Datamancer: Bimanual Gesture Interaction in Multi-Display Ubiquitous Analytics Environments

Biswaksen Patnaik
Department of Computer Science
University of Maryland College Park
College Park, Maryland, USA
bpatnaik@umd.edu

Marcel Borowski
Department of Computer Science
Aarhus University
Aarhus, Denmark
marcel.borowski@cs.au.dk

Huaishu Peng
Computer Science
University of Maryland
College Park, Maryland, USA
huaishu@umd.edu

Clemens Nylandsted Klokmose
Department of Computer Science
Aarhus University
Aarhus, Denmark
clemens@cs.au.dk

Niklas Elmqvist
Department of Computer Science
Aarhus University
Aarhus, Denmark
elm@cs.au.dk

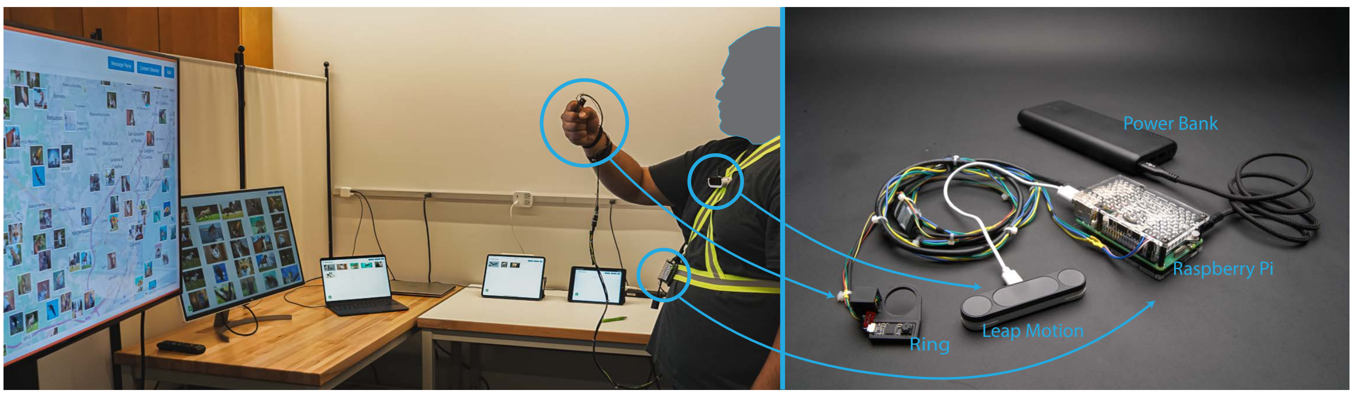

Figure 1: Datamancer. A person using Datamancer to interact with a large geographic map (indicated as the focused display by the orange border) with overlaid geolocated photographs (left image). The physical environment contains multiple screens, some of them being powered by separate personal computers or tablet devices. Each device can be acquired using a pointing gesture with the ring-mounted pinhole camera and then interacted with using bimanual gestures. The components of the Datamancer device (right image): the ring-mounted pinhole camera, the Leap Motion 2, the Raspberry Pi, and the power bank.

Abstract

We introduce Datamancer, a wearable device enabling bimanual gesture interaction across multi-display ubiquitous analytics environments. Datamancer addresses the gap in gesture-based interaction within data visualization settings, where current methods are often constrained by limited interaction spaces or the need for installing bulky tracking setups. Datamancer integrates a finger-mounted pinhole camera and a chest-mounted gesture sensor, allowing seamless selection and manipulation of visualizations on distributed displays. By pointing to a display, users can acquire the display and engage in various interactions, such as panning, zooming, and selection, using both hands. Our contributions include (1) an investigation of the design space of gestural interaction for physical ubiquitous analytics environments; (2) a prototype implementation of the Datamancer system that realizes this model; and (3) an evaluation of the prototype through demonstration of application scenarios, an expert review, and a user study.

CCS Concepts

• **Human-centered computing → Visualization systems and tools**; **Visualization application domains**; **Gestural input**.

Keywords

Gestural interaction, ubiquitous analytics, immersive analytics, Augmented Reality, visualizations, situated analytics.

ACM Reference Format:
Biswaksen Patnaik, Marcel Borowski, Huaishu Peng, Clemens Nylandsted Klokmose, and Niklas Elmqvist. 2025. Datamancer: Bimanual Gesture Interaction in Multi-Display Ubiquitous Analytics Environments. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3706598.3713123

ACM ISBN 979-8-4007-1394-1/25/04
https://doi.org/10.1145/3706598.3713123

1 Introduction

John Anderton in *Minority Report* (2002); Tony Stark in *Iron Man 3* (2013); Johnny Mnemonic in the eponymous film (1995)—all examples of fictional characters from Hollywood science fiction movies using gestural interaction to navigate and manipulate data. Clearly, the idea of using your arms and body to interact with visual interfaces is a popular, fascinating, even futuristic prospect for many [18]. However, while pointing and gesturing is a popular and much-studied interaction metaphor for general HCI [10, 26], it has not been nearly as closely investigated in the data visualization field [5]. In particular, current approaches for gestural interaction with visualizations are either limited to small interaction areas for a personal computer [37], fully instrumented large-display environments [2], or optical gesture tracking using a head-mounted display in immersive space [6]. To the best of our knowledge, no research has explored mobile untethered gesture detection for general multi-display analytics environments in the wild, such as a conference room full of laptops or a tiled display wall powered by separate computers.

We present Datamancer: a wearable device for tracking both pointing actions and bimanual gestures for universal interaction with distributed visualization display. The device consists of a finger-mounted pinhole camera and a chest-mounted gesture sensor powered by a Raspberry Pi. Devices involved in the data analysis session will run a visualization dashboard workspace accessed using a standard web browser and based on the web-based Webstrates software stack [35]. Each visualization display will be uniquely identified using a dynamic ArUco marker [56]. Given a room of displays, the user can select one by pointing the pinhole camera at the display and pressing a finger-mounted button, thus focusing it as the interaction target. Pressing the button again unfocuses it. While the display is focused, the user's dominant hand is used to switch between interaction types whereas the non-dominant hand is used for continuous six-degrees-of-freedom 3D input. Our current implementation of Datamancer supports 2D panning, zooming, selection, dragging, and dropping. Additional gestures, even user-defined ones, can easily be added.

We validate Datamancer by demonstrating its use in three application scenarios within multi-display environments. Furthermore, we present findings from an expert review with a senior data analyst on the suitability of Datamancer in a large-scale collaborative data analysis environment where analysts work together on real-time transportation management and presentations. Finally, we report on a qualitative user study of using Datamancer for interacting with and organizing spatial data across multiple screens in a ubiquitous analytics environment.

Our contributions in this paper can be summarized as follows: (1) an investigation of the design space of gestural interaction for physical ubiquitous analytics environments; (2) a prototype implementation of the Datamancer system that realizes this model; (3) an evaluation of the prototype through demonstration of application scenarios, an expert review, and a user study.

2 Background

The field of data visualization and analytics is increasingly moving towards more immersive [16] and ubiquitous [18, 20] environments where users can interact with data across multiple displays and spaces. This shift necessitates new interaction paradigms that can seamlessly bridge the physical and digital worlds. Our work on Datamancer builds upon several key areas of research in human-computer interaction and data visualization.

2.1 Gesture-based Interaction

Gesture-based interfaces have long been a cornerstone of natural user interaction by offering intuitive ways to manipulate digital content. The foundations of this field were laid by pioneering works such as Hauptmann and McAvinney's research on combining gestures with speech for graphic manipulation [26], and Bolt's "Put-that-there" system [10], which demonstrated the power of multimodal interfaces leveraging both voice and gesture for spatial commands. These early efforts highlighted the potential of gestures to provide natural and intuitive control in computer interfaces, setting the stage for more sophisticated developments in the field.

As gesture-based interaction matured, researchers developed taxonomies and frameworks to better understand and categorize different types of gestures. Karam and schraefel's comprehensive classification of gestures [32] provided a structured approach to the growing body of research in this area. More recent systematic reviews by Vuletic et al. [62] and Koutsabasis and Vogiatzidakis [38] have further refined our understanding of hand gestures in HCI, offering insights into the most effective types of gestures for different tasks and contexts.

The application of gesture-based interaction to data visualization and analytics has opened up new possibilities for more engaging and intuitive data exploration. Lee et al. [41] highlighted the need to move beyond traditional input methods in visualization contexts, while Roberts et al. [58] envisioned a future where data visualization extends beyond traditional desktop environments, emphasizing the role of natural interactions, including gestures. Practical applications of these ideas can be seen in works such as that of Badam et al. [2], which explored how gesture-based interactions can facilitate collaborative data analysis in multi-display settings.

Recent years have seen increasingly sophisticated applications of gesture-based interaction in data visualization and analytics [31]. Filho et al. [21] demonstrated how gestures can enhance the exploration of complex spatiotemporal data in immersive environments. SketchStory [42] showed how gesture-based sketching can be used to create more engaging data narratives, blending the intuitiveness of hand-drawn sketches with the power of data visualization. The potential of gesture-based interaction in remote collaboration scenarios was explored by Hall et al. [24], proposing novel ways of using gestures to enhance data presentations in distributed settings. Additionally, He et al. [27] demonstrated how gesture-based natural user interfaces can be effectively integrated into sophisticated data analysis tools, pointing towards a future where gestures play a central role in how we interact with and understand complex data.

