

# “I Feel Like Iron Man”: Authoring, Exploring, and Presenting Data Visualizations in Immersive AR

Marcel Borowski, Peter W. S. Butcher, Shannon J. Jones, Panagiotis D. Ritsos, *Member, IEEE*,  
Clemens N. Klokrose, and Niklas Elmquist, *Fellow, IEEE*

**Abstract**—A truly anytime, anywhere approach to immersive analytics is predicated on the ability for users to author and explore visualizations in an immersive fashion and without a desktop computer. In this paper, we present a generalizable visual programming technique for immersive authoring, exploration, and presentation of data visualizations through direct manipulation, and implement it in a web-based augmented reality (AR) system. In a preregistered user study, we use this implementation to explore how analysts author and use data visualizations in immersive AR accessed using a head-mounted display (HMD). The study focuses on the authoring technique, spatial organization strategies, and sensemaking and presentation practices in such immersive analytics environments for HMD-based AR. Our findings suggest that participants particularly liked the direct manipulation of our approach, and that participants' spatial organization is not based on physical landmarks in the room.

**Index Terms**—Augmented reality, immersive analytics, data visualization, visualization authoring, organizational strategies, sensemaking.

---

◆

## 1 INTRODUCTION

THE burgeoning fields of ubiquitous [1], immersive [2], and situated analytics [3] (UA/IA/SA) are purporting to revolutionize the way we interact with data by providing a mobile [4], immersive [5], and spatially embedded [6] space for analysis that can be accessed anytime, anywhere [7]. One of the persistent visions of such “anytime anywhere analytics” is to enable analysts to author their own visualizations in this mobile and immersive space, allowing them to explore complex data relationships and uncover new insights while on the go and with no desktop computer in sight. However, with a few exceptions [8], [9], [10], current UA systems still mostly rely on pre-authored visualizations, workflows, and datasets, thereby limiting the flexibility and creativity of data exploration. There is no generalized approach to immersive authoring and limited empirical results on their practical usage.

In this paper, we address this gap by proposing and evaluating a visual programming technique for authoring declarative visualization specifications through 3D direct manipulation and voice input. It provides a highly flexible and expressive method for constructing and manipulating visualizations by dragging and dropping name-value code fragments into hierarchical groups and connecting them to data. We implemented it in our web-based augmented reality system DashSpace [11], that runs on consumer-level HMDs such as the Quest 3 and uses the Vega-Lite visual grammar [12].

To validate our approach, we conducted a qualitative preregistered study in which participants used our system implementation to author and explore data visualizations. More specifically, beyond assessing the utility of our authoring technique, we designed our study to also investigate spatial organization strategies for

immersive analytics in 3D as well as specific strategies employed for analysts when performing sensemaking as well as presentation.

Our work contributes (1) a visual programming technique for authoring and exploring data for immersive visualizations using direct manipulation that we implemented in the DashSpace [11] platform; and (2) insights into authoring, spatial organization, sensemaking, and presentation in immersive analytics.

## 2 RELATED WORK

Our research on immersive authoring and exploration of ubiquitous data visualizations builds upon several areas of prior work. Here we review immersive and ubiquitous analytics, spatial organization in 3D environments, and authoring tools for immersive visualizations.

### 2.1 Immersive and Ubiquitous Analytics

The field of data visualization has evolved significantly with the advent of novel computing technologies. Ubiquitous analytics (UA) [1] emerged in the early 2010s as a paradigm that leverages these technologies—including mobile devices, large displays, and various interaction modalities—to enable effective data analysis and visualization across diverse environments. Within this broader field, immersive analytics (IA) [2], [5] and situated analytics (SA) [3] focus specifically on utilizing immersive technologies, such as augmented reality (AR), virtual reality (VR), and extended/mixed reality (XR), to analyze data within specific contexts or locations.

The development of platforms for these analytics approaches has seen significant progress over the years [7], [13], [14], [15]. Early systems primarily targeted traditional computing environments, but the landscape has since expanded to encompass a wide range of immersive and ubiquitous technologies. As the field has matured, researchers have identified key challenges and opportunities in immersive analytics. Ens et al. [16] conducted a comprehensive review of the state of immersive analytics, outlining grand challenges that need to be addressed for the field to reach its full potential. These challenges span areas such as spatially situated

- 
- Marcel Borowski, Clemens N. Klokrose, and Niklas Elmquist are with the Department of Computer Science at Aarhus University, Aarhus, Denmark. E-mail: {marcel.borowski, clemens, elm}@cs.au.dk
  - Peter W. S. Butcher, Shannon J. Jones, and Panagiotis D. Ritsos are with the School of Computer Science and Engineering at Bangor University, Bangor, United Kingdom. E-mail: {p.butcher, shannon.jones, p.ritsos}@bangor.ac.uk

Manuscript received XXX XX, 2024; revised XXX XX, 2024.



Fig. 1. **Participant analyzing data in immersive AR.** Participant P7 (anonymized through creative license) during their “I feel like Iron Man” moment, after modifying three visualizations by adding a single piece that connects to all three and sorts their X-axis.

data visualization, collaborative analytics, and the development of standards and best practices for immersive analytics systems.

Researchers have explored the potential of both VR and AR for immersive analytics. For VR-based systems, Cordeil et al. [17] conducted a comparative study on immersive collaborative analysis of network connectivity, examining the efficacy of CAVE-style environments versus head-mounted displays. Their findings highlighted the potential of VR for collaborative data analysis tasks while also identifying challenges related to user interaction and spatial awareness. Likewise, Lee et al. [18] explored the use of space and surfaces in VR for collaborative visualization tasks performed in a room-sized immersive environment, indicating, among other findings, the need for reconfigurable immersive spaces. Finally, several efforts explore novel interaction mechanism. Gesslein et al. [19] explore the use of pen-based input to interact with spreadsheets in VR. Biener et al. [20] explore the joint interaction space between touchscreens and VR-based displays when interacting with data.

Focusing specifically on AR-based systems, recent developments have shown promising advancements on interaction mechanisms, in situated and immersive analytics settings. Bach et al. [21] investigated the effectiveness of tangible AR for interactive exploration of 3D visualizations, comparing it with traditional desktop and touch-based interfaces. Their work also demonstrates the potential of AR to enhance spatial understanding and direct manipulation of data in immersive environments. MARVisT [22] introduces a mobile AR glyph-based visualization authoring tool for non-expert users, enabling them to bind data to real world objects and demonstrating the potential of AR to bridge the gap between digital data and physical contexts. Büschel et al. [23] investigate the use of mobile devices as controllers, to provide pan and zoom techniques for interacting with 3D data spaces in AR.

## 2.2 Spatial Organization in 3D Environments

Intelligent use of physical space can simplify choice, perception, and internal computation [24], playing a crucial role in data analysis and sensemaking. Building on this, Andrews et al. [25]

demonstrated the importance of spatial arrangements in 2D analysis, coining the term “space to think” to describe how analysts leverage large, high-resolution displays for complex sensemaking tasks.

The benefits of spatial organization trivially extend to 3D environments. For example, space has been shown to support memory and recall [26]. Lisle et al. investigated sensemaking strategies in immersive “space to think” environments, examining how analysts can leverage 3D space for complex analytical tasks [27]. Davidson et al. extended this research by evaluating how immersive analytic systems support sensemaking tasks over time [28]. HydrogenAR [29] combines AR with 2D depictions, providing data-driven storytelling, spatially situated close to points of interest to provide additional contextualization. Likewise, research on immersive economics analysis using the ImAxes [8] system has provided valuable insights into how analysts utilize 3D space for data exploration and pattern discovery [10]. These studies contribute to our understanding of spatial organization strategies in immersive environments and inform the design of data visualization and analysis in AR/VR.

