





Talk to the Wall: The Role of Speech Interaction in Collaborative Visual Analytics

Gabriela Molina León , Anastasia Bezerianos , Olivier Gladin , and Petra Isenberg 




Fig. 1: Collaborative sensemaking system on the wall-sized display: (a) Participants working together on solving the task. (b) Selecting a document via touch interaction. (c) Searching for a keyword via speech interaction.

Abstract—We present the results of an exploratory study on how pairs interact with speech commands and touch gestures on a wall-sized display during a collaborative sensemaking task. Previous work has shown that speech commands, alone or in combination with other input modalities, can support visual data exploration by individuals. However, it is still unknown whether and how speech commands can be used in collaboration, and for what tasks. To answer these questions, we developed a functioning prototype that we used as a technology probe. We conducted an in-depth exploratory study with 10 participant pairs to analyze their interaction choices, the interplay between the input modalities, and their collaboration. While touch was the most used modality, we found that participants preferred speech commands for global operations, used them for distant interaction, and that speech interaction contributed to the awareness of the partner’s actions. Furthermore, the likelihood of using speech commands during collaboration was related to the personality trait of agreeableness. Regarding collaboration styles, participants interacted with speech equally often whether they were in loosely or closely coupled collaboration. While the partners stood closer to each other during close collaboration, they did not distance themselves to use speech commands. From our findings, we derive and contribute a set of design considerations for collaborative and multimodal interactive data analysis systems. All supplemental materials are available at <https://osf.io/8gpfv>.

Index Terms—Speech interaction, wall display, collaborative sensemaking, multimodal interaction, collaboration styles.

1 INTRODUCTION

Wall displays come with a set of benefits and challenges for the visualization and analysis of large datasets. They provide more screen real-estate and offer people “space to think” [2]. The large physical size promotes physical navigation for better performance [5] and facilitates

co-located collaboration [25]. As data grows in size and complexity, combining the knowledge and skills of multiple people can facilitate sensemaking tasks [43]. When two or more people collaborate in front of a wall display, their interactions should not interfere with the coordination of tasks and activities between team members.

Interacting with large displays requires techniques that may be unlike the standard mouse and keyboard interactions usually coupled with windows, icons, menus, and pointers (WIMP) interfaces. New technologies now allow us to interact with data using all our physical senses. Ten years ago, already, Roberts et al. [46] pointed to the importance of future research on how these technologies could be fluidly integrated. On large displays, touch input is a classic support for close-up interaction via direct manipulation [49]. Touch, however, cannot be the only input modality for wall displays because people often stand far away from the wall to look at overviews of the shown data. For interaction distant from the wall, researchers have proposed techniques that do not require touching the display such as laser pointers [37], mobile phones/smartwatches [11, 22], gestural interfaces [38], and speech [47].

- Gabriela Molina León is with the University of Bremen. E-mail: molina@uni-bremen.de.
- Anastasia Bezerianos is with the Université Paris-Saclay, CNRS, Inria. E-mail: anastasia.bezerianos@universite-paris-saclay.fr.
- Olivier Gladin is with Inria. E-mail: olivier.gladin@inria.fr.
- Petra Isenberg is with Université Paris-Saclay, CNRS, Inria. E-mail: petra.isenberg@inria.fr.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

Integrating different input modalities is a research challenge in this context because the close-up and distant interaction techniques need to be seamlessly combined to support sensemaking activities. In addition, it is not only important how input modalities integrate for one analyst but also how they support multiple analysts working together.

Our work focuses on exploring two specific input modalities during collaborative sensemaking: touch and speech. We chose touch as the most common type of close-up modality for wall displays. Our choice of speech is motivated by the outcome of our recent collaborative elicitation study [36], where participants preferred speech commands over touch, pen, and mid-air gestures to explore data on a large vertical display. Speech (or natural language) interaction is also an interesting candidate to study as it is not yet clear how its use may influence the communication between collaborators, their awareness of each others' activities, and whether personal characteristics may influence the choice of using speech at all. As little is still known about how speech and touch may be used in practice during collaborative analysis, we conducted an exploratory study, focused on how pairs solve a sense-making task. More specifically, we were interested in their modality choices, their movements in front of the display, and how multimodal interaction and collaboration styles intertwine.

To investigate these aspects, we designed and implemented TOUCHTALKINTERACTIVE, a wall-based interactive system supporting touch and speech commands for co-located collaborative work. We conducted a study with 20 participants (10 pairs) who solved a fictitious mystery by interacting with a document collection on the wall display. We included interaction techniques facilitated by both modalities and allowed participants to freely choose between them. We analyzed their choices, movements, and collaboration styles and report on our findings regarding how the interaction modalities were used in collaboration.

In summary, we make the following contributions:

- We present insights on the interaction choices, movements, and collaboration styles of the participants and the interplay with speech interaction.
- We examine how personality traits relate to the use of input modalities, such as speech commands.
- We derive a set of design considerations on how to support co-located collaborative work leveraging speech interaction on wall-based interactive systems.
- We release TOUCHTALKINTERACTIVE as an open-source tool to support reproducibility.

2 RELATED WORK

We present relevant work related to our research on the topics of interaction with large vertical displays, natural language interfaces, collaborative visual analytics, and personality traits.

2.1 Interaction with Large Vertical Displays

Large vertical displays provide more pixels and more space to interact with [7]. As the space in front of the display allows physical navigation, which can lead to better performance [5], wall-based systems should provide different ways of interacting from varying distances, to make the most of the surface and space. Touch is the standard input modality to support direct manipulation, but it comes with the disadvantage that the display areas beyond the arm's reach become inaccessible. Thus, researchers have come up with interaction techniques to virtually move the display content towards the user [45]. Interacting with a pen or stylus is another type of direct manipulation that people leverage to write annotations [36] and author visualizations [58]. While touch and pen can be a powerful combination [20], we decided to leverage touch only as it is the most common modality supported and it requires no additional hardware. However, direct manipulation requires participants to stand close to the wall display. Thus, we need other input modalities to facilitate interaction from a distance. For instance, Nancel et al. [38] evaluated and proposed mid-air gestures to navigate through pan and zoom interactions on wall displays. Another possibility is to use additional devices, such as mobile phones [31] and smartwatches [22]

to interact from a distance. However, our work focuses on supporting distant interaction without needing additional handheld devices.

Previous work has studied the combination of multiple input modalities on large vertical displays. Srinivasan and Stasko [53] proposed combining direct manipulation and speech commands to interact with network visualizations and found that participants chose different modalities across tasks, e.g., speech for filtering and touch for highlighting connections between nodes. DataBreeze [51] combines touch, pen, and speech to interact with unit visualizations. While interacting, participants tended to use speech for global actions and touch-based context menus for local actions. Accordingly, we incorporated examples of such interactions in our system, to examine whether the same happens in the context of collaborative work. Saktheeswaran et al. [47] found that when comparing touch, speech, and touch-and-speech interactions, participants preferred the multimodal interactions. However, Molina León et al. [36] found that people preferred speech commands, over touch, pen, and mid-air gestures, alone or in combination, for a series of exploratory tasks. Therefore, we examine the use of touch and speech in a real-world system to study how participants choose to interact.

2.2 Natural Language Interfaces

Natural language interfaces (NLI) facilitate interaction through written or spoken language. They allow users to express their intent without having to translate questions into data attributes [15] but suffer from discoverability issues and depend on feedback [52]. Eviza [48] is an example of such an interface, supporting visual data analysis through written text, and giving the user a sense of dialogue by supporting follow-up queries. Gao et al. [14] designed DataTone, which supports both written and spoken text and addresses the issue of language ambiguity by generating widgets for users to refine their queries. Recently, Guo et al. [18] proposed a deep-learning solution that generates visualizations to answer multiple aspects considered relevant to a question posed by the user. While these solutions are all desktop-based, there are also NLIs designed for other devices. Kassel and Rohs [30] designed Valetto, a tablet-based interface that supports visualization authoring through speech and touch. The authors assigned each modality to specific actions, but their evaluation revealed that redundancy would be preferable. Tabalba et al. [54] designed an NLI that allows pairs to author charts on a large display, using their group conversation as additional query context. In contrast, we study interaction techniques facilitated through both touch and speech. While we do not focus on conversation or query refinement, we seek to support diverse low-level actions to give the option of choosing between two modalities.