2.2 Multi-Display Environments and Ubiquitous Analytics

Ubiquitous analytics extends data analysis beyond desktop computers to multiple diverse devices and displays so that the analytical process is embedded into the physical environment [18, 20].

These environments present unique challenges and opportunities for interaction design [14] and data visualization. Grudin's seminal work [22] on partitioning digital worlds through multiple monitors laid the foundation for understanding how users interact with information spread across multiple displays.

Therefore, a key focus of research has been on pointing and targeting for multi-display environments. Nacenta et al. [51] introduced the Perspective Cursor, an interaction technique that provides perspective-based feedback for selecting targets across multiple displays. This work was extended by Xiao et al. [65], who compared direct and indirect pointing feedback in multi-display environments. Benko and Feiner [8] proposed pointer warping as a method to improve target acquisition across heterogeneous displays, while Waldner et al. [63] further refined this technique to bridge gaps between displays. The challenges of interaction across "displayless space"—the areas between physically separated displays—were addressed by Nacenta et al. [49], who proposed and evaluated several techniques for targeting across these gaps. In related work, Nacenta et al. [50] introduced E-conic, a perspective-aware interface for multi-display environments that adjusts visual feedback based on the user's viewing angle.

Recent developments in mobile and wearable technologies have opened new possibilities for ubiquitous analytics. Langner et al. [40] proposed VisTiles, a system that coordinates multiple mobile devices for visual data exploration, demonstrating how small, portable displays can be combined for complex analytics tasks. Horak et al. [28] proposed an approach for automatically distributing visualizations and their corresponding interactions across a dynamically changing display environment. Batch et al. [6] evaluated view management techniques for situated visualizations in web-based handheld augmented reality, providing insights into the effectiveness of different strategies for placing and adapting visualizations in AR environments. Finally, Batch et al. very recently also studied a "magic"-inspired visualization system designed for immersive and ubiquitous analytics [5], which leverages augmented reality to create interactive data visualizations in the physical world.

2.3 Bimanual Interaction Techniques

Bimanual interaction leverages the coordinated use of both hands, potentially offering more natural and efficient ways to manipulate data. The theoretical foundation for understanding bimanual interaction was laid by Guiard's Kinematic Chain model [23], which describes the asymmetric division of labor between the dominant and non-dominant hands in skilled bimanual actions. Building on this foundation, researchers have explored various aspects of bimanual interaction in human-computer interfaces. Leganchuk et al. [44] conducted experimental studies to quantify the manual and cognitive benefits of two-handed input, providing empirical evidence for the advantages of bimanual interaction in certain tasks. Balakrishnan and Hinckley [3] further investigated symmetric bimanual interaction, expanding our understanding beyond the asymmetric model proposed by Guiard.

Bimanual interaction has seen increasing use for specialized computing platforms. For large display environments, Nancel et al. [52] developed and evaluated mid-air pan-and-zoom techniques for wall-sized displays, demonstrating how bimanual gestures can be effectively used for navigation in large information spaces. Banerjee et al. [4] have presented an in-air pointing technique to manipulate out-of-reach targets while showing no loss of performance. Hough et al. [29] investigated the fidelity and plausibility of bimanual interaction in mixed reality environments, providing insights into the design of natural and effective bimanual interfaces in immersive space. Building on this, Peng et al. [54] explored freehand bimanual gestures for cross-workspace interaction in virtual reality, demonstrating how two-handed interactions can facilitate navigation and manipulation across multiple virtual workspaces. Talvas et al. [61] provided a comprehensive survey of bimanual haptic interaction, highlighting the importance of tactile feedback in two-handed interfaces. Building on all these prior research projects, our work on Datamancer uses bare-hand bimanual gesture interaction tracked using a chest-mounted Leap Motion device.

2.4 Wearable Devices for Gesture Interaction

Wearable technologies offer new possibilities for gesture recognition, allowing for more natural and unobtrusive interaction in ubiquitous computing environments. Early work in this field focused on creating compact devices. Rekimoto [57] introduced GestureWrist and GesturePad, pioneering the concept of unobtrusive wearable interaction. As technology progressed, researchers explored novel ways to leverage the human body as an input surface. Skinput [25] appropriates the human body for acoustic input, allowing for touch input on the skin. This work expanded the potential interaction space beyond traditional wearable devices. The concept of on-skin input has been further refined for smartwatch interactions. For example, TapSkin [67] recognizes on-skin tap gestures near a smartwatch, effectively extending the device's interaction space well beyond its physical dimensions.

Recent developments have explored new sensing modalities for gesture recognition. Soli [45] is a gesture sensing technology using millimeter-wave radar. This approach allows for high-resolution gesture tracking in a compact form factor, suitable for integration into various wearable devices. Kim et al. [33] introduced Digits, a wrist-worn gloveless sensor that enables freehand 3D interactions. This device demonstrates how wearable technology can capture complex hand movements without the need for external cameras or markers. The combination of wearable devices with other interaction modalities has also been explored. Pfeuffer et al. [55] investigated the combination of gaze and pinch gestures in VR, demonstrating how wearable sensors can complement other input methods in immersive environments. This approach has since been adopted in current-generation HMDs such as the Apple Vision Pro and the Meta Quest Pro.

Researchers have also explored wearable projection systems as a means of expanding the interaction space. Mistry et al. [48] introduced WUW (Wear Ur World), a wearable gestural interface that combines a camera and a projector to turn any surface into an interactive display. Similarly, Beardsley et al. [7] explored interaction techniques using a handheld projector, which could be adapted for wearable contexts. Recently, Patnaik et al. [53] proposed a handheld projector device with camera tracking to enable flashlight interaction with situated data [60]. Our approach builds on these approaches: we use a chest-mounted camera to track bimanual gestures, as in Mistry et al.'s WUW, and a ring-mounted device for hand-held display acquisition, similar to Beardsley et al. and Patnaik et al.

3 Design Space: Unified Interaction for Physical Ubiquitous Analytics Spaces

Ubiquitous analytics is the idea of embedding the analytical process into the physical environment to facilitate sensemaking anywhere and anytime [18, 20]. We envision a *ubiquitous analytics workspace* as a physical environment consisting of *multiple devices*, *multiple users*, and *multiple data representations* that can be created *ad-hoc* to suit specific tasks and contexts. The goal of a unified interaction paradigm for these spaces is to shift from device-proprietary inputs to device-agnostic, content-specific inputs creating a homogenized interaction workspace. This approach allows users to seamlessly acquire devices, select interaction modes, and engage with the data representations displayed within the environment. Below we outline the design requirements for this interaction paradigm, the modes we can identify, and the specific interaction types that we foresee being useful to interact with the data representations in the environment.

3.1 Design Requirements

We highlight several design requirements for unified interaction in ubiquitous and immersive analytics spaces:

- DR1 **Multiple displays and devices.** The ubiquitous and immersive analytics environments we envision consist of multiple physical displays and devices providing a seamless ad hoc sensemaking workspace [1, 34], all which are potential interaction targets.
- DR2 **Minimal instrumentation.** There should be a minimum of costly and centralized infrastructure in the environment to lower the barrier against entry and participation.
- DR3 **Mobility.** The solution should prioritize a mobile form factor, including no tethering with wireless communication offering portability. This also means that environments cannot be spatially mapped in advance.
- DR4 **Multimodal and Distant interaction.** Devices should be interacted with from a distance to minimize physical navigation as well as support Post-WIMP (windows, icons, menus, pointer) interaction, which has shown to support more natural, flexible, and engaging user experiences when going beyond a desktop [43].
- DR5 **Multiple Interaction Modes.** Data analysis requires multiple types of interactions, including selection, annotation, and navigation [66].
- DR6 **Multi-user.** While not a focus in this paper, the approach should also support collaborative data analytics [9, 30].

3.2 Interaction Modes

Regardless of the technical solution chosen to support the above design requirements, we can identify three distinct interaction modes. Figure 2 gives an overview of these modes and their transitions.

- Device Acquisition: Since our target environments contain multiple devices and displays (DR1) as well as multiple users (DR6), there will need to be an acquisition mode where a specific user selects a *device focus* for which further interaction will be directed. We call this acquisition "focusing" vs. "unfocusing" a device. For multiple users, we will have to consider whether a device that is already focused can be acquired by other users, and whether this will automatically unfocus earlier users or not.
- Mode Selection: Similarly, because we envision multiple interaction modes (DR5) for typical data analytics scenarios, the interaction model will need to support a selection mode where the user can switch between the different types of interaction. These types will depend on the device and the visual representation currently focused—for example, while all data representations support navigation, only parallel coordinate plots support axis reordering. We discuss data interaction types in Section 3.3.
- Data Interaction: Finally, once the device has been focused and the interaction mode chosen, we can enter into a direct interaction mode that is directly sent to the focused device and visual representation. This input is often continuous, such as 2D or 3D relative or absolutely position, but can also be discrete commands, such as specific gestures or device actions (clicks or taps).

3.3 Ubiquitous Data Interactions

Here we synthesize common interactions for data visualization, such as from Yi et al.'s taxonomy [66], cross-cutting interaction for visualization [17], and prior surveys on interaction for immersive analytics [15]. These interactions are all defined for 2D visual representations shown on regular displays within the physical analytics environment. For volumetric 3D visualizations, such as in VR or AR, many of these interactions also have 3D equivalents.