Although in our present work we do not implement specific spatial organization mechanisms, or recommendations for spatial arrangements, we were inspired by these insights to develop our immersive authoring and exploration techniques that leverage the spatial affordances of AR environments accessed through head-mounted displays. Moreover, our evaluation mirrors much of the methodology in Batch et al. [10], including preregistration with predictions regarding spatial organization, in addition to sensemaking, and presentation. With this, our goal was to see how our users would utilize the space available to them, unprompted, seeking to see if our predictions would be met and to what degree.

## 2.3 Visualization Authoring

The development of authoring tools for data visualization spans a wide spectrum, from general-purpose tools to specialized immersive environments. Grammel et al. [30] distinguish between different types of *visualization construction user interfaces*, providing a

framework for categorizing and analyzing these tools. This classification helps in understanding the diverse approaches to visualization authoring, including template-based galleries in spreadsheets such as Microsoft Excel, shelf-based authoring systems in tools such as Polestar and Tableau (Polaris [31]), visualization grammars such as Vega-Lite [12] and ggplot2 [32], and no-code authoring systems such as iVisDesigner [33], Lyra [34], Lyra 2 [35], Data Illustrator [36], and Charticulator [37]. Satyanarayan et al. [38] review such systems in terms of expressivity and learnability, which are challenges encountered in most authoring systems.

In addition, commercial systems, such as Microsoft Power BI, Tableau and Google Sheets make limited use of natural language processing for authoring, whereas research efforts such as from Wang et al. [39] explore related usage patterns. Building on this, and beyond relying on typed statements, other systems [40], [41] make use of voice input, and more recently large language models [42]. In this work, we integrate voice authoring using the OpenAI Whisper speech-to-text model and the OpenAI GPT-4o LLM. Last but not least, our work is influenced by the concept of block-based authoring, encountered in environments such as Scratch [43] and Alice [44], widely used in education, as well as more specialized applications such as haptic visualization authoring [45]. Block-based authoring has become the standard way to introduce learners to programming [46], [47], with several researchers exploring its use in immersive settings [48].

## 2.4 Authoring Tools for Immersive Visualizations

In the realm of immersive analytics, authoring tools are still in their infancy, with few general-purpose systems available. ImAxes [8], an early IA system, allows for authoring visualizations using direct manipulation, albeit with limitations on visualization types. IA tools and toolkits such as VRIA [49], IATK [50], DXR [51], and RagRug [52] provide frameworks for creating immersive visualizations. However, these toolkits are all limited to desktop-based authoring. DXR and VRIA both support changing the mapping between data dimensions and visual channels while in immersive space, but lack the ability to construct full Vega-Lite and 3D visualization specifications from scratch that our tool supports.

Recent advancements in authoring for immersive and ubiquitous analytics include Wizualization [53], which incorporates gestures and voice commands [54], [55], and NCARvis [56], which explores novel interaction techniques for the purpose. Bridging the gap between mobile and immersive authoring, Satkowski et al. [57] present an approach for in-situ authoring of AR visualizations using mobile devices. Their work demonstrates the potential for creating immersive visualizations in context, leveraging the ubiquity and familiarity of mobile interfaces while targeting AR environments.

Despite these advancements, current immersive authoring tools often struggle to match the flexibility and ease of use found in their 2D counterparts, particularly for non-expert users. Additionally, many existing systems fail to fully leverage the unique affordances of immersive environments, such as spatial organization and embodied interaction. Our present work builds upon these advances by focusing on creating intuitive and flexible tools for immersive and ubiquitous data visualizations in AR environments using declarative visualization grammars as the foundation.

## 3 IMMERSIVE VISUALIZATION AUTHORING

We designed and implemented a visual programming system for immersive and declarative authoring of data visualizations through

direct manipulation (see the accompanying video figure). Our system is realized in DashSpace [11], a live collaborative platform for immersive and ubiquitous analytics built on the web. It is accessible through the web browser and runs on desktop using a collaborative 3D space analysts can navigate, in handheld AR, or in immersive AR and VR. Borowski et al. [11] report on a series of proof-of-concept systems in DashSpace for authoring visualizations based on D3 [58] and Vega-Lite [12]. In this paper, we iteratively redesigned and refined the immersive visualization authoring technique based on the Vega-Lite grammar and added support for all four levels of the sensemaking process [14], [59].

### 3.1 Design Goals

We aim to enable analysts to create advanced visualizations entirely within an immersive analytics 3D environment. We define three design goals that aided our authoring design and implementation:

**DG 1—Enable Visual Programming with Declarative Grammars.** Analysts should be able to author visualizations using only immersive input devices, e.g., motion controllers or hand input, using direct manipulation. They should be able to stay in the flow [60], rather than working around syntax problems in a code editor. Declarative visualization grammars such as Vega-Lite [12] offer an extensive set of “puzzle pieces” that can be assembled to create expressive visualizations, and can be leveraged to allow analysts to reuse their existing knowledge. By employing a proximity-based design, these pieces can be combined by merely bringing them into proximity or dropping them onto each other.

**DG 2—Use Simple Composable Primitives and Nested Hierarchical Structures.** Elements in immersive space need to be simple, yet expressive enough to author meaningful visualization. Experiences from block-based programming languages suggest that visually simple, structurally clear, and directly manipulable primitives can reduce cognitive load and support learnability [47]. Particularly in AR, blocks should be easily movable and simple in representation to prevent clutter. This is especially important due to the lower pixels-per-degree resolution of HMDs compared to high-resolution display, especially at larger distances. This requires elements to be larger in size, cluttering the scene more easily. For more complex visualizations, we employ hierarchical structures inspired by Boxer [61]. The latter uses a *spatial metaphor* where boxes can be placed in a 2D canvas and nested into each other and *naive realism* where users interact with visual representations of computational objects. The spatial nesting mechanism allows analysts to reduce complexity in immersive space by hiding nested elements when they are not needed, while retaining the ability to unpack them if required.

**DG 3—Support All Sensemaking Loop Levels.** Another key requirement of an immersive authoring system for visualizations is to support all four levels of the sensemaking loop: (a) read, (b) explore, (c) schematize, and (d) report [14], [59]. Analysts should be able to conduct all four levels within immersive space.

### 3.2 DashSpace’s Proof-of-Concept Authoring

DashSpace [11] included a proof-of-concept authoring mechanism. This mechanism used four primitives including visualizations, specs, datasets, and snippets. Snippets—similar to *pieces* in this paper—were not mutable, requiring creation and deletion of multiple snippets to simply explore different variations of a visualization which led to a complex authoring experience. The

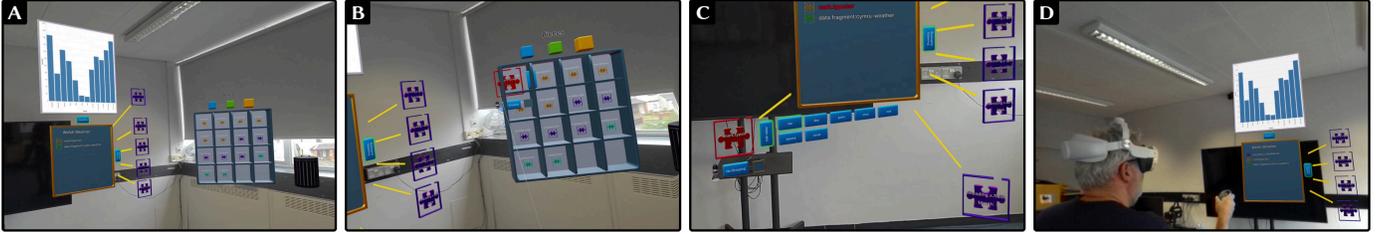


Fig. 2. **Authoring experience.** (A) The training scene of the user study with a visualization and a bookshelf. (B) An analyst creates a piece (red) by dragging it out of the bookshelf. (C) Once connected to a visualization, editors allow to modify the piece. (D) The view from the facilitator of an analyst editing the visualization.