A common challenge to facilitate natural language interaction is the technological capabilities to recognize natural language. Accordingly, Narechania et al. [39] created a toolkit to author visualizations based on a written query and related systems made use of HTML5 speech recognition APIs to translate from spoken language to text [50]. To circumvent common accuracy issues, we trained a deep-learning model with the vocabulary of our dataset and integrated it into our system.

2.3 Collaborative Visual Analytics

Researchers have proposed different ways of supporting multi-user interaction on large displays. Langner et al. [31] proposed using touch and touch-enabled mobile devices to allow participants to interact from different distances from the display. Badam et al. [4] designed interaction techniques combining proxemics with mid-air gestures to interact with lenses on wall displays. James et al. [29] proposed to leverage augmented reality to extend the interaction possibilities in the space in front of the display. We designed a system to leverage touch and speech interaction between two users in parallel if they want.

When people do not only interact with a system, but also with each other, social aspects need to be considered, such as collaborative distance, awareness, and communication among team members [32]. We examine how the use of speech may influence these aspects, e.g., encouraging people to move away from each other before using a speech command. We studied how speech interaction may influence awareness and its interplay with human communication. Territoriality and privacy are other aspects to consider. For example, Reipschläger et

al. [44] proposed using augmented reality to extend the display surface by creating a personal virtual space for each collaborator. While these aspects are out of the scope of our work, we designed our system to provide a set of interface elements (e.g., tag list, virtual keyboard) that would allow participants to work alone if desired.

Regarding the use of multiple input modalities for collaborative work, Tse et al. [56] explored how combining speech and hand gestures affected collaborative work on tabletops. The authors identified speech use in parallel commands as a design issue, i.e., if speech commands are used often for actions that people need to perform in parallel, one person may decide to work sequentially to avoid voice overlapping. Their findings inspired us to study the effects of speech commands on wall displays. Furthermore, the way people work together can qualitatively vary depending on how closely or loosely they collaborate. Tang et al. [55] proposed a set of codes to classify collaboration styles according to how strong the relationship between the activities of each person was. Isenberg et al. [27] extended the set for analyzing collaborative work on tabletops to take multiple data views into account. Other researchers have extended or adapted the styles proposed by Tang et al. and Isenberg et al. for scenarios with more than one device [9] or for hybrid collaborations [40]. However, those characteristics do not apply to our scenario, thus, we use the code set of Isenberg et al. to analyze the collaborative work of our participants. We selected the dataset and designed our experiment based on their work and follow-up studies [1, 28] to facilitate comparison and validation of our findings.

2.4 Personality Traits and Visualization

Designing one-size-fits-all interfaces for visualization systems is problematic because the individual differences among people can affect how a person uses a visualization [33]. Thus, researchers have looked at personality traits to analyze whether and how they may affect the experience of using visualizations. The most known personality research model is the Five-Factor Model. As its name suggests, it refers to five personality traits that describe individual differences: extraversion, openness, agreeableness, conscientiousness, and neuroticism. We focus on the first three traits, as we consider them relevant for speech interaction and collaborative work. *Extroverted* people are outgoing and sociable [57]. In a study on task performance, Green et al. [16] found that extroverts solved search tasks faster than introverts. However, introverts gained more insights. Ziemkiewicz et al. [59] found that introverts took more time to analyze the task and, thus, solved more tasks accurately than introverts. We expect that extroverted participants will use speech interaction more often than introverts because they may feel more comfortable speaking publicly.

Openness to experience is the second trait that could influence the likelihood of using speech interaction, as it refers to curiosity and the wish to seek new experiences [13, 57]. We expect the willingness of the participants to experiment with less common modalities (such as speech) may influence their choices. *Agreeableness* is the third trait we examine, as it refers to being cooperative [57] and is considered a positive trait for people to succeed in collaborative work [33]. We expect that the use of speech interaction negatively correlates with agreeableness. In other words, we suspect that participants will be reluctant to use speech so as not to disturb their colleagues.

3 EXPERIMENT DESIGN

To explore the synergy between speech and touch interaction, we designed an in-person experiment where pairs of participants had to solve a sensemaking task collaboratively. Our goal was to investigate the following research questions:

RQ1 How do partners use touch and speech individually while working as part of a team?

RQ2 How are touch and speech used during collaboration? When is each modality chosen and how does it relate to in-team communication and coordination?

We asked each pair to solve a fictitious mystery described in a collection of text documents that could be interacted with on a touch and speech-enabled wall display. We designed the experiment inspired by

two prior collaborative studies on wall displays [27, 28]. Following a prior VAST challenge [17], the researchers asked participants to find a weapons-smuggling plot hidden among a large number of text documents. Teamwork can be helpful to solve the task as more documents can be read together and participants can discuss their hypotheses. In our system, participants were free to use any supported touch gestures and speech commands. To understand participants' choice of modalities, we ensured that all commands were available in both speech and touch; see Sec. 4 for details on the interaction techniques. We video-recorded the sessions, tracked the position of the participants, and logged their interactions with the system.

Before the main experiment, we conducted two pilot studies. The first helped us finalize the list of interaction techniques the system should support to solve the task. The second pilot study with another participant pair helped us identify usability issues that we then solved before the main study, and to refine the protocol and the setup regarding audio capture when multiple people were talking simultaneously.

3.1 Dataset and Task

We created an adapted version of the dataset used in the "Stegosaurus" scenario from the interactive session of the 2006 VAST challenge [17]. Our dataset contained 61 text files and four images. The text files were news articles from a fictitious town called Springfield¹ and other miscellaneous documents, such as a list of diseases and a list of terrorist organizations. From the 61 text files, 10 news articles contained relevant information about a hidden weapon-smuggling plot, seven provided background information, and the miscellaneous documents gave some important clues about the evidence (e.g., the characteristics of a chemical involved). Similar to previous work [27, 28], the pairs had one main task: generating hypotheses about what happened. The task encourages participants to collaborate loosely and closely on different parts of the task [28] (e.g., reading documents and discussing hypotheses), which allowed us to study various modes of collaboration.

3.2 Personality Traits

We measured three personality traits per participant: extraversion, openness to experience, and agreeableness (see Sec. 2.4). We hypothesized that participants with higher extraversion and openness to experience would be more at ease to use speech commands, as they may feel more comfortable speaking publicly and using new technology; whereas participants with high agreeableness may be reluctant to use speech so as to not distract others. We decided to examine personality traits because, in the first pilot study, both participants suggested that personality was an important factor regarding the choice to use speech interaction. For instance, one participant indicated that "*I found that I do not want to use speech in front of him (...) I am just an introverted person.*"

We used the IPIP-NEO-60 questionnaire by Maples-Keller et al. [34] of the NEO PI-R domains from the instrument of Costa and McCrae [12]. Based on the answers, we calculated three scores per participant that we could operationalize to describe each trait and compare them to the total of speech commands. While dozens of personality assessment instruments exist, many are copyrighted and paid-only. Therefore, we chose the freely available IPIP-NEO-60, which is easily accessible for future replications. It includes 12 questions per trait, and it is, thus, more practical than others available with over 300 questions.

3.3 Procedure

We started the experiment by explaining our research motivation and handing participants a consent form to read and sign. If both participants gave consent, we gave each of them a pre-questionnaire to fill out. The questionnaire included five questions about their previous experience with interactive technologies, two questions about their collaboration experience with their partner, two demographic questions, and 36 questions from the personality traits questionnaire. Afterward, we proceeded to a 10-minute training session. The experimenter introduced the system and then participants went through a list of all

¹We modified the town name (previously named Alderwood) and other proper nouns in the dataset to ensure the speech engine would understand them.