- *Selection:* Choosing specific data points or elements within the visualization. These selected data points can then be used for further interaction.
- *Panning:* Translating the viewport within the visualization. Panning is typically conducted in both horizontal and vertical dimensions simultaneously.

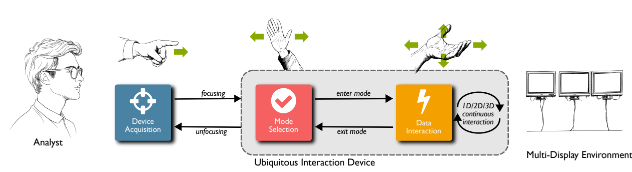

Figure 2: Main interaction modes. We identify three interaction modes in our approach for universal interaction for ubiquitous analytics spaces: device acquisition, mode selection, and data interaction.

- *Zooming:* Adjusting the magnification level of the view to focus on specific details or to see a broader overview. Sometimes zooming and panning can be combined with 3D relative input—left, right, up, and down controls translation, whereas backwards and forwards controls zoom.
- *Flipping (left/right):* Switching from the current view to the previous (left) or next (right) one, such as stepping backwards and forwards among views in a dashboard.
- *Filtering:* Applying criteria to display only the relevant data or elements within the visual representation. Filter operations are typically preceded by choosing which data dimension to filter on; this could, for example, be accomplished by flipping left or right between dimensions.
- *Sorting:* Controlling the order of presentation for data items, such as in a list or table. The sort order may depend on different data dimensions, which could potentially be flipped through, and could also be reversed.
- *Drill down (roll up):* Expand (drill down) a group of nodes into its children, or collapse them (roll up) into their higher-order grouping. This is common for clustering representations, especially for hierarchical clustering [19].
- *Grab (place):* Picking up a visualization to drop (place) it on a different display in the analytics environment.
- *Text input:* Entering textual information or commands using a keyboard or other input method. Textual input is often challenging in immersive or physical environments, where there is no easy way to place a physical keyboard. Soft keyboards or voice input may be good alternatives.
- *Pointer emulation:* Some legacy visualizations designed for mouse or touch input may require directly controlling a virtual pointer using gestural input.
- *Details-on-demand:* Displaying additional information or details about selected data points upon request.

Note also that these are generic definitions of interactions for visualization. Each visualization will have a concrete implementation with a meaning specific to that chart type. For example, flipping left and right may switch between adjacent axes in a parallel coordinate plot, whereas it could be used for switching between time-series in a line chart. In fact, some visualizations may not implement all interactions; for example, a scatterplot may not provide a meaningful flipping interaction.

4 Datamancer: Bimanual Gestures for Ubiquitous Analytics

Datamancer is a novel wearable system in our universal ubiquitous interaction design space (Section 3). It is designed to facilitate seamless bimanual gesture interactions across (ad-hoc) multi-display ubiquitous analytics environments. This section details the system's components, functionality, and architecture, as well as its implementation and capabilities.

4.1 System Overview

Datamancer comprises three main components (Figure 3): (1) a wearable hardware unit with two lightweight and compact sensors (DR2, DR3), (2) a software framework for gesture recognition and screen interaction, and (3) a distributed application layer for multi-display environments (DR1). The hardware provides a ring-mounted pinhole camera for acquiring displays, even at a distance (DR4), a chest-mounted gesture tracker to support bimanual data interaction (DR4, DR5) with a focused display, and a Raspberry Pi providing the computation.

4.2 Basic Use

The Datamancer system is designed for intuitive use in various ubiquitous analytics scenarios. A typical interaction session with our implementation of Datamancer proceeds as follows (Figure 4).

- **Setup:** Any screen/device that runs a browser and is connected to the internet can open up the Datamancer interaction workspace using a specific predefined web address. The user dons the wearable component, adjusting the chest-mounted harness and finger-mounted ring for optimal fit.
- **Hand Tracking Indicator:** All connected devices display a hand tracking icon in the bottom left corner, providing real-time feedback on the system's tracking status (green for active tracking, red for inactive).
- **Screen Acquisition:** To initiate interaction with a specific display, the user employs the finger-mounted ring equipped with a camera. Pressing and holding the ring's button triggers the display of fiducial markers (ArUco) on all connected screens. The user then points onto the fiducial marker at the desired screen, and the ring's camera detects the corresponding marker. Successful selection is confirmed by a red border appearing on the chosen display.
- **Data Interaction:** Once a screen is selected, the user can perform various hand gestures to interact with the content. For example, a right-hand grab/fist gesture enables panning of a map in any direction (left, right, up, down) or selecting between tiles in a gallery, a left hand grab/fist enables zooming in/out of a map, and a right hand pinch enables placing content whereas a left hand pinch enables removing content.

Figure 3: Datamancer hardware. Datamancer hardware primarily consists of a Leap Motion Controller 2 for recognizing hand gestures, a finger mounted camera for screen selection via pointing, and a Raspberry Pi as the computer.

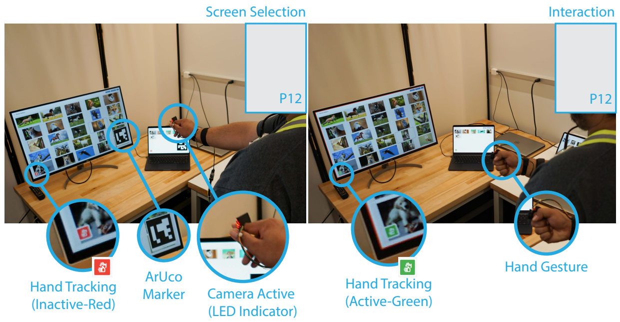

Figure 4: Basic usage. We identify three interaction modes in our approach for universal interaction for ubiquitous analytics spaces: device acquisition, mode selection, and data interaction.

Here, the input is activated upon detecting a gesture (such as a right-hand fist or a pinch etc.), which facilitates continuous interaction, and is turned off when the gesture is no longer detected. This mechanism helps avoid accidental input due to an unintentional gesture.

4.3 Hardware

The hardware component of Datamancer is a carefully integrated system of sensing, computing, and interaction devices designed for wearability and precision (Figure 3). It consists of the following key elements:

- **Hand Gesture Recognition:** At the core of Datamancer's gesture sensing capability is a Leap Motion Controller 2 camera. This high-precision optical tracking device with a tracking range of 10–110 cm and a horizontal-vertical field of view of 160°–160° is mounted on the chest harness, providing continuous and accurate hand pose estimation for both hands of the user.
- **Screen Recognition:** A custom-designed 3D-printed ring, worn on the index finger, houses an Adafruit Ultra Tiny GC0307 Camera. This ultra-compact camera performs fiducial marker detection, enabling precise screen selection. Datamancer comfortably works in detecting fiducials on small screens such as a 10.2-inch iPad from 2 m and larger screens such as a 60-inch TV from 7 m. This distance can be extended by adjusting the size of the displayed fiducial. The ring incorporates a tactile switch that, when pressed, activates the fiducial detection system. An LED indicator provides visual feedback to the user when the detection system is active.
- **Computation:** The system's computational needs are met by a Raspberry Pi 5, featuring a Broadcom BCM2712 2.4 GHz quad-core 64-bit Arm Cortex-A76 CPU with 8 GB RAM. An active cooler is employed to maintain optimal operating temperatures, ensuring consistent performance.
- **Power:** The system is powered by an Anker 337 portable power bank with a capacity of 26800 mAh , offering a runtime of over 10+ hours.
- **Wearable Harness:** The hardware components are integrated into a vest-style harness. The hand gesture sensor is positioned on the chest for an unobstructed view of the user's hands, while the computational unit is mounted at the waist to balance the system's weight. The ring is directly worn on the finger. Putting on the device takes about a minute; it is akin to putting on a safety harness.

4.4 Software Architecture

The Datamancer software architecture runs on the displays in a ubiquitous analytics space. It builds on the MyWebstrates [36] platform and the Codestrates [11] development environment. Each display accesses the system using a web browser. The interface is implemented using React, providing a responsive and modular front-end structure. Visualizations are generated using Vega-Lite [59] and interactive geographic maps using Leaflet.[1] Which content is shown on which display can be changed dynamically, e.g., if one display fails.

The connection between the Raspberry Pi and the screen software is facilitated using Automerge,[2] a conflict-free replicated datatype (CRDT) and a network layer for data synchronization. MyWebstrates uses Automerge as its data substrate and exposes its document to the client. The Raspberry Pi connects to the same document through a simple Node.js application. Using Automerge, a broadcast message channel can then be opened between all clients connected to the document—each screen and the Raspberry Pi. The Raspberry Pi uses this channel to send messages to all screens like, e.g., showing or hiding the marker on the screen. Similarly, when gestures are detected on the Raspberry Pi, they are sent as gesture messages to all screens. Screens then react to these gestures by, for example, panning or zooming a visualization. However, they only do so when they are active, i.e., were acquired using the pinhole camera. This architecture also preserves the native input of the device screens allowing for surface level input (say) on a touch screen or mouse based input on a desktop or projected screen.

5 Application Scenarios

Here we explore three scenarios that demonstrate how Datamancer can be used (Figure 5): individual use in a personal office setting, presentation in a conference room, and collaborative analysis. We note that these are conceptual scenarios; while we have not implemented them, they are all possible with the current Datamancer implementation.

