 11, 1 (April 2018), 1–31. DOI:https://doi.org/10.1145/3178855
[218] Kristen Shinohara and Josh Tenenberg. 2007. Observing Sara: A Case Study of a Blind Person's Interactions with Technology. In *Proceedings of the 9th International ACM SIGACCESS Conference on Computers and Accessibility* (Assets '07), ACM, New York, NY, USA, 171–178. DOI:https://doi.org/10.1145/1296843.1296873
[219] Ben Shneiderman. 2000. Universal Usability. *Communications of the ACM*. Retrieved January 12, 2020 from https://link.galegroup.com/apps/doc/A62355861/AONE?sid=lms
[220] Britta Shoot. 2017. The 1977 Disability Rights Protest That Broke Records and Changed Laws. *Atlas Obscura*. Retrieved May 20, 2023 from http://www.atlasobscura.com/articles/504-sit-in-san-francisco-1977-disability-rights-advocacy
[221] Rachel N. Simons, Danna Gurari, and Kenneth R. Fleischmann. 2020. "I Hope This Is Helpful": Understanding Crowdworkers' Challenges and Motivations for an Image Description Task. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2 (October 2020), 105:1-105:26. DOI:https://doi.org/10.1145/3415176
[222] Olle Sköld. 2017. Getting-to-know: Inquiries, sources, methods, and the production of knowledge on a videogame wiki. *JD* 73, 6 (October 2017), 1299–1321. DOI:https://doi.org/10.1108/JD-11-2016-0145
[223] Dr Stacy L Smith, Marc Choueiti, and Dr Katherine Pieper. 2020. *Inequality in 1,300 Popular Films: Examining Portrayals of Gender, Race/Ethnicity, LGBTQ & Disability from 2007 to 2019*. Annenburg Inclusion Initiative, USC.
[224] Andrii Soviak, Anatoliy Borodin, Vikas Ashok, Yevgen Borodin, Yury Puzis, and I.V. Ramakrishnan. 2016. Tactile Accessibility: Does Anyone Need a Haptic Glove? In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '16), ACM, New York, NY, USA, 101–109. DOI:https://doi.org/10.1145/2982142.2982175
[225] Tumblr Staff. About Tumblr. *What is Tumblr? Well, it's a website.* Retrieved April 7, 2023 from https://about.tumblr.com/#quick-facts
[226] Abigale Stangl, Jeeeun Kim, and Tom Yeh. 2014. 3D Printed Tactile Picture Books for Children with Visual Impairments: A Design Probe. In *Proceedings of the 2014 Conference on Interaction Design and Children* (IDC '14), ACM, New York, NY, USA, 321–324. DOI:https://doi.org/10.1145/2593968.2610482
[227] 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. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, ACM, Honolulu HI USA, 1–13. DOI:https://doi.org/10.1145/3313831.3376404
[228] Abigale Stangl, Kristina Shiroma, Bo Xie, Kenneth R. Fleischmann, and Danna Gurari. 2020. Visual Content Considered Private by People Who are Blind. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Virtual Event Greece, 1–12. DOI:https://doi.org/10.1145/3373625.3417014
[229] 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. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ACM, Virtual Event USA, 1–15. DOI:https://doi.org/10.1145/3441852.3471233

[230] Linn Steen-Hansen and Siri Fagernes. 2016. The Importance of Process-Oriented Accessibility Guidelines for Web Developers. (2016), 11.

[231] Miriah Steiger, Timir J Bharucha, Sukrit Venkatagiri, Martin J. Riedl, and Matthew Lease. 2021. The Psychological Well-Being of Content Moderators: The Emotional Labor of Commercial Moderation and Avenues for Improving Support. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (CHI '21), Association for Computing Machinery, New York, NY, USA, 1–14. DOI:https://doi.org/10.1145/3411764.3445092

[232] Kevin M Storer, Harini Sampath, and M. Alice Alice Merrick. 2021. "It's Just Everything Outside of the IDE that's the Problem":: Information Seeking by Software Developers with Visual Impairments. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ACM, Yokohama Japan, 1–12. DOI:https://doi.org/10.1145/3411764.3445090

[233] Morgan Sung. 2022. "Goncharov" isn't a real Martin Scorsese movie, but Tumblr convinced the internet it's a classic. *NBC News*. Retrieved April 7, 2023 from https://www.nbcnews.com/pop-culture/viral/goncharov-mafia-movie-martin-scorsese-fake-tumblr-meme-rcna58777

[234] Kayla S. Sweet, Jennifer K. LeBlanc, Laura M. Stough, and Noelle W. Sweany. 2020. Community building and knowledge sharing by individuals with disabilities using social media. *Journal of Computer Assisted Learning* 36, 1 (2020), 1–11. DOI:https://doi.org/10.1111/jcal.12377

[235] Jeanna Sybert. 2022. Tumblr's enduring appeal reveals the potency of the web's cultural memory. *The Conversation*. Retrieved April 5, 2023 from http://theconversation.com/tumblrs-enduring-appeal-reveals-the-potency-of-the-webs-cultural-memory-183788

[236] Hironobu Takagi, Shinya Kawanaka, Masatomo Kobayashi, Takashi Itoh, and Chieko Asakawa. 2008. Social accessibility: achieving accessibility through collaborative metadata authoring. In *Proceedings of the 10th international ACM SIGACCESS conference on Computers and accessibility - Assets '08*, ACM Press, Halifax, Nova Scotia, Canada, 193. DOI:https://doi.org/10.1145/1414471.1414507

[237] Hironobu Takagi, Shinya Kawanaka, Masatomo Kobayashi, Daisuke Sato, and Chieko Asakawa. 2009. Collaborative web accessibility improvement: challenges and possibilities. In *Proceedings of the 11th international ACM SIGACCESS conference on Computers and accessibility* (Assets '09), Association for Computing Machinery, New York, NY, USA, 195–202. DOI:https://doi.org/10.1145/1639642.1639677

[238] Danielle M. Taylor. 2018. *Americans With Disabilities: 2014*. U.S. Census Bureau, United States. Retrieved from https://www.census.gov/library/publications/2018/demo/p70-152.html

[239] Kaitlyn Tiffany. 2022. How the Snowflakes Won. *The Atlantic*. Retrieved April 5, 2023 from https://www.theatlantic.com/technology/archive/2022/02/tumblr-internet-legacy-survival/621419/

[240] Carlos Toxtli, Siddharth Suri, and Saiph Savage. 2021. Quantifying the Invisible Labor in Crowd Work. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 1–26. DOI:https://doi.org/10.1145/3476060

[241] Shari Trewin, Sara Basson, Michael Muller, Stacy Branham, Jutta Treviranus, Daniel Gruen, Daniel Hebert, Natalia Lyckowski, and Erich Manser. 2019. Considerations for AI fairness for people with disabilities. *AI Matters* 5, 3 (December 2019), 40–63. DOI:https://doi.org/10.1145/3362077.3362086

[242] Shari Trewin, Brian Cragun, Cal Swart, Jonathan Brezin, and John Richards. 2010. Accessibility challenges and tool features: an IBM Web developer perspective. (2010), 10.

[243] Shari Trewin, Cal Swart, and Donna Pettick. 2013. Physical Accessibility of Touchscreen Smartphones. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (ASSETS '13), ACM, New York, NY, USA, 19:1-19:8. DOI:https://doi.org/10.1145/2513383.2513446