*piece editors* in the new version make it possible to stay in XR while editing pieces. Furthermore, it did not have any notion of grouping or hierarchy, which made moving visualizations with multiple connected specs, datasets, and snippets in the scene more cumbersome and quickly cluttered the scene, as all these elements had to be visible all the time. The addition of groups and grouping overcome this challenge. Lastly, the initial version added friction by requiring users to switch to a desktop computer and use a code editor to make more complex changes, such as adding a filter to a visualization. Again, the updated *voice authoring* that is compatible with pieces, enables users to also author these more complex pieces in immersive space.

For our authoring system, we reused DashSpace’s overall architecture and technology stack to bootstrap our system, including the sticky note, bookshelf, and trashcan elements. We kept sticky notes, as they can help to organize content during the analysis task (DG 3b). The immersive visualization authoring mechanism, however, was redesigned based on the above design goals.

### 3.3 Our Authoring Experience in Brief

Our system, first, requires setting up a new space<sup>1</sup> and uploading data. This is most conveniently done on a desktop computer, but from thereon all authoring is conducted in immersive space.

In immersive space, analysts have access to a virtual bookshelf (Figure 2A) containing the available components for building visualizations from datasets, specs, and pieces, that can be combined in arbitrary ways as detailed further below. Authoring happens through analysts dragging out components from the bookshelf (Figure 2B) and spatially combining them either by grouping them by dropping them on one another, or through proximity of components (if enabled). The latter allows analysts to, e.g., toggle on and off a tweak to a visualization by moving components in and out of proximity of each other. Simple changes, such as changing the mark type, are done in immersive space using editors (Figure 2C). To change more complex parts of a Vega-Lite spec that traditionally would require code editors, analysts can use voice commands that are passed through an LLM (large language model) capable of editing Vega-Lite specifications.

### 3.4 Authoring Concepts

Our immersive authoring approach revolves around only two key concepts: *pieces* and *groups* (DG 2). Taking inspiration from Boxer [61], pieces can be merged into groups, added into existing

<sup>1</sup>. DashSpace is based on Webstrates [62] to author visualizations. Creating a new space is done by copying an existing (empty) space, that serves as a prototype for new spaces.

groups, and removed from groups. Groups can also be nested into each other, allowing for more complex specifications (DG 2).

We decided to use Vega-Lite as the visualization grammar (DG 1). We chose Vega-Lite over Vega and D3, as it is the highest-level language of the three and requires the least building blocks to generate visualizations, thereby reducing the complexity of the authoring experience. Hence, we removed the D3 capabilities of DashSpace, as D3 is too low-level for our needs.

#### 3.4.1 Pieces

A *piece* is a path-value pair that represents a part of a Vega-Lite specification. For instance, paths that a piece can store are the mark type (`mark.type`) of a visualization or the encoding field of the Y-axis (`encoding.y.field`). The path describes the path in the JSON of the Vega-Lite specification. The value part of a piece describes the value for this field of the Vega-Lite specification. The path and value of a piece are displayed on the front of the 3D icon (Figure 3A). The color of a piece depends on its path.

Pieces are created by dragging them out of a *bookshelf* (Figure 4A), causing a new copy of the piece to be instantiated. The bookshelf has three categories: specs, datasets, and pieces. Specs in the bookshelf are derived from Vega-Lite specifications that were loaded into the system. Datasets are derived from JSON datasets that analysts can upload. The pieces category, lastly, contains a hard-coded list of essential pieces required to create visualization in our system (marks, encodings, datasets, transforms). The system supports also other types of pieces, i.e., other specification paths; however, these will not support editors (see below) but only voice authoring. These other types of pieces can be created when importing and decomposing a Vega-Lite specification (Section 3.6.1), or when a group is modified using voice. Pieces and other elements can be removed from the scene by bringing them close to a *trashcan*.

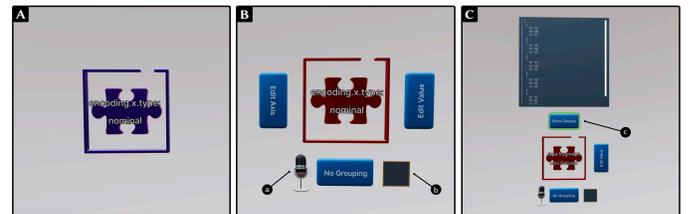


Fig. 3. **Pieces.** (A) An unselected piece that shows its path in the first line and value in the second. (B) Once selected, the piece shows additional options like buttons to edit it, prevent grouping, activate voice authoring (a) or to convert it to a group (b). (C) Pieces, which path includes data, have an extra button to preview the first five rows of the dataset (c).

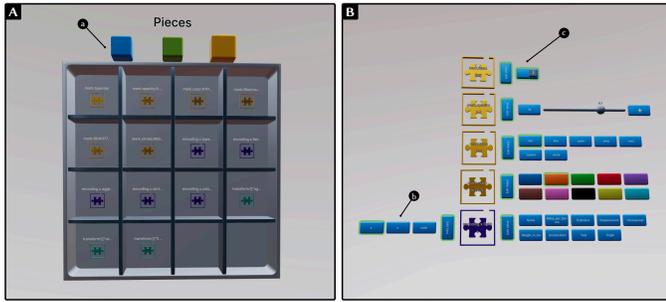


Fig. 4. **Bookshelf and editors.** (A) The bookshelf allows to instantiate pieces and groups in the scene. It comprises three categories (a): Specs, Datasets, and Pieces (the last being selected here). As seen in the figure, pieces have different colors depending on their path. (B) Five pieces with five value and one axis editor enabled. A button on the left allows to create an editor for the axis of encodings (b) and a button on the right allows to create editors (c) for different types (top to bottom): boolean, number, enumerated string, enumerated color, enumerated data field.

Pieces can be edited (DG 3b) using *editors* (Figure 4B). A button on the piece opens an editor for the value and, if the piece is an encoding, the button on the left opens an editor for its encoding axis (Figure 3B). We created a variety of editors for different types of values: boolean toggle, number slider, string option buttons, color option buttons, field option buttons, dataset option buttons. If the piece is a dataset, it has an option to inspect the first five rows (DG 3a) in an editor panel on top of the piece (Figure 3C).

### 3.4.2 Groups

A *group* is a container for pieces and other nested groups (Figure 5A). It can be named by analysts to keep track of what each group is doing (DG 3c). A paginated list inside the group element displays the pieces and nested groups that are contained within the group. Each entry in this list can be disabled using a toggle on the list entry; we will discuss why this might be useful in Section 3.5.1. A group can show a visualization on top of it and its currently composed specification (DG 3a) in a panel on the left side (Figure 5B)—these can be activated using the buttons on the top (Visualization) and left (Spec).