5.1 Personal Office: Multi-Screen Data Manipulation

Overview: In a personal office environment, Datamancer can empower an individual user to seamlessly interact with data visualizations across multiple screens (Figure 5 (a)).

Scenario: Consider David, a data scientist working in his office with a complex dataset spanning several displays. By pointing to a specific display, David can acquire it and engage in various interactions using both hands. For instance, he can use bimanual gestures

[1] Leaflet: https://leafletjs.com/
[2] Automerge: https://automerge.org/

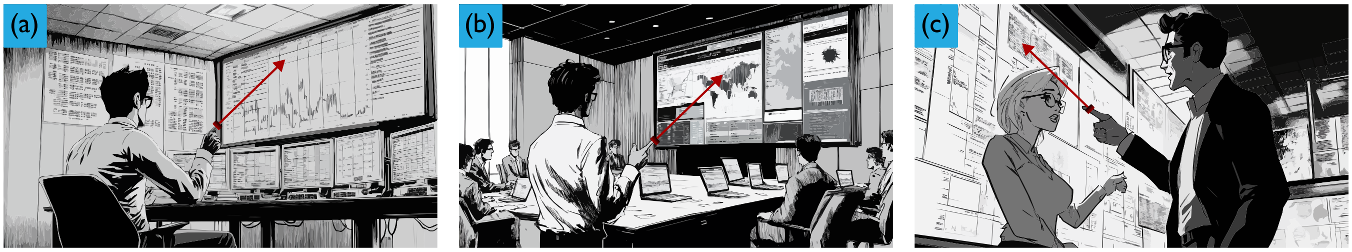

Figure 5: Application scenarios. Three conceptual examples of Datamancer being used for (a) single-user office work, (b) large-scale presentations, and (c) coupled collaborative work. (Images by MidJourney v6.1.)

to pan across a large dataset on one screen and then zooming into details on another. The system's selection capabilities allow for easy comparison of data points across different visualizations. The key benefit is that Datamancer binds together all of David's devices without the need to download or install specialized software. This allows David to make the most of available resources on the fly—such as wall displays, projectors, tablets, or monitors—without being confined to a specific room or setting, whether it's his home, a huddle room, a personal workspace, or a fully equipped data analysis room. He can now assemble a set of visualizations in preparation for a sales meeting with the board. Even if David would be using an HMD, the Datamancer system would facilitate blending virtual and physical content by allowing him to acquire and interact with any physical screen (such as in his office).

5.2 Conference Room: Dynamic Sales Data Presentation

Overview: Datamancer's capabilities are particularly powerful in a conference room setting, where a presenter shares sales data across a company's geographical regions with board members using their own laptops (Figure 5 (b)).

Scenario: After permitting David access to their devices, David, equipped with Datamancer, can use the laptops of board room attendees as additional display surfaces, creating an expansive, interactive presentation environment. Using gesture controls, he can swiftly distribute different regional sales charts to various screens by simply pointing and performing selection gestures. Bimanual gestures enable real-time data filtering and aggregation, allowing him to respond dynamically to audience questions. For example, a sweeping gesture could aggregate sales data from multiple regions, while a pinch-and-zoom action could drill down into specific product categories. At the end of the meeting, Jeanine, the CEO of the company, asks David and his colleague Anna to use the sales data to plan the launch of the company's new product line.

5.3 Collaborative Analysis: Multi-Display Product Launch Strategy

Overview: Here a team of business analysts uses Datamancer in a multi-display environment to collaboratively analyze data for deciding the location of their next product line launch (Figure 5 (c)).

Scenario: The room is equipped with several large displays and touch-enabled surfaces, allowing multiple team members to interact with the data simultaneously using their Datamancer devices. Anna and David can share and manipulate visualizations across displays by pointing and using bimanual gestures, fostering a truly collaborative analysis process. For instance, Anna might use gestures to overlay demographic data onto a map on one screen, while David combines market trend graphs on another display. The team can use common gestures to compare different datasets, such as consumer behavior and economic indicators, across multiple screens. Datamancer's interface allows for rapid iteration of ideas by allowing team members able to quickly acquire a screen, adjust its visualizations, and thereby explore various scenarios. In this setting, if Anna and David are wearing HMDs, Datamancer would facilitate them also coordinating the physical screens in the room with the virtual content shown in their personal displays.

6 Expert Review: Multi-Display Collaborative Decision Making

We conducted an expert review on the use of Datamancer for collaborative decision-making in a large-scale multi-display environment. We engaged with a professional from a transportation laboratory to gather insights on how Datamancer could be applied in their workflow. This expert review aimed to understand the practical implications and potential benefits of our system in a complex, data-driven decision-making process.

6.1 Method

Our expert review involved a senior data analyst from a prominent urban transportation laboratory with over 15 years of experience in analyzing urban mobility data. The study session consisted of three parts: a demonstration of Datamancer (with an animal image dataset as described in Section 7.2), hands-on experience for the expert participant, and a semi-structured interview. We showcased Datamancer's capabilities in a simulated multi-display environment, then allowed the expert to interact with the system while thinking aloud. The concluding interview focused on potential integration into their workflow, perceived benefits and challenges, and ideas for future enhancements. We audio-recorded the session and took notes for subsequent analysis.

6.2 Results

We spent a total of 60 minutes interviewing our analyst participant, demonstrating Datamancer, and letting them use the tool in their own collaborative decision-making space (Figure 6).

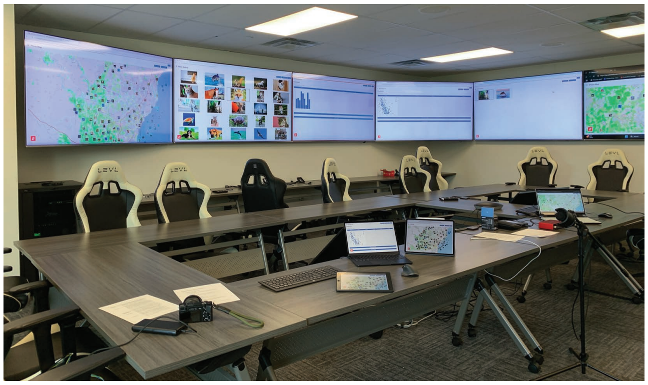

Figure 6: Collaborative decision-making space. A conference room used by our expert review participant that is used for presentations to stakeholders and collaborative decision-makers. The screens around the room is showing our Datamancer prototype in action.

6.2.1 Lab Setup and Workflow. The analyst described how they worked with data analytics and data visualization tools primarily for the transportation sector using geographic and big data visualizations. The analyst described how the six displays in the room would be used to display different data visualizations and analyses created for their clients. During a typical presentation, clients are seated in the room and an analyst will have the different screens pre-loaded with information, which allows them to walk around the room. However, lacking a better alternative, they would control the mouse cursor with a conventional cordless mouse using their own leg as a surface. The analyst described the workflow essentially as a storytelling session where a story is lined up and they go from left to right across the screens.

6.2.2 Current Challenges. The analyst describes three core challenges in their use of their setup: (1) when they lose track of the mouse on the six large screens; (2) the considerable time it takes to move content between monitors that are far apart; and finally (3) the "*catastrophic*" situation where a computer does not boot, making it extremely difficult to rearrange content across the remaining screens.

6.2.3 Potential of Datamancer. The expert review revealed several promising aspects of Datamancer for multi-display collaborative decision making in transportation analysis. The participant expressed enthusiasm about the system's potential to enhance real-time collaboration among team members, envisioning scenarios where multiple analysts could simultaneously interact with different datasets across various displays, which is in line with the scenario described in Section 5.2. The analyst said, "*it would be very impressive if I could walk into a room where you got 20 people with laptops and all of a sudden I'm able to throw things around onto their machine. It'd be pretty awesome.*" and "*I think the wearable (device) be used for shocking on a big presentation, just kind of do impressive fast things.*" This capability, they noted, could lead to more efficient decision-making processes, particularly when dealing with complex urban mobility data. "*I imagine it's got to be quicker for me to just be able to point at that screen, say, boom, over here, as opposed like grabbing my mouse. Okay, now I gotta move it all the way, because you gotta keep picking up the mouse and moving it to get to where you want to be. So I imagine it quicker.*"

The analyst also suggested using multiple Datamancer devices to allow different team members independent control over specific displays during collaboration. While they found the gesture-based interactions intuitive, the participant noted a potential learning curve for some team members and suggested a training program for widespread adoption.

We also see Datamancer as being a useful tool in cases of system going down—which the participant has described as being "*catastrophic*"—as Datamancer provides an ad-hoc way to construct an analytical workspace, building an ecosystem using any computers that run a browser, as well as facilitates easy interaction with content across screens. Furthermore, the current Datamancer implementation is based on MyWebstrates [36], which uses a local-first federated architecture and thus yields better robustness against failure than a pure client/server system.

Overall, the expert review participant received Datamancer positively, viewing it as a promising tool for enhancing collaborative decision-making in their field. They emphasized that the system's success would largely depend on its ability to integrate smoothly with existing workflows and its adaptability to specific domain needs.

7 User Study

We conducted a user study to investigate the effectiveness and usability of Datamancer. The goal was to understand how users acquire different displays, interact with digital objects, and move content between screens in a setting that resembles a multi-display ubiquitous analytics environment. The tasks in our user study required acquiring (focusing) devices, selecting from a limited set of interaction modes, and finally performing gestural interactions to successfully complete them (Section 3).