[244] Phil Turner and Susan Turner. 2011. Is stereotyping inevitable when designing with personas? *Design Studies* 32, 1 (January 2011), 30–44. DOI:https://doi.org/10.1016/j.destud.2010.06.002
[245] Stephen Uzor, Jason T. Jacques, John J Dudley, and Per Ola Kristensson. 2021. Investigating the Accessibility of Crowdwork Tasks on Mechanical Turk. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ACM, Yokohama Japan, 1–14. DOI:https://doi.org/10.1145/3411764.3445291
[246] Gregg Vanderheiden and Jim Tobias. 2000. Universal Design of Consumer Products: Current Industry Practice and Perceptions. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 44, 32 (July 2000), 6–19. DOI:https://doi.org/10.1177/154193120004403206
[247] W3C Web Accessibility Initiative. 2019. Image Concepts • WAI Web Accessibility Tutorials. Retrieved August 8, 2020 from https://www.w3.org/WAI/tutorials/images/
[248] Emily Q. Wang and Anne Marie Piper. 2022. The Invisible Labor of Access in Academic Writing Practices: A Case Analysis with Dyslexic Adults. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1 (March 2022), 1–25. DOI:https://doi.org/10.1145/3512967
[249] Web Accessibility In Mind (WebAIM). 2022. The WebAIM Million - The 2022 report on the accessibility of the top 1,000,000 home pages. Retrieved June 17, 2022 from https://webaim.org/projects/million/
[250] Web Accessibility Initiative - W3C. 2019. How to Meet WCAG (Quickref Reference). Retrieved April 9, 2021 from https://www.w3.org/WAI/WCAG21/quickref/#top
[251] WebAIM. 2019. Testing with Screen Readers - Questions and Answers. Retrieved May 20, 2023 from https://webaim.org/articles/screenreader_testing/
[252] WebAIM. 2021. Survey of Web Accessibility Practitioners #3 Results. Retrieved August 18, 2023 from https://webaim.org/projects/practitionersurvey3/#srproficiency
[253] Etienne Wenger. 2010. Communities of practice and social learning systems. In *Social learning systems and communities of practice*, Chris Blackmore (ed.). Springer.
[254] Wikipedia. 2023. List of disability-related terms with negative connotations. *Wikipedia*. Retrieved February 5, 2023 from https://en.wikipedia.org/w/index.php?title=List_of_disability-related_terms_with_negative_connotations&oldid=1135461664
[255] Candace Williams, Lilian de Greef, Ed Harris, Leah Findlater, Amy Pavel, and Cynthia Bennett. 2022. Toward supporting quality alt text in computing publications. In *Proceedings of the 19th International Web for All Conference*, ACM, Lyon France, 1–12. DOI:https://doi.org/10.1145/3493612.3520449
[256] Michele A. Williams, Caroline Galbraith, Shaun K. Kane, and Amy Hurst. 2014. "Just Let the Cane Hit It": How the Blind and Sighted See Navigation Differently. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '14), ACM, New York, NY, USA, 217–224. DOI:https://doi.org/10.1145/2661334.2661380
[257] Michele A. Williams, Amy Hurst, and Shaun K. Kane. 2013. "Pray before you step out": describing personal and situational blind navigation behaviors. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility - ASSETS '13*, ACM Press, Bellevue, Washington, 1–8. DOI:https://doi.org/10.1145/2513383.2513449
[258] Jacob O. Wobbrock, Shaun K. Kane, Krzysztof Z. Gajos, Susumu Harada, and Jon Froehlich. 2011. Ability-Based Design: Concept, Principles and Examples. *ACM Trans. Access. Comput.* 3, 3 (April 2011), 9:1-9:27. DOI:https://doi.org/10.1145/1952383.1952384
[259] 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*, ACM, Portland Oregon USA, 1180–1192. DOI:https://doi.org/10.1145/2998181.2998364

[260] Yeliz Yesilada, Giorgio Brajnik, Markel Vigo, and Simon Harper. 2012. Understanding web accessibility and its drivers. In *Proceedings of the International Cross-Disciplinary Conference on Web Accessibility*, ACM, Lyon France, 1–9. DOI:https://doi.org/10.1145/2207016.2207027
[261] Dora Zhao, Angelina Wang, and Olga Russakovsky. 2021. Understanding and Evaluating Racial Biases in Image Captioning. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 14810–14820. DOI:https://doi.org/10.1109/ICCV48922.2021.01456
[262] Liming Zhu, Xiwei Xu, Qinghua Lu, Guido Governatori, and Jon Whittle. 2022. AI and Ethics—Operationalizing Responsible AI. In *Humanity Driven AI: Productivity, Well-being, Sustainability and Partnership*, Fang Chen and Jianlong Zhou (eds.). Springer International Publishing, Cham, 15–33. DOI:https://doi.org/10.1007/978-3-030-72188-6_2
[263] Kathryn Zyskowski, Meredith Ringel Morris, Jeffrey P. Bigham, Mary L. Gray, and Shaun K. Kane. 2015. Accessible Crowdwork?: Understanding the Value in and Challenge of Microtask Employment for People with Disabilities. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, ACM, Vancouver BC Canada, 1682–1693. DOI:https://doi.org/10.1145/2675133.2675158
[264] 2020. Facebook: active users worldwide. *Statista*. Retrieved August 28, 2020 from https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/
[265] 2021. WebAIM: Alternative Text. *WebAIM: Web Accessibility In Mind*. Retrieved April 25, 2023 from https://webaim.org/techniques/alttext/
[266] 2023. Tumblr – Fandom, Art, Chaos. *App Store*. Retrieved April 5, 2023 from https://apps.apple.com/us/app/tumblr-fandom-art-chaos/id305343404
[267] Vision impairment and blindness. Retrieved January 31, 2022 from https://www.who.int/news-room/fact-sheets/detail/blindness-and-visual-impairment
[268] reader-response theory. *Oxford Reference*. DOI:https://doi.org/10.1093/oi/authority.20110803100406762
[269] Discord Servers - Home. *Discord*. Retrieved April 8, 2023 from https://discord.com/servers/undefined
[270] About Discord | Our Mission and Values. Retrieved April 8, 2023 from https://discord.com/company
[271] Alt-Text as Poetry. Retrieved April 13, 2021 from https://alt-text-as-poetry.net
[272] The ADHD Analog Brain Tool. *The ADHD Analog Brain Tool*. Retrieved May 22, 2023 from https://adhdanalogbrain.blogspot.com/p/introduction-to-analog-brain-sometimes.html

APPENDIX A: Contributions

This appendix provides three contributions derived from my research as a whole: a guide for creating inclusive image descriptions in different contexts, an updated set of models describing accessible imagery creation processes, and an "owner's manual" describing the elements of a functioning volunteer-based accessibility crowdsourcing project.

Guidelines for Creating Inclusive Image Descriptions

The first contribution coming from my dissertation is a compiled resource for creating inclusive image descriptions. I focused on addressing as many possible scenarios as possible, in a similar decision-making branching logic as my Study 3 participants indicated they used. The format of the guidelines as they appear here is unideal. The intention was to mimic resources such as an "ADHD Analogue Brain Tool" which makes heavy use of hyperlinks [272]. However, for the sake of the dissertation document it has been rendered more linearly. I hope that this comprehensive guide can be helpful for a wide range of scenarios that researchers or accessibility activists such as my Study 3 participants might come across.

Another approach I considered was going more general rather than more specific, describing Principles that should inform alt text creation, or a sort of Style Guide for considering how to balance elements of image descriptions discussed in Study 1. I believe there is space and need for such a resource as well: something less prescriptive that instead prompts describers to use their good judgement and personal expertise when applying them. However, as discussed in Chapter 4, many accessibility advocates working on image

accessibility are already overwhelmed with the cognitive load and judgement calls they have to make as describers. An “Elements of Alt Text Style” or “Principles of UX for Image Accessibility” may be suited to and welcomed by designers and other accessibility practitioners who are used to exercising their professional expertise, but for frazzled academics and advocates busy with an entirely different day job, I believe a straightforward guide was more suitable.

Initially I planned to create a set of unified guidelines, but I realized that trying to fit all the findings into a single set of guidelines would make the resource overly complicated. Instead, I start the resource with general guidelines for image descriptions, derived from both secondary and original research I conducted. Then, to introduce readers to the topic, I define and explain the stakes of describing social identities inclusively, as well as giving some examples of social identities and ways they may come up in the language of an image description. After that, I give tips for various types of images without subjects (or with real but non-human subjects such as animals) separately, as they generally will have a lower likelihood of referencing social identity. Then I discuss photographs of human subjects, both known and unknown to the describer. By focusing on known and unknown subjects, I cover what to do when there are various types of uncertainty or barriers to finding the ideal identity language for that subject. I then discuss illustrations or artistic depictions of real, fictional, and ambiguous subjects. Artistic depictions are different to photographs as the subject cannot be consulted directly (unless the subject is a real person), and instead the most accurate information about the identities of the subject will likely come from the artist or the source fiction. In the case of ambiguity, I provide guidelines covering both the

inclusive ideal as well as the pragmatic scenario in which making assumptions cannot be practically avoided. All sections of the Inclusive Image Description Guide are below.