Groups can be created and deleted using the grouping authoring mechanism. The button on the right side of a group enables proximity authoring. Both authoring mechanisms will be explained in the next section.

## 3.5 Authoring Mechanisms

Our authoring system includes three main mechanisms to author visualizations: *grouping*, *proximity authoring*, and *voice authoring*. Each of these is designed for a different phase of the analytics process: grouping is designed for persistent changes to a visualization (DG 3c), while proximity authoring enables rapid experimentation during the exploration phase for basic changes (e.g., changing the mark type) to a visualization (DG 3b). The latter also enabling more interactive and embodied authoring. Voice authoring, lastly, can be used in both phases and, while also handy for simple changes, is essential for complex changes such as adding filter transforms, as these cannot be added or changed using editors.

### 3.5.1 Grouping

The *grouping* mechanism enables moving pieces and groups in and out of other groups. Groups can be created by moving one piece on

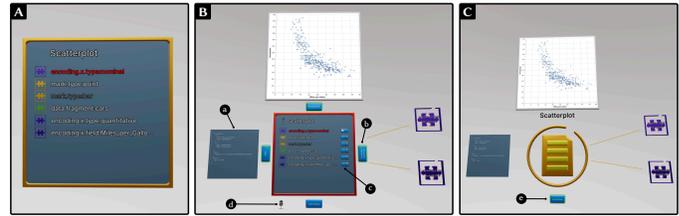


Fig. 5. **Groups.** (A) An unselected group that contains six pieces. The first piece is disabled, which is indicated by the red color. The third piece is grayed out, indicating that it is overwritten by the second piece in the list, which has the same path. (B) A selected group. It shows a preview of the composed spec on the left (a), a visualization on the top, and has proximity authoring activated (b), which connects two pieces outside the group. Inside the group list entries now show a toggle to disable them (c) and a microphone allows to rename the group by transcribing audio. The microphone below the group (d) allows to modify the whole group using voice. (C) A group that has been locked (e). It is not possible to add or remove pieces to it using grouping. The proximity authoring is still active.

top of another (i.e., moving them close together in the 3D scene), which will create a group with two pieces. Alternatively, a button in the bottom right of a piece allows to create a group with one piece (Figure 3B). More pieces or other groups can be added in the same way: by moving them on top of the target group.

Pieces or nested groups can be removed from a group by grabbing their handle—the icon on the left side of the list entries in the group—and dragging them out of a group. If only one list entry is left within a group, the group collapses and converts back to the piece or nested group that was within it.

Grouping can be disabled on individual pieces or groups so they are not dropped into one another, such as when moving multiple pieces with editors close to one another. For groups, this done by locking them, which also replaces the group body with an icon that prevents changing its contents (Figure 5C). Locking a group also visually cleans up the scene, hiding details that are not required for, e.g., a presentation of results (DG 3d).

### 3.5.2 Proximity Authoring

When exploring a visualization, *proximity authoring* allows for rapid experimentation and changes to a visualization (DG 3b). The basic design of proximity authoring was derived from DashSpace [11], however, it now can be toggled on or off on a group using the button on the right side. Once active, other pieces and groups within a 75 cm radius around the group are connected to it using lines (Figure 5B/C). The composed specification and visualization of the group will behave as if these connected pieces or groups are part of itself. Moving them in and out of this radius will update the visualization live.

Proximity authoring is disabled by default to prevent false positives when moving pieces in proximity that should not be connected. It can be toggled on during an experimentation phase, and toggled off once a visualization is finalized or almost finished. Analysts might also want to disable proximity authoring when they want to compare multiple visualizations in close proximity, as these would otherwise connect to each other and break the visualizations.

### 3.5.3 Voice Authoring

Our authoring system also provides *voice authoring* to support more complex changes to visualizations without having to write Vega-Lite code in immersive 3D space. Voice authoring is activated by selecting the microphone icon underneath a piece or group

(Figure 3B), which will then record and transcribe five seconds of audio. The transcribed audio is then passed to a LLM, along with information about the current path and value, if a piece was selected, or the current specification and dataset, if a group was selected. The model will then return an updated path and value for pieces, or an updated Vega-Lite specification.

Voice authoring is mainly intended for pieces like filter transforms which are highly flexible and can take many forms, and, thus, would be impossible to author using a simple GUI editor like the one used for numbers or boolean values. Typing these more complex queries, e.g., a filter, is cumbersome in immersive space. Voice, on the other hand, can be employed more easily than typing in XR (DG 3b); nevertheless, voice is not a full substitute.

### 3.6 Implementation

Our authoring system builds on DashSpace [11]. To convert between pieces and groups to Vega-Lite specifications, we created a *composition* and *decomposition* process.

#### 3.6.1 Specification Composition and Decomposition

We call the conversion of a Vega-Lite JSON specification into pieces *decomposition*, and the reverse process *composition*. These processes are required to generate pieces from imported specifications, and to generate composed specifications in groups.

The decomposition process splits a Vega-Lite JSON specification into one piece per each value in the JSON file. The path required to reach the value is combined into a dot-notated child—similar to how values in a JavaScript object are accessed. When decomposing a specification, we remove certain meta and styling properties that are not required or overwritten by our system. For instance, the size of a visualization is fixed in our system and we do not currently support multi-view visualizations.

The composition process, on the other hand, takes a list of pieces—path-value pairs—and combines them into a JSON specification. It, first, converts each piece into a JSON object and then deep merges these into a specification. If the same path appears in multiple pieces, the latter will overwrite the earlier ones, similar to how CSS style rules can be overwritten. The exception for this behavior are array properties like transforms, where array elements are merged into one array instead of overwriting each other.

#### 3.6.2 Editors and Voice Authoring

Editors of pieces are rendered based on the type of the property of the specification. To determine the types of properties, we combine retrieving the type and enumeration options from the JSON schema of Vega-Lite<sup>2</sup> with hard-coded values for some types.

Voice authoring is implemented using the OpenAI Whisper text-to-speech model and the OpenAI GPT-4o LLM. When used on a piece, the LLM receives a system prompt stating that it is an assistant to set the path and value of a Vega-Lite property. When used on a group, it receives a system prompt stating that it helps with modifying Vega-Lite specifications. In both cases, the LLM then receives a list of all fields in all datasets that are currently existing in the scene to provide it with knowledge about the datasets. Additionally, when used on a group, it receives the first data point of the dataset loaded in a group, if there is any. Lastly, the LLM receives the current path-value pair of a piece, or the current composed specification of a group, and the transcribed

prompt of analysts. It then proceeds to generate a new path and/or value for a piece, or a new specification for a group, which will overwrite the current values.

## 4 USER STUDY

To assess our immersive visualization authoring system, we conducted a preregistered user-based evaluation study.<sup>3</sup> The university where our user study was conducted requires all human-subjects research to be reviewed by a college-level ethics committee. We engaged fully with the ethics procedures of our institution and our study received ethical clearance (favorable opinion).

Our evaluation approach was driven by the following research questions: **RQ 1** (Authoring): What are the key design principles and guidelines for an authoring system that allows analysts to author interactive visualizations for IA/UA in HMD-based AR? **RQ 2** (Spatial Organization): What spatial organization strategies and layouts do analysts use to effectively integrate IA/UA visualizations in HMD-based AR? **RQ 3** (Sensemaking): How do analysts leverage IA/UA visualizations in HMD-based AR to support their sensemaking and visual exploration processes? **RQ 4** (Presentation): What presentation strategies do analysts use to effectively communicate insights gained from IA/UA visualizations in HMD-based AR?

