

Pointing at Data Together: Establishing Common Ground with Multi-modal Data Representations in Blind/Sighted Collaboration

Jonathan Zong  and Arvind Satyanarayan 

Abstract—Research on collaborative analytics has long recognized that data representations enable collaborative sensemaking by establishing a common point of reference. Yet research on visualization accessibility has largely approached data representations in terms of individual comprehension—whether they enable a blind reader to independently extract equivalent information. In this paper, we draw on linguistic anthropology to understand data representations’ central role in mediating communication in collaborative, mixed-ability settings that include both blind and sighted data analysts. Through a contextual inquiry with Bower Lab, an oceanography lab led by blind PI Amy Bower, we show how multi-modal representations provide shared linguistic affordances for data reference and social signaling that enable collaborators to establish common ground and manage joint attention. These findings suggest a reframing of accessibility in visualization as effective bidirectional communication—for instance, tactile representation not only makes data accessible non-visually but also makes a blind analyst’s analysis process accessible to sighted collaborators.

Index Terms—accessibility, data visualization, multimodal data representation, collaborative analytics, contextual inquiry

1 INTRODUCTION

When data analysts collaborate to make sense of data, a central challenge is supporting *common ground*: the shared understanding between participants that enables them to refer to the same things, direct each other’s attention, and build on each other’s interpretations [9, 17]. Establishing common ground depends heavily on shared reference and *deixis*—the ability to point to and direct attention toward entities in a shared context [19]. When an analyst points at a chart and says to their collaborator, “this outlier in the top right,” the visualization is providing a shared representation that enables both collaborators to reference and build a mutual understanding of the data.

Challenges around supporting common ground are amplified when teams of data analysts include both blind / low-vision (BLV) and sighted members. In such cases, collaboration typically spans multiple sensory modalities—one person may examine a visualization while another listens to a textual summary or reads a tactile chart of the same data. Sighted data analysts looking at the same visualization might be able to establish common ground by directly pointing things out. In contrast, mixed-ability collaborators must establish common ground across representations—e.g. a visualization and a tactile chart—that do not share a common perceptual modality or points of reference.

Unfortunately, little existing work has examined the role multi-modal representations play in helping mixed-ability collaborators establish common ground. Although recent work in accessible visualization increasingly emphasizes the value of multi-modal data representations [8, 23, 35, 36, 43], researchers have primarily evaluated these representations in terms of whether an individual BLV user can independently comprehend the data, extracting information equivalent to what a sighted reader gets from a chart.

This individual framing surfaces two tensions across the fields of visualization and accessibility. The first is the observation, central to collaborative analytics, that the social use of visualization is the norm rather than the exception [17, 19]. Analysts gather around visualizations in meetings, refer to them as they interpret the data, and present them to audiences to convey a message. Focusing on whether one person can extract the right facts from a chart overlooks the social and collaborative aspects of visualization. The second is a longstanding argument in disability studies that access is fundamentally relational—produced

jointly by the people who share a social setting rather than delivered by one person to another. Reading this argument alongside collaborative analytics suggests that the relational work of access and the analytic work of establishing common ground are intertwined. A mixed-ability team trying to achieve shared reference and mutual understanding for data analysis must also continually attend to one another’s access needs as they go. Understanding how multi-modal representations mediate this joint work is therefore necessary both for supporting collaborative sensemaking and for treating access as something a team accomplishes together.

If access in collaborative settings is accomplished through communication, then we need to understand how that communication works across modalities. While collaborative analytics has a rich literature on reference—pointing at and describing data—it has comparatively little for the other communicative acts that keep a multi-person, multi-modal conversation aligned, like confirming you have understood one another or signaling that you are still engaged. To analyze this fuller range of communication, we draw on linguistic anthropology, a field which studies how people communicate by combining speech, gesture, touch, and other semiotic resources. Linguistic anthropology gives us a vocabulary for communicative work that does not presuppose a specific modality, which lets us examine communication that leverages multiple modalities and kinds of data representations. We build on two frameworks in particular—Jakobson’s functions of language [34] and Keating and Sunakawa’s participation cues [25]—which we introduce more fully in Sec. 2.

Through a contextual inquiry with Bower Lab, an oceanography lab led by blind PI Amy Bower, we investigate how mixed-ability collaborators use shared visual and tactile representations to establish common ground across modalities. Our contextual inquiry surfaces three sets of findings.

First, we show that tactile data representations provide non-verbal affordances for three of Jakobson’s functions of language: the *referential* function, the work of pointing at and describing the data under discussion; the *metalingual* function, the work of checking that collaborators are understanding the same thing; and the *phatic* function, the work of signaling that collaborators are still actively engaged in the conversation. The referential function is the one collaborative analytics already studies closely, though not necessarily by this name. For instance, work on design considerations for shared reference and deixis in collaborative visualization [16, 17, 19, 31] discusses how analysts point to and name elements of a shared display. However, we show that the metalingual and phatic functions, which this prior work leaves largely implicit, become more prominent in mixed-ability collaboration. Verifying a reference requires cross-modal translation across representations, and signaling engagement relies on non-verbal cues like a hand

• Jonathan Zong is with University of Colorado Boulder (work done at MIT CSAIL). E-mail: jzong@colorado.edu

• Arvind Satyanarayan is with MIT CSAIL. E-mail: arvindsatya@mit.edu

Manuscript received xx xxx. 202x; accepted xx xxx. 202x. Date of Publication xx xxx. 202x; date of current version xx xxx. 202x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.202x.xxxxxxx

resting on a tactile map.

Second, we show how co-operative computer use between a blind and a sighted analyst—taking turns at a shared keyboard and mouse—depends on verbal participation cues and a shared cursor. Because neither collaborator can rely on simply watching the other reach for the mouse, they develop an explicit, spoken protocol for handing off control. This requires a cursor that both collaborators can place, and which is perceptible through both screen reader narration and visually on a computer monitor. Because the cursor is legible across multiple modalities, it can be used to resolve ambiguous references like “this next sentence.”

Third, we show that these collaborative practices are made possible by durable relational and institutional structure. Most notably, the lab has created a dedicated Access Assistant role—a half-research, half-accessibility position—that redistributes the labor of producing tactile representations away from the blind PI and onto shared infrastructure, an arrangement that makes the lab’s multi-modal workflow possible. The lab also models a collectively held, evolving body of access know-how, and a working culture in which collaborative sensemaking is treated as the preferred mode of analysis rather than a fallback from independent exploration.

Across our findings, we observe that common ground in mixed-ability collaboration is most effectively built and sustained through communicative signals that collaborators implicitly produce while doing their own work. For example, as a blind analyst reads a tactile chart, their hand placement also shows their sighted collaborator what data they’re attending to without needing to verbally describe it. We observe, however, that this information currently flows asymmetrically. A sighted analyst gains evidence of a blind analyst’s attention without much effort, simply by watching their hands move across a tactile map or by following the screen reader cursor across a document. There is no comparably effortless channel in the other direction. To learn where a sighted collaborator is looking, the blind analyst must rely on costlier, deliberate acts—i.e., asking, or waiting to be told. This suggests a promising direction for future research, in which designers support common ground through affordances that multi-modally supply ambient information about collaborator state, especially in the modalities that currently lack such affordances.

2 RELATED WORK

Our work is informed by interdisciplinary research threads across collaborative analytics, accessible multi-modal data representations, disability studies, and linguistic anthropology. In this section, we situate our contributions within this prior work, and describe the conceptual background we are drawing from in disability studies and linguistic anthropology.

2.1 Collaborative Data Analysis and Visualization

Research on collaborative visualization has studied visualizations as shared communicative artifacts [38]. Researchers have drawn distinctions between co-located and distributed settings, identified design dimensions such as synchronous vs asynchronous interaction [21]. Common ground has emerged as an important design consideration for collaborative visual analytics [17], with researchers identifying the crucial role of deixis [19] (the ability to reference specific elements of a shared representation) in supporting common ground. More recently, researchers have explored visual affordances for common ground and cursor-based deixis in synchronous collaboration [16, 31].

While this body of work richly characterizes collaborative visualization among sighted users, it largely assumes a shared visual modality. Our work extends these conversations by examining how common ground, deixis, and collaborative sensemaking function when collaborators work across different sensory modalities. Specifically, we examine the case of blind and sighted analysts collaborating using visual and tactile representations of the same data.