1. General Guidelines
    1.1. Be succinct (150 characters or less for alt text, can have more detailed description either linked or on a different part of the page so that people can access the longer version if they're interested)
    1.2. Consider the purpose of the image – is it intended to educate, illustrate an idea, or visually reflect a key component of the written content. The purpose of the image defines what type of information is relevant to include.
    1.3. Consider the image in context – an image of someone posing in front of a fountain might have very different details included if it is, for example, on a clothing product page, on the individual's tourism blog, or it's a stock image in a presentation demonstrating modeling poses. Context helps you pinpoint the purpose of the image and decide which details to include based on that.
    1.4. Consider who the potential audience of the image is and what they might care about based on where they are going to encounter it – fashion magazine readers will have different informational needs than elementary school students, for example. While determining a specific intended audience can sometimes be difficult, context can be a huge hint. If there are multiple possible audiences, consider how different likely audiences will understand your description and tailor or broaden your language appropriately.
    1.5. Avoid repetition of surrounding text – images have a direct context (the text that accompanies that image in the place it is being posted) and extended context

(information gleaned from multiple sources related to the subject, artist, topic, etc.), when you are writing image descriptions or alt text do not repeat text from the direct context (as repetition is highly unpleasant to many screen reader users considering the already tedious process of navigating with a screen reader much of the time) but you can use information from the extended context to inform and lend increased accuracy to your description

1.6. Do not start alt text, or a labeled image description, with a generic phrase like "image of" or "graphic of" as screen readers already read out "graphic" before speaking the alt text of an image, and the image description label will already cue someone that it's referring to the image. If the specific type of image (photo vs screenshot vs digital painting, etc) is relevant, it can be included, along with a potential short description of the style (e.g., black and white photo, watercolor portrait, lineless art, etc.)

1.7. Repeat any text from the image (including specific numbers, dates, etc.) and include details if they may be relevant (e.g., in a research publication the details of the graph values may be important, even if it leads to very long alt text).

1.8. For alt text (or short initial version of image description) use specific names rather than sensory details, particularly when describing commonly known images. For example, saying "US flag" or "the Mona Lisa" is more appropriate than going into detail about colors and shapes and artistic details.

1.9. Be consistent with names and labels, particularly when describing the same element multiple times (such as a repeated logo, the same character in a sequence of images, etc.)

1.10. If the image is not adding any additional information besides what is already covered in the direct context, leave the alt text blank or give a short description labeling the image as decorative/repetitive. Do not repeat information unnecessarily, but ideally do not leave users guessing if the image was simply not made accessible and might in fact be important.

2. Definition and (Non-Exhaustive) List of Social Identities

2.1. Social identities are any socially constructed identities or any identities that have a personal or culturally situated meaning. How you describe social identities matter and should, whenever possible, not be assumed about a subject in an image, as the language around social identity can be very fraught. Using inclusive language to reference social identities is an indication, although it does not guarantee, that the imagery or at least the image description is made with care towards marginalized identities that are depicted.

2.2. Social Identities include (and take the following form in image descriptions):

2.2.1. Gender identity (e.g., he, she, they, or neo-pronouns, which gendered terms to use such as boy/girl, actor/actress, blond/blonde, niece/nephew/nibling, etc.)

2.2.2. Race, ethnicity, or nationality (e.g., Black, Latino, person of color, South Asian, Mexican, Asian-American, etc.)

2.2.3. Ethno-religion, religious or cultural / ethnic group (e.g., Māori, Jewish, Sikh, Choctaw, etc.)

2.2.4. Disability status / identity, medical diagnosis (e.g., Autistic person, wheelchair user, legally blind, person with cerebral palsy, person with a disability, etc.)

2.2.5. Sex, sexuality, or sexual / romantic orientation under a split attraction model (e.g., same-sex couple, mlm (man loving man), asexual, lesbian, biromantic homosexual, queer, etc.)

3. Images Without Subjects or with Real, Non-human Subjects

3.1. If the image has a lot of text or is only text (e.g., the screenshot of a tweet, a passage from a book), run the image through an Optical Character Recognition (OCR) program such as https://www.onlineocr.net/ or your preferred OCR software. Remove text that isn't relevant to the context or audience (e.g., timestamp or @ from a tweet, or page number of a book is not, in most circumstances, important or interesting)

3.2. If the image is complex (e.g., a diagram, table, graph, or a complex and highly detailed piece of art) there should be a short description in alt text or initial image description and a longer version (see 1.1.). The shorter version should feature a high-level summary (e.g., for a graph include the type of graph/diagram, axes or labels, and trends) or description of key elements.

3.3. Generally, describe the main content, including details only as necessary given context. If you don't know the correct term for visual details or wording for what you are looking at, try consulting with peers, looking for specifically tailored guides online, or reaching out to an expert (e.g., if you don't know the correct term for a

clothing style, finding a YouTube video discussing style trends in the given time period may be necessary).

3.4. However, do not use overly specific terms that many readers won't be able to understand or interpret (e.g., don't call something "sea glass green" or other precise terminology, when "light grey-green" might be clearer).

3.5. If you cannot find the correct way to describe the image or you can't figure out what you are looking at, describe what you believe is in the image to the best of your ability, with additional language to denote your uncertainty (e.g., "what appears to be a very long poodle skirt" or "something that looks like a sea urchin").

3.6. For photos of non-human subjects (e.g., an animal, or a robot who looks humanoid or appears to be demonstrating agency), use the correct name given to the non-human and/or the correct pronouns if you know them or can find them out, otherwise use gender neutral (they/them or it/its) pronouns and focus on describing what you see (see 7.1.)

4. Photographs with Human Subject(s) Known to the Describer

4.1. If the subject(s) can be consulted directly, ask them to describe the photo or themselves, or ask if there are any social identities / visual details they would like to have mentioned. If certain identities are relevant to the image's purpose in context, ask subject(s) specifically how or if they would like the relevant identities to be described and defer to their preferences.

4.2. If the subject(s) cannot be consulted directly, refer to any documented language preferences for this specific subject around any relevant social identities (e.g., look up the Wikipedia page for a famous person, or find prior bios written by a deceased

subject). If you can't find any documented language preferences, consult with someone else who may know the subject's preferred social identity terminology (e.g., someone who is closer to them or more up to date on any shifting identities). If language preferences cannot be found out through either of these methods, see section 5.

5. Photographs with Human Subject(s) Unknown to the Describer
   5.1. If social identities are necessary in the alt text (i.e., relevant given audience, context, and purpose), and unclear from the context of the image: consult with other people or sources, specifically seek out information directly from or written by people who have the same apparent social identity regarding most inclusive or correct language in general. If no clear answer is available after researching the identity group, refer to 5.3. below.
   5.2. If social identities are not necessary to the alt text and identity language can be avoided (e.g., there are no gendered terms necessary, race is not relevant or doesn't add anything to the image), then use neutral language / do not refer to the identity.
   5.3. If social identities are not necessary to the alt text but identity language cannot be avoided (e.g., contextually it's important to describe someone as blond/blonde), use your best judgment or guess about the subject's identity while indicating uncertainty (see 3.5.)
6. Illustrations or Artistic Depictions with Subject(s)
   6.1. If the art depicts a real person who the describer can identify, see section 4.
   6.2. If the art depicts a fictional subject from an established fictional work, and the art visually depicts a canonical social identity, use the terms from the source fiction (if

relevant to the image in context) or don't mention the social identity as it will already be known to those familiar with the character.

6.3. If the art depicts a fictional subject from an established work, but the art visually depicts or signals a social identity differently than it is established in the canon, use the terms the artist uses or states are correct. Or describe the identities using the preferred language by that social group (see 5.1.).

6.4. If the art depicts a fictional subject from an established work, but it is unclear what identity is being depicted, see section 7. Guidelines for Ambiguity.

6.5. If the art depicts a fictional subject, but you do not know what established work it is from, then try consulting with others, reverse image searching, looking at extended context, etc. to find the source fiction. If the source is found, see 6.2.-6.4. (above). If the source is not found, see 6.6.-6.7. (below).

6.6. If the art depicts a subject who may or may not be fictional / real (e.g., illustration of a person where you don't know if it is based on a real person, or established character, or neither) and the identity is relevant to the image description, see 7. Guidelines for Ambiguity.

6.7. If the art depicts a subject who may or may not be fictional / real, but social identities are not relevant to the alt text or image description, see 5.2.-5.3.