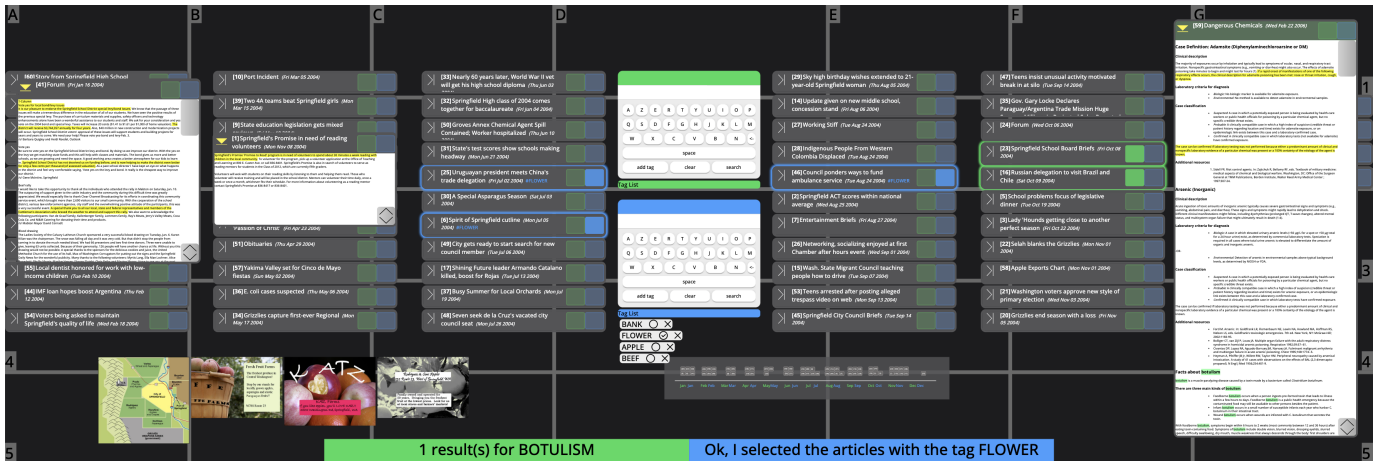


Fig. 2: Screenshot of TOUCHTALKINTERACTIVE. The blue and green colors identify the interface elements associated with each participant. The bars on the bottom of the screen give feedback to the user regarding their last speech command and any search-related tasks with either touch or speech.

supported interactions and performed every interaction using speech and touch gestures at least once. We proceeded to the main phase of the study once participants felt confident enough to interact with the system independently. We also provided participants with a cheat sheet (one page) of all interaction techniques supported by the system, describing the necessary actions for using them with touch and speech. Most of them kept the sheet in their hand while solving the task.

The main phase of the study took 45 minutes. Participants started by reading a one-page background document introducing the task scenario. This document told the story of a mysterious incident in the fictional city of Springfield that the police wished to investigate further. To generate hypotheses, explore the dataset, and solve the mystery participants could interact with the system as they wished. Participants could ask questions to the experimenter about the system, but they did not receive additional help to solve the mystery. Once the 45 minutes were over, participants filled out a post-questionnaire about their experience in terms of collaboration and input modalities. Afterward, we conducted a short interview where we asked them to elaborate on what they found out about the mystery and their impressions of the input modalities, as well as of the collaborative work with their partner. We did not control the initial position of the participants, they were free to move in the space in front of the wall display as they wished. We also did not measure their performance because our goal was to capture holistic, qualitative, observational data to understand how they interacted with each other and the system—and not to compare accuracy or speed that in this task is highly dependent on text reading speed and understanding.

3.4 Data Collection and Analysis

To answer **RQ1**, we logged interactions to know how often each participant used each input modality and for which actions. To measure the distance to the screen, we used a Vicon motion tracking system to document the movements of the participants in front of the wall display.

We video-recorded the main study phase with a camera positioned at the back of the room and a microphone above the display. We used two wireless microphones to perform the speech commands. A speech-to-intent engine translated the commands directly into machine-readable text in JSON format, including only the information necessary to process the command. We analyzed the videos qualitatively to identify the interaction and collaboration styles of the participants to help us answer **RQ2**. One author of this paper coded the videos in two passes with the help of a commercial coding tool. The first pass was meant to validate and complete the interaction logs and the second was focused on coding the collaboration styles. We applied the codes associated with collaboration styles proposed by Isenberg et al. [27], except for their SIDV code that refers to multiple copies of the same document, as our system showed each document only once.

The post-questionnaire and the interview helped to obtain subjective feedback about the collaborative experience and the interactions. We

generated the first transcription of the interviews with an automatic tool that we then revised manually. We analyzed the interviews through an open coding scheme based on grounded theory.

4 SYSTEM AND INTERACTION DESIGN



We designed and implemented a visual analytics system called TOUCHTALKINTERACTIVE to support co-located collaborative sense-making on wall-sized displays, using touch and speech interaction. Its functionality is inspired by previous research-oriented systems created to support collaborative tasks [1, 26, 28] and its speech support by studies on interaction with large vertical displays [36, 47]. While prior wall display systems support speech and other modalities for a single user [47, 51] or speech for visualization authoring in a group context [54], ours enables the use of speech interaction for collaborative sensemaking. Our goal was to use the system as a *technology probe* [23] with two goals: a) to help understand how well the proposed solution addresses interaction and collaboration needs and b) to inspire users and researchers to reflect on speech interaction in collaboration.

Our prototype is designed to assist two people in exploring and making sense of a large collection of documents through touch gestures and speech commands. Each document is visually represented as a unit, that can be opened or extended to read the content. The system includes a timeline to visualize and select the documents according to the temporal dimension and supports images. By default, the images are positioned at the bottom of the screen, as previous work suggests that the lower area of wall displays should not be used for data representations [8]. Each person can interact with one or multiple documents at a time, and both persons can interact simultaneously. Speech recognition is activated with a click on a wireless mouse. Thus, using a wake word, such as “Hey Siri”, is not required, and participants can be confident that the system only listens to them when they wish to. Each person has a designated color (blue or green) and a set of interaction widgets in their color, e.g., a keyboard to search for keywords. That way, participants are free to work individually if they wish to.

4.1 Interaction Techniques

Our goal was to support a comprehensive list of interactions with the documents, both through touch gestures and speech commands, so that each of the two participants could interact in parallel while choosing any modality. Although single multimodal interactions would have been possible (i.e., combing speech and touch in one command), we kept the modalities separate to compare our results to the elicitation study of our related work [36], and to examine when and in what contexts participants would choose either modality. We present the 20 interaction techniques the system supports in **Tab. 1** (for the speech commands, we show only the main examples because synonyms were supported). We carefully chose them to facilitate co-located and synchronous work,

Table 1: Interaction techniques supported by the system. Each interaction can be performed with touch gestures and speech commands.

Action	Touch gesture 	Speech command 
Open document	Tap on the left corner of header (toggle).	"Open document 12"
Close document	Tap on the left corner of header (toggle).	"Close document 12"
Select document	Tap on radio box on the document header.	"Select document 12"
Select documents per month	Tap on the month name on the timeline.	"Select documents of the month <i>May</i> "
Select documents per tag	On the tag list, tap on the tag name	"Select documents with the tag <i>money</i> "
Deselect document	Tap on the header of the selected document	"Deselect 12"
Deselect all	Long press on canvas to open the context menu. Tap "Deselect."	"Deselect all"
Sort	Long press on canvas to open the context menu. Tap "Sort." Then, tap on sort parameter.	"Sort by <i>date</i> "
Resize document	Drag bottom-right corner of an opened document.	"Make document 12 smaller"
Move a document	Drag document by its header.	"Move document 12 to column A, row 3"
Move multiple documents	Drag selected documents by dragging one of them.	"Move selected to column A, row 3"
Navigate in document	Drag scrolling bar inside document.	"Scroll <i>down</i> in document 12"
Mark document	Tap on sentence (toggle).	"Mark [sentence 1 in] document 12" (first by default)
Clear sentence	Tap on sentence (toggle).	"Clear sentence 1 in document 12"
Clear document	Tap on sentences (toggle).	"Clear document 12"
Search	Tap phrase on keyboard. Then, tap on "search."	"Search for <i>flowers</i> "
Clear search	Tap on "clear" on keyboard.	"Clear search"
Add tag	Type tag name on keyboard. Then, tap on "add tag."	"Add the tag <i>flowers</i> "
Assign tag to document	Select document. Tap on O icon next to the tag.	"Assign tag <i>flowers</i> to the selected documents"
Remove tag from document	Select tagged document. Tap on O icon next to the tag.	"Remove tag <i>flowers</i> from the selected documents"
Delete tag	On the tag list, tap on X icon next to the tag.	"Delete the tag <i>flowers</i> "

based on prior work on multimodal and collaborative interaction [27, 28, 36, 47]. When the system recognized a speech command, it provided feedback by showing a message on the corresponding bar at the bottom of the screen (e.g., "Ok, I selected document 6"). If the recognition failed, the message "Sorry, I did not understand that" appeared.