2.2 Multi-Modal Data Representations

Accessibility researchers have developed a range of systems that represent data through non-visual modalities, including structured textual

description [4, 13, 28, 42], sonification [26], and tactile graphics [7]. Multi-modal systems such as Chart Reader [36] and MAIDR [35] combine multiple modalities into a single reading interface. Across this body of work, however, multi-modal representations are predominantly designed and evaluated in terms of individual comprehension: whether a single BLV reader can independently extract information equivalent to that of a visualization.

Among these systems, Umwelt [43] is a particularly relevant example in that it explicitly names common ground between mixed-ability collaborators as a design goal. Umwelt links visual, textual, and sonified representations through synchronized state, such that a screen reader user’s navigation is reflected as visual highlighting—making the focus of their attention legible to sighted collaborators. Participants in Umwelt’s evaluation envisioned using these affordances to present to mixed audiences and contribute to data-driven discussions with sighted colleagues. These collaborative uses, however, were speculative.

Our work complements this design-led approach with an empirical account of how common ground is accomplished in a working mixed-ability team. Our observations affirm some of Umwelt’s design intuitions—in particular, the value of continuous, passively produced evidence of a collaborator’s attention. But they also surface communicative functions that synchronized state alone does not capture, including metalingual work to verify understanding across modalities and phatic signaling of ongoing engagement, as well as an asymmetry in which information supporting common grounding flows more readily from blind to sighted collaborators (Sec. 4.1).

2.3 Collaboration and Interdependence in Accessibility and Disability

Accessibility research in HCI has increasingly moved beyond individual-focused accommodation to examine how access is negotiated in social and collaborative contexts. A recurring finding is that access in mixed-ability settings depends on ongoing, often-invisible coordination work. Branham and Kane, studying blind employees in sighted workplaces, found that blind and sighted coworkers actively negotiate shared practices, and blind workers perform substantial behind-the-scenes labor to make collaboration run smoothly—labor that went largely unrecognized by their sighted colleagues [5]. Subsequent work has examined the specific strategies mixed-visual-ability teams use to share information, such as managing the transformations a document undergoes as it moves between visual and non-visual representations [39], and the collaboration tools teams adopt for videoconferencing and project management [1]. Parallel findings on Deaf-hearing teams similarly show collaboration is shaped by the continuous work of bridging modalities [37]. Across this literature, a consistent theme is that access is produced through interaction between collaborators. This empirical pattern is given theoretical grounding in disability studies. Where a medical model locates disability in an individual’s body, Kafer advances a political/relational model in which disability is produced through social interactions and structural arrangements [22]. Two concepts extend this relational view in ways directly relevant to collaborative analysis. *Interdependence* [3] is the idea that access is produced jointly by the people who share a social setting. *Access intimacy* [30] describes the felt sense that others understand and anticipate one’s access needs.

Our work brings these relational concepts to bear on collaborative data analysis. Concretely, we use interdependence and access intimacy to interpret collaboration in a mixed-ability lab (Sec. 4.3). Our contextual inquiry argues that the team’s data analysis depends on jointly produced access and that the lab has restructured its labor and roles to sustain that interdependence over time. In doing so we connect a body of work in disability studies to the everyday processes by which blind and sighted analysts establish common ground.

2.4 Multi-modal Participation in Linguistic Anthropology

Establishing common ground is fundamentally a communicative problem. Linguistic anthropology—which studies how people communicate by combining speech, gesture, touch, and other modalities—offers us two frameworks for understanding multi-modal communication that

operate at complementary levels of analysis. Jakobson’s functions of language [34] is a taxonomy of the distinct kinds of work an utterance can do—referring to things, checking mutual understanding, maintaining a conversation, and more. Crucially, Jakobson defines these functions without assuming a specific modality, which lets us ask how each is accomplished across speech, touch, and vision. Keating and Sunakawa’s participation cues [25], developed to explain how people coordinate joint action when their interaction spans physical co-presence and a shared digital world, gives us a complementary interaction-level vocabulary for how collaborators hand off control and direct one another’s attention moment to moment. In our analysis, Jakobson identifies what mixed-ability grounding requires, and participation cues describe how a blind and sighted analyst accomplish it in practice.

3 METHODS

In this section, we introduce the goals of our contextual inquiry method, explain how we engaged our participants, and discuss why our field site is particularly suited for studying collaborative data analysis on mixed-ability teams.

3.1 Method: Contextual Inquiry

To understand how teams that include both blind and sighted people collaborate on data analysis, we conducted a contextual inquiry with an oceanography lab led by a blind principal investigator. Contextual inquiry is a field research method in HCI where researchers observe and interview participants in their actual work environment as they perform real tasks [24]. By balancing observation and interviewing, we aimed both to gather observational data that reflected realistic, context-specific work practices, and also to collaboratively interpret observed actions in follow-up conversations between the researchers and participants. Our approach is based on ethnographic traditions, using observations and qualitative data as the basis for developing grounded theories [6] about social and cultural practices. This means that our qualitative results are not claiming to be representative of a hypothetical average mixed-ability team of analysts; rather, observations are meant to provide an initial framework that can juxtapose and relate to practices in other contexts [11].

Research ethics. In our research, we followed standard university ethics review processes and implemented proposed best practices around citational justice in accessibility research. This research was reviewed by MIT COUHES, the university’s institutional review board, and received an Exempt Determination (#6371). We were approved to request consent from participants to reference them by name. Although qualitative research norms favor hiding identities with pseudonyms, accessibility researchers practicing citational justice have argued that doing so without intentionality can erase the contributions of disabled people from research [41]. Because our project involves deep engagement with a small group of people with unique accessibility expertise, we felt it was important to publicly credit them. We asked all participants for consent and preferences about how they would like to be referenced, and have followed those preferences in this paper.

3.2 Research Setting: Bower Lab

We chose Bower Lab as the field site of our contextual inquiry because they are an established mixed-ability team conducting scientific data analysis using multi-modal data representations, exemplifying the type of collaboration we seek to understand. The lab is led by Dr. Amy Bower, a Senior Scientist in Physical Oceanography. Lab members include Access Assistant Anna Pinckney, Postdoctoral Investigator Dr. Ali Exley, and Senior Research Specialist Dr. Heather Furey. Bower is blind with low partial vision, which means that she has some light perception but cannot read computer screens, relying on assistive technology including screen readers. She uses a variety of strategies to collaborate with sighted colleagues in her lab. For example, the lab has developed expertise in creating 3D models and tactile graphics to support their oceanography research, and in using sonification to communicate their work to the public.

Beyond technical approaches, the lab has redistributed access labor at the organizational level by creating a dedicated Access Assistant role that now serves as essential collaborative infrastructure. This role, which Bower advocated for her institution to establish, is a half-time research assistant and half-time accessibility assistant. That means that responsibilities include both contributing intellectually to research projects and also developing accessibility expertise by, for example, creating tactile representations of figures and other artifacts of the research process. The lab’s first and current Access Assistant, Pinckney, combines accessibility expertise with domain-specific scientific knowledge to design and produce accessible data representations such as tactile charts. Bower credits the idea for this type of role to Dr. Mona Minkara, a blind Assistant Professor of Bioengineering at Northeastern University. Given that the labor of accessibility is often treated as “invisible work” [5] that disproportionately falls on disabled people, the existence of this type of role is notable because it reflects a rare redistribution of accessibility-related labor.

The lab’s extensive public engagement around accessibility also makes them well-suited as participants in a contextual inquiry. Bower has co-authored numerous research papers on accessibility in science and founded OceanInsight, an outreach program through which the lab shares practical guides and resources for creating accessible data representations—for example, a how-to guide on creating tactile graphics with a Picture-in-a-Flash (PIAF) machine [2]. OceanInsight exemplifies *crip technoscience* [15], or “disabled knowing and making,” because it is created by people who develop expertise through navigating inaccessibility and then share that knowledge within disabled communities. Because the Bower Lab actively reflect on and share their approach to access with others, this also enhances their participation in contextual inquiry because they are able to reflexively articulate their own practices as part of a collaborative interpretation process.

Because Bower Lab represents an unusually experienced practice, our findings characterize what mixed-ability collaborative analysis can look like under favorable conditions. We treat the lab as an exemplary case whose practices surface communicative functions and design opportunities that may be less apparent in other settings.

3.3 Data Analysis and Interpretation

We conducted our contextual inquiry with the Bower Lab at Woods Hole Oceanographic Institution. Author Zong spent 2 full days on-site at the lab’s offices in Woods Hole, Massachusetts. During this site visit, Zong participated in multiple routine working meetings on various projects in the lab, during which he observed, took notes, and asked occasional questions for participants to explain their actions and thought processes. Zong also met with every lab member for individual semi-structured interviews to discuss observations and understand collaborative processes in the lab. In total, the site visit included 16 hours of in person contact with the lab, including 9 scheduled meetings.