7. Guidelines for Ambiguity

7.1. Ideally do not assume anyone's social identity: instead describe visually apparent details (e.g., describe visible skin tone rather than racial or ethnic identity, assistive aids rather than disability identity, etc.)

7.2. If describing visually apparent details (such as potential signals for gender expression that a sighted person uses to draw conclusions about gender identity) would create a significantly longer or harder to parse description, you can make some assumptions but indicate hedging language (e.g., appears to be feminine presenting, might be a trans masculine person)

7.3. Even when making assumptions, do research to try and understand inclusive language and/or the language preferred by members of that group (as seen in examples above). Avoid overt bias or stereotypical / othering language such as describing a wheelchair user as "confined to a wheelchair" or describing a person of color as "ethnic."

Models of Accessible Imagery Creation Processes

For the models of accessible imagery creation processes, I will start with descriptions of each method including key players, when they're best used or are most likely to apply, and advantages or disadvantages to each method. I then present a table with relative pros and cons of each method so that processes can be compared directly to each other.

The Direct Hand-off Method

The Direct Hand-off Method is what I'm calling the Lone Writer Process discussed in Study 2. I have shifted the name slightly to cover the fact that the process, at least taken to its logical conclusion, is about more than just the writer working alone, but all the actors working alone. The roles involved in this process are fairly elastic. At the bare minimum you must have: the subject, the artist, the describer, and the end user. However, there may

be additional roles depending on the context, such as the content writer or the QA tester, as seen in the Formal Iteration Method. Under this process, as the original and new name suggests, the subject is depicted by a single artist, who then hands off the images to a single describer. The single describer works alone to produce the alt text, and then hands them off the accessible imagery completely to the content writer or QA tester or puts them directly into the content management system and pushes them to the end users.

In some ways the Direct Hand-off Method is different from the other processes because it is more of a paradigm or a modifier than it is a completely different process. Additionally, the method is defined by a negative: it is the process without collaboration and feedback. However, despite this there are a few advantages to this approach. Because gathering and propagating feedback between many parties is logistically difficult, and because extended iteration is labor intensive, the Direct Hand-off Method is in many ways the most streamlined process. However, because there is no collaboration between roles or between different people within a role, the largest drawback of this method is the quality of the final product may suffer, particularly for complex images.

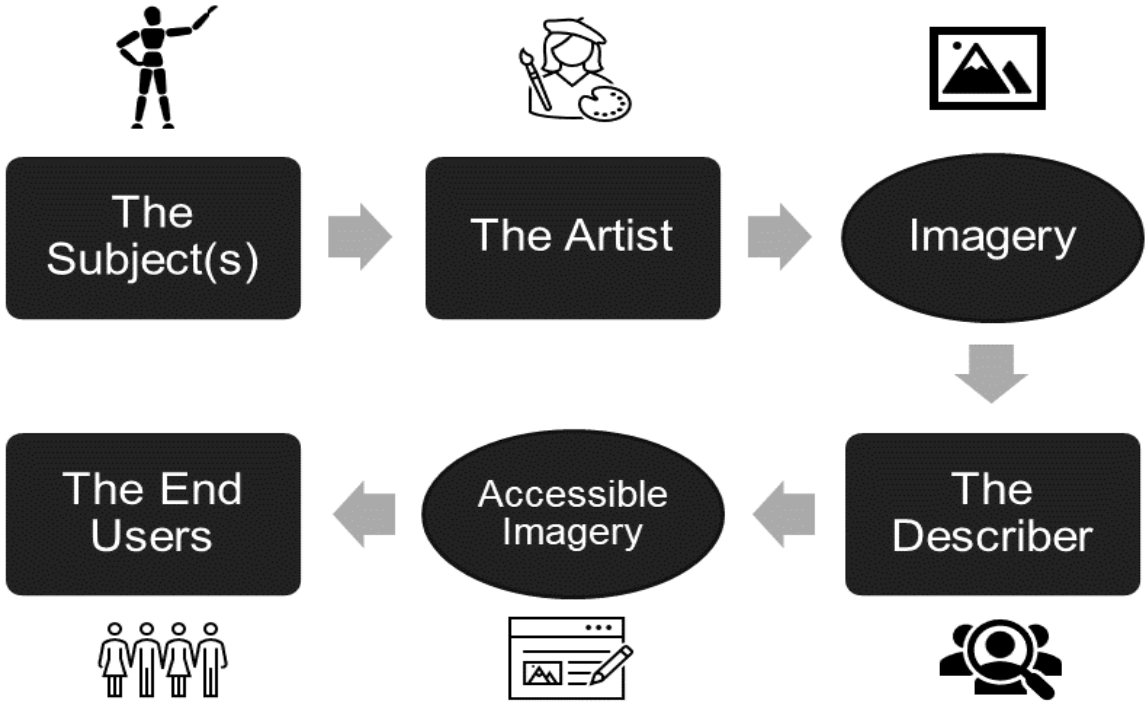

Figure A0.1: Actors and Products in the Direct Hand-Off Process

The Collaborative Describing Method

Collaborative Describing is most similar to the process I identified in Study 2 as the Team Write-A-Thon Process. However, I'm including in this version of the process some of the elements I observed in how the accessibility community in Study 3 conducted their practice. Collaborative Describing is characterized by having these roles: the subject, the artist, the team of describers, and the end users. The primary mechanism that differs in this process from other processes discussed is that the describers collaborate to create the alt text. Because there are multiple describers, there can be a broader range of identities represented on the team, potentially removing some of the need to consult with content experts through an Inclusivity Evaluation. However, the mechanism to reach consensus is complicated.

What I didn't originally describe in the Team Write-A-Thon Process, but which became clear from watching how image descriptions were made collaboratively through the PDDS, is that for a collaborative process to truly reach consensus, rather than just a lack of continued discussion, there needs to be a leader within the team of describers. If the team of describers is too large or is without a leader, then it is very likely that some describers will be paying less attention to learning about accessibility or will simply remain silent rather than actively disagree with the group. In these cases, consensus appears to be reached, but in fact isn't, because not all describers will go on to use the same terms consistently or to understand their role as describers fully. Thus, I could potentially add one additional role to the set that makes up this process: a lead describer. The lead describer's role would be to make sure everyone on the team of describers was actually on the same page. They would also function as somewhat of a mediator because their role would be important for crafting the collaboration in such a way that all describers felt comfortable voicing their thoughts, rather than being socially pressured in any way to keep silent.