4.2 Setup

Our wall display was 5.91×1.96 meters in size, composed of 75 liquid-crystal displays (LCDs) with a total resolution of 14400×4800 pixels (at 60 ppi), controlled by a cluster of 10 computers. The wall was placed in a room of 31.5 m^2 , and it recognized touch input via infrared light. For recording the position of the participants, each of them wore a headband tracked by a VICON motion tracking system. Each participant had a wireless mouse and a microphone. They could activate the speech recognition by clicking on the mouse.

TOUCHTALKINTERACTIVE is a web-based application, implemented in Typescript and Svelte. The documents are provided to the system as JSON files. We implemented the speech recognition on Rhino, a deep learning speech-to-intent engine by Picovoice [41]. We trained an English-based model according to the vocabulary used in the dataset documents to enable custom speech commands, with a reported accuracy above other commercial solutions [42]. The open-source code is available at <https://gitlab.inria.fr/aviz/TouchTalkInteractive>.

5 RESULTS

We recruited 20 participants (five identified as women, and fifteen as men) distributed into 10 pairs. We recruited the pairs from within a research organization, as the goal was to recruit people who would interact with data as part of their work. Participants were required to be at least 18 years old and fluent English speakers. The study was approved by the corresponding IRB (Inria COERLE, opinion 2023-39). Participants worked or studied at the institution, but had diverse backgrounds. We had participants from Europe, the Americas, and Asia. However, cultural background was not a factor we controlled.

The average age was 27 years. All participants reported interacting with touch-based systems daily. Nine persons used speech interaction at least monthly, while four had never used it before. For 14 participants, this was the first time interacting with a wall display. As the task was similar to solving a puzzle, we asked how often participants played such games. Eighteen participants had played at least once, seven did so monthly or daily. We did not require that participants knew each other before the experiment to get a diverse sample of collaborators. Six pairs

were familiar with each other before participating and four were not. Sixteen participants reported to work closely with others at least once a week. All pairs were able to find important clues about the mystery. At the end of the interaction phase, several groups expressed that they would have liked to have more time with the task. This suggests that people engaged with the game-like setting.

5.1 Input Modalities per Action

In total, we registered 4,020 interactions from the 20 participants. Of those, 3,633 were touch gestures and 387 were speech commands. Given that some actions happened more often than others (e.g., moving documents), we analyzed the interactions per action **RQ1**. The system supported 20 actions. However, we excluded interactions for the action *select documents per month* from our analysis (77 interactions), as we had technical problems with that feature in the first two sessions.

Participants were free to choose how to interact to solve the sense-making task, after having tried all the interaction techniques with both modalities in the training session. Overall, participants made use of all the interaction techniques provided. The most used ones were searching for keywords, opening, closing, and moving documents, as well as highlighting and clearing sentences of interest in the documents. Most groups created and assigned tags to documents, but only a few removed or deleted those tags. Participants selected documents one by one more often than per tag or month. Similarly, they deselected individually more often than *deselecting all*.

In Fig. 3a, we present the recorded interactions distributed across actions, with the proportion of modalities used. We show only the successful interactions (e.g., without speech recognition errors) where the system interpreted the action correctly. Participants mainly used touch interaction, but there were exceptions among the actions. Participants sorted the documents, cleared highlighted documents, and used *deselect all* more often with speech interaction. To sort, nine participants used speech only, and two used touch (once in both cases). To deselect all, seven participants used only speech, and only one person used touch. To search for keywords, 18 participants used the touch-based keyboard, while six used speech commands more often than touch. To open and close documents, all except one participant used touch gestures more often than speech. To select single documents, one person used only speech while everyone else used touch more often (or only). The diversity of the choices suggests that participants were able to use both modalities, but made different personal choices. For instance, while one participant (P12) interacted with speech commands throughout the




Fig. 3: Use and preferences of input modalities per action.

whole session, another participant (P20) only used speech once.

In Fig. 3b, we present the modality preferences of our participants after interacting with the system. Most participants preferred speech interaction to sort the documents. For keyword search, more than half of the participants preferred speech or favored touch and speech equally. To clear search results, half of the participants preferred speech or both. For the remaining actions, touch was preferred. To select documents with a given tag, seven participants preferred touch, but six wished to do it with either modality. Some participants made remarks about their reasoning when they chose the *I prefer them equally* answer. For example, P2 pointed out that choosing between touch and speech for closing a document depended on their distance from the document. Similarly, for moving documents, the choice would depend on the distance between the current and the desired position of the document on the wall display. Another participant also referred to the distance to the wall and to the document as a decisive factor in choosing how to open a document. For two participants, speech would be the first choice to search for a keyword, and if there would be a recognition error, they would switch to the touch-based keyboard. Another participant considered that searching by touch took too much time.

5.2 Personality Traits

Based on their answers to the personality questions in the pre-questionnaire, we divided participants into three groups per trait: those with low, high, or average scores, in comparison to the sample mean and standard deviation, as suggested in the literature [24, 59]. Seven participants scored high on extraversion, six were average, and seven scored low. Seven participants scored high on openness to new experiences, eight were average, and five were low. Nine participants scored high on agreeableness, three were average, and eight scored low. We compared the scores of the three personality traits against the number




Fig. 4: Scatterplots visualizing the scores from each personality trait in relation to the total count of speech commands per participant.

of times each participant attempted to use speech interaction (regardless of the action and outcome), to examine whether those traits could relate to the tendency of using speech interaction. We visualize the relations in Fig. 4. We found a significant relation between agreeableness and speech using Spearman’s correlation ($\rho = -0.445, p = .0495$), suggesting a moderately strong relationship between how agreeable a participant was and how often they used speech interaction. The more cooperative a person was, the fewer speech commands they used, potentially to avoid interrupting their partner, given that eight participants reported being afraid to distract their partner in the post-questionnaire. In contrast, the correlation between extraversion and speech use was not significant ($\rho = 0.016, p = .947$). Furthermore, contrary to our expectations, speech use and openness to experience correlated negatively, although not significantly ($\rho = -0.351, p = .129$).

5.3 Speech Interaction from a Distance

We compared the frequency of speech commands with the distance of the participants from the wall display, based on the interaction logs and the tracking data we collected during the study. We present the distribution of the participant positions per modality in Fig. 5.

While touch interaction required standing at a close distance, participants could interact with speech from any position. The data we collected about the participants’ movements suggest they tended to stay away from the wall while using speech commands. On average, they stood at 1.52 meters from the display when interacting with speech and at 0.59 meters when interacting with touch. Furthermore, when a participant used a speech command, the average distance to their partner was 1.88 meters ($std = 0.86$) and the median was 1.80 m. In contrast, when someone interacted with touch gestures, the distance between partners varied more, with an average of 1.94 meters ($std = 1.25$) and a median of 1.52. While the larger variation of the distances associated with touch may be due to the higher frequency of touch gestures overall, the distances indicate that participants did not intentionally go away from their partner to use speech commands.

5.4 Collaboration Styles

To answer RQ2, we qualitatively coded the collaboration styles of the participants throughout the sessions according to the codes proposed by Isenberg et al. [27]. Overall, participants spent more time in closely coupled collaboration than in loosely coupled collaboration. On average, the groups spent 54% of the session collaborating closely and 36% loosely. The remaining 10% of the time corresponds to periods of no clear collaboration, e.g., when participants talked to the experimenter.

All pairs reported to have divided tasks among the two persons while solving the mystery. In the post-questionnaire, we asked each participant to estimate the proportion of time they spent working together with

their partner, alone researching a shared question, and alone researching their own question. On average, participants reported to have spent 31% of the time working with their partner ($std = 21.05$), 31% alone researching a shared question ($std = 19.43$), and 38% researching their own question ($std = 22.85$). The 31% of alone-time does not fully match our impressions from the coding, but that may relate to having talked to their partner before focusing on one question (which we coded as loose collaboration). When asked about their effectiveness as a team, nine out of the 10 pairs considered that they worked together effectively or very effectively to solve the mystery.