Author Zong conducted all observations and interviews during the site visit. Data took the form of field notes, including written observations and sketches of artifacts such as tactile charts and the lab’s desk configurations. Notes were taken in-situ during meetings and interviews, then reviewed and expanded the same day during breaks between sessions.

Zong analyzed the field notes through an iterative process of thematic interpretation, identifying recurring patterns in how lab members coordinated attention and understanding across modalities. The theoretical frameworks described in Sec. 2.4 were brought into the analysis after these initial patterns emerged: for example, repeated observations of non-verbal confirmation and engagement signaling motivated the turn to Jakobson’s Functions of Language as an interpretive lens.

Consistent with contextual inquiry’s emphasis on collaborative interpretation [24], participants were involved in interpreting observations at multiple points: clarifying questions during observation sessions, individual semi-structured interviews in which Zong discussed emerging observations with each lab member, and follow-up conversations after the site visit. Additionally, all lab members reviewed a complete draft of our findings prior to submission and confirmed the accuracy of our accounts of their practices.

Because data collection and analysis were conducted by a single researcher, we note our approach to validity. In the interpretivist tradition this study draws on, validity comes from being transparent about our process and showing how our interpretations are grounded in fieldwork, rather than from measures like inter-coder agreement [11, 27]. We accordingly rely on multiple data sources and on participant involvement in interpretation. We also note Zong’s position relative to the field site: Zong has a research background in accessible data visualization and an existing relationship with the lab. Bower participated in a user study for a previous system [43] and has provided feedback on prior research. This familiarity aided rapport and access during the site visit, while also shaping the inquiry’s orientation toward data representations as objects of analysis.

4 CONTEXTUAL INQUIRY RESULTS: ESTABLISHING COMMON GROUND IN BLIND/SIGHTED COLLABORATION

In this section, we present qualitative findings from our contextual inquiry. In Sec. 4.1, we focus on the mixed-ability team’s use of tactile data representations, arguing that multi-modal representations provide shared affordances supporting fundamental linguistic functions required for common grounding. In Sec. 4.2, we turn to observations around co-operative computer use scenarios, where Bower and sighted collaborators managed control flow and joint attention while taking turns navigating a shared digital workspace. Finally, in Sec. 4.3, we discuss how Bower Lab instantiates disability scholars’ concepts of interdependence and access intimacy in practice.

4.1 Tactile Map: Shared Affordances for the Functions of Language

In the majority of our observation sessions, Bower and collaborators used tactile representations to aid their discussions. These included a 3D printed tactile map of a region of the ocean floor that much of their work focuses on, as well as embossed tactile charts of figure prototypes from draft research publications. In this section, we first describe how tactile representations provide affordances for communication functions that are typically assumed to be verbal. Then, we explain how these affordances play a crucial role in multi-modal common grounding.

4.1.1 Tactile representations provide non-verbal affordances for the Functions of Language

In our observations, tactile representations enabled communicative functions that are often assumed to be primarily verbal. We observed evidence for three of Jakobson’s functions of language that we expect to be most relevant to data analysis: the *referential*, *metalingual*, and *phatic* functions.

Referential function. The referential function concerns how a speaker points to or describes things in the *extra-linguistic context* [12]—the objects, places, and ideas outside language that the conversation is about. This is the function visualization researchers already know best, even if not by this name. When two analysts lean toward a scatterplot and one says “this outlier in the top right,” the phrase “this outlier” is doing referential work by picking out a specific mark in the shared visual context. Using pointing words like “this,” “that,” “here” (also known as deixis) helps to populate a collaborative analytical context with shared references.

Bower and collaborators made extensive referential use of tactile representations. When I asked Bower about their upcoming expedition to the Grand Banks area of the North Atlantic, she read the 3D map with her hands and pointed out where the team would go. The map provided context for deictic gestures and statements: pointing to “this ridge,” or saying that the expedition would travel “over here.” It also provided a common frame for verbal references. For example, in discussions with Exley, Bower used relative statements like “west of Reykjanes Ridge,” or “near Charlie Gibbs” (referring to the Charlie Gibbs Fracture Zone). Shared landmarks on the map, together with the team’s oceanography expertise, enabled them to describe relative locations in space.

Spatial reference and deixis are already known to be fundamental to collaborative analytics, so their presence here is not necessarily surprising. However, the referential function rarely operates alone. As

we will discuss, the combination of reference with the less-discussed metalingual and phatic functions provide important building blocks for common ground. For visualization researchers, this suggests that a shared representation does not only enable data reference, but also supports verifying and sustaining a conversation.

Metalingual function. The metalingual function is for checking that people in a conversation actually understand each other—aligning on the meaning of what is communicated. In a visualization setting, metalingual work is what happens when references are ambiguous. Imagine one analyst says “the spike in Q3” and the other replies “the tall bar, third from the left?” That reply adds no new information about the data; instead, it confirms that both people are looking at the same mark. Asking “wait, do you mean the orange line or the red one?” is also metalingual. These questions re-align a shared understanding of the representation. The function becomes especially important when collaborators are looking at different representations of the same data, because then verification also requires translating between the representations—confirming that what one person reads in a chart matches what the other reads in a table, a description, or a tactile graphic.

We observed Bower and her collaborators use tactile representations metalingually. In one session, Exley was describing a figure from a draft paper to get Bower’s feedback. Reading from her laptop, Exley described a chart of the Iceland Basin with latitude and longitude axes and arrow annotations representing flow through various landmarks. Bower followed along on the 3D map. She interjected with clarifying questions—“Over here?”—and Exley visually confirmed where Bower was pointing: “Exactly.” Bower then asked, “Do you have an arrow here going toward Greenland?”, and Exley described the corresponding annotations on the visual chart.

This exchange shows clarification and verification, which are typical of metalingual work; however, it also highlights a challenge specific to the mixed-ability setting. Because the two analysts were reading different representations, this required confirming a shared message specifically by reconciling two representations of the same data, functioning like translation between modalities. We also observed that the tactile graphic enabled Bower to initiate metalingual exchanges that would not have been possible without a tactile representation. In one instance, Bower noticed a feature on the map she was unsure about, saying, “I was looking at this feature—I don’t know if it’s real. There’s an interesting channel, I wonder if it’s for water.” This prompted Exley to check the visualization, which opened up a new discussion. The tactile representation enabled her to raise the question, and the cross-check with Exley gave them a way to resolve it. If solely relying on verbal description, Bower would have had no comparable reference against which to notice something was amiss, needing to rely on Exley’s description rather than her own reading.

Phatic function. The phatic function is about maintaining the channel of communication, making sure that communication is still ongoing (i.e. a conversation has not ended and everyone is still paying attention). Where the metalingual function asks “do you understand the same thing I do?”, the phatic function asks the more basic question, “are you still with me?” In a visualization context, phatic signals are the easily-overlooked behaviors that let an analyst keep talking without stopping to check whether anyone is receiving. For instance, this can include a collaborator’s “mm-hm” while you walk through a dashboard or nodding along as you trace a line. None of this conveys information about the data, but it reassures the speaker that their partner is listening and that it is safe to continue.

Tactile reading supports this function too, but in an unexpected way. During data discussions, Bower kept both hands in near-constant contact with the map, running them across the model and exploring it even while talking. Even when she was not looking for anything specific, she would feel around familiar landmarks to hold her place in the conversation. When I later asked about this, she said the repeated contact helps her “stay familiar” with the map and that she is “following along” while others talk. Bower’s sustained contact with the map is a phatic signal, because the hands’ presence on the map conveys that Bower is still engaged in the conversation (as opposed to distracted or

twiddling her thumbs).

The phatic signal here is interesting because it is an ambient action that is legible to collaborators without needing to directly address them or interrupt the flow of discussion. As we discuss in [Sec. 4.1.2](#), this property of conveying information to a collaborator without a dedicated communicative act is important to what makes the tactile map effective as communication infrastructure, and it is something that future design work can take inspiration from.

4.1.2 Multi-modality creates new sources of evidence during common grounding, but asymmetrically

Adding a tactile representation for the blind analyst makes establishing common ground easier by introducing a visual way for the sighted analyst to understand the blind analyst’s tactile reading process.