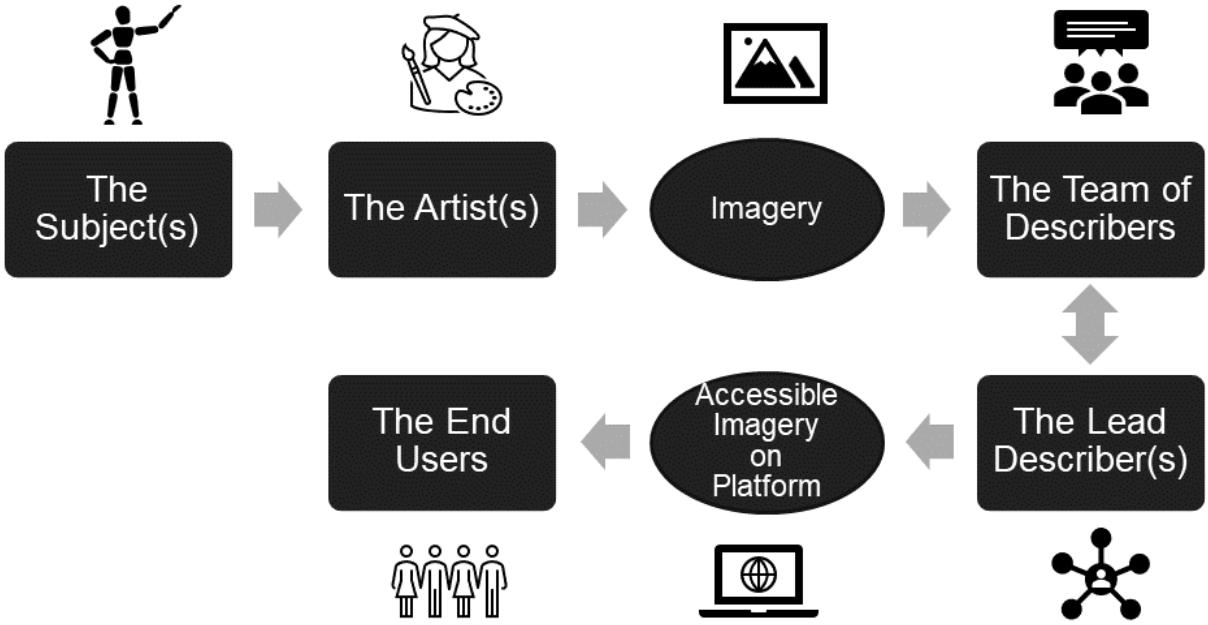

Figure A0.2: Actors and Products in the Collaborative Describing Process

The Community Generated Method

The Community Generated method is based most closely on the process used by members of the accessibility community observed in Study 3. It is characterized most strongly by the fact that roles overlap, or individuals serve in different roles at different times. In an attempt to retain clarity, I'll first present the roles in their most simple form: the subject, the artist, the poster, the describer, the team of volunteers, and the end user. The process involves the artist depicting the subject and then either having their art (in the form of a movie screenshot or other type of common social media content) posted by another person or playing the role of the poster themselves. After that, the post reaches the describer, either directly after or with many end users consuming the post before it reaches someone who will describe it. The describer generally attempts to write a description, or at least has the intention of writing a description, but goes to the team of volunteers to help.

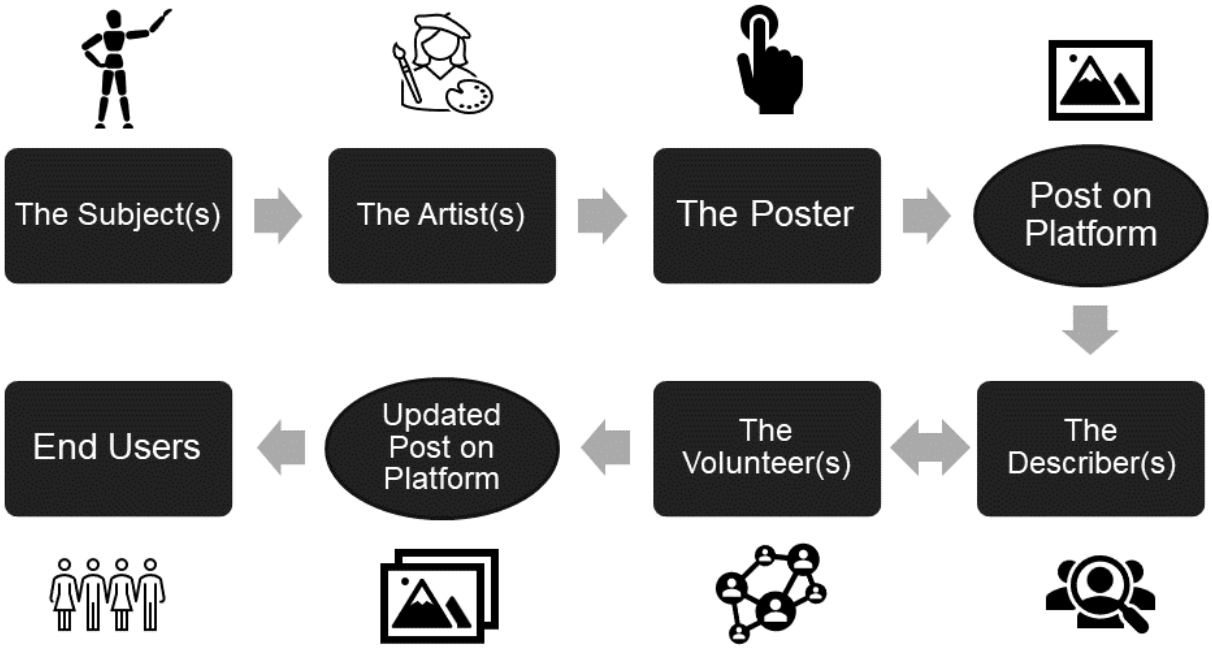

Figure A0.3: Actors and Products in the Community Generated Process

The volunteers function as co-describers who generate the description in its entirety or suggest edits to the description. Once the team of volunteers and the describer reach some kind of agreement that the description is sufficient, the post then gets passed along to more end users, now in an accessible format. The tricky thing about the team of volunteers, however, is that they function simultaneously as describers and users. They often present their comments or critiques based on their own reading of the image description or personal issues they may have with it. This is not necessarily a bad thing, indeed if the volunteers themselves have disabilities or other marginalized identities it means that they functionally combine the two Evaluation Methods (discussed below) within the process of collaboratively describing. However, it does make it hard to clearly define the role that the volunteers play, particularly when from one moment to the next a describer may turn around and volunteer to provide feedback to another describer's request.

Ultimately, there are a lot of exciting potentials to the community generated method, because it has the potential to be both highly distributed and efficient for a wide variety of image types. However, at least in the PDDS I still observed that users were taking on a higher labor and emotional burden than they had to spare. Additionally, while community norms and distributed learning allowed for a certain degree of agreement and education to be reached by most members, there was not a formal process for either checking understanding, assessing quality, or integrating new findings. Thus, the actual quality of the image descriptions produced is variable, with the amount of practice as volunteers apparently predicting how well they describe.

The Accessibility Evaluation Method

An Accessibility Evaluation is a process for creating imagery that focuses on making sure the images and descriptions are fully understood by and properly prepared for consumption by disabled audiences. The focus of an Accessibility Evaluation Process is to create a positive user experience of the imagery in context. Thus, it would include usability testing with screen reader users, as well as confirming the comprehensibility of the imagery. For example, an Accessibility Evaluation Process would involve determining the ideal combination of wording, capitalization, and punctuation given different screen readers and verbosity settings. But it would also deal with the overall comprehensibility of abstract or unfamiliar language, written tone, and other stylistic or content elements. Importantly, a full Accessibility Evaluation would go beyond testing with screen reader users, and test the accessibility of the images and alt text with a wide variety of disabilities. For example, imagery and image descriptions (or the experience of consuming it) may have to be customized for people with cognitive or learning disabilities, people with various

types of color blindness, low vision users who may use a combination of visual and written information to interpret the image, or users who are prone to migraines and may have

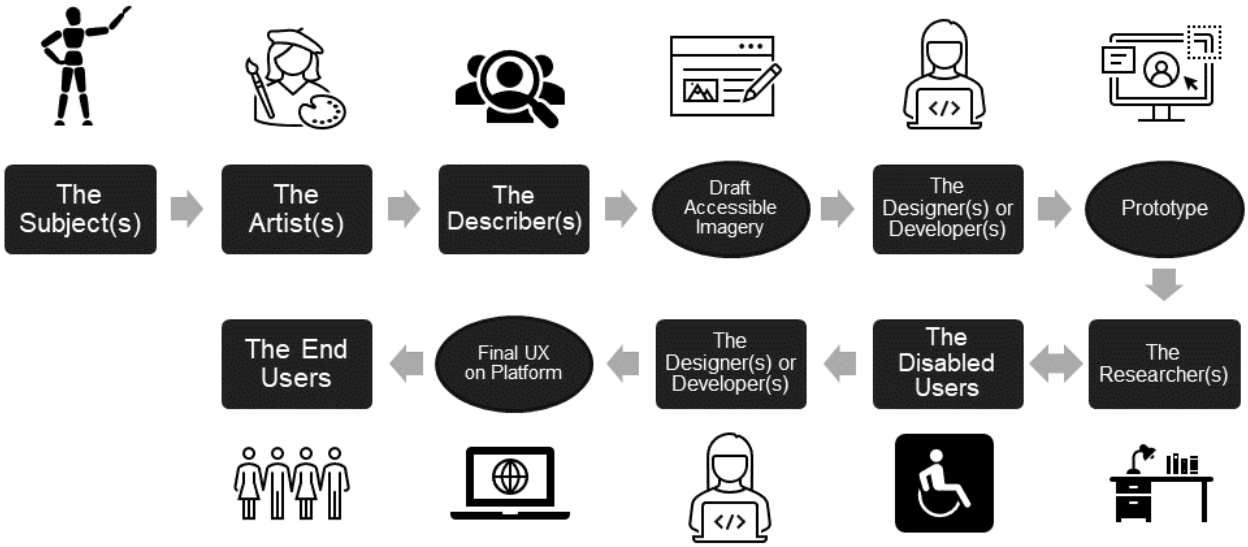

Figure A0.4: Actors and Products in the Accessibility Evaluation Process

sensitivities to certain colors on digital displays.