Figure 6 (a) shows the average distance to the wall of each participant per collaboration style. Most participants stood further away from the wall while collaborating closely than while collaborating loosely. However, the difference was small, given that, on average, participants tended to stand 1.02 m. away from the display in close collaboration, and 0.98 m. away while collaborating loosely. Most groups discussed and worked on the same specific problems while standing slightly further from the screen, potentially to look at the overview together (they did not say it explicitly). The most common type of close collaboration was *active discussion*, which refers to when participants discussed their findings and hypotheses about the mystery they were trying to solve. Meanwhile, the most common form of loose collaboration was working on different problems. Participants worked independently while standing closer to the screen, probably reading a document, while their partner interacted with another document related to a different problem.

We also examined how the distance between the two participants changed across collaboration styles. Figure 6 (b) compares the average distance between closely and loosely coupled collaboration per pair. We found that the teams tended to be closer to each other during closely coupled collaboration than during loosely coupled collaboration. That difference was clearer and larger than in the case of the distance to the wall display. On average, participants stood at 1.45 meters from each other while collaborating closely and at 1.90 m. from each other while collaborating loosely. Therefore, they came closer or farther by half a meter. The largest difference corresponds to the pair of participants P9 and P10. Overall, they collaborated closely most of the time, but when they collaborated loosely — mainly at the end of the session — they tended to go to opposite ends of the display.

5.4.1 Modality Choices During Collaboration

We examined the relationship between input modalities and collaboration styles to find out whether participants tended to choose a specific modality while collaborating. Surprisingly, participants interacted similarly across collaboration styles. From the interactions that overlapped with a collaboration style, 53% of the touch gestures happened during close collaboration and 47% during loose collaboration. Regarding speech commands, it was 49% and 51%. Although speech interactions constituted less than 10% of all interactions, they were evenly performed across collaboration styles, suggesting that speech commands were deemed appropriate in both loose and close collaboration.

5.4.2 Awareness

In the post-questionnaire, all but one participant claimed to be aware of what their partner was doing, and 18 thought their partner was aware of their actions. Nine participants reported to have looked across the wall often to find out what their partner was doing. However, six participants from six different pairs had the impression that their partner




Fig. 5: Distribution of the participant positions in the space in front of the wall display while interacting, separated per modality.




Fig. 6: Slope charts comparing: (a) the distance to the wall of each participant across collaboration styles. (b) The distance between participants of each pair across collaboration styles.

often found important information that would have helped them, but they only learned about it later in the session.

We asked participants whether they were afraid to use speech commands as they may annoy their partner. The answers suggest that it was a concern, although not shared by the majority: Eight participants were concerned, while 11 disagreed and one was undecided. When on the other side participants were asked about being annoyed by the speech commands of their partner, 14 participants disagreed, and only three were actually annoyed by the commands of their partner during the study. Furthermore, 14 participants commented positively about having learned about what their partner was doing from their speech commands, suggesting that speech interaction helped them be aware of the actions of their partner. When we asked about the different ways they learned about what their partner was doing, participants learned by hearing the speech commands 5% of the time on average. They knew by noticing what the partner was touching on the screen 8% of the time. The two most common answers were that the partner explicitly said it (24%) and that they knew about the action because they performed it together (25%). These answers corresponded to the groups spending most of the time in close collaboration, actively discussing.

5.5 Interviews

After the participants had completed their main task on the wall display, we conducted a semi-structured interview with each pair to discuss their hypotheses about the mystery, their impressions on the input modalities, and how those may have affected teamwork and awareness.

Impressions on input modalities: Of the 20 participants, 16 found touch more effective than speech interaction, three found both equally effective and one person found speech more effective than touch. Three participants mentioned that touch felt more familiar, and two referred to direct manipulation regarding using their fingers and getting immediate visual feedback (e.g., marking sentences in a document). Participants especially highlighted that moving documents and scrolling on them was faster and more precise with touch gestures.

Although most participants found touch more suitable overall, nine mentioned specific actions for which they found speech more appropriate. While participant P18 said “*it’s something about local versus global*”, P20 used the term *abstract actions* to describe operations that would benefit from speech interaction, such as sorting. Two participants found speech best for actions involving multiple elements, and another

one found it helpful for moving documents all the way across the screen. Others considered speech useful for keeping an overview by interacting from a distance. P17 indicated that “*if you make a search command, you can already see everything without having to move around*”. Unsurprisingly, accuracy was the most common issue participants mentioned concerning speech. Interestingly, seven of them mentioned that not being native English speakers played a role. At the same time, five participants explicitly separated the concept of speech interaction from the recognition accuracy they experienced. P19 expressed:

“I’m not that confident in the fact that it will recognize what I want (...) but my feeling is that in an ideal setup, like, maybe if the voice commands were in French, I would have used a lot more of voice commands.”

According to P18, “*maybe just the quality of speech recognition was not bad. In fact, quite good. But still (...) it’s also about accents.*” Participants also suggested additional interactions and other modalities. To search, some wished for the virtual keyboard to move with them through proxemics and others wished for a physical keyboard to type more quickly. Two participants wished for an interaction technique to cluster documents of interest by automatically rearranging them in a screen region and moving the remaining documents to other areas.

Teamwork: We asked participants whether and how the input modalities affected their teamwork. Participants discussed the advantages and disadvantages of both modalities. Some found touch helpful for highlighting sentences to silently indicate what they found relevant in a document to their partner. Yet, two participants found touch problematic because it required possibly interrupting their partner by getting close to them to interact with a document nearby. Thus, one of them preferred using speech to move documents away. However, two participants avoided using speech so as not to bother their partner, which fits the post-questionnaire answers where eight participants expressed such a concern. In contrast, four participants explained that speech was not an obstacle because they felt the difference between talking to each other and talking to the system was clear. The reasons included that speech commands required a click with the wireless mouse and each person had a microphone, so there was no need to speak loudly. P13 said “*the difference between you saying a command and you addressing to me was quite clear. So it was okay if you were mumbling on your side.*” While some participants did not notice the speech commands of their partners, others said that they could not avoid listening. Two persons associated touch interaction with working on their own, e.g., “*because the interaction makes us look at the wall and not at each other.*”

Awareness: Listening to the speech commands of the other person was often associated with awareness. Five participants mentioned that listening to their partner talking to the wall helped them to have an idea of what keywords or documents the person was focusing on, which, at the same time, could be distracting:

“The advantage of the speech interaction might be we can know what my partner is looking for because I cannot avoid to listen (...) so when he says something and I heard that, I will realize something about what he is now looking for”.

P19 mentioned that they sometimes looked at the feedback bar at the bottom of the screen to see what keywords their partner searched for. Two participants considered touch and body movements helpful as the position of their partner would give them a hint of what documents they were working on. Eight participants expressed that talking to their partner helped them to know what they were doing.

5.6 Design Considerations

Based on our quantitative and qualitative findings, we draw the following design considerations for multimodal collaborative systems:

1. **Speech interaction for global operations.** Speech was preferred for interactions that influenced the complete dataset, such as sorting and deselecting all documents. The interaction choices of our participants indicate that the findings of Srinivasan et al. [51] about individuals leveraging speech commands for global actions also apply to collaborative scenarios.

2. **Speech interaction supports closely and loosely coupled collaboration.** A likely concern about speech interaction is that it may disrupt collaborative work. However, speech was not an obstacle to collaborating closely or loosely, and participants did not walk away from each other to use speech commands. Moreover, participants used half of their speech commands during loosely coupled collaboration and the other half during close collaboration, suggesting that speech was appropriate in both cases.
3. **Speech interaction contributes to the awareness of the partner.** According to the post-questionnaire answers of our study, 14 of 20 participants became aware of the actions of their partner through speech commands. While some participants agreed with being concerned about annoying their partner by using speech, only three of them actually reported being annoyed by the speech commands. Therefore, including speech interaction may have a positive effect on collaborative work by increasing awareness. Given that some participants found it distracting, we should attempt to minimize distraction by considering what interactions to associate voice with, but in practice, our results indicate that likely the distraction is not high.
4. **Personal characteristics influence interaction choices.** From the three personality traits we examined, we found a significant relation between how agreeable a person was and the number of speech commands they used. Thus, it is important to provide alternatives so that people who do not feel comfortable with speech can still use touch. Furthermore, two of our participants pointed out that being non-native speakers made them hesitant to use speech, which suggests that in a sample of native speakers only, participants may be more likely to leverage speech. Given the different factors that may influence a person’s choice to use speech commands, speech interaction should not be the only input modality to support collaborative work.
5. **Speech interaction technologies can support probe-based research.** Our study indicated that it is possible with current technologies to create speech-based interaction prototypes for collaborative work. For the exploratory purpose of our study, we used a deep-learning solution that allowed us to recognize custom commands according to our dataset. The speech engine [41] processed the speech locally, which allowed us to avoid sending the audio to an external server in favor of local data privacy protection regulations. However, recognition errors still happened, so other solutions should be tested. Using a wireless device to activate the speech recognition was well received by the participants.