### 4.1 Participants

We recruited 13 unpaid participants from within one of the author’s institution’s college of science and engineering using email and mailing lists. We specifically recruited participants who either had experience of Vega-Lite or some coding experience.

Overall, our participants had a varied educational level, as well as experience with programming, data analysis and analytical software (Figure 6). All reported normal or corrected to normal vision and none reported any color deficiencies.

### 4.2 Experimental Setup

The study took place in a 5.8 m × 5.6 m open-plan room with a roughly 3 m × 3.6 m usable space in the center (Figure 7A). Each study was conducted by two facilitators, F1 and F2, who assumed the same role in every session to ensure consistency. Participants, first, used a desktop computer with two screens (bottom-left of Figure 7A) for a screening test, later used by F2 for session observations. Upon completing the screening test, participants moved to a second table with a laptop, study information, and consent form (bottom-center of Figure 7A). The laptop was used to fill out the demographic questionnaire, the pre- and post-experiment Simulator Sickness Questionnaires (SSQ) [63], and the System Usability Scale (SUS) [64] questionnaire. Participants used a Meta Quest 3 and F1 a Meta Quest Pro HMD.

During the training, exploration task, and presentation, both participants and F1 used their HMD and shared the same DashSpace scene. F2 was seated at the PC desk to which the participants’ first-person view from the HMD was streamed. This enabled F2 to monitor the view of participants, give aid in case of problems, and observe their use of the system during the study. F2 was also responsible for starting and stopping recording and logging during the exploration task and presentation.

3. OSF Preregistration (created August 12, 2024): [https://osf.io/w5emj/?view\\_only=5a35f43d65f843ba9ee953ebe138c15c](https://osf.io/w5emj/?view_only=5a35f43d65f843ba9ee953ebe138c15c)

2. Vega-Lite JSON Schema: <https://vega.github.io/schema/vega-lite/v5.json>

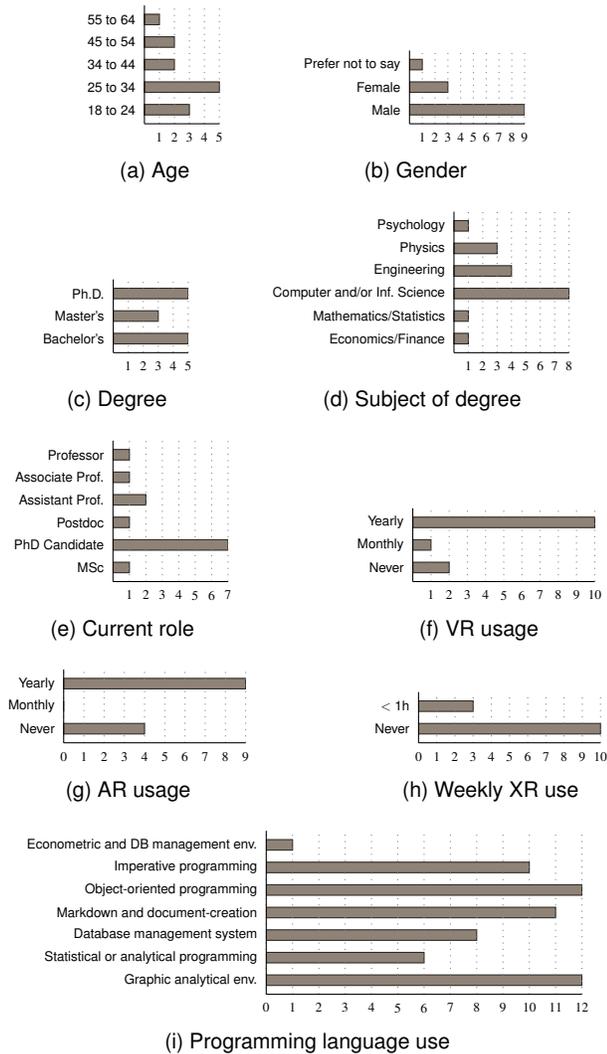


Fig. 6. **Participant profiles.** We recruited participants from the college of science and engineering of one of the authors' institution. Although most had experience with programming and analytical software and tools, their experience with XR was fairly limited.

While DashSpace [11] and our authoring mechanism support 3D visualizations using the Optomancy [53] toolkit, we decided to disable the feature for this user study to reduce complexity (Vega-Lite does not support 3D) and to focus on spatial organization and direct manipulation of elements in the virtual space instead.

### 4.3 Procedure

**Screening Test.** Before attending a study session, each participant was sent basic training material using a tutorial based on the Vega-Lite “Getting Started” guide.<sup>4</sup> On the day of the study, we used a simple screening test based on this tutorial to gauge whether participants had a sufficient understanding of Vega-Lite to reduce the potential for any learning effects associated with the grammar. They used a dataset on monthly average rainfall in Wales in 2018 for this test. Although we intended to disqualify any participants who did not complete our screening test within ten minutes, none of our participants needed more time to do so.

**Consent and Pre-Questionnaires.** After passing the screening test, participants were asked to re-read the information sheet and sign

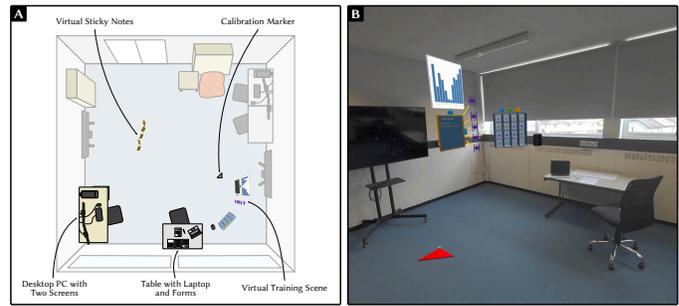


Fig. 7. **Room setup.** (A) A top-down view of the room used in the study. The starting layout of elements for the study are overlaid on the 3D scan. (B) The same setup from the first person perspective of the facilitator, standing on the left in front of the sticky notes.

the consent form. They then were asked to fill out the demographic questionnaire and the pre-experiment SSQ.

**AR Setup.** Participants were then asked to stand up,<sup>5</sup> put on their HMD, and start our system with a training scene. A physical triangle marker was placed on the floor of the room onto which participants were asked to align a virtual, movable marker within the AR scene. This would ensure the scenes of participants and F1 were aligned. Once participants indicated they were comfortable, F1 joined them in the scene using their HMD.

**Training.** During the training (Figure 7B), F1 explained all the supported features of the authoring system using an example chart in AR, using the same dataset used in the screening task. This way, participants were already familiar with the dataset and could focus on the authoring system's features and interactions. Once the guided training was completed, F1 instructed participants to independently recreate the example chart within the scene using the authoring system. Participants were allowed to ask questions while recreating the chart. Training was concluded once participants indicated they felt comfortable using the authoring system, and F1 and F2 were satisfied that all features had been demonstrated and that the training task had been completed satisfactorily.

**Exploration Task.** Participants remained in the same AR scene they used during training. F1 instructed participants to select a previously unseen dataset (monthly average rainfall in Wales from 2014–2023) from the bookshelf. Participants were asked to explore the new dataset and create charts using the authoring system, i.e., generate insights on the dataset. At this point F1 would move close to the participant's working area, moving textual prompts in the form of virtual sticky notes into view, e.g., “What's the wettest/driest month?” These prompts were designed to stimulate analysis, however, participants were free to ignore them and explore the dataset on their own. There was no time limit and participants were allowed to interact with the facilitators.