The general roles involved in an Accessibility Evaluation process are the subject, the artist, the describer, the designer and/or developer, the researcher, the disabled users, and the end users. The subject is depicted by the artist, which is then translated into language by the describer. The designer and/or developer takes the image and alt text and places it into a prototype or real user experience. The researcher then tests the user experience with the disabled users. After feedback is received and the developer and/or designer tweaks the user experience, the final images and alt text are delivered to the end users. This model only describes changes made at the design level, however obviously there is some feedback that could be received from the users with disabilities which would require going back to the describer or even the artist to change fundamental parts of the imagery and alt text to make it accessible. However, such a process could be more accurately described as a

combination of an Accessibility Evaluation and the Formal Iteration Method described below.

The Inclusivity Evaluation Method

An Inclusivity Evaluation, in contrast to an Accessibility Evaluation, focuses on determining the most appropriate, accurate, and welcoming ways to depict social identities in imagery and image descriptions. To do this, the imagery creation process focuses on consulting with content experts, i.e., people who have similar identities to those planning to be depicted in the imagery, who can provide their feedback on how their identities should be depicted. The roles involved in an inclusivity evaluation are the subject, the artist, the describer and/or researcher, the content experts, and the end users.

As shown in the co-design I conducted in Study 1—an early version of an Inclusivity Evaluation—, the person conducting the evaluation can be, and may well best function as, the writer of the alt text. This is theoretically also true for the Accessibility Evaluation Method described above, but logistically finding a user researcher who is experienced with a range of disabilities and assistive technology, capable of pinpointing usability issues for the designer/developer to address, and adept at writing and re-writing alt text as needed may be difficult. In both methods it is also just as possible for the describer and researcher to be separate people who are collaborating.

Essentially the process for an Inclusivity Evaluation has the subject being depicted by the artist. Then the describer translates the image into an image description, at which point the researcher uses focus groups or other qualitative methods to gather feedback from the content experts on the image description, image, or both. The writer (and ideally

the artist) then iterates based on the user feedback before the images are passed on to the

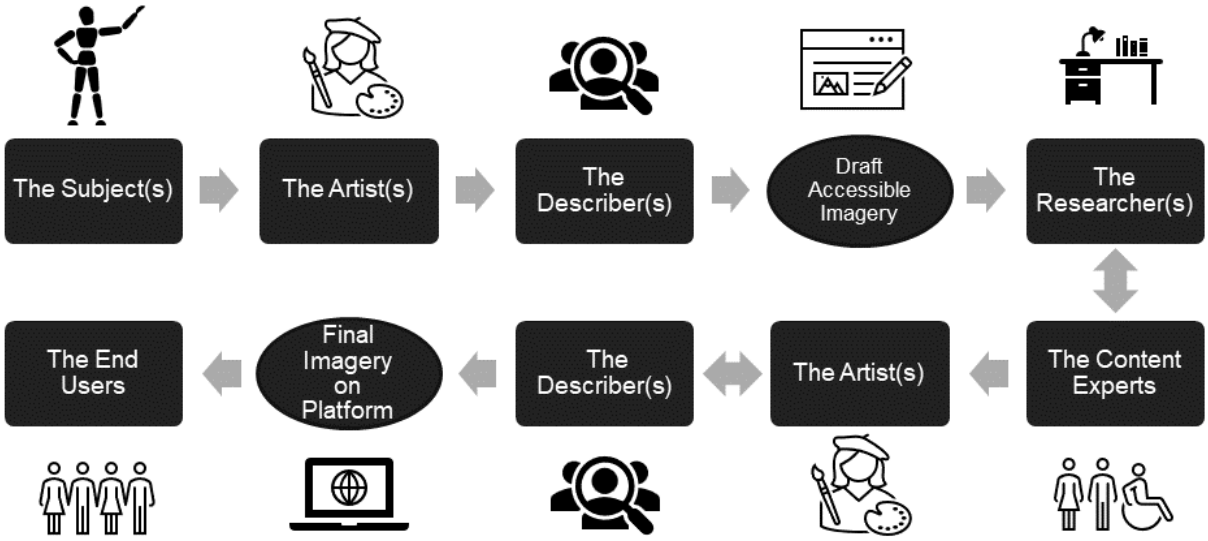

Figure A0.5: Actors and Products in the Inclusivity Evaluation Process

final end users.

The primary difficulty with this process is the pragmatic need to recruit enough people with the specific identities being represented. Because identities and identity intersections in particular [73] are so complicated to understand in-depth, theoretically an Inclusivity Evaluation would need to be conducted anytime a new combination of identities was being depicted. This is logistically difficult and would take a long time depending on how many identities were depicted in a specific image set. The additional challenge is what I identified as a drawback in Study 1, namely the fact that designing by committee, i.e., trying to find one solution that pleases everyone, can lead to an overly complicated or patchwork solution that ultimately pleases no-one completely.

The Formal Iteration Method

The Formal Iteration Method is most similar to the Artist-Writer Process from Study 2 with the addition of the other roles found in the study of this method. This type of process

is likely to be seen within large matrixed organizations where imagery creation is part of a product or service the organization provides. Specific variations on how the Formal Iteration Process is enacted are highly dependent on circumstances and pre-existing job roles in place. The model as I describe it in general has 7 roles: the subject, the commissioner, the content writer, the artist, the describer, the quality assurance tester, and the user. The subject (be they real or an amalgam of inspirations formed into a new fictionalized subject) is decided on between the commissioner and content writer. The commissioner gets the project approved by their organization and commissions the artist. In the ideal version of this process, the describer joins the project at this time as well, and

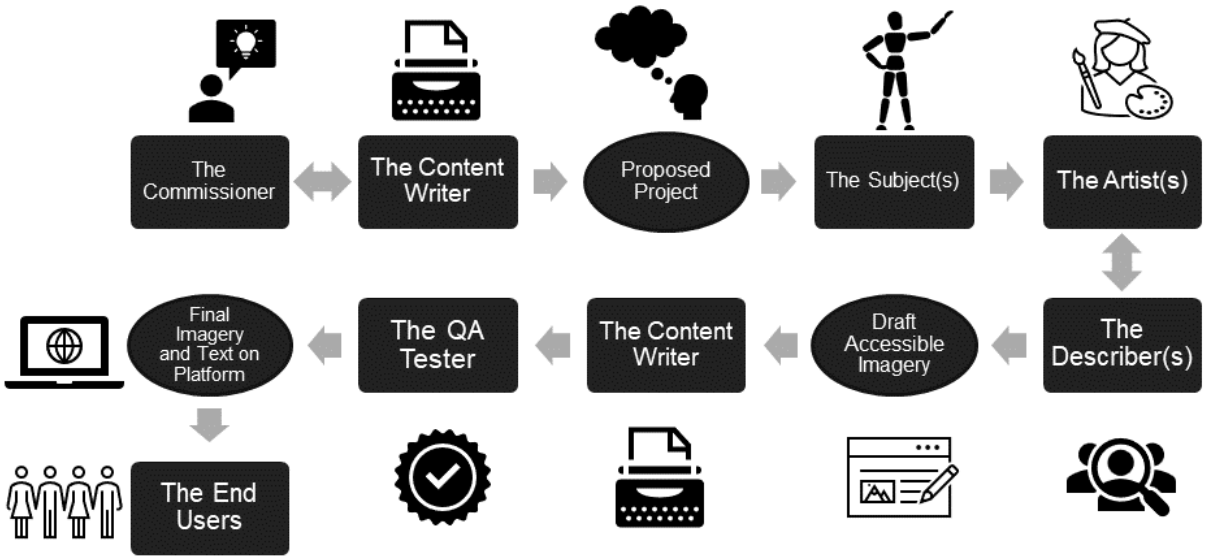

Figure A0.6: Actors and Products in the Formal Iteration Process

together the artist and describer iterate on depictions until the commissioner is happy.

The final imagery is then accompanied by some written content which was crafted by the content writer. This can sometimes be the same person as the commissioner, depending on if the imagery or the content was conceived of first. If the imagery was the main deliverable, as with the Avatar Project, then the content writer may only come into

the process at this later stage. But the accessibility practitioners in Study 2 also explained that sometimes the content writer would commission small pieces just to visually fit with the main content. If this latter version of the process is used, then the content writer may also, theoretically, iterate on the writing during the same time that the describer and the artist are iterating on the other components. Once all three components–the imagery, description, and primary written content–is finalized, the entire package is passed along to the Quality Assurance (QA) tester for final checks before being sent out to the end users.