6 DISCUSSION AND FUTURE WORK

Speech interaction is considered a promising input modality because it allows people to express what they wish for, instead of taking the time to translate it into code, formulas, or widget input, and has the potential to make interactive visualizations more accessible [21]. However, it often suffers from a lack of trust due to recognition errors [6] and past work has shown that it can discourage parallel work [56]. Our findings show that speech commands can support distant interaction during collaboration without major issues but tend to be used only for specific actions and still suffer from technological challenges.

We used up-to-date tools powered with deep learning to enable speech interaction, but recognition errors were still a common problem, and two participants seemed to give up on speech after some errors. For instance, P17 said in the interview: “*I believe I stopped using speech because I couldn’t get it to understand*”. Compared with the strong preference for speech command interaction on large displays in a previous study [36], it seems that the recognition accuracy of current solutions still needs to catch up to support the use of speech commands as often as people actually may want to use them. As P8 indicated, “*for example, [commercial voice assistant], they worked on that many years and they’re, they still have problems*”. In addition, people may take speech recognition issues personally. In our study, a common concern among participants was that a foreign accent may be the cause of the errors. That may be plausible because the only native speaker had only

four errors during the whole session (the maximum was 36), but a more diverse sample is needed to assess the effect. Overall, there were 150 speech errors. In case of an error, participants switched to touch in 35 cases. In 61 cases, the person insisted on using speech until being successful. In 14 cases, the person tried using speech up to four more times before switching to touch. Another unexpected challenge we faced during the system implementation was that the speech engine did not recognize multiple names as English words. For example, the town name in the original dataset was Alderwood, but the engine did not recognize it. Thus, we had to modify some terms from the dataset.

6.1 Facilitating Specific Actions

Despite the challenges, participants clearly favored speech interaction for global operations such as sorting and *deselect all*. Moreover, search commands via speech were common and were even part of the conversation within some teams. P19 indicated that they started working on the task by discussing what to focus on and used speech: “*because we were together, it was like implicit that we were talking to the wall (...) then we went apart and it was not that clear and we focused on touch*”. The sensemaking task promoted collaborative work but relied heavily on browsing and reading multiple documents. Some participants argued that using a speech command to open or move a document was cumbersome because it required first identifying the document by its number. That suggests that speech interaction is more suitable for group or set-related actions, such as in prior work about data exploration by individuals, e.g., encoding an attribute with color [53]. Some participants indicated that speech commands helped them interact with multiple documents that were either far from the person or far from each other on the screen. In the preceding elicitation study [36], participants wished to filter data, change the visual encoding, toggle brushing & linking, and merge and split multiple views via speech. These operations involved the whole dataset or all views, corresponding to global operations and reaching far-away elements. In the context of visual analytics, advanced operations such as invoking automated analysis methods on the dataset may be a good fit for speech interaction.

While most participants were not disturbed by the speech commands of their partner, these commands represent less than 10% of the logged interactions. To put that into context, 15% of the interactions correspond to participants moving documents via touch. Maybe there was little conflict because the speech commands were so few in contrast to touch gestures, and, therefore, it was unlikely that both persons would try to talk to the wall simultaneously. Global operations, such as sorting, may indeed happen less often than local operations. Thus, speech may be a modality used less often than others that suit more common actions, e.g., touch for scrolling. A follow-up question would be whether there is a collaborative scenario where speech commands would prevail. Then, the likelihood of collaboration conflicts may rise.

6.2 Supporting Co-located Collaboration

Regarding the characteristics of collaborative work, our study focused on examining how speech interaction may affect group communication, the distance between team members, and awareness. In contrast to the findings of Tse et al. [56] on tabletops, our participants did not find speech interaction problematic for working in parallel. One reason may be the larger space to interact in front of the wall display. Another one may be the speaking volume, given that some participants referred to not having to speak loudly with the wireless microphone. Given the concern of some participants about distracting others through speech interaction, one may consider including dedicated awareness features. For example, Cernea et al. [10] proposed visualizing emotions through *emotion-prints* to increase emotional awareness among collaborators. Such techniques may help assess whether specific actions negatively affect any collaborator. There are also other social aspects relevant to collaboration that may be affected by speech interaction. Future work should investigate how speech may affect privacy and territoriality [3]. While our participants did not report any inconveniences, working separately and keeping interactions private may be more relevant in other collaborative scenarios. For instance, P3 said that touch may be more appropriate for “*privacy, maybe. Not particularly in this context, but*

(...) if you work on sensitive data.” Participants appreciated that the speech recognition had to be manually activated to avoid unintended interactions with the system while talking to each other. Regarding territoriality, previous work found that the size of a personal space changes across cultures [19]. Therefore, the distance between our participants may have been influenced by their cultural background. However, we did not control cultural backgrounds and can, therefore, not provide supporting data. Moreover, we had participants with similar cultural backgrounds and different behaviors. Thus, personality differences may be more important. Still, speech allows distant interaction with elements in the personal space of others. Future work could investigate the relationship between speech interaction and social norms.

While the interaction techniques that our system supported used one modality at a time, future work should consider the combination of two modalities in one action. Our participants suggested combining touch and speech for some actions, such as tapping on a document and then saying “*Tag this with police*” or tapping on an empty screen space followed by “*Bring [document ID] over here*” to move documents. These interactions could also be collaborative, i.e., performed jointly by two persons. However, such interactions require good coordination and may come with technical challenges, such as distinguishing which action comes from which person. Moreover, the chance of having overlapping voices increases with the number of collaborators, which may make speech recognition more challenging.

Finally, we did not focus on performance, as our study was rather exploratory. An important future research direction would be to investigate how the use of speech commands may influence performance in the context of collaborative work. In the interview, participants often favored touch over speech due to speed and accuracy. However, the participant with the highest number of successful speech commands argued that getting used to speech helped to use it more. Familiarity may thus play an important role in ensuring performance via speech.

Which other visual analytics tasks would benefit from leveraging speech? Future work should systematically explore interaction taxonomies to determine when is speech suitable. Moreover, with larger groups, it would be possible for two persons to talk while a third one interacts via speech. Would speech become a significant obstacle to communication then? With our findings, we hope to encourage researchers and practitioners to consider speech interaction for interactive systems not only for individuals but also for collaborative scenarios.

7 CONCLUSION

We conducted an in-depth exploratory study on combining touch gestures and speech commands for collaborative sensemaking to extend our understanding of whether and how speech interaction can be leveraged for co-located collaborative work. Our findings provide evidence that speech commands are a viable option to support global operations and distant interaction from the wall display. The differences between modality preferences and actual use suggest that in concordance with previous work [36], people are interested in using speech interaction, but other aspects, such as speech recognition errors and the lack of familiarity, may influence their interaction choices on implemented systems. Furthermore, we found the tendency to use speech commands correlated with the personality trait of agreeableness. Future work can investigate whether the choice to leverage other input modalities may be influenced by personal characteristics. Regarding the communication between team members, we did not find enough evidence that speech interaction was an obstacle to collaboration. On the contrary, 14 of 20 participants reported that speech commands contributed to becoming aware of their partner’s actions. Based on our findings, we introduced design considerations for incorporating speech commands to collaborative systems and highlight opportunities for future research. Our collaborative visual analytics prototype, TOUCHTALKINTERACTIVE, which combines touch and speech interaction for collaborative sensemaking, is openly available together with our supplemental material.