**Presentation.** At the end of the exploration task, participants were asked to give a short presentation in AR to F1, with a focus on interesting points they had identified about the data.

**Post-Questionnaires and Interview.** After the presentation, participants took off their HMD, and completed a post-experiment SSQ immediately followed by a SUS questionnaire and a semi-structured interview between the participant and F1.

4. Tutorial: [https://vega.github.io/vega-lite/tutorials/getting\\_started.html](https://vega.github.io/vega-lite/tutorials/getting_started.html)

5. Participants were allowed to sit back down whenever they wanted.

#### 4.4 Data Collection and Analysis

We collected questionnaire data on demographics, two SSQs, and one SUS questionnaire. During training, exploration task, and presentation, we recorded the video of participants’ first-person view, as well as audio in the room. During the exploration task, we also recorded telemetry and interaction logs, capturing the participant’s HMD position and rotation, system interactions, and the positions of objects in the scene over time. Finally, we recorded the audio of the interviews. The audio recorded while participants were in AR was transcribed using Whisper [65]. The audio from the interviews was recorded and transcribed using Otter.ai.

Telemetry and interaction logs provided us with a rich and multifaceted dataset to analyze how participants interacted with our system. We used the data to create a series of interactive visualizations using Vega-Lite and combined them in a web-based log analysis tool, which is accessible online.<sup>6</sup> The SSQ and SUS questionnaires were analyzed following their scoring methods.

The interview transcripts were coded and the utterances of participants were categorized top-down into four main themes related to the research questions: Authoring, Spatial Organization, Sensemaking, and Presentation. One independent coder conducted the analysis, with three co-authors independently calibrating on a subset of codes to ensure consistency. Within each research question we then created sub-themes bottom-up by clustering similar utterances and comments. For instance, “Drag and Drop Interaction” and “Voice Authoring” were sub-themes of Authoring, and “One Focused Workspace” and “Place Elements wherever there is Space” were sub-themes of Spatial Organization. We used the log analysis tool and the first person video recordings to supplement the emerging sub-themes. These sub-themes were then ordered and some of them combined into the results section below.

## 5 RESULTS

Our results are structured by our four research questions (Section 4) related to authoring visualizations with our system (RQ 1), how participants organized the space spatially (RQ 2), how our system supported participants in their sensemaking process (RQ 3), and how it supported presenting results in immersive AR (RQ 4).

Table 1 shows the varying amounts of time participants spent using the system. P10 and P11 did not complete the exploration task due to motion sickness after training, we excluded their data from the SUS results, as well as the log data analysis. However, as both P10 and P11 spent 32 and 30 minutes respectively using the system, we decided to include their qualitative feedback from the interview. The log data of P4, P7, and P13 was partially/completely lost due to technical issues and excluded from the log data analysis.

Except for only having eight participants with full log data instead of a target minimum of ten participants, we report no deviations from the preregistration.

### 5.1 Questionnaire Results

Our participant group completed the SSQ questionnaire [63] before their training session and after the exploration task. The calculated scores for the three symptom clusters, Nausea, Oculomotor, and Disorientation, along with the Total Severity, are summarized in Table 2. In the pre-questionnaire, all participants reported no symptoms, and, in the post-questionnaire, the majority of participants reported only slight symptoms. The exception were

TABLE 1

**Participant study overview.** The number of minutes participants spent in each phase, the number of visualizations they created during the exploration task, whether participants sat down during the exploration task, and whether their log data was included in the analysis.

|                     | P1 | P2 | P3 | P4 | P5 | P6 | P7 | P8 | P9 | P10 | P11 | P12 | P13 |
|---------------------|----|----|----|----|----|----|----|----|----|-----|-----|-----|-----|
| <b>Training</b>     | 21 | 11 | 21 | 38 | 19 | 22 | 26 | 26 | 25 | 20  | 30  | 23  | 27  |
| <b>Explor. Task</b> | 28 | 16 | 19 | 15 | 27 | 21 | 37 | 33 | 21 | 12  | –   | 10  | 13  |
| <b>Presentation</b> | 1  | 4  | 1  | 3  | 2  | 1  | 3  | 2  | 3  | –   | –   | 4   | 3   |
| <b>Vis. Created</b> | 1  | 1  | 1  | 1  | 2  | 2  | 3  | 1  | 1  | 1   | 0   | 1   | 1   |
| <b>Sat Down</b>     | ✓  | –  | –  | –  | –  | –  | –  | ✓  | –  | –   | –   | –   | –   |
| <b>Log Data</b>     | ✓  | ✓  | ✓  | –  | ✓  | ✓  | –  | ✓  | ✓  | –   | –   | ✓   | –   |

TABLE 2

**Simulator Sickness Questionnaire.** The resulting scores are labeled as *slight*, *moderate*, and *severe* (none recorded).

|             | Nausea |       | Oculomotor |       | Disorientation |        | Total Severity |        |
|-------------|--------|-------|------------|-------|----------------|--------|----------------|--------|
|             | Pre    | Post  | Pre        | Post  | Pre            | Post   | Pre            | Post   |
| <b>Mean</b> | 2.20   | 20.55 | 3.50       | 22.74 | 2.14           | 31.05  | 3.16           | 27.62  |
| <b>SD</b>   | 7.63   | 28.21 | 8.77       | 28.98 | 5.37           | 51.55  | 8.59           | 37.96  |
| <b>Min</b>  | 0.00   | 0.00  | 0.00       | 0.00  | 0.00           | 0.00   | 0.00           | 0.00   |
| <b>Max</b>  | 28.62  | 95.4  | 30.32      | 83.38 | 13.92          | 153.12 | 29.92          | 115.94 |

P10 and P11, who reported moderate sickness symptoms after the training/exploration task and did not complete the exploration.

A total of 11 participants (all except for P10 and P11) completed the SUS questionnaire, resulting in an average SUS score of 73.64 (SD = 12.37), which indicates a good level of usability, according to Bangor et al. [66]. Notably, one of the participants gave the system an overall score of 45, one of 62.5, and the rest scored it over 70. In particular, the lowest score from P1 can be traced to not feeling confident using the system (SUS Q9: “disagree”) and that they would need training before using it (SUS Q10: “agree”), with both responses markedly different than other participants. Due to the data anonymity we could not follow up with them after data analysis.

### 5.2 RQ1 — Authoring

The first research question focused on the piece-based authoring system we introduce in this paper. In the preregistration, we predicted that the authoring system is sufficiently expressive for participants to generate visualizations, that they would use proximity authoring for rapid and iterative exploration, while using grouping for persisted changes and schematizing. Lastly, we expected participants to use editors for simple changes to pieces and voice authoring for more complex queries. All of these predictions held up to varying degrees with most of the participants.

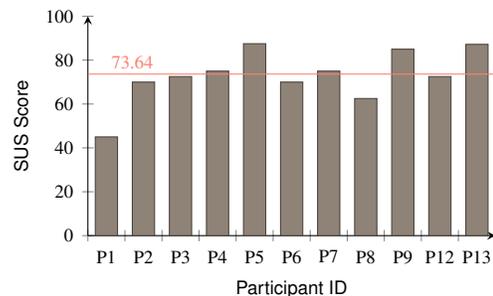


Fig. 8. **System Usability Scale plot.** The average score is 73.64 (SD=12.37), indicating “good” usability [66].