Pros and Cons

Table A0.1: Comparative Pros and Cons of Different Imagery Creation Processes

	Accessibility Evaluation	Inclusivity Evaluation	Formal Iteration	Direct Hand-Off	Collaborative Describing	Community Generated
Feedback	Pro: Direct feedback from users with disabilities	Pro: Direct feedback from marginalized users	Mix: QA testing but no feedback from users	Con: No user or writer feedback	Pro: Direct feedback from writers	Pro: Writers act as users to give feedback
Buy-in	Pro: Buy-in from writers and a11y advocates	Pro: High level of buy-in from users and writers	Pro: Buy-in for project from management	Con: No buy-in beyond few people involved	Pro: Buy-in from the many writers involved	Pro: Shaped and approved by norms of community
Distribution of Expertise	Mix: Writers need to consult many PWD but can apply broadly	Con: Writers need to learn rules for every identity intersection	Pro: No need to teach anyone; roles are played by experts	Pro: No need to teach anyone; roles are played by experts	Pro: All writers learn about alt text and how to make it	Con: Different community members have differing expertise
Use Case	Pro: Good for images served to many users	Pro: Good for images with high visibility	Pro: Avoids unintended ambiguity for identities	Pro: Good for large, simple image sets	Con: Hard to reach consensus on images	Pro: Good for wide variety of image types

	Accessibility Evaluation	Inclusivity Evaluation	Formal Iteration	Direct Hand-Off	Collaborative Describing	Community Generated
Integration of Artist	Con: Complex; hard to integrate artist	Con: Complex; hard to integrate artist	Pro: Means downstream uses reflect original intent	Con: Does not attempt to integrate feedback with artist	Con: Complex; hard to integrate artist	Mix: Inconsistent; sometimes artist is integrated
Timeline	Mix: Can take a long time, but fits in with other user research	Con: Can take a long time and be logistically difficult	Con: May take many iterations to reach correct depictions	Pro: Quick turnaround, both per-image and per-set	Con: Reaching consensus can take a long time	Pro: Quick turnaround per-image and wide reach
Consistency	Con: Might be impossible to find one solution to all access needs	Con: Can suffer from "design by committee" issues	Pro: Consistency within a set plus potential for personality	Pro: Consistency in form / language	Pro: Once consensus is reached, ongoing consistency	Con: High variability depending on how involved community members are
Labor Burden	Con: Requires labor from users with disabilities	Mix: Requires labor from users but they also feel respected	Con: Requires a lot of labor from artist to iterate	Con: High labor and emotional burden on individuals	Mix: Can lead to distributed labor after initial time investment	Mix: Highly distributed, but still undue burden on participants
Complex Images	Pro: Works well for complex images and experiences	Pro: Works well for complex images and depictions	Pro: Complex imagery / concepts can be simplified	Con: Potential bottleneck for complex images	Con: Consensus may be difficult with complex images	Pro: Works well for large volume of complex images

Crowdsourcing Community for Accessibility: An Owner's Manual

The owner's manual is set apart from my other contributions by both tone and intent. While all contributions are designed to be practically beneficial in some ways, the owner's manual was originally conceived as a toolbox or even a how-to manual. It is

informed deeply by my research on PDDS, as well as by the literature around accessibility communities, peer production communities, and communities of practice. However, due to the format and the intertwined nature of PDDS with the theoretical accessibility community I'm describing, direct citations do not feel appropriate. Instead, whenever possible, I have used examples from or recognizable references to content in Chapter 4.

Getting Started

To start an online accessibility community or start using volunteer-based crowdsourcing to increase accessibility, there are a few things to keep in mind.

Key Components

The key components of an accessibility community can be divided into three categories: goals or purpose, technical or social prerequisites, and necessary roles. In some senses your goal or purpose with an accessibility community will always be the same–encourage people to learn about and practice accessibility. However, having a more specific purpose or starting point can be very important.

When starting an accessibility community, it can be helpful to tap into and work within an existing source community or target platform. If you can tie your specific accessibility goals to an existing source, something that people are already passionate about, then you're more likely to have a steady stream of people joining and investing time in your community. Specifying what you want your community to do in terms of a target platform, a specific body of work that needs to be made accessible, or a type of accessible accommodation that needs support, can be valuable. When deciding what the purpose of your community is, I encourage you to set your sights on a specific but always expanding set of images, websites, or other documents. Making user-generated content accessible within the bounds of a specific platform, topic, or type is often a good central purpose for a community. Although keep in mind the bounds of that purpose may change and expand organically over time.

After deciding the goals and bounds of your community, you are already part of the way to reaching the technical and social prerequisites that make up the second key

component of an accessibility community. Another thing you will need is a medium or host platform to work within or organize on. This may be the same platform that you are seeking to improve accessibility for, but it doesn't have to be. However, when deciding what platform to use for organizing, consider that there will be a certain amount of drop-off when moving platforms. So, if it's possible to host your community organizing on a platform that the existing community is already using, then that will lead to fewer potential members failing to join in your accessibility community because they are not comfortable shifting platforms.

In addition to a prerequisite digital space for organizing and hosting your community, it can be valuable, although not required, for the target platform you are looking to change to have some basic accessibility features. This is helpful for two reasons. First, if your accessibility community follows the pattern of in-person communities, people who have disabilities themselves are more likely to become engaged in supporting your accessibility efforts. Thus, having basic accessibility features already in place on the host and target platform, even if they're not perfect or widely used, can be important for allowing the disabled members of your community to effectively participate. The second reason having basic versions of accessibility features present on the target platform is important is because there needs to be something to build off of and advocate for. Your community's actual interventions do not always have to use the given accessibility features, and you may instead find it more advantageous to bootstrap a better, more accessible solution from other features of the platform. But importantly, if there are no people with disabilities currently able to use the target platform, particularly people with the disabilities your interventions are trying to help, then convincing

potential members of the source community that it is worthwhile and important to make the platform accessible will be significantly harder to do. As we'll discuss later in this manual, members want to feel like their work matters and makes a difference for people. Starting an accessibility community with the goal of hopefully attracting people who need those accessibility accommodations to the platform means that at least to start there is no obvious or direct beneficiary to your community's efforts. This therefore forms the first social prerequisite for your accessibility community: the presence of actual users on the target platform that could benefit from increased accessibility.

Beyond these three prerequisites, the additional social prerequisite, mentioned earlier when talking about goals, is that there ideally needs to be an existing source community, or group of people who use and care about the target platform. These will form your pool of potential accessibility community members. There are some characteristics that, if you find them within your source community, may indicate that that group of people are a good fit for joining an accessibility community. Some of these characteristics are: passion (both in general and specifically directed towards the target platform), concern for social issues or equity (if this already includes disabilities that is wonderful, but concern for other social issues are still a good indicator that consideration of accessibility may catch on), and time or energy to spare (this is not an absolute requirement if you have the other characteristics, but it can definitely help). If you have found a source community and understand it well enough to engage with that community without being seen as an interloper, then you can move on to the final set of prerequisites: necessary roles and players.

The key roles that an accessibility community needs are: a leader or leaders (you are hopefully in this group), moderators or experts (i.e. those informed about accessibility and/or community rules who can guide and act as points of contact for new community members), visible community members (i.e. members who are visible enough on the target platform or within the source community that they can successfully conduct outreach), inaccessibility ameliorating members (i.e. the volunteers that are going to actually be doing the bulk of the accessibility work), and then inaccessibility pinpointing members (who may or may not also be beneficiaries of the community's accessibility efforts). These roles are likely to overlap, with some users acting as pinpointing members and visible members at the same time, or as moderators and ameliorating members. There will probably be at least a few members, however, that have a primary alignment with each role. It would be paradoxical to claim you have to have these members before beginning your community, as most or all members of the community will join after the leader(s) have completed the setup (see next section). Nevertheless, it is a prerequisite to at least understand what types of source community or target platform users might fit into these roles. You have to sufficiently understand the context you are building your community in to have allies, or at least understand how friends and allies are made on the target/host platform, and to have a clear vision for what types of source community participants might be able to fall into each of the necessary roles.