SUPPLEMENTAL MATERIALS

All supplemental materials are available on OSF at <https://osf.io/8gpv2>, released under a CC BY 4.0 license. In particular, they include

(1) the questionnaires, (2) the answers to the questionnaires, and (3) the interaction logs. The open-source code of the system is available at <https://gitlab.inria.fr/aviz/TouchTalkInteractive>.

ACKNOWLEDGMENTS

This work was supported by a fellowship of the German Academic Exchange Service (DAAD) as well as by the French National Research Agency (ANR) under the Investments for the Future program (PIA) grant ANR-21-ESRE-0030 (CONTINUUM). An earlier version of this paper is part of the dissertation of the first author [35].

REFERENCES

- [1] C. Andrews, A. Endert, and C. North. Space to think: Large high-resolution displays for sensemaking. In *Proc. CHI*, pp. 55–64. ACM, New York, 2010. doi: 10.1145/1753326.1753336 3, 4
- [2] C. Andrews, A. Endert, B. Yost, and C. North. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization*, 10(4):341–355, Oct. 2011. doi: 10/fbmqct 1
- [3] A. Azad, J. Ruiz, D. Vogel, M. Hancock, and E. Lank. Territoriality and behaviour on and around large vertical publicly-shared displays. In *Proc. DIS*, pp. 468–477. ACM, New York, 2012. doi: 10.1145/2317956.2318025 9
- [4] S. K. Badam, F. Amini, N. Elmqvist, and P. Irani. Supporting visual exploration for multiple users in large display environments. In *Proc. IEEE VAST*, pp. 1–10. IEEE Computer Society, Los Alamitos, 2016. doi: 10.1109/VAST.2016.7883506 2
- [5] R. Ball, C. North, and D. A. Bowman. Move to improve: Promoting physical navigation to increase user performance with large displays. In *Proc. CHI*, pp. 191–200. ACM, New York, 2007. doi: 10.1145/1240624.1240656 1, 2
- [6] A. Baughan, X. Wang, A. Liu, A. Mercurio, J. Chen, and X. Ma. A mixed-methods approach to understanding user trust after voice assistant failures. In *Proc. CHI*, article no. 7, 16 pages. ACM, New York, 2023. doi: 10.1145/3544548.3581152 8
- [7] I. Belkacem, C. Tominski, N. Médoc, S. Knudsen, R. Dachsel, and M. Ghoniem. Interactive visualization on large high-resolution displays: A survey. *Computer Graphics Forum*, n/a:e15001, 2024. To appear. doi: 10.1111/cgf.15001 2
- [8] A. Bezerianos and P. Isenberg. Perception of visual variables on tiled wall-sized displays for information visualization applications. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2516–2525, Oct. 2012. doi: 10.1109/TVCG.2012.251 4
- [9] F. Brudy, J. K. Budiman, S. Houben, and N. Marquardt. Investigating the role of an overview device in multi-device collaboration. In *Proc. CHI*, article no. 300, 13 pages. ACM, New York, 2018. doi: 10.1145/3173574.3173874 3
- [10] D. Cernea, C. Weber, A. Ebert, and A. Kerren. Emotion-prints: interaction-driven emotion visualization on multi-touch interfaces. In D. L. Kao, M. C. Hao, M. A. Livingston, and T. Wischgoll, eds., *Proc. Visualization and Data Analysis*, vol. 9397, p. 93970A. International Society for Optics and Photonics, SPIE, Bellingham, USA, 2015. doi: 10.1117/12.2076473 9
- [11] O. Chapuis, A. Bezerianos, and S. Frantzeskakis. Smarties: An input system for wall display development. In *Proc. CHI*, pp. 2763–2772. ACM, New York, 2014. doi: 10.1145/2556288.2556956 1
- [12] P. T. Costa and R. R. McCrae. Manual for the revised neo personality inventory (NEO-PI-R) and neo five-factor inventory (NEO-FFI). *Psychological Assessment Resources*, 396, 1992. 3
- [13] C. G. DeYoung, R. G. Grazioplene, and J. B. Peterson. From madness to genius: The openness/intellect trait domain as a paradoxical simplex. *Journal of Research in Personality*, 46(1):63–78, Feb. 2012. doi: 10.1016/j.jrp.2011.12.003 3
- [14] T. Gao, M. Dontcheva, E. Adar, Z. Liu, and K. G. Karahalios. DataTone: Managing ambiguity in natural language interfaces for data visualization. In *Proc. UIST*, pp. 489–500. ACM, New York, 2015. doi: 10.1145/2807442.2807478 2
- [15] L. Grammel, M. Tory, and M.-A. Storey. How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):943–952, Oct. 2010. doi: 10.1109/TVCG.2010.164 2
- [16] T. M. Green and B. Fisher. Towards the personal equation of interaction: The impact of personality factors on visual analytics interface interaction. In *Proc. IEEE VAST*, pp. 203–210. IEEE, Los Alamitos, 2010. doi: 10.1109/VAST.2010.5653587 3
- [17] G. Grinstein, T. O’Connell, S. Laskowski, C. Plaisant, J. Scholtz, and M. Whiting. VAST 2006 contest - a tale of alderwood. In *Proc. IEEE VAST*, pp. 215–216. IEEE, Baltimore, 2006. doi: 10.1109/VAST.2006.261420 3
- [18] Y. Guo, D. Shi, M. Guo, Y. Wu, N. Cao, and Q. Chen. Talk2Data: A natural language interface for exploratory visual analysis via question decomposition. *ACM Trans. Interact. Intell. Syst.*, 14(2), article no. 8, 24 pages, Apr. 2024. doi: 10.1145/3643894 2
- [19] E. T. Hall. *The Hidden Dimension*, vol. 609. Anchor, New York, 1966. 9
- [20] K. Hinckley, K. Yatani, M. Pahud, N. Coddington, J. Rodenhouse, A. Wilson, H. Benko, and B. Buxton. Pen + touch = new tools. In *Proc. UIST*, pp. 27–36. ACM, New York, 2010. doi: 10.1145/1866029.1866036 2
- [21] E. Hoque, V. Setlur, M. Tory, and I. Dykeman. Applying pragmatics principles for interaction with visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):309–318, 2018. doi: 10.1109/TVCG.2017.2744684 8
- [22] T. Horak, S. K. Badam, N. Elmqvist, and R. Dachsel. When David meets Goliath: Combining smartwatches with a large vertical display for visual data exploration. In *Proc. CHI*, article no. 19, 13 pages. ACM, New York, 2018. doi: 10.1145/3173574.3173593 1, 2
- [23] H. Hutchinson, W. Mackay, B. Westerlund, B. B. Bederson, A. Druin, C. Plaisant, M. Beaudouin-Lafon, S. Conversy, H. Evans, H. Hansen, N. Roussel, and B. Eiderbäck. Technology probes: Inspiring design for and with families. In *Proc. CHI*, pp. 17–24. ACM, New York, 2003. doi: 10.1145/642611.642616 4
- [24] IPIP Project: ORI Community and Evaluation Services. Interpreting individual IPIP scale scores, 2023. <https://iip.ori.org/InterpretingIndividualIPIPScaleScores.htm>, last accessed February 2, 2024. 6
- [25] P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization Journal*, 10(4):310–326, Oct. 2011. doi: 10.1177/1473871611412817 1
- [26] P. Isenberg and D. Fisher. Collaborative brushing and linking for co-located visual analytics of document collections. *Computer Graphics Forum*, 28(3):1031–1038, 2009. doi: 10.1111/j.1467-8659.2009.01444.x 4
- [27] P. Isenberg, D. Fisher, M. Ringel Morris, K. Inkpen, and M. Czerwinski. An Exploratory Study of Co-located Collaborative Visual Analytics around a Tabletop Display. In *Proc. VAST*, pp. 179–186. IEEE, Los Alamitos, 2010. doi: 10.1109/VAST.2010.5652880 3, 4, 5, 6
- [28] M. R. Jakobsen and K. Hornbæk. Up close and personal: Collaborative work on a high-resolution multitouch wall display. *ACM Transactions on Computer-Human Interaction*, 21(2), article no. 11, 34 pages, Feb. 2014. doi: 10.1145/2576099 3, 4, 5
- [29] R. James, A. Bezerianos, and O. Chapuis. Evaluating the extension of wall displays with ar for collaborative work. In *Proc. CHI*, article no. 99, 17 pages. ACM, New York, 2023. doi: 10.1145/3544548.3580752 2
- [30] J.-F. Kassel and M. Rohs. Valletto: A multimodal interface for ubiquitous visual analytics. In *Extended Abstracts of CHI*, article no. LBW005, 6 pages. ACM, New York, 2018. doi: 10.1145/3170427.3188445 2
- [31] R. Langner, U. Kister, and R. Dachsel. Multiple coordinated views at large displays for multiple users: Empirical findings on user behavior, movements, and distances. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):608–618, 2019. doi: 10.1109/TVCG.2018.2865235 2
- [32] B. Lee, P. Isenberg, N. H. Riche, and S. Carpendale. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2689–2698, 2012. doi: 10.1109/TVCG.2012.204 2
- [33] Z. Liu, R. J. Crouser, and A. Ottley. Survey on individual differences in visualization. *Computer Graphics Forum*, 39(3):693–712, 2020. doi: 10.1111/cgf.14033 3
- [34] J. L. Maples-Keller, R. L. Williamson, C. E. Sleep, N. T. Carter, W. K. Campbell, and J. D. Miller. Using item response theory to develop a 60-item representation of the neo pi-r using the international personality item pool: Development of the ipip-neo-60. *Journal of Personality Assessment*, 101(1):4–15, 2019. doi: 10.1080/00223891.2017.1381968 3
- [35] G. Molina León. *Multimodal and Collaborative Interaction for Visual Data Exploration*. PhD thesis, University of Bremen, Germany, Apr. 2024. doi: 10.26092/elib/2967 10