7.1 Participants

We recruited 12 paid participants (9 male, 3 female) from a diverse pool of university students and local professionals. The participants were selected based on their familiarity with mobile devices and basic computer skills, as well as to represent a range of experience levels with gestural interfaces. The demographics of our participant pool are as follows:

- **Age:** 22–37 years old (average 29.5).
- **Academic/Professional background:** Computer Science (6), Research Scientist in XR (1), Robotics (1), Environmental Science (1), Real Estate Analysis (1), Mechanical Engineering (1), Quantum Physics (1).
- **Prior experience with gestural interfaces:** 4 participants had extensive experience, 5 had some experience, and 3 were novice users.
- **Prior experience with analytics and visualization:** 3 participants were experts, 2 had a passing knowledge, and 7 had good experience.

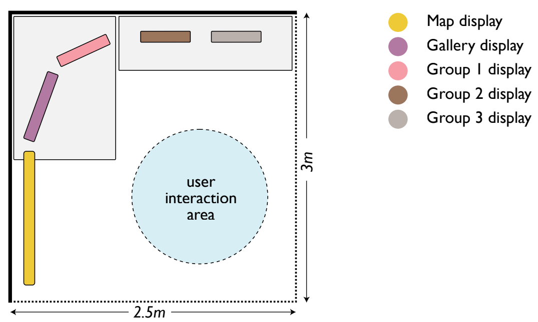

Figure 7: Overhead view of physical setup. The physical configuration of screens in our user study. Photo map shown on Map Display, Photo Gallery shown on Gallery Display, and Photo Groups 1, 2, 3 were shown on Group 1, 2, 3 displays.

7.2 Apparatus and Dataset

We used the Datamancer prototype, which consists of a finger-mounted pinhole camera for precise pointing and selection, and a chest-mounted gesture sensor for hand and finger tracking. The physical layout of screens included can be seen in Figure 1. We give an overhead view of the space in Figure 7. Each screen was connected to a separate computer or laptop, utilizing the Webstrates environment for seamless data replication across devices.

The setup (in clockwise order as seen in Figure 7 and with their respective contents shown in Figure 8) consisted of a Samsung 55-inch TV run on a Razer Blade Windows laptop with Chrome; a LG 32-inch monitor run on a Dell XPS Windows laptop with Chrome; an Apple 13-inch M3 MacBook Air with Safari; an Apple 12.9-inch iPad Pro with Safari; and an Apple 10.2-inch iPad (9th gen) with Safari. The Datamancer content on each of these screens (Figure 8) and their supported interaction were the following:

- **Photo Map:** (i) *Pan:* Lateral movement of right fist, (ii) *Zoom:* Back-forth movement of left fist.
- **Photo Gallery:** *Selection (and Grab):* Lateral movement of right fist (current selection was grabbed automatically).
- **Photo Group (1,2,3):** (i) *Place:* Right-hand pinch (after a selection), (ii) *Remove:* Left-hand pinch (after a selection), (iii) *Selection (and Grab):* Lateral movement of right-hand fist.

Given the small set of interactions, we did not include dedicated visual feedback on the exact gesture being performed, instead relying only on the tracking indicator. However, we discuss the need for effective in-situ guidance system as a potential enhancement in Section 9.3 and Section 10. This setup allowed us to test Datamancer's capabilities across various display sizes and orientations, mimicking real-world ubiquitous analytics environments.

The Photo Map showed a set of 100 photographs representing 10 different animal categories placed onto a geographic map of Sicily (Figure 8). The photographs were taken from the Animal Image Dataset.[3] The dataset was created by manually placing 10 photographs for each of 10 animal categories on a geographic map. We ensured that each animal was placed in an appropriate geographic location based on its natural habitat. For example, forest-dwelling animals were placed in wooded areas, while aquatic animals were positioned near bodies of water. This approach required participants to consider both visual and spatial information in their decision-making process, simulating real-world data analysis.

[3]Animal Image Dataset: https://www.kaggle.com/datasets/iamsouravbanerjee/animal-image-dataset-90-different-animals

7.3 Task

We designed our study around a geospatial sensemaking task where participants were asked to organize photographs of animals located in a geographic region into categories. Participants were given three tasks related to different animal categories. These tasks were designed to require multiple screens with different roles (Figure 7), screens where panning and zooming were necessary, and to move content between screens.

- T1 (a) **Search animals on map**: Catania is a municipality in Sicily and is close to the coast. What are the top 3 animals (land and aquatic) that are found closest/within the city center of Catania?
- T1 (b) **Place images in group**: Place 3 unique pictures of each of these in Group 1.
- T2 (a) **Search region on map**: Find the region on the map that has mainly eagles, goats, and foxes.
- T2 (b) **Place images in group**: Place (any) 3 of these animals (1 picture each) into Group 2.
- T3 **Search animals on map and place in group**: Find (any) 3 pictures of owls and place them into Group 3.

All tasks were performed sequentially. The tasks required acquiring/focusing on the respective display to begin with. Search tasks on the map (T1 (a), T2 (a), T3) required using panning and zooming gestures on the Photo Map on the large Map display (Figure 7). Placement tasks (T1 (b), T2 (b), T3) were done by, first, selecting pictures in the Photo Gallery on the medium-sized Gallery display, and, second, placing the selected picture in one of the Photo Groups on the smaller displays using a pinch gesture. All gestures are demonstrated in the accompanying video figure.

7.4 Procedure

The study was conducted in a controlled laboratory setting designed to simulate a realistic multi-display work environment and lasted approximately 55–60 minutes per participant. The procedure was as follows:

1. **Introduction and Consent** (7 minutes): Participants were briefed on the study's purpose and gave informed consent. They completed a questionnaire on demographics and prior experience with gestural interfaces.
2. **Training Session** (12 minutes): Participants underwent a short training phase with the Datamancer system. They practiced the acquisition of screens using the camera, interacting with the map and gallery using gestures, and placing images. After the training session the facilitator emptied the groups before starting the main task.
3. **Main Study** (25-30 minutes): Participants were presented with the three sensemaking tasks described above.
4. **Post-Task Evaluation** (10 minutes): Participants completed a post-task questionnaire with custom Likert-scale ratings

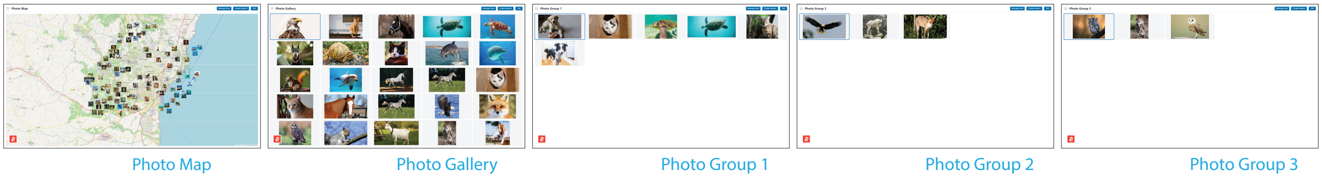

Figure 8: User study task screens. (From left) Map, Gallery and Photo Groups in the Datamancer user study.

on specific aspects of Datamancer. A brief semi-structured interview was conducted to gather qualitative feedback on the participants' experience, challenges faced, and suggestions for improvement.

(5) **Debriefing** (1 minute): Participants were debriefed on the study and given a chance to ask final questions.

7.5 Data Analysis

We employed a mixed-methods approach to analyze the collected data. On the qualitative side, we performed a direct observation of the participants as they performed the tasks accompanied by a thematic analysis of interview responses to identify recurring themes and user perceptions. The semi-structured interview responses were transcribed using Otter[4] and open coded by two researchers; after which the codes were merged into a code book which included rephrasing codes as well as adding and removing codes. A total of 15 codes were used across 156 excerpts from the interview data of 12 participants.

On the quantitative side, we analyzed the interaction logs cross-referencing with the direct observation and created an event diagram of active screens and participant interactions. Lastly, the SUS questionnaire was analyzed according to its scoring method [13].

8 Results

Here we report the results from our user study, including an analysis of the participant event logs, the verbal feedback on the system, and the system usability score. All 12 participants were able to complete the tasks within the time frame of the study. Since our study is not comparative, we do not perform an analysis of completion times or accuracy.

8.1 Event Logs and Inferences

Figure 9 shows an overview of the interaction events for all 12 participants in the study. It shows discrete interaction events, which display was focused, and the current task participants were working on. A clear pattern emerged during the individual tasks, with participants frequently transitioning between screens and interacting with each. Typically, participants began by consulting the main map, followed by the gallery display, and finally the individual groups.

Several participants seem to alternate between the gallery and the groups, presumably when selecting and placing photos. We also find a relatively large variance in completion time across participants. During the search task on the Map display, a partial reason for this is that some participants were more lucky than others in finding the target location. Additionally, some participants forgot to release the gesture, resulting in unintended gesture detections that caused the map to pan too far. This required them to manually readjust the map to its desired position. We also see participants getting better with gestures over time (Figure 9) such as P3 between T1(b) and T3, and P12 between T2(b) and T3.

8.2 Thematic Analysis of Qualitative Data

Here we describe the broad themes and user perceptions obtained through thematic analysis of interview responses.