6. Log Analysis Tool: <https://demo.webstrates.net/tvcg-log-analysis/sub/>

Overall, most participants enjoyed the piece-based direct manipulation afforded by the authoring system. They used proximity authoring for quick experimentation and grouping for more finished visualizations. They appreciated the benefits of the live updating visualizations, yet raised concerns about the potential challenges of using Vega-Lite as the underlying grammar of the system.

### Visual Programming using Drag and Drop Largely Appreciated.

Most participants liked the piece-based visual programming direct manipulation interaction to author visualizations. They found it “*easier to drag and drop*” (P1) than to write code, and that the system took care of the syntax (P2, P3, P12). During the task, P7 said “*I feel like Iron Man*” when arranging elements in the scene.

Still, they also experienced challenges: While P11 appreciated that pieces have different colors depending on their path, P6 found it hard to differentiate pieces from each other, especially encodings where only a single letter decided whether it applies to the X- or Y-axis. P10 questioned whether the visual programming approach is the right one, as they found dragging and dropping elements “*tiring*” and a “*quite physical activity*.” P11 and P13, on the other hand, liked to “*physically touch the scene*” (P11) to grab and move pieces around, having only used HMDs for video games before.

P6 felt that it was “*like Lego blocks, to be able to add a piece in and then have it change the graph*.” P7, P9, and P12 highlighted how the piece-based interaction was a playful experience.

**Proximity Authoring, Grouping, and Voice Authoring.** Proximity authoring, as predicted, was used for both experimentation and exploration. Participants could make changes to visualizations using editors. For instance, P2 liked that it “*helped you to quickly put in pieces*” without having to be overly precise. P5 preferred proximity authoring over the grouping as “*it made it more intuitive*.” Still, arranging pieces around a group can quickly cause clutter, as P8 put it: “*eventually you’re going to have a large circle of sort of elements surrounding it, and things are going to get confusing*.”

This resonated with P9, who said that it got hard to remember what all these surrounding pieces were doing, which was exaggerated by them being hard to distinguish. P4 preferred to display the pieces in the group as a list, but disliked the need to remove pieces from the group in order to edit them using the editors.

Voice authoring was well received overall: P2 mentioned it helped to keep focus on the visualization, P4 sees potential in it to save time, and P5 thinks it is a good solution for XR environments where using keyboards “*is beyond a nightmare*.” Yet, P7 felt voice authoring was “*not so easy to use*,” as they sometimes did not know whether it was activated. This happened to multiple participants, which caused them to accidentally rename groups or talk too early or late when activating the voice authoring feature. P2 and P7 both suggested mapping a button on the controller to activate voice authoring instead. P9 also would have liked to use voice authoring more but felt it lacked better support to do so.

**Real-Time Editing and Vega-Lite Requirements.** In our authoring system, any change to a piece or a group immediately updates the visualizations to which it is connected. P5, P11, P12, and P13 mentioned liking this reactive way of instant updates. The editors of pieces were praised by P8 and P11, who liked the clear layout.

P9, however, felt they needed more guidance on how to combine pieces, stating that something like “*a prompt that basically suggested, right, you need fields and types for all your encodings*” would have helped them in getting started. This guidance would especially be useful for participants with little or no Vega-Lite knowledge. Multiple participants, including P5 and P10, noted that

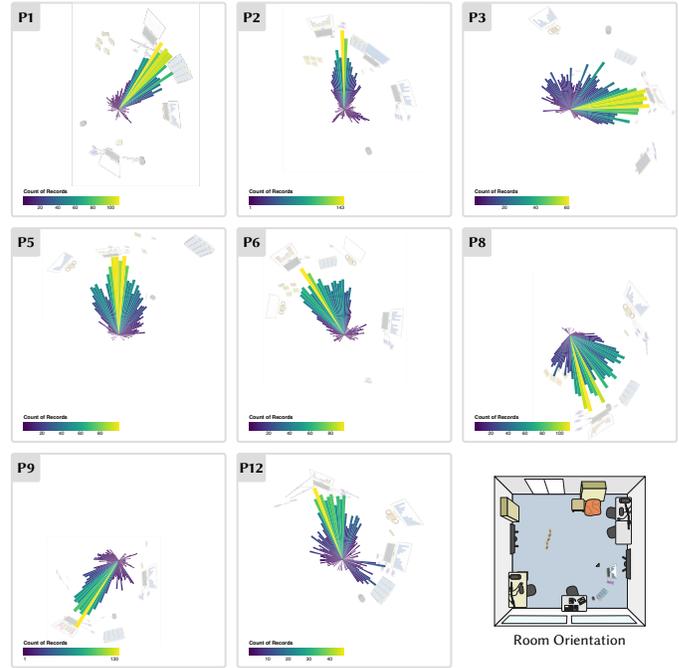


Fig. 9. **Participant rotation telemetry data.** The aggregated telemetry data of the Y-axis rotation of participants during the task and presentation phase. Most participants face one predominant direction. (The total count varies between participants due to different lengths of their study duration.) We superimposed a top-down view of their final 3D scene.

an understanding of the Vega-Lite semantics is required; in fact, P4 mentioned that “*you can’t use DashSpace to learn Vega-Lite*.” P5 stated that “*if I’d spent like, a little bit more time with Vega before, then, I would have found the whole thing very intuitive*.” However, once someone has Vega-Lite knowledge, P10 acknowledged using it for this authoring system and that they thought “*it was quite intuitive*.” Also, multiple participants commented that they enjoyed not having to deal with syntax. P4 stated that “*it took care of the syntax for me, just by moving the blocks around and checking the different menu items within that particular box. I did not have to worry about whether I got the syntax right or not*.”

### 5.3 RQ2 — Spatial Organization

The second research question focused on how participants organize the 3D space spatially and where they place elements. We can confirm the first three of our predictions: that people arrange elements in an egocentric way, that they place most elements between waist and head height, and that they used a default orientation in which the main work happened. However, the other two predictions were proven wrong, as participants neither drew semantic meaning from their physical environment, nor used flat surfaces as anchors.

We report that most participants used a central workspace to work on their visualizations, and that for most of them the background or location where they did it did not matter.

**Creating a Central Workspace.** Most participants created a central workspace where they did most of their visualization authoring. This is also supported by the telemetry data, which shows that all participants have a roughly 45-degree cone they were predominantly facing (Figure 9), and that most elements in the scene are placed at waist to head height (Figure 10), echoing findings from [10]. Multiple participants (P1, P2, P3, P5, P7, P9)

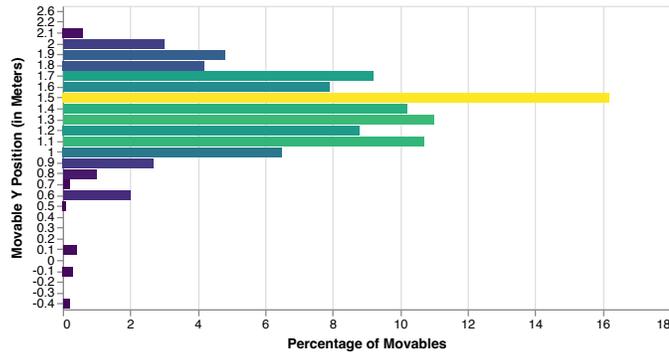


Fig. 10. **Element height data.** The aggregated height of elements in the scene over all participants during the task and presentation phase. Most elements are placed from waist to head height (1–2 m).

moved the bookshelf closer to the visualization they were working on. P3 explained they “*wanted just to have it easy access,*” while P8 explained they wanted to work in their “*foveal area, because you want to look at multiple things at the same time.*”

P13, on the other hand, felt the bookshelf is something of “*a fixed object*” and did not move it. P7 employed a model where they would have one group in front of them that they are actively working on, and once finished, move that group “*to the side, ideally such that [they] could line them up.*” One potential other reason for trying to keep everything in a central workspace is motion sickness induced by poor system performance (P10).