In Clark and Brennan’s model [9], *grounding* is the process by which participants accumulate *evidence* that they have understood one another. In this paper, we distinguish *active evidence*, produced through a dedicated communicative act (e.g., a verbal confirmation), from *passive evidence*, produced as an incidental byproduct of a collaborator’s analysis work (e.g., hand position while reading a tactile map). We use “passive” to characterize communicative intent rather than effort: a collaborator may be highly active in their own task while passively emitting grounding evidence.

According to Clark and Brennan, grounding varies by medium, and different media impose different costs—for example, the effort to formulate a message, or to receive and understand one, differs across speech and writing. One might expect that combining two modalities, tactile and visual, would raise these costs, since the analysts cannot simply point at the same display. Instead, we observe that adding the tactile representation also introduces visual grounding evidence for the sighted analyst. In [Fig. 1](#), we show a diagram that places the blind and sighted analyst on either side, with the tactile chart and the visualization respectively. Each arrow represents communicative actions labeled by sensory modality. Consider this concrete example of a flow depicted by the diagram. Exley describes a feature on the visualization—for example, “there are particles over there.” On the diagram, this is the *description* arrow from sighted analyst to blind analyst, traveling verbally. Bower’s hands then move across the tactile map to the corresponding place. For Bower, this is the *touch: read* arrow on the diagram—she is locating the feature in her own representation. But the same motion is visible to Exley as the *observe hand position* arrow, a visual trace of Bower’s tactile reading.

In the example we walked through, Exley could have confirmed an aligned understanding with verbal metalingual questions, but Bower’s hands already supplied evidence of understanding with no verbal exchange needed. Rather than adding overhead, the tactile map provided a continuous source of passive evidence about Bower’s attention. This parallels how shared cursors in synchronous collaborative visualization passively convey a remote user’s location as a byproduct of their interaction [31], providing ambient awareness without disrupting either collaborator’s independent activity. What is distinctive in the mixed-ability setting is that this passive evidence crosses modalities, because Bower’s tactile reading produces visual evidence legible to Exley.

However, one key observation is that the benefit of the map’s combined visual/tactile affordances benefit the two analysts unequally. The tactile map introduces passive visual evidence for the sighted analyst, but there is no corresponding passive tactile evidence for the blind analyst. One could imagine that the sighted analyst moves the blind analyst’s hand on the tactile surface to indicate visual attention. However, this feels higher-cost (more active), requiring the sighted analyst to perform a dedicated action rather than being a byproduct of their own reading. Social norms around touch may make this strategy less commonly used—we did not observe anyone doing this in our setting.

Nonetheless, we observed that sighted analysts also benefit from pointing on a visual representation as a way to follow along with their collaborator and provide evidence of attention. In an exchange with Bower, Furey described a figure from a paper draft, then Bower took over and described her own reading of the data. As she spoke, Furey used a pen to point at a printout, tracking her own understanding. I later




Figure 1: Communication channels between a blind analyst and sighted analyst collaborating across a tactile chart and visualization. The solid line crossing the middle shows non-verbal visual grounding evidence made available by the tactile chart. The dashed line shows a hypothetical design opportunity: non-verbal tactile evidence produced by a sighted analyst’s interaction with a visualization.

confirmed with Bower that despite some light perception, she could not see Furey’s pointing. Furey’s gesture mirrored the work a blind analyst does to align a verbal description with a tactile chart. But because it was invisible to Bower, it did no grounding work. Furey would still need an additional verbal follow-up to convey her understanding to Bower. This could suggest potential future design directions around making a collaborator’s visual pointing tactually accessible in a continuous, passive way. We discuss this idea further as a design opportunity in [Sec. 5.2](#).

4.1.3 Aligning understanding across modalities shifts discussion toward analysis and design

As analysts successfully align their understanding across modalities, the focus on their conversation shifts from describing representations to discussing the underlying data and science. In the exchanges above, much of the early work went into matching what one analyst saw against what the other felt. As common ground accumulated, that work receded and the analysts were freed to focus on more substantive discussions.

We observed verbal exchanges moved from the appearance of figures to oceanographic phenomena. In one discussion, Bower referred to a chart she and Exley had both encountered (as an embossed tactile graphic for Bower and as a visual figure for Exley). Exley pulled it up to check a depth value, and Bower noted that “all the particles go south,” reasoning about flow behavior rather than describing an image. In another instance, Bower asked a question about particle behavior “at 1800” meters of depth, to which Furey responded “the plot only goes to 1000.” Here, Bower was reasoning about the oceanography beyond the figure, treating it as a launching point for analysis rather than the object of the discussion.

Bower also took part in design decisions about the figures themselves. When Exley pitched a new figure, Bower judged it “would be a nice summary,” reasoning about the rhetorical work it would do. She also voiced preferences about encoding conventions, like “red should be anticyclonic.”

These observations have implications for how the visualization community defines success for accessible design. Once the blind analyst can successfully read the chart, this opens up opportunities for the analyst can fully participate as an analytical and design partner—reasoning about phenomena the figure doesn’t show, and helping decide how findings are presented. In other words, the right success metric for an accessible representation is not equivalent information access in isolation, but whether it lets a blind analyst participate in the parts of analysis that happen after and around reading. This reframes accessibility from a property of a single artifact to a property of a collaboration.

4.1.4 Charts commonly used by oceanographers pose challenges for tactile design

The Bower Lab has deep experience producing tactile representations, but some charts common in oceanography resist tactile translation entirely. Their workflow is well-developed—for 3D maps, they use MATLAB to generate STL files for printing and keep track of what smoothing parameter produces the best result. For 2D charts, Pinckney generates plots in MATLAB and then exports them to Illustrator to adjust lines manually. While this may sound tedious, it aligns with existing workflows used by professional tactile designers [7]. According to Pinckney, she usually spends about 15 minutes adjusting each chart in Illustrator before embossing.

Some charts, though, have features that are, in Exley’s words, “kind of impossible to make accessible.” Pinckney and Exley showed me one such example, a bathymetric map, which is a topographic map representing underwater depth. The map had latitude and longitude, contour lines for depth, and an additional layer with a color gradient representing data values. The tactile version needed the contours so Bower could orient herself, but it also needed to overlay the data through some tactile analogue of the color gradient. Pinckney manually traced the contours and turned the gradients into discrete textures by binning the data. The result was that the data textures covered up the contours. This made it hard for Bower to wayfind, because she could not tactually separate the data texture from the contour line.

This difficulty points to a more general limitation that layering does not translate well from visual to tactile. A sighted reader can visually separate overlaid layers because color and line weight are contrasting enough that the reader can attend to one layer at a time, so the contour lines and the color gradient remain perceptually distinct even where they overlap. Touch has no equivalent mechanism for this kind of selective attention. Where two tactile features coincide, the hand encounters a single combined texture and cannot easily pull them apart, so a tactile chart that overlays data textures with contour lines ends up feeling perceptually indistinct.

Unfortunately, visualization practice leans heavily on layering, which suggests that designers must attend to modality-specific differences when designing complex tactile charts. Simply directly replicating the visual appearance of a chart in tactile form will often not be enough. This is why the lab is excited about spatial audiotactile interfaces, which separate the data layer from the wayfinding layer by moving data into the auditory channel and reserving the tactile surface for spatial structure. More broadly, it suggests that porting a visualization to touch is less about faithful reproduction than about deciding what information is well-suited to tactile representation or not.

4.2 Co-op Computer Use: Participation Cues for Screen Reader Users

When sessions involved shared documents, such as draft research papers containing figures discussed in [Sec. 4.1](#), Bower and collaborators engaged in co-operative computer use. They took turns using the input devices and used the cursor to direct both collaborators’ joint attention. Bower’s desk was specifically laid out to support co-operative use. In [Fig. 2](#), we illustrate Bower’s computer setup. Bower uses a Windows laptop running the JAWS screen reader, which she uses with keyboard input. The laptop is connected to a monitor that is facing across her desk, toward where a collaborator would sit, allowing sighted lab members to see what she is doing. There is also a mouse connected to the laptop and placed with the monitor, so sighted collaborators can also provide input when needed. This unique layout enables co-operative use by providing multiple displays (i.e., visual and speech outputs) and input devices (i.e., mouse and keyboard) according to each collaborator’s abilities. Co-operative use in a mixed-ability collaboration also introduces unique challenges around explicitly managing control and initiative to avoid input conflicts. Bower said the setup was originally motivated by inaccessible institutional software (e.g., required finance software), which she copilots with a sighted person who navigates visually and verbally passes control for steps like entering a password. However, we observed that the setup is fruitful beyond this




Figure 2: Diagram of Bower’s desk, set up for co-op computer use with sighted collaborators. A) Bower’s chair. B) Collaborator’s chair. C) Laptop and keyboard facing Bower. D) Monitor and mouse facing Collaborator. E) 3D printed tactile map. *Disclosure: Diagram created with Gemini using hand-drawn sketch from field notes as input, labels manually added.*

initial intention, providing essential collaborative infrastructure in the mixed-ability setting.