- [36] G. Molina León, P. Isenberg, and A. Breiter. Eliciting multimodal and collaborative interactions for data exploration on large vertical displays. *IEEE Transactions on Visualization and Computer Graphics*, 30(2):1624–1637, Feb. 2024. doi: 10.1109/TVCG.2023.3323150 2, 4, 5, 8, 9
- [37] B. A. Myers, R. Bhatnagar, J. Nichols, C. H. Peck, D. Kong, R. Miller, and A. C. Long. Interacting at a distance: Measuring the performance of laser pointers and other devices. In *Proc. CHI*, pp. 33–40. ACM, New York, 2002. doi: 10.1145/503376.503383 1
- [38] M. Nancel, J. Wagner, E. Pietriga, O. Chapuis, and W. Mackay. Mid-air pan-and-zoom on wall-sized displays. In *Proc. CHI*, pp. 177–186. ACM, New York, 2011. doi: 10.1145/1978942.1978969 1, 2
- [39] A. Narechania, A. Srinivasan, and J. Stasko. NL4DV: A toolkit for generating analytic specifications for data visualization from natural language queries. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):369–379, 2021. doi: 10.1109/TVCG.2020.3030378 2
- [40] T. Neumayr, H.-C. Jetter, M. Augstein, J. Friedl, and T. Luger. Domino: A descriptive framework for hybrid collaboration and coupling styles in partially distributed teams. *Proc. of the ACM on Human-Computer Interaction*, 2(CSCW), article no. 128, 24 pages, Nov. 2018. doi: 10.1145/3274397 3
- [41] Picovoice Inc. Train, develop and deploy custom voice features - Picovoice, July 2023. <https://picovoice.ai/>, accessed July, 2023. 5, 8
- [42] Picovoice Inc. Natural language understanding benchmark, Mar. 2024. <https://picovoice.ai/docs/benchmark/nlu/>, accessed: March, 2024. 5
- [43] Y. Qu and D. L. Hansen. Building shared understanding in collaborative sensemaking. In *Proc. CHI Sensemaking Workshop*, 2008. 1
- [44] P. Reipschläger, T. Flemisch, and R. Dachsel. Personal augmented reality for information visualization on large interactive displays. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1182–1192, 2021. doi: 10.1109/TVCG.2020.3030460 3
- [45] P. Riehmman, G. Molina León, J. Reibert, F. Ehtler, and B. Froehlich. Short-contact touch-manipulation of scatterplot matrices on wall displays. *Computer Graphics Forum*, 39(3):265–276, 2020. doi: 10.1111/cgfm.13979 2
- [46] J. C. Roberts, P. D. Ritsos, S. K. Badam, D. Brodbeck, J. Kennedy, and N. Elmqvist. Visualization beyond the desktop—the next big thing. *IEEE Computer Graphics and Applications*, 34(6):26–34, Nov. 2014. doi: 10.1109/MCG.2014.82 1
- [47] A. Saktheeswaran, A. Srinivasan, and J. Stasko. Touch? Speech? Or touch and speech? Investigating multimodal interaction for visual network exploration and analysis. *IEEE Transactions on Visualization and Computer Graphics*, 26(6):2168–2179, June 2020. doi: 10.1109/TVCG.2020.2970512 1, 2, 4, 5
- [48] V. Setlur, S. E. Battersby, M. Tory, R. Gossweiler, and A. X. Chang. Eviza: A natural language interface for visual analysis. In *Proc. UIST*, pp. 365–377. ACM, New York, 2016. doi: 10.1145/2984511.2984588 2
- [49] Shneiderman. Direct manipulation: A step beyond programming languages. *Computer*, 16(8):57–69, Aug. 1983. doi: 10.1109/MC.1983.1654471 1
- [50] A. Srinivasan, B. Lee, N. Henry Riche, S. M. Drucker, and K. Hinckley. Inchorus: Designing consistent multimodal interactions for data visualization on tablet devices. In *Proc. CHI*, article no. 653, 13 pages. ACM, New York, 2020. doi: 10.1145/3313831.3376782 2
- [51] A. Srinivasan, B. Lee, and J. Stasko. Interweaving multimodal interaction with flexible unit visualizations for data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 27(8):3519–3533, Aug. 2021. doi: 10.1109/TVCG.2020.2978050 2, 4, 8
- [52] A. Srinivasan and J. Stasko. Natural language interfaces for data analysis with visualization: considering what has and could be asked. In *Proc. EuroVis*, pp. 55–59. Eurographics Association, Goslar, Germany, 2017. doi: 10.2312/eurovisshort.20171133 2
- [53] A. Srinivasan and J. Stasko. Orko: Facilitating multimodal interaction for visual exploration and analysis of networks. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):511–521, 2018. doi: 10.1109/TVCG.2017.2745219 2, 9
- [54] R. S. Tabalba, N. Kirshenbaum, J. Leigh, A. Bhattacharya, V. Grosso, B. Di Eugenio, A. E. Johnson, and M. Zellner. An investigation into an always listening interface to support data exploration. In *Proc. IUI*, pp. 128–141. ACM, New York, 2023. doi: 10.1145/3581641.3584079 2, 4
- [55] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *Proc. CHI*, pp. 1181–1190. ACM, New York, 2006. doi: 10.1145/1124772.1124950 3
- [56] E. Tse, S. Greenberg, C. Shen, C. Forlines, and R. Kodama. Exploring true multi-user multimodal interaction over a digital table. In *Proc. DIS*, pp. 109–118. ACM, New York, 2008. doi: 10.1145/1394445.1394457 3, 8, 9
- [57] S. T. Völkel, R. Schödel, D. Buschek, C. Stachl, V. Winterhalter, M. Bühner, and H. Hussmann. Developing a personality model for speech-based conversational agents using the psycholexical approach. In *Proc. CHI*, article no. 83, 14 pages. ACM, New York, 2020. doi: 10.1145/3313831.3376210 3
- [58] J. Walny, B. Lee, P. Johns, N. Henry Riche, and S. Carpendale. Understanding pen and touch interaction for data exploration on interactive whiteboards. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2779–2788, 2012. doi: 10.1109/TVCG.2012.275 2
- [59] C. Ziemkiewicz, A. Ottley, R. J. Crouser, A. R. Yauilla, S. L. Su, W. Ribarsky, and R. Chang. How visualization layout relates to locus of control and other personality factors. *IEEE Transactions on Visualization and Computer Graphics*, 19(7):1109–1121, July 2013. doi: 10.1109/TVCG.2012.180 3, 6