8.2.1 User Experience. Overall the participants felt that their experience was positive, the interaction was intuitive, and they were mostly in control except for occasional incidents. Participants said they felt a sense of accomplishment after the tasks (P1), and added it was "*fun*" (P8, P11), "*easy to use*" (P2, P3, P9), and to "*understand*" (P2, P3, P9). All participants described the Datamancer system as being intuitive—said P9, "*I thought it was intuitive, like actually pinching the button to select a screen felt awesome as an idea, and I had fun doing that.*" and P12 said "*very intuitive to use, because it borrows from zooming, panning, all of these pre-existing interactions.*" P1 and P2 described their experience as being "*smooth*" and "*natural,*" respectively. While some participants described feeling confident and being in control using the system (P3, P10, P11), others noted adapting to the sensor sensitivity over time "*I think it was. I was in control of things that were happening. There were a few jitters on zooming and panning, but I think with time, I could get over them.*" (P7), "*I'll say it has a little bit learning curve at the beginning for few minutes, but once you get a hang of it, it's pretty seamless.*" (P10).

8.2.2 Novel Experience. Participants described their experience as being novel (P1, P2, P11). P12 said "*It was pretty unique. I have never interacted with screens using gestures before, well, arguably, virtual reality systems that I sort of tested had something similar, but it still didn't feel exactly the same, because I was still interacting with multiple (physical) screens.*" P7 loved the idea of "*having a variable on [my] body*" without obstructing their hands that facilitated cross-device interaction. Participants also preferred Datamancer to existing technologies such as an XR controller or a mouse which works on certain ecosystems (only) and blocks the hand (P2), as well as the need for a desk space/surface to use a mouse on (P6). P6 also said AR headsets could be comparable but are bulky and inconvenient to wear.

8.2.3 Features and Improvement. Participants liked being able to quickly select and deselect screens that enabled them to quickly

[4]Otter: https://otter.ai/

Figure 9: Study Event Logs. The vertical lines represent discrete input events, the horizontal lines inside charts represent the currently focused display, and the horizontal lines underneath charts represent the current task participants were working on. The horizontal scale shows elapsed time in minutes and seconds.

switch between them (P2, P4, P11). P4 said this would enable "*cross referencing data [...] with visual attributes*" while describing an analytical task in a cross-device setup as a potential scenario. Critical comments about the features were around hand tracking and the screen selection feature using fiducial markers. Some participants found the hand tracking to be slow and occasionally inaccurate (P9)—"*I found it hard for it to detect any pinches or any hand gestures*"—but also added that they might get better over time—"*but maybe if I had more time with this.*" We also noted that the ring ergonomics turned out to be a challenge where the participants that had a perfect fit had an easier time selecting screens as opposed to ones who did not. Only one participant, P1, found it difficult to memorize gestures. Overall participants commented on improving the tracking accuracy and screen selection.

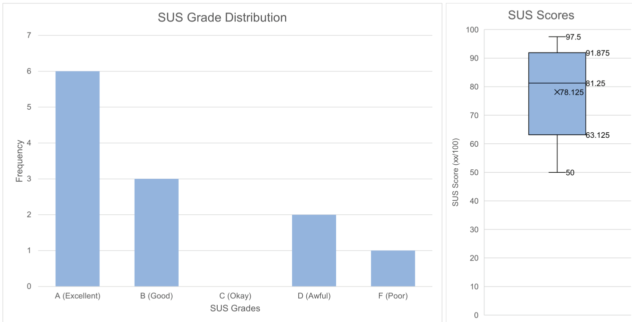

Figure 10: SUS Scores. System Usability Scale score for Datamancer during our qualitative evaluation.

8.2.4 User Adaptability. We were positively surprised to see participants quickly adapt to the system constraints and tracking inaccuracies to make successful gestures. We also observed that all participants learned to get better over time, where the later tasks were carried out with much ease as opposed to the opening ones. Said P10 said, "*It's pretty smooth, actually, (and) after a few minutes of human calibration,*" whereas said P12, "*but I still felt more in control because I was able to overcome those issues by just maybe learning to adapt my gesturing to the system.*"

8.2.5 Application Scenarios. Participants primarily discussed two broad categories of scenarios when asked about potential uses of Datamancer: (i) moving content across screens in multi-display environments either to better support workflow or due to hands being busy causing a situational impairment, and (ii) in collaborative settings for presentation and analysis. P1 spoke about managing a group people and P2 described potential use cases in meetings. P4 and P10 speculated on using it in a lab settings to move data across machines when hands are busy, such as when wearing gloves. P5, P9, and P11 spoke about using Datamancer to make meetings and presentations more engaging: "*You're not stuck at a desk or at a podium trying to, like, move things like, you can actually walk around and stuff, which I think makes people more engaging*" (P11). P6 felt the tool could potentially be used in classrooms.

8.3 System Usability Scale (SUS)

A standardized System Usability Scale [13] (SUS) test was administered at the end of the user study to determine the perceived usability of Datamancer. The average SUS score (xx/100) across 12 participants was 78.125 with a min of 50, max of 97.5 and a standard deviation of 16.135 as shown in Figure 10 (right). The distribution of SUS grades is shown in Figure 10 (left), where we see 75% of the participants rate usability above average.

9 Discussion

Here we critically examine the results of our user study and expert review on Datamancer. We interpret our findings in the context of existing research, discuss the implications and generalizability of our results, and address its limitations.

9.1 Explaining the Results

The results from our user study and expert review provide valuable insights into the effectiveness and potential of Datamancer in multi-display ubiquitous analytics environments. These findings, while encouraging, are not particularly surprising given the growing body of research on gesture-based interactions and multi-display environments.

Our evaluations confirmed that remote pointing and gestures are relatively natural for users, aligning with previous research on embodied interactions [14]. Participants in our user study quickly adapted to using Datamancer, describing the interaction as "intuitive" and "smooth." This supports the notion that gesture-based interactions can provide a more natural and immersive experience compared to traditional input methods—and in form factors more amenable to ubiquitous use than a cordless mouse used on your own thigh!

Our expert review with a senior data analyst from a transportation laboratory validated the utility of universal interaction in real-world ubiquitous analytics settings. The analyst's positive response to Datamancer's potential for enhancing real-time collaboration and streamlining complex data presentation workflows underscores the system's relevance in professional analytical environments. The user study, on the other hand, demonstrated that Datamancer is effective for acquiring different displays and moving content between them. This functionality addresses a key challenge in multi-display environments, where users often struggle with managing content across multiple screens.

9.2 Generalizing the Results

While our results are promising, we must approach generalization with caution. The expert review, involving only one participant, provides valuable insights but is limited in its ability to represent diverse use cases and preferences. Our 12-person user study, on the other hand, provides a stronger foundation for generalization. The consistently positive feedback across participants, coupled with the high average SUS score, suggests that Datamancer has broad appeal and usability. The diverse range of potential application scenarios proposed by participants, from collaborative meetings to classroom settings, indicates that the system's benefits may extend beyond our limited study use case.

However, the study also highlighted that the adaptability and training of gestures remains a challenge. This is a generally known fact for all gesture-based interfaces [39]. While most participants found the system intuitive, there was a learning curve, particularly in adapting to the sensor sensitivity and memorizing gestures. This suggests that while gesture-based interactions can be powerful, careful consideration must be given to the design and introduction of gesture vocabularies to ensure widespread adoption.

9.3 Limitations

Despite overall positive outcomes, our study revealed several limitations that warrant further investigation. Foremost among these is gesture learnability and adaptability. While participants generally found Datamancer intuitive once they understood the gestures, the initial learnability of available interactions posed a challenge. This highlights the need for effective onboarding strategies and possibly in-situ guidance systems to help users learn and remember gesture commands.

Although we did not observe significant participant fatigue during the study—possibly due to the structure of the study comprising multiple smaller tasks—prolonged use of Datamancer may still induce fatigue over time. Notably, P4 reported experiencing some discomfort when repeatedly performing a pointing gesture to acquire a display. A promising direction for future research could be exploring sensing architectures that offer indirect mid-air inputs [12] and hip-level sensing [46] to reduce physical strain and improve overall user comfort.

Another limitation is that we have not yet studied collaborative scenarios, despite their clear potential as highlighted in our expert review. Future research should explore how multiple users can interact simultaneously with Datamancer, addressing potential issues of interference and developing protocols for coordination between multiple devices.

Our expert review involving a single analyst participant would have benefited from using datasets and tasks specific to the application rather than our generic geospatial sensemaking task. Unfortunately, the sensitive nature of the data studied in this transportation laboratory made it impractical for us to create a customized application scenario. Future work should study this application in more detail.

The current prototype of our ring device is tethered, which may limit mobility and natural interaction in some scenarios. Developing a wireless version of the ring would enhance the system's flexibility and potentially its adoption.

Finally, a technical limitation of our current implementation is its reliance on web-based content for the displays. While this approach offers advantages in terms of cross-platform compatibility, it may not fully address the needs of professionals who rely on native applications. To bridge this gap, we could explore technologies like WebRTC and application sharing to cast native applications to a web browser, thereby expanding Datamancer's compatibility with existing software ecosystems.

10 Conclusion and Future Work

We have presented Datamancer, a device that supports bimanual interaction for acquiring and interacting with ubiquitous multi-display environments. The device was specifically designed for data visualization and analytics tasks, providing a comprehensive set of two-handed gestures for selecting, filtering, grabbing, placing, and drilling into data visualizations. We have validated Datamancer in three different ways: three application scenarios, an in-depth qualitative demonstration and interview session with a transport professional, and a qualitative user study involving 12 participants conducting a spatial sensemaking task in a five-screen analytics space. Our qualitative results indicate an overall positive potential for using Datamancer in ubiquitous analytics settings, including for presentations, collaboration, and multi-display analysis settings.