In this section, we first describe how mixed-ability collaborators use participation cues to signal when they’d like to transfer control of the computer between them. Then, we discuss the role of the cursor, shared between the mouse and screen reader, in directing joint attention and resolving deictic references. Finally, we discuss how the screen reader plays a role in collaboration akin to a conversational participant.

4.2.1 Collaborators use participation cues to transfer control

When Bower and her collaborators reviewed documents together on the computer, they traded off who was operating the computer situationally, based both on who was presenting and on accessibility considerations. As a result, it was important for them to explicitly designate who was currently operating the computer, since a lack of clarity would cause clashing inputs. Keating and Sunakawa’s *participation cues* describe how people “orient to, plan, and execute collaborative actions” that span physical and digital space [25]. In their original study, co-present participants could draw on physical cues like body language to coordinate digital actions. In our setting, blind and sighted analysts perceive different views of a shared digital document (via spoken narration and visual display respectively), and lack symmetric access to physical, visual cues. As a result, they rely even further on verbal participation cues—such as verbal narration of what they’re seeing and doing on screen, and using deixis to point into shared digital space.

An observation session with Bower and Furey provided an illustrative example of how participation cues made collaboration possible. Working through Excel figures for a paper, Bower walked through the document while narrating her screen-reader actions so Furey could follow. Furey requested control of Bower’s computer by saying, “Can I use your mouse?” She then took control of the cursor and began describing what she was doing. She navigated to a directory on the computer and said, “I’m going to open this file.” Then, she returned control to Bower, saying, “You’ve got the steering wheel.” Later in the discussion, Furey once again requested control, saying “Can I drive?” Bower acknowledged and accepted this request by replying “I relinquish control.” When Bower decided to pass control over to Furey, she used keyboard hotkeys to turn off JAWS so that the screen reader doesn’t interfere with mouse input.

The use of participation cues here is significant due to unique features of the mixed-ability setting. First, managing control is entirely verbal and explicit. Phrases like “Can I drive?” and “I relinquish control” are requests and acknowledgments that also establish a shared understanding of the current state of the digital workspace. Fully sighted collaborators would simply reach for the mouse, but here verbal cues

do the work that might otherwise be done by gestures and body language. Second, Bower and Furey use metaphors like “steering wheel” that frame the computer as a vehicle with a single driver, reinforcing a turn-taking model with unambiguous handoffs. Future research could explore design opportunities for structured, perceivable signals of who holds control, potentially drawing on something like the steering wheel metaphor.

4.2.2 A shared mouse/screen reader cursor helps direct attention and resolve deictics

As collaborators pass control back and forth, the shared cursor helps maintain common ground. We observed that the cursor functions like pointing (another example of the referential function), serving both as a proxy for a collaborator’s attention and as a way to direct attention to a location in the workspace. Previous work has argued that cursors facilitate joint attention between a user and a computer [40]. In this mixed-ability setting, the cursor also facilitates joint attention between collaborators, because each participant can both control it (with mouse or keyboard) and read its state through whichever modality is more useful to them (monitor or screen reader).

As an example of how the collaborators used the cursor to manage reference and joint attention, consider this exchange. While reviewing a document with Furey, Bower wanted to return to her email. Furey placed the cursor on Bower’s inbox with the mouse and said “go ahead” to hand back control. Bower began reading an email by screen reader, then paused: “This is worrisome. This next sentence.” She moved the cursor to the spot with the screen reader, and Furey, watching the monitor, could resolve the deixis and confirm where Bower was reading—aligning their understanding and grounding the reference.

What makes this exchange work is that the cursor is a single referent that both analysts can address and perceive, each through their own modality. Bower positions it with the screen reader, Furey reads it by sight, and is able to understand the deictic reference in “this next sentence” without anyone describing the location. This is the same deixis phenomenon that collaborative visualization has long studied for shared cursors. However, in the multi-modal setting, the two collaborators are sharing a cursor that renders in two different modalities (the monitor and the screen reader narration). The design implication is that a shared cursor that can be both moved and read through multiple modalities can support deixis multi-modally.

Beyond pointing in the moment, collaborators also built durable structure into the document. Reviewing a Word file with Exley, Bower used comments she had added, jumping between them with the screen reader menu and talking through each. Each jump also showed Exley the surrounding text on the monitor. Bower sometimes added a comment purely as a “bookmark” to return to later. For Bower, these bookmarks created a navigation structure—named anchor points the screen reader can jump between, providing quick traversal over a document that is otherwise linear and slow to scan non-visually. For Exley, the same jumping externalizes Bower’s reading path: each landing scrolls the monitor to where Bower is, so Exley passively sees the route Bower takes through the document. The design implication is that non-visual navigation aids—comment-jumping, heading navigation, landmarks—double as ambient awareness channels for sighted collaborators when they are surfaced on a shared display.

4.2.3 Screen reader narration is a shared resource that sometimes helps and sometimes disrupts

When the screen reader is speaking out loud, it takes up the verbal channel—if one of the analysts wants to speak, they either need to pause the screen reader or wait for it to finish. As a result, we observed how Bower and her collaborators needed to manage the screen reader’s narration in conjunction with their discussion.

Sometimes, screen reader narration was a helpful additional source of grounding evidence. For instance, Bower often took meeting notes in a plaintext file, narrating as she typed. The screen reader would then read back what she wrote, helping her confirm that she typed it correctly. The sighted collaborator could also confirm visually on the

monitor. In this case, the screen reader provided grounding evidence for both collaborators.

Other times, collaborators needed to actively manage the screen reader to ensure it was not disruptive. In a meeting about the lab’s OceanInsight website, the group needed to look something up. Bower noted that having JAWS read it aloud would make it hard for others to listen, so they had someone read it to her instead. On another occasion, she and Furey both needed to read a bulleted list in an email, and chose to have Furey read aloud, since Bower noted that “with JAWS I have trouble [with bulleted lists] sometimes.” But Furey stumbled over words, possibly because of the small on-screen font, and they switched: “Let’s have JAWS read it.” Because JAWS was now reading for both of them, Bower slowed its speech rate for Furey—an accessibility accommodation from the blind collaborator to the sighted collaborator—then sped it back up afterward while Furey followed on the monitor.

The takeaway is that assistive output can sometimes compete with human conversation for space in the auditory perceptual channel. Accessible visualization sometimes treats adding a modality (sonification, narration) as strictly additive—more channels, more access. Instead, our observations reflect the fact that the auditory channel can hold either screen-reader speech or collaborative talk, not both at once, so adding narration can crowd out conversation. The design consequence is that auditory representations intended for collaborative use should consider how their output shares a channel with human speech, and ideally offer ways to background or redirect that output.

4.3 Interdependence and Access Intimacy in Practice

Disability studies and assistive technology scholars have articulated *interdependence* as a frame for accessibility [3, 14]. Interdependence rejects the idea of the self-sufficient individual and instead recognizes that all people rely on others. Therefore, understanding access means attending to social relationships rather than only focusing on individual users. Bower Lab’s practices exemplify interdependence as a foundational way of working. Their processes exemplify *access intimacy*, which disability activist and scholar Mia Mingus defines as “that elusive, hard to describe feeling when someone else ‘gets’ your access needs” [29]. For Mingus, “access intimacy is interdependence in action” because it centers the sustained relationships that, over time, produce accessible spaces with an equitable distribution of labor. Mingus further articulates *liberatory access* in which “demands that the responsibility for access shifts from being an individual responsibility to a collective responsibility” [30]. Together these frames help explain something we observed repeatedly: accessible data analysis in the lab depends not only on tools and representations, but on how expertise, labor, and workflow are arranged.

In this section, we discuss three aspects of how interdependence and access intimacy operate in the Bower Lab. First, we discuss how institutionalizing access labor through a dedicated support role made multi-modal data analysis processes more possible. Second, we describe how access knowledge circulates as shared infrastructure rather than residing in one person. Finally, we reflect on how practicing interdependence requires challenging the conventional goals of accessible data visualization research.