While not necessary before starting, it's also important in the early days of your community to establish expectations around some of these roles. In particular, having open discussions around what constitutes a moderator or expert in your community (even if originally you are the only moderator) is important. Additionally, you must have

a clear understanding of what everyday members are being asked to do and how they will do it, including what combination of target/host platform communication tools is your community going to use, how will accessibility issues get ameliorated or information shared with community members, and how membership can be signaled. It can be beneficial, although not strictly necessary, in the early stages to also get at least a few members who can work as inaccessibility pin pointers. Be they actual users who benefit from accessibility accommodations, or simply users of the target platform with an understanding of what the community's purpose is, finding people who can point out where the problem lies in existing practices can be hugely beneficial when first getting started.

Setup

Once you have all the prerequisites, and an idea or plan for all your components, actually starting your community is still a little tricky. While there may be many ways to jumpstart an accessibility community, the method I'm familiar with is education and outreach through modeling behavior. Essentially, at the beginning of your community you (and any other community leaders you already have on-board) will have to take on all the necessary roles. You will be the experts, or at least the avenue for expertise to flow through. You will be the key visible community member, attaching your accessibility efforts to a popular or up-and-coming topic or trend in your source community. And you will be both pinpointing inaccessibility and doing the work to solve it. The leader(s) of an accessibility community have a hard role at first, and this is when your community is most likely to falter. However, with an understanding of the prerequisites and a hopefully deep dedication to accessibility and/or the target platform, you will hopefully find that

visibly demonstrating the problem and its solution for other source community members will naturally draw allies and eventual members into your community.

Once your community has enough members that other moderators, experts, and key members can share the workload, or once the pinpointing members and ameliorating members can function as a stable feedback loop for increasing accessibility. Then you will know that your community has officially completed the startup process. Any shifts to host platform, target platform, or purpose after this stage should come only organically or if supported by the bulk of your existing members.

Ongoing Maintenance

Keeping your community running smoothly may involve extra work from you as the community leader, as well as investments from other moderators or key members. However, it is also important to rest once your community has passed the fledgling stage. Burnout is a clear and present danger to all accessibility community members, so part of community maintenance is setting aside time to replenish personal energy stores. As a community leader you did a lot of work to start up the community, allowing others to carry it forward while you take a step back can be valuable. It also shows your fellow community members that participation is valuable only to the degree that you don't burn out your long-term energy.

Encouraging Participation

Communities are likely to go through natural ebbs and flows of activity. Depending on the source community you are drawing most of your members from, these fluctuations may by minor shifts that align with certain times of the year such as holidays, or it's just as likely your community could be impacted profoundly by world events,

changes in the target or host platform, personal upheavals of highly sociable members, and other unforeseeable occurrences. Therefore, when community participation initially lulls, don't immediately panic, it may be a natural or cyclical downturn and participation will bounce back over time. However, as you gain more experience with your community, you may begin to notice an overall downward trend, particularly from ameliorating members who often have the most important job in the community and invest the most time and effort into practicing accessibility. If this is indeed something you're noticing, I present a few suggestions for encouraging more participation.

First, understand that the source of motivation in a volunteer accessibility community is not likely to be a sense of competition. No one is trying to be the most active member or the loudest voice in the community. Instead, participation is valuable to members, particularly ameliorating members, because it generates a sense of camaraderie and satisfaction. So, one way to encourage participation is by setting a precedent or supporting members who provide positive feedback for participation. A simple "thank you! <3" from a pinpointing member when an ameliorating member supports accessibility is incredibly valuable. All types of positivity and generosity within the community should be uplifted and modeled by the leaders and moderators. Casually complimenting members on particularly good work, thanking them for quick replies or thoughtful suggestions, general compliments even on unrelated topics such as a member's icon or aesthetic, all of these things make the space feel welcoming, supportive, and enjoyable to participate in. The more goodwill and positive associations you can create when members participate, the more likely they are to come back and participate more.

The second valuable element to community participation that can be supported is the sense of community itself. Ideally, members of your community would get a sense of social connectedness or belonging from interacting with other members. This can take the form of inside jokes, personal conversations and comments, making connections based on shared interests outside accessibility. The same ways you might support a social club or set of colleagues–through social events, crafting casual opportunities to get to know each other, and fun shared activities–can be helpful for crafting a more tight-knit accessibility community. The only challenge with this method is balancing the natural fact that not everyone will have the time, energy, or inclination to socially engage with the fact that supporting those who do can be valuable. Ultimately, as a community leader, you must consider if the core group of active members is becoming too small and cliquish in a way that discourages more widespread participation, or if the camaraderie demonstrated by core members makes the overall space feel fun, supportive, and collegial.

Growing Your Community

Another key element of community maintenance is encouraging new members to join. While continuous or exponential growth of your community is not the goal, recruiting enough new members to replenish those who have to step back from active participation is a good idea. And of course, the ultimate goal is for the target platform itself to change, which would mean enough people either officially joining the community or unofficially adopting accessibility practices that a new standard is set.

The most important way to grow an accessibility community is essentially outreach through example. For all members, but particularly your visible members, engaging often and actively on the target platform is important. As long as your members

can function as valuable members of the source community, while still practicing and educating around accessibility, they are functionally recruiting users to the community. Connecting to and nurturing visible members or members engaged in generally popular activities or trends within the source community pays dividends in terms of education.

Part of the process of educating is demonstrating what it looks like to practice accessibility: demystifying the process so casual members of the source community can feel comfortable beginning to ameliorate inaccessibility as well. Another element is more direct education or outreach: answering questions, sharing resources, explaining many times and in many ways why accessibility is important. Some users may join right away, but the majority of target platform users will need extended interactions with and discussion of accessibility before they fully understand its value. Even after understanding how valuable accessibility can be, depending on the complexity of inaccessibility amelioration in your community's context, or depending on how difficult it is to join or participate in the host platform, getting users to take the leap from wanting to help to actively practicing accessibility can be another long-term project.

Troubleshooting

This is the final portion of the manual, where we'll cover troubleshooting some potential issues that may arise in your community and how you can address or preempt them.

Minimizing the Bystander Effect

Related to encouraging participation, as discussed above, one issue you may have while running your community is that there are "loafers" or community members who join but don't actively participate. In most cases this is a natural element of community

make up: there will always be many more people comfortable "lurking" than actively engaging in discussion. You may notice a problem in your community in particular if there is a mismatch between the number of members pinpointing inaccessibility and the number of members ameliorating them. If many more issues are being raised than there are solutions being offered, you may find that community members become discouraged.

One way to address this mismatch is supporting positive feelings from community participation in general, as discussed above. Another solution may be to gently push members into different or more active roles. If more pinpointing members also take the time to ameliorate, the ratio of issues to solutions will hopefully stabilize. The most valuable tool for encouraging this is in fact members' own self-set expectations for reciprocity. If there are many more problems being raised than solutions, then suggesting or demonstrating to members that they should potentially offer a solution to the previous issue before raising their own problem can immediately level the playing field. While all members are volunteers, ameliorating members generally are spending the most time and effort on their labor. It's about shifting users into the mindset that investing a little more energy for others, by ameliorating for them, results in others supplying labor for the issues they have pinpointed.

Withstanding Burnout

A long-term issue for accessibility communities is the potential for individual and collective burnout. It can help to address these two types of burnout separately. For individual burnout, even when it is happening to a number of members at once, is essentially an issue of investing more energy than the volunteers have to spare. To address this, allow and support members in taking breaks, building their energy stores

back up. Or, decrease the amount of energy individuals need to invest, either by making the labor more fulfilling or by distributing tasks more evenly according to member's willingness and ability. Remember that the ultimate goal is for enough people to adopt accessible practices that it becomes second nature, habit rather than active effort.

For collective burnout, that is a community-wide feeling that their time and labor is not being valued or is not making a difference, there are a few suggestions I can make. First, allow volunteers to openly vent their frustrations to one another. Even if it's not pretty or productive, it's important to have a shared space where members can be bitter or upset when their efforts to support accessibility are rebuffed. That said, don't build up spaces just for venting or complaining as much as possible redirect members' frustration with others back into dedication to and respect for their fellow members. Emphasizing the shared values, community practices, and emotional payout that arises from participation ultimately turns the chore of pushing the boulder of inaccessibility ever upwards toward the peak, into a shared responsibility and point of satisfaction in energy well-spent.