This work is part of a larger effort on physical computing devices to support ubiquitous [18, 20] and immersive analytics [16, 47]. In the future, we would be interested in understanding how multiple users, each with a Datamancer device, can work together as well as integrate other available input devices—such as a clicker, pointer, controller—while minimizing interference. Another avenue of future research would be to develop our gesture recognition further to enable more advanced, even personalized, gestures for interacting with data. Finally, like all gestural interfaces, Datamancer suffers from discoverability concerns [64]; we are keen on investigating visual feedback mechanisms to instruct the user or even using haptic feedback to guide them.

Acknowledgments

This work was supported partly by Villum Investigator grant VL-54492 by Villum Fonden. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency.

References

[1] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to Think: Large, High-Resolution Displays for Sensemaking. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 55–64. doi:10.1145/1753326.1753336

[2] Sriram Karthik Badam, Fereshteh Amini, Niklas Elmqvist, and Pourang Irani. 2016. Supporting Visual Exploration for Multiple Users in Large Display Environments. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. IEEE Computer Society, Los Alamitos, CA, USA, 1–10. doi:10.1109/vast.2016.7883506

[3] Ravin Balakrishnan and Ken Hinckley. 2000. Symmetric bimanual interaction. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 33–40. doi:10.1145/332040.332404

[4] Amartya Banerjee, Jesse Burstyn, Audrey Girouard, and Roel Vertegaal. 2011. Pointable: an in-air pointing technique to manipulate out-of-reach targets on tabletops. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. ACM, New York, NY, USA, 11–20. doi:10.1145/2076354.2076357

[5] Andrea Batch, Peter W. S. Butcher, Panagiotis D. Ritsos, and Niklas Elmqvist. 2024. Wizualization: A "Hard Magic" Visualization System for Immersive and Ubiquitous Analytics. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 507–517. doi:10.1109/TVCG.2023.3326580

[6] Andrea Batch, Sungbok Shin, Julia Liu, Peter W. S. Butcher, Panagiotis D. Ritsos, and Niklas Elmqvist. 2023. Evaluating View Management for Situated Visualization in Web-based Handheld AR. *Computer Graphics Forum* 42, 3 (2023), 349–360. doi:10.1111/cgf.14835

[7] Paul A. Beardsley, Jeroen van Baar, Ramesh Raskar, and Clifton Forlines. 2005. Interaction Using a Handheld Projector. *IEEE Computer Graphics and Applications* 25, 1 (2005), 39–43. doi:10.1109/MCG.2005.12

[8] Hrvoje Benko and Steven Feiner. 2007. Pointer warping in heterogeneous multi-monitor environments. In *Proceedings of the Graphics Interface Conference*, Vol. 234. ACM, New York, NY, USA, 111–117. doi:10.1145/1268517.1268537

[9] Mark Billinghurst, Maxime Cordeil, Anastasia Bezerianos, and Todd Margolis. 2018. Collaborative Immersive Analytics. In *Immersive Analytics*. Lecture Notes in Computer Science, Vol. 11190. Springer Publishing, New York, NY, USA, 221–257. doi:10.1007/978-3-030-01388-2_8

[10] Richard A. Bolt. 1980. "Put-That-There": Voice and gesture at the graphics interface. *Computer Graphics* 14, 3 (1980), 262–270. doi:10.1145/965105.807503

[11] Marcel Borowski, Janus Bager Kristensen, Rolf Bagge, and Clemens N. Klokmose. 2021. *Codestrates v2: A Development Platform for Webstrates*. Technical Report. Aarhus University. https://pure.au.dk/portal/en/publications/codestrates-v2-a-development-platform-for-webstrates(66e1d4d9-27da-4f6b-85b3-19b0993caf22).html

[12] Eugenie Brasier, Olivier Chapuis, Nicolas Férey, Jean-Marc Vézien, and Caroline Appert. 2020. ARPads: Mid-air Indirect Input for Augmented Reality. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, Los Alamitos, CA, USA, 332–343. doi:10.1109/ISMAR50242.2020.00060

[13] John Brooke. 1996. SUS - A 'Quick and Dirty' Usability Scale. In *Usability Evaluation In Industry*. CRC Press, Boca Raton, FL, USA, 189–194.

[14] Frederik Brudy, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmose, and Nicolai Marquardt. 2019. Cross-Device Taxonomy: Survey, Opportunities and Challenges of Interactions Spanning Across Multiple Devices. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 562. doi:10.1145/3290605.3300792

[15] Wolfgang Büschel, Jian Chen, Raimund Dachselt, Steven Mark Drucker, Tim Dwyer, Carsten Görg, Tobias Isenberg, Andreas Kerren, Chris North, and Wolfgang Stuerzlinger. 2018. Interaction for Immersive Analytics. In *Immersive*

Analytics. Lecture Notes in Computer Science, Vol. 11190. Springer Publishing, New York, NY, USA, 95–138. doi:10.1007/978-3-030-01388-2_4

[16] Tom Chandler, Maxime Cordeil, Tobias Czauderna, Tim Dwyer, Jaroslaw Glowacki, Cagatay Goncu, Matthias Klapperstueck, Karsten Klein, Falk Schreiber, and Elliot Wilson. 2015. Immersive Analytics. In *Proceedings of the International Symposium on Big Data Visual Analytics*. IEEE Computer Society, Los Alamitos, CA, USA, 1–8. doi:10.1109/BDVA.2015.7314296

[17] Jungu Choi, Deok Gun Park, Yuet Ling Wong, Eli Raymond Fisher, and Niklas Elmqvist. 2015. VisDock: A Toolkit for Cross-Cutting Interactions in Visualization. *IEEE Transactions on Visualization and Computer Graphics* 21, 9 (2015), 1087–1100. doi:10.1109/TVCG.2015.2414454

[18] Niklas Elmqvist. 2023. Data Analytics Anywhere and Everywhere. *Commun. ACM* 66, 12 (2023), 52–63. doi:10.1145/3584858

[19] Niklas Elmqvist and Jean-Daniel Fekete. 2010. Hierarchical Aggregation for Information Visualization: Overview, Techniques, and Design Guidelines. *IEEE Transactions on Visualization and Computer Graphics* 16, 3 (2010), 439–454. doi:10.1109/TVCG.2009.84

[20] Niklas Elmqvist and Pourang Irani. 2013. Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime. *IEEE Computer* 46, 4 (2013), 86–89. doi:10.1109/mc.2013.147

[21] Jorge A. Wagner Filho, Wolfgang Stuerzlinger, and Luciana Nedel. 2020. Evaluating an Immersive Space-Time Cube Geovisualization for Intuitive Trajectory Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 514–524. doi:10.1109/TVCG.2019.2934415

[22] Jonathan Grudin. 2001. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 458–465. doi:10.1145/365024.365312

[23] Yves Guiard. 1987. Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. *Journal of Motor Behavior* 19, 4 (1987), 486–517. doi:10.1080/00222895.1987.10735426

[24] Brian D. Hall, Lyn Bartram, and Matthew Brehmer. 2022. Augmented Chironomia for Presenting Data to Remote Audiences. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 18:1–18:14. doi:10.1145/3526113.3545614

[25] Chris Harrison, Desney S. Tan, and Dan Morris. 2010. Skinput: appropriating the body as an input surface. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 453–462. doi:10.1145/1753326.1753394

[26] Alexander G. Hauptmann and Paul McAvinney. 1993. Gestures with Speech for Graphic Manipulation. *International Journal of Man-Machine Studies* 38, 2 (1993), 231–249. doi:10.1006/IMMS.1993.1011

[27] Shizhe He, Mao Lin Huang, and Lin Zhu. 2015. Natural User Interface Design in DA-TU: An Interactive Clustered Data Visualization System. In *Proceedings of the International Conference on Information Visualisation*. IEEE Computer Society, Los Alamitos, CA, USA, 83–88. doi:10.1109/iV.2015.25

[28] Tom Horak, Andreas Mathisen, Clemens Nylandsted Klokmose, Raimund Dachselt, and Niklas Elmqvist. 2019. Vistribute: Distributing Interactive Visualizations in Dynamic Multi-Device Setups. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 616:1–616:13. doi:10.1145/3290605.3300846

[29] Gregory Hough, Ian Williams, and Cham Athwal. 2015. Fidelity and plausibility of bimanual interaction in mixed reality. *IEEE Transactions on Visualization and Computer Graphics* 21, 12 (2015), 1377–1389. doi:10.1109/TVCG.2015.2480060

[30] Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma, and Hans Hagen. 2011. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization* 10, 4 (2011), 310–326. doi:10.1177/1473871611412817

[31] Sujin Jang, Niklas Elmqvist, and Karthik Ramani. 2014. GestureAnalyzer: visual analytics for pattern analysis of mid-air hand gestures. In *Proceedings of the ACM Symposium on Spatial User Interaction*. ACM, New York, NY, USA, 30–39. doi:10.1145/2659766.2659772

[32] Maria Karam and m. c. schraefel. 2005. *A Taxonomy of Gestures in Human Computer Interactions*. Project Report. University of Southampton. https://eprints.soton.ac.uk/261149/