4.3.1 Institutionalizing access labor shifts agency in collaborative analysis

Formalizing access work through a dedicated Access Assistant role redistributes labor so that experienced researchers can spend more time on research rather than managing access. Furey shared that the “Access Assistant has been a game changer” for her because she wants to focus on science, and does not consider herself an access expert. Furey has over 20 years of oceanography experience, and wants to lean into that strength, which is why “everyone is happier with Anna around.”

As we discussed in Sec. 4.1, tactile charts and maps are central to the lab’s workflow. Institutionalizing the labor required to create those charts and maps makes that workflow possible. The Access Assistant role is skilled, specialized work that is defined by the lab’s data analysis needs. Pinckney fits it well precisely because she combines scientific knowledge with tactile design skill. Before the lab had someone who

could make tactile charts, Furey spent considerable time describing figures aloud. This was worse for Bower because she had to just trust the descriptions, which previous work has also surfaced as a problem for BLV user agency [42]. Furey said that “tactiles brought [Bower] back into being able to engage in the data more fully.”

Increasing agency for blind analysts is a socio-technical problem, not just a technical one. For Bower Lab, the shift from verbal description to tactile representation was made possible by an institutional decision to fund the labor of producing tactile representations. Accessible visualization research increasingly recognizes user agency as an important design goal, and Bower Lab’s experience supports the idea that a representation is most successful when it lets the reader make their own interpretations rather than relying on a collaborator’s. Our observations further point toward relevant factors outside of the design of a representation that contribute to interpretive agency. This matters because, without institutionalized support, disabled academics often face assumptions that frame access as individual preferences instead of structural needs. Even though “disabled workers possess specific knowledge of their workplace and its barriers,” many academics experience pushback in the form of questions like “Are you sure you really need that?” when asking for institutional accommodations—especially when the asks are perceived as nonstandard or expensive [33]. Bower Lab’s success suggests that comparable structural interventions adapted to other contexts could support other mixed-ability teams.

4.3.2 Access knowledge is collectively learned and shared as infrastructure for data work

Bower Lab members develop and share practical access knowledge through collaboratively solving everyday problems. This knowledge then circulates interpersonally in the lab such that expertise is held communally rather than by one person.

This collective approach is visible in how the lab shares paper drafts with Bower. Members typically write in LaTeX on Overleaf but need to share drafts as Word documents for accessibility, and converting between the two reliably is a recurring problem. People have tried different methods. Bower notes that JAWS has OCR, but line numbers disrupt reading. In one meeting, Exley described using an online Adobe conversion tool and then cleaning up the result by hand. Bower asked Exley to show that process to Pinckney—an instance of best practices circulating between members. In a later conversation, Exley said: “we help each other” learn about access.

Document conversion may seem peripheral to data analysis, but for a research lab, drafting publications is the culmination of analysis. Papers are where visualization designs are decided, and sensemaking models likewise treat presentation as a constitutive stage of analysis [32]. The example illustrates a dynamic that applies across the lab’s whole workflow. Each stage of data analysis, from exploring data to sharing and communicating results, involves access challenges that benefit from collective problem-solving. This expectation around knowledge sharing and mutual learning is an example of *liberatory access* [30] because the responsibility shifts from the individual to the collective.

It is tempting to think of accessibility as a property of an artifact (i.e. a chart is accessible or it is not); instead, what we instead observed is that the lab’s most valuable accessibility resource beyond any individual data representation is a shared, evolving body of know-how. This includes knowing how to convert a document, which smoothing parameter renders a map well, or how to bin a gradient into textures. Because that knowledge is held collectively, it survives turnover and accumulates as the lab’s institutional memory, rather than vanishing when one person leaves. This suggests that supporting accessible collaboration is not only a matter of producing better artifacts or tools, but of supporting the processes through which a team builds and retains access expertise over time. Current visualization accessibility research, which primarily focuses on individual artifacts and single-session use, largely does not address these socio-technical factors of long-term collaboration.

4.3.3 Emphasizing collaborative sensemaking suggests design goals for analysis tools

For Bower, accessibility tools for independent analysis are usually not the most efficient approach. Instead, she thinks of collaborative sensemaking as her most efficient workflow—looking at tactile charts together with Pinckney and asking questions. This framing challenges a common assumption in accessible visualization research that the goal is to give individuals tools for independent data exploration.

Relying on collaborative sensemaking introduces slower processes, but in the Bower Lab, that slowness can also be productive. It can certainly be frustrating—Exley finds that it’s harder to get quick feedback and rapidly prototype figures, and Bower notes that “charts come in too fast to get tactiles of all of them.” Yet, Pinckney says that access “productively slows things down,” because she needs to spend more time with each chart. Similarly, Bower takes time to read tactile charts and understand them in order to explain them back to collaborators. The additional preparation and synthesis adds time, but deepens engagement with the data.

Embracing this model meant rejecting academia’s culture of independence. As Bower puts it, “I had to turn my back on that. That doesn’t work for me, this does,” referring to sustained collaborative and interdependent relationships over time. For instance, Bower credits her decades-long collaboration with Furey for many of her accomplishments. As disabled researcher Cindy Bennett reflects, she has observed blind people “[doing] access work creatively and intimately, in that access work and friendship were indistinguishable” [18]. If we think of access as understanding and responding to needs, long-term relationships involving collective responsibility for access are crucial for successful collaborations.

Taking interdependence seriously as accessible visualization researchers suggests that tools and workflows should be designed not only for individual exploration, but to support long-term collaborative analytical processes that evolve over time like the Bower Lab’s. For Bower, independent exploration is the inefficient workflow, and understanding emerges through slower, collaborative engagement with data. Designing for this type of work could mean designing for collaborations that evolve over time.

5 DISCUSSION

In this section, we first synthesize our findings from the contextual inquiry into a set of key insights for visualization accessibility researchers. Then, we build on these insights to present directions for future work designing multi-modal affordances for common ground. Finally, we argue that our findings motivate reframing access in collaborative data analysis as effective bidirectional communication.

5.1 Synthesis: Key Insights on Mixed-Ability Common Grounding

Across the tactile map and co-operative computer use scenarios, our observations converge on a set of insights about how common ground is established and maintained in mixed-ability collaborative data analysis:

- **Grounding requires more than reference.** Collaborative analytics research has primarily considered representations’ referential affordances—pointing at and describing data. Our observations show that mixed-ability grounding equally depends on metalingual work (verifying understanding across modalities) and phatic work (signaling continued engagement). The referential and metalingual functions recurred across both the tactile map and the co-operative computer scenarios (Sec. 4.1 and Sec. 4.2); the phatic function was most pronounced in tactile reading, where continuous hand contact passively signaled engagement. Together this suggests all three are requirements of mixed-ability grounding, met through different affordances in different settings.
- **A single representation can serve different collaborators through different channels simultaneously.** The tactile map is a tactile reading surface for the blind analyst and, at the same time, a visual display of the blind analyst’s attention for the sighted analyst. Likewise, the shared cursor is read through speech by

the screen reader user and visually by the sighted collaborator. Representations in mixed-ability settings are best thought of as communicative resources that convey information in multiple perceptual channels.

- **The most valuable grounding evidence is passive, produced as a byproduct of a collaborator's task.** A sighted analyst can see where a blind analyst is attending as a byproduct of their hands moving across the tactile map, or by watching the screen reader cursor traverse a document. This kind of passive evidence helped collaborators maintain common ground while reducing verbal turn-taking overhead.
- **Common ground enables fuller participation in analysis.** Once understanding aligned across modalities, discussion shifted from the appearance of figures to the underlying science and design decisions (Sec. 4.1), suggesting that successful accessible representations must not only convey equivalent information but also enable collaborative participation.
- **Mixed-ability collaboration requires explicit control handoffs when sharing a digital workspace.** Co-operative computer use depends on a spoken protocol ("Can I drive?" / "I relinquish control"). These verbal participation cues do the coordination work that gesture and body language do for sighted teams, and they double as a way of establishing shared understanding of a collaborative workspace's current state.
- **Accessible collaboration is supported by relational and institutional structure.** The lab is developing an evolving body of know-how supported by an institutionalized Access Assistant role that redistributes access labor. Supporting accessible collaboration also includes attending to non-technical factors that help a team build and retain access expertise over time.

Limitations. Several features of our setting limit the transferability of specific observations. The lab's extensive referential use of the tactile map depends on the geographic nature of oceanographic data and on long-lived artifacts that collaborators have repeatedly used over years. Analytical domains with abstract or rapidly-changing data may not support equivalent tactile reference. Likewise, the lab's grounding practices rest on decades-long working relationships and an institutionalized Access Assistant role that most teams lack. And although we previously noted the lab's expertise and interest in sonification, we did not actually observe any instances of lab members using sonification in their workflows. This could be due to limitations in their existing tooling that made sonification inappropriate for everyday work, or due to limitations in our ability to conduct longer-term inquiry. We therefore are limited to offering evidence that the lab's specific practices are successfully used in this context. Nonetheless, we hypothesize that our evidence about the communicative functions they serve are requirements of mixed-ability collaboration more generally, that other settings will satisfy through different site-specific practices.

5.2 Future Work: Designing Multi-Modal Affordances for Common Ground

In collaborative visual analytics, common grounding for analysts sharing a single visual display can often be accomplished primarily visually, with some supporting verbal and gestural communication. Mixed-ability common grounding, in contrast, requires analysts to resourcefully combine multiple modalities to establish and maintain alignment. Our analysis of three communicative requirements—referential, metalingual, and phatic—showed that while collaborative analytics tools frequently consider the referential function, future research might explore dedicated affordances for metalingual and phatic uses of representations to support common ground.

Across both the tactile map and co-operative computer use scenarios, we observed that the most valuable grounding evidence was the continuous, low-cost trace that was an incidental byproduct of what the collaborator was already doing. This suggests a guiding principle for design: representations should passively provide grounding evidence across multiple modalities as a byproduct of interaction, rather

than requiring dedicated communicative acts that interrupt the flow of analysis.

In reflecting on our observations, we found it useful to think about grounding evidence in terms of the active/passive distinction we introduced in Sec. 4.1.2, together with the question of which perceptual channel the evidence travels through. For grounding, we specifically focus on the receiver's perceptual channel. For example, hand position on a map is visual to an observing collaborator even though the act of reading is tactile for the person producing it, so what determines whether a collaborator can use a piece of evidence is the channel through which it reaches them. The passive evidence available in our setting was almost entirely visual. Screen reader narration offers some auditory evidence, though it required more active management to not interrupt the surrounding conversation. But passive evidence in the tactile channel was essentially absent.

This asymmetry points toward a direction for future work examining how tactile or other non-visual grounding evidence can be produced passively, as a byproduct of interaction rather than a dedicated act. Several existing technologies suggest starting points. Refreshable tactile displays (RTDs), increasingly used to present charts non-visually [20], could render a sighted collaborator's attention tactually to convert visual pointing into passive tactile evidence—for example, by raising or lowering pins to mirror the position of a mouse cursor. Ambient auditory approaches offer another possibility, such as continuous sonification that provides soundscape-like signaling of attention.

These directions concern how a collaborator might produce grounding evidence passively, but our observations suggest that just making evidence available doesn't mean it will be easy to receive. Evidence has to be perceived through a channel, and channels vary in how much backgrounded information they can carry. In Sec. 4.2.3, we saw that screen reader narration takes over the verbal/auditory channel, requiring active management to avoid disrupting conversation. A channel already carrying speech cannot easily carry continuous backgrounded signal as well. So designs aiming to fill the tactile gap should look for evidence that is not only passively produced but also peripheral to attend to, perceivable without displacing the receiver's primary task.

5.3 Reframing Access as Effective Bidirectional Communication

Our findings complicate a common assumption that access is something sighted people provide and blind people receive. DeafBlind poet and activist John Lee Clark identifies this assumption directly in his essay "Against Access," asking why access is always a question of whether the sighted world includes disabled people, rather than whether disabled people include the sighted [10]. Clark's provocation reverses the usual direction of the question, and our observations suggest the reversal is not just rhetorical. In the Bower Lab, the blind analyst is continuously producing signals—phatic hand contact, cursor movements, spoken verification—that her sighted collaborators depend on to follow her analysis. This shows that the work of making a shared analysis legible needs to happen in both directions.

By understanding access as bidirectional, we are encouraged to pay more attention to multi-modal representations' communicative functions. The tactile map might be conventionally thought of as a non-visual substitute that lets Bower read data she cannot see. But as we showed in Sec. 4.1, the same artifact simultaneously renders her attention and engagement visible to her sighted collaborators, doing communicative work for them that the visual chart alone did not. An assistive representation, in other words, can be designed for two audiences. It can both serve the blind reader and, as a byproduct, also serve the sighted collaborators in establishing common ground. As such, we argue for rethinking accessible design beyond delivering information to an individual user, toward supporting mutual intelligibility.

ACKNOWLEDGMENTS

The authors wish to thank Bower Lab members for their time and hospitality. We also thank Graham Jones, and our anonymous reviewers. This work was supported in part by NSF Award #2341748 and MIT-SERC.