[33] David Kim, Otmar Hilliges, Shahram Izadi, Alex Butler, Jiawen Chen, Iason Oikonomidis, and Patrick Olivier. 2012. Digits: freehand 3D interactions anywhere using a wrist-worn gloveless sensor. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 167–176. doi:10.1145/2380116.2380139

[34] David Kirsh. 1995. The Intelligent Use of Space. *Artificial Intelligence* 73, 1-2 (1995), 31–68. doi:10.1016/0004-3702(94)00017-U

[35] Clemens Nylandsted Klokmose, James R. Eagan, Siemen Baader, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2015. Webstrates: Shareable Dynamic Media. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 280–290. doi:10.1145/2807442.2807446

[36] Clemens Nylandsted Klokmose, James R. Eagan, and Peter van Hardenberg. 2024. MyWebstrates: Webstrates as Local-first Software. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 42:1–42:12. doi:10.1145/3654777.3676445

[37] Zona Kostic, Catherine Dumas, Sarah Pratt, and Johanna Beyer. 2023. Exploring Mid-Air Hand Interaction in Data Visualization. *CoRR* abs/2311.15372 (2023), 18 pages. doi:10.48550/ARXIV.2311.15372 arXiv:2311.15372

[38] Panayiotis Koutsabasis and Panagiotis Vogiatzidakis. 2019. Empirical Research in Mid-Air Interaction: A Systematic Review. *International Journal of Human-Computer Interaction* 35, 18 (2019), 1747–1768. doi:10.1080/10447318.2019.1572352

[39] Gordon Kurtenbach, Thomas P. Moran, and William Buxton. 1994. Contextual Animation of Gestural Commands. *Computer Graphics Forum* 13, 5 (1994), 305–314. doi:10.1111/1467-8659.1350305

[40] Ricardo Langner, Tom Horak, and Raimund Dachselt. 2018. VisTiles: Coordinating and Combining Co-located Mobile Devices for Visual Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 626–636. doi:10.1109/TVCG.2017.2744019

[41] Bongshin Lee, Petra Isenberg, Nathalie Henry Riche, and Sheelagh Carpendale. 2012. Beyond Mouse and Keyboard: Expanding Design Considerations for Information Visualization Interactions. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2689–2698. doi:10.1109/TVCG.2012.204

[42] Bongshin Lee, Rubaiat Habib Kazi, and Greg Smith. 2013. SketchStory: Telling more engaging stories with data through freeform sketching. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2416–2425. doi:10.1109/TVCG.2013.191

[43] Bongshin Lee, Arjun Srinivasan, Petra Isenberg, and John T. Stasko. 2021. Post-WIMP Interaction for Information Visualization. *Foundations and Trends in Human-Computer Interaction* 14, 1 (2021), 1–95. doi:10.1561/1100000081

[44] Andrea Leganchuk, Shumin Zhai, and William Buxton. 1998. Manual and Cognitive Benefits of Two-Handed Input: An Experimental Study. *ACM Transactions on Computer-Human Interaction* 5, 4 (1998), 326–359. doi:10.1145/300520.300522

[45] Jaime Lien, Nicholas Gillian, Mustafa Emre Karagozler, Patrick Amihood, Carsten Schwesig, Erik Olson, Hakim Raja, and Ivan Poupyrev. 2016. Soli: ubiquitous gesture sensing with millimeter wave radar. *ACM Transactions on Graphics* 35, 4 (2016), 142:1–142:19. doi:10.1145/2897824.2925953

[46] Mingyu Liu, Mathieu Nancel, and Daniel Vogel. 2015. Gunslinger: Subtle Arms-down Mid-air Interaction. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 63–71. doi:10.1145/2807442.2807489

[47] Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas (Eds.). 2018. *Immersive Analytics*. Lecture Notes in Computer Science, Vol. 11190. Springer Publishing, New York, NY, USA. doi:10.1007/978-3-030-01388-2

[48] Pranav Mistry, Pattie Maes, and Liyan Chang. 2009. WUW - Wear Ur World: a wearable gestural interface. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 4111–4116. doi:10.1145/1520340.1520626

[49] Miguel A. Nacenta, Regan L. Mandryk, and Carl Gutwin. 2008. Targeting across displayless space. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 777–786. doi:10.1145/1357054.1357178

[50] Miguel A. Nacenta, Satoshi Sakurai, Tokuo Yamaguchi, Yohei Miki, Yuichi Itoh, Yoshifumi Kitamura, Sriram Subramanian, and Carl Gutwin. 2007. E-conic: a perspective-aware interface for multi-display environments. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 279–288. doi:10.1145/1294211.1294260

[51] Miguel A. Nacenta, Samer Sallam, Bernard Champoux, Sriram Subramanian, and Carl Gutwin. 2006. Perspective cursor: perspective-based interaction for multi-display environments. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 289–298. doi:10.1145/1124772.1124817

[52] Mathieu Nancel, Julie Wagner, Emmanuel Pietriga, Olivier Chapuis, and Wendy E. Mackay. 2011. Mid-air pan-and-zoom on wall-sized displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 177–186. doi:10.1145/1978942.1978969

[53] Biswaksen Patnaik, Huaishu Peng, and Niklas Elmqvist. 2024. VisTorch: Interacting with Situated Visualizations using Handheld Projectors. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 208:1–208:13. doi:10.1145/3613904.3642857

[54] Chao Peng, Yangzi Dong, and Lizhou Cao. 2021. Freehand Interaction in Virtual Reality: Bimanual Gestures for Cross-Workspace Interaction. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*. ACM, New York, NY, USA, 73:1–73:2. doi:10.1145/3489849.3489900

[55] Ken Pfeuffer, Benedikt Mayer, Diako Mardanbegi, and Hans Gellersen. 2017. Gaze + pinch interaction in virtual reality. In *Proceedings of the ACM Symposium on Spatial User Interaction*. ACM, New York, NY, USA, 99–108. doi:10.1145/3131277.3132180

[56] Francisco J. Romero Ramírez, Rafael Muñoz-Salinas, and Rafael Medina Carnicer. 2018. Speeded up detection of squared fiducial markers. *Image and Vision Computing* 76 (2018), 38–47. doi:10.1016/J.IMAVIS.2018.05.004

[57] Jun Rekimoto. 2001. GestureWrist and GesturePad: unobtrusive wearable interaction devices. In *Proceedings of the IEEE International Symposium on Wearable Computers*. IEEE Computer Society, Los Alamitos, CA, USA, 21–27. doi:10.1109/iswc.2001.962092

[58] J. Roberts, P. Ritsos, S. K. Badam, D. Brodbeck, J. Kennedy, and N. Elmqvist. 2014. Visualization beyond the Desktop—the Next Big Thing. *IEEE Computer Graphics and Applications* 34, 6 (Nov 2014), 26–34. doi:10.1109/MCG.2014.82

[59] Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2017. Vega-Lite: A grammar of interactive graphics. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 341–350. doi:10.1109/TVCG.2016.2599030

[60] Sungbok Shin, Andrea Batch, Peter W. S. Butcher, Panagiotis D. Ritsos, and Niklas Elmqvist. 2024. The Reality of the Situation: A Survey of Situated Analytics. *IEEE Transactions on Visualization and Computer Graphics* 30, 8 (2024), 5147–5164. doi:10.1109/TVCG.2023.3285546

[61] Anthony Talvas, Maud Marchal, and Anatole Lécuyer. 2014. A Survey on Bimanual Haptic Interaction. *IEEE Transactions on Haptics* 7, 3 (2014), 285–300. doi:10.1109/TOH.2014.2314456

[62] Tijana Vuletic, Alex H. B. Duffy, Laura Hay, Chris McTeague, Gerard Campbell, and Madeleine Grealy. 2019. Systematic literature review of hand gestures used in human computer interaction interfaces. *International Journal of Human-Computer Studies* 129 (2019), 74–94. doi:10.1016/J.IJHCS.2019.03.011

[63] Manuela Waldner, Ernst Kruijff, and Dieter Schmalstieg. 2010. Bridging gaps with pointer warping in multi-display environments. In *Proceedings of the ACM Nordic Conference on Human-Computer Interaction*. ACM, New York, NY, USA, 813–816. doi:10.1145/1868914.1869036

[64] Haijun Xia, Michael Glueck, Michelle Annett, Michael Wang, and Daniel Wigdor. 2022. Iteratively Designing Gesture Vocabularies: A Survey and Analysis of Best Practices in the HCI Literature. *ACM Transactions on Computer-Human Interaction* 29, 4 (2022), 37:1–37:54. doi:10.1145/3503537

[65] Robert Xiao, Miguel A. Nacenta, Regan L. Mandryk, Andy Cockburn, and Carl Gutwin. 2011. Ubiquitous cursor: a comparison of direct and indirect pointing feedback in multi-display environments. In *Proceedings of the Graphics Interface Conference*. ACM, New York, NY, USA, 135–142. doi:10.5555/1992917.1992939

[66] Ji Soo Yi, Youn ah Kang, and John Stasko. 2007. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1224–1231. doi:10.1109/tvcg.2007.70515

[67] Cheng Zhang, Abdelkareem Bedri, Gabriel Reyes, Bailey Bercik, Omer T. Inan, Thad E. Starner, and Gregory D. Abowd. 2016. TapSkin: Recognizing On-Skin Input for Smartwatches. In *Proceedings of the ACM Conference on Interactive Surfaces and Spaces*. ACM, New York, NY, USA, 13–22. doi:10.1145/2992154.2992187