REFERENCES

- [1] T. Akter, A. S. Marathe, D. Gergle, and A. M. Piper. Beyond Accessibility: Understanding the Ease of Use and Impacts of Digital Collaboration Tools for Blind and Low Vision Workers. In *Proceedings of the 27th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '25, pp. 1–17. Association for Computing Machinery, New York, NY, USA, Oct. 2025. doi: [10.1145/3663547.3746332](https://doi.org/10.1145/3663547.3746332) 2
- [2] Anna Pinckney. How to make accessible graphics using a PIAF machine, Jan. 2025. 3
- [3] C. L. Bennett, E. Brady, and S. M. Branham. Interdependence as a Frame for Assistive Technology Research and Design. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*, pp. 161–173. ACM, Galway Ireland, Oct. 2018. doi: [10.1145/3234695.3236348](https://doi.org/10.1145/3234695.3236348) 2, 7
- [4] M. Blanco, J. Zong, and A. Satyanarayan. Olli: An Extensible Visualization Library for Screen Reader Accessibility. In *IEEE VIS Posters*, vol. 6, p. 65, 2022. 2
- [5] S. M. Branham and S. K. Kane. 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, pp. 163–171. Association for Computing Machinery, New York, NY, USA, Oct. 2015. doi: [10.1145/2700648.2809864](https://doi.org/10.1145/2700648.2809864) 2, 3
- [6] K. Charmaz. *Constructing Grounded Theory*. Sage Publications, London ; Thousand Oaks, Calif, 2006. 3
- [7] M. K. Chen, I. Pedraza Pineros, A. Satyanarayan, and J. Zong. Tactile Vega-Lite: Rapidly Prototyping Tactile Charts with Smart Defaults. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, pp. 1–23. Association for Computing Machinery, New York, NY, USA, Apr. 2025. doi: [10.1145/3706598.3714132](https://doi.org/10.1145/3706598.3714132) 2, 6
- [8] P. Chundury, J. B. Jordan, Y. Reyazuddin, N. Elmqvist, and J. Lazar. Sound, Touch, or the Full Monty? A Comparative Study of Accessible Data Exploration Systems for Blind Users. *ACM Transactions on Accessible Computing*, 19(1):1:1–1:41, Mar. 2026. doi: [10.1145/3798100](https://doi.org/10.1145/3798100) 1
- [9] H. H. Clark and S. E. Brennan. Grounding in communication. In *Perspectives on socially shared cognition*, pp. 127–149. American Psychological Association, Washington, DC, US, 1991. doi: [10.1037/10096-006](https://doi.org/10.1037/10096-006) 1, 5
- [10] J. L. Clark. Against Access. *McSweeney's Quarterly Concern* 64: *The Audio Issue*, Oct. 2021. 9
- [11] P. Dourish. Reading and Interpreting Ethnography. In J. S. Olson and W. A. Kellogg, eds., *Ways of Knowing in HCI*, pp. 1–23. Springer, New York, NY, 2014. doi: [10.1007/978-1-4939-0378-8_1](https://doi.org/10.1007/978-1-4939-0378-8_1) 3, 4
- [12] T. Edwards. *Going Tactile: Life at the Limits of Language*. Oxford Studies in the Anthropology of Language. Oxford University Press, Oxford, New York, July 2024. 4
- [13] F. Elavsky, L. Nadolskis, and D. Moritz. Data Navigator: An Accessibility-Centered Data Navigation Toolkit. *IEEE Transactions on Visualization & Computer Graphics*, 30(01):803–813, Jan. 2024. doi: [10.1109/TVCG.2023.3327393](https://doi.org/10.1109/TVCG.2023.3327393) 2
- [14] A. Hamraie. Designing Collective Access: A Feminist Disability Theory of Universal Design. *Disability Studies Quarterly*, 33(4), Sept. 2013. doi: [10.18061/dsq.v33i4.3871](https://doi.org/10.18061/dsq.v33i4.3871) 7
- [15] A. Hamraie and K. Fritsch. Crip Technoscience Manifesto. *Catalyst: Feminism, Theory, Technoscience*, 5(1):1–33, Apr. 2019. Number: 1. doi: [10.28968/cftt.v5i1.29607](https://doi.org/10.28968/cftt.v5i1.29607) 3
- [16] C. Han and K. E. Isaacs. A Deixis-Centered Approach for Documenting Remote Synchronous Communication Around Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):930–940, Jan. 2025. doi: [10.1109/TVCG.2024.3456351](https://doi.org/10.1109/TVCG.2024.3456351) 1, 2
- [17] J. Heer and M. Agrawala. Design Considerations for Collaborative Visual Analytics. *Information Visualization*, 7(1):49–62, Mar. 2008. doi: [10.1057/palgrave.ifs.9500167](https://doi.org/10.1057/palgrave.ifs.9500167) 1, 2
- [18] L. Hickman and C. Bennett. Access Work: Laboring with Non-Innocent Authorization. *interactions*, 30(4):60–64, June 2023. doi: [10.1145/3603494](https://doi.org/10.1145/3603494) 8
- [19] W. Hill and J. Hollan. Deixis and the future of visualization excellence. In *Proceeding Visualization '91*, pp. 314–320. IEEE Comput. Soc. Press, San Diego, CA, USA, 1991. doi: [10.1109/VISUAL.1991.175820](https://doi.org/10.1109/VISUAL.1991.175820) 1, 2
- [20] L. Holloway, P. Cracknell, K. Stephens, M. Fanshawe, S. Reinders, K. Marriott et al. Refreshable Tactile Displays for Accessible Data Visualisation, Jan. 2024. arXiv:2401.15836 [cs.HC]. doi: [10.48550/arXiv.2401.15836](https://doi.org/10.48550/arXiv.2401.15836) 9
- [21] P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Ha-gen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, Oct. 2011. doi: [10.1177/1473871611412817](https://doi.org/10.1177/1473871611412817) 2
- [22] A. Kafer. *Feminist, Queer, Crip*. Indiana University Press, 2013. 2
- [23] S. S. Kamath, A. N. Zeidieh, V. Potluri, S. O'Modhrain, K. Perry, and J. Seo. Three Modalities, Two Design Probes, One Prototype, and No Vision: Experience-Based Co-Design of a Multi-modal 3D Data Visualization Tool. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26, pp. 1–23. Association for Computing Machinery, New York, NY, USA, Apr. 2026. doi: [10.1145/3772318.3791272](https://doi.org/10.1145/3772318.3791272) 1
- [24] H. Karen and J. Sandra. Contextual Inquiry: A Participatory Technique for System Design. In *Participatory Design*. CRC Press, 1993. Num Pages: 34. 3
- [25] E. Keating and C. Sunakawa. Participation cues: Coordinating activity and collaboration in complex online gaming worlds. *Language in Society*, 39(3):331–356, June 2010. doi: [10.1017/S0047404510000217](https://doi.org/10.1017/S0047404510000217) 1, 3, 6
- [26] H. Kim, Y.-S. Kim, and J. Hullman. Erie: A Declarative Grammar for Data Sonification. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, pp. 1–19. Association for Computing Machinery, New York, NY, USA, May 2024. doi: [10.1145/3613904.3642442](https://doi.org/10.1145/3613904.3642442) 2
- [27] N. McDonald, S. Schoenebeck, and A. Forte. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):72:1–72:23, Nov. 2019. doi: [10.1145/3359174](https://doi.org/10.1145/3359174) 4
- [28] C. Mei, J. Pollock, D. Hajas, J. Zong, and A. Satyanarayan. Benthic: Perceptually Congruent Structures for Accessible Charts and Diagrams. In *Proceedings of the 27th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '25, pp. 1–17. Association for Computing Machinery, New York, NY, USA, Oct. 2025. doi: [10.1145/3663547.3746342](https://doi.org/10.1145/3663547.3746342) 2
- [29] Mia Mingus. Access Intimacy: The Missing Link, May 2011. 7
- [30] Mia Mingus. Access Intimacy, Interdependence and Disability Justice, Apr. 2017. 2, 7, 8
- [31] R. Neogy, J. Zong, and A. Satyanarayan. Representing Real-Time Multi-User Collaboration in Visualizations. In *2020 IEEE Visualization Conference (VIS)*, pp. 146–150. IEEE, Salt Lake City, UT, USA, Oct. 2020. doi: [10.1109/VIS47514.2020.00036](https://doi.org/10.1109/VIS47514.2020.00036) 1, 2, 5
- [32] P. Pirolli and S. Card. The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis. In *Proceedings of the International Conference on Intelligence Analysis*, vol. 5, pp. 2–4, 2005. 8
- [33] M. Price. *Crip Spacetime: Access, Failure, and Accountability in Academic Life*. Duke University Press, May 2024. doi: [10.1515/9781478093992](https://doi.org/10.1515/9781478093992) 8
- [34] Roman Jakobson. Linguistics and Poetics. In Thomas Sebeok, ed., *Style in Language*, pp. 350–377. MIT Press, Cambridge, 1960. 1, 3
- [35] J. Seo, Y. Xia, B. Lee, S. Mccurry, and Y. J. Yam. MAIDR: Making Statistical Visualizations Accessible with Multimodal Data Representation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, pp. 1–22. Association for Computing Machinery, New York, NY, USA, May 2024. doi: [10.1145/3613904.3642730](https://doi.org/10.1145/3613904.3642730) 1, 2
- [36] J. R. Thompson, J. J. Martinez, A. Sarikaya, E. Cutrell, and B. Lee. Chart Reader: Accessible Visualization Experiences Designed with Screen Reader Users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pp. 1–18. Association for Computing Machinery, New York, NY, USA, Apr. 2023. doi: [10.1145/3544548.3581186](https://doi.org/10.1145/3544548.3581186) 1, 2
- [37] E. Q. Wang and A. M. Piper. Accessibility in Action: Co-Located Collaboration among Deaf and Hearing Professionals. *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW):180:1–180:25, Nov. 2018. doi: [10.1145/3274449](https://doi.org/10.1145/3274449) 2
- [38] J. Wood, H. Wright, and K. Brodie. Collaborative visualization. In *Proceedings. Visualization '97*, pp. 253–259. IEEE, Phoenix, AZ, USA, Oct. 1997. doi: [10.1109/VISUAL.1997.663890](https://doi.org/10.1109/VISUAL.1997.663890) 2
- [39] Y. Zhao, M. A. Nacenta, M. A. Sukhai, and S. Somanath. Accessibility-Driven Information Transformations in Mixed-Visual Ability Work Teams. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26, pp. 1–14. Association for Computing Machinery, New York, NY, USA, Apr. 2026. doi: [10.1145/3772318.3790872](https://doi.org/10.1145/3772318.3790872) 2
- [40] J. Zong. Pointing, Mutual Intelligibility, and the Seeing Subject in HCI. *Interface Critique*, (4):95–105, Apr. 2023. Number: 4. doi: [10.11588/ic.2023.4.93413](https://doi.org/10.11588/ic.2023.4.93413) 7

- [41] J. Zong. Using Real Names of Disabled Participant-Contributors to Practice Citational Justice in Accessibility. In *2025 IEEE Workshop on Accessible Data Visualization (AccessViz)*, pp. 30–33. IEEE, Vienna, Austria, Nov. 2025. doi: [10.1109/AccessViz68666.2025.00011](https://doi.org/10.1109/AccessViz68666.2025.00011) 3
- [42] J. Zong, C. Lee, A. Lundgard, J. Jang, D. Hajas, and A. Satyanarayan. Rich Screen Reader Experiences for Accessible Data Visualization. *Computer Graphics Forum*, 41(3):15–27, Aug. 2022. doi: [10.1111/cgf.14519](https://doi.org/10.1111/cgf.14519) 2, 8
- [43] J. Zong, I. Pedraza Pineros, M. K. Chen, D. Hajas, and A. Satyanarayan. Umwelt: Accessible Structured Editing of Multi-Modal Data Representations. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, pp. 1–20. Association for Computing Machinery, New York, NY, USA, May 2024. doi: [10.1145/3613904.3641996](https://doi.org/10.1145/3613904.3641996) 1, 2, 4