In contrast, P13 explained that they enjoyed “*having everything around you*” and compared it to having many tabs open in a web browser, and also did not mind turning to reach the bookshelf.

**Does the Physical Environment Matter?** In our predictions, we expected participants to use the physical environment as landmarks for placing elements; for instance, placing them on tables or shelves. However, to our surprise, none of the participants placed elements in this way but instead were just looking for an area of empty space to put their visualization. Many participants (P2, P6, P8, P9, P10) stated that they did not care about the background, with P8 putting it as: “*I was not focusing on the background at all. I focused only on foreground objects.*”

Only P11 mentioned that the background was important to them, but only when it was disturbing them: they moved pieces away from the table and window area as they found the background “*too busy.*” During the interview, P8 was the only participant to mention the possibility of using physical objects as an “*organizational tool.*”

P9, who was working mostly on the window side of the room (Figure 9) facing the seated facilitator, noted that they did not move towards the facilitator as they considered the motion controller a laser gun and “*you don’t point lasers at people.*”

#### 5.4 RQ3 — Sensemaking

This research question focuses on how working with visualizations in HMD-based AR supported their sensemaking process. For some participants, we can partially confirm the first prediction, in that they moved from inspecting the dataset (read), to exploring the visualizations using proximity authoring and grouping (explore), to creating copies of visualizations to compare them to answer different questions about the data (schematize), to presenting them in the end (report). While all participants that started the task did the first two and last phases, only some schematized their findings

through creating different copies of their visualizations and making small modifications to them. We, however, cannot meaningfully confirm or disprove the other two predictions, as—to our surprise—the task was too simple for participants to reach increasing levels of abstraction or iterate enough through different sensemaking levels.

We report on how participants used our system for prototyping and exploration of the dataset, and how they also saw value for it in teaching scenarios.

**Prototyping and Quick Experimentation.** The participants appreciated that the system allowed them to do prototyping and explore datasets. P4 found it “*quite useful for prototyping*” as one could “*quickly change the type of a chart or play around with some additional data or add filters.*”

Also, P5 saw future use of the system for “*experimenting with different visualizations*” to get “*an initial insight into what the data is telling you.*” P6 described how they could see the system especially useful for “*non-technical people*” who might prefer it over coding the Vega-Lite specification in a code editor.

**Potential for Teaching.** Multiple participants commented on potential use of our system for teaching and its use in schools. P1 and P12 compared our drag-and-drop authoring mechanism with the Scratch platform<sup>7</sup> and thought it is easier than coding, which might help in learning to create visualizations. P3 said they would be interested in seeing the system used in schools, also because they perceived the collaboration to work very well, and P11 and P12 considered it a useful tool for more visual learners who “*learn better by doing physical things than just sitting down at a desk*” (P11). Taking this further, P6 was confident that even “*eight-year-olds*” could use our system and “*would be able to get to grips with that a lot easier than they would with Vega.*”

#### 5.5 RQ4 — Presentation

The last research question focused on how participants were presenting results in the immersive space. We can confirm the first prediction that participants create exocentric views for presentations, for participants that created more than one visualization: Those usually placed them side-by-side and not, e.g., in a cylinder around them. We can also confirm the second prediction, as most participants placed elements between head and waist height during the task and presentation phase (Figure 10).

We report that participants were liking the collaborative features of our system in AR, and that they saw potential in giving presentations in a lecture context using our system.

**Collaborative Use and Presentations.** Multiple participants (P3, P6, P7, P8, P11, P13) mentioned the collaborative features of our system positively and as a potential future use case. P8, for instance, praised the possibility to collaborate together and “*rather than having two people stuffed over a screen, you can both, you have your own space, but you are able to see what’s going on.*”

P3 found it “*quite seamless*” to collaborate in the same space using the calibration marker to calibrate the 3D scenes. P11 said they could imagine teams using the system in remote scenarios. While generally mentioning collaboration as something positive, P13 also raised concerns, as moving pieces might appear sudden for other collaborators, causing disorientation.

**Teaching and Lecture Presentations.** Another use case that participants mentioned was to use our system in lecture presentations. The lecturer could show visualizations and modify it live while

7. Scratch: <https://scratch.mit.edu/>

giving the presentation (P3). P9 also raised that the space in such a context is greatly increased, “*rather than just having one screen, we’ve suddenly got the whole room to use as a screen.*” P13 echoed this by comparing lecturing in DashSpace as having “*more of like a gallery instead of screen.*”

## 5.6 Other Findings

A recurring issue that many participants (P1, P3, P5–P11) mentioned, was lag and stutter due to performance issues on the HMD. This caused various problems like increased motion sickness for some participants, and that some participants tried to move and look around less in the scene to reduce lag and keep or more stable frame rate. We were aware that our system is pushing the limits of the Meta Quest 3 HMD. Furthermore, streaming the scene from the participant’s perspective exacerbated the problem further.

Another issue some participants mentioned, was that legibility of some of the smaller text in the scene, for example those on pieces within the bookshelf. For P1, P6, P9, P10, and P13 much of the text was too small to comfortably read, and P4 had difficulty reading some text when looking at it at an angle. Though text was clearer when looked at close-up, participants were reluctant to move too close to objects.

## 6 DISCUSSION

We discuss our findings on authoring, exploring, and presenting visualizations in immersive 3D. Our results provide insights into how analysts interact with data visualizations in AR, highlighting both the potential and challenges of true immersive analytics.

### 6.1 Summary of Findings

Our predictions were largely confirmed, with some notable exceptions. Most interestingly, our spatial organization predictions regarding using the physical environment as landmarks to place virtual elements were disproved. While not expecting all participants to use the physical world for their spatial organization, we observed no one using these cues and instead acting mostly like they were in VR. Our expectation was based on how we, as developers of the system, used space during data analysis. Hence, one potential difference is that our participants were all new users of the system. More experienced users might use space differently.

One key finding was the flexibility of our tool, as evidenced by the diverse sensemaking, presentation, and organizational strategies employed by participants. Put differently, we were consistently amazed at the many and varied directions our participants took in analyzing data using our tool. This versatility suggests that our technique to immersive authoring and exploration can accommodate various individual preferences and working styles.

Using visual programming through direct manipulation for authoring visualizations was well-received by participants. This suggests that such interaction paradigms may be particularly suited to immersive AR environments. The proximity-based authoring was distinctly well received. While it, at first, may seem as a gimmick compared to the conventional mechanisms for authoring that are also available, it seemed to support an embodied user experience. This embodiment is key in what mixed reality can offer compared to conventional desktop computer use.

The voice authoring was well received and, as predicted, used for interactions that would otherwise have required a keyboard. We see it as a good fit for immersive analytics scenarios, as

typing on virtual keyboards is a hassle in XR. However, future work is required on more advanced feedback and error recovery mechanisms when using AI-assisted voice support. It is also not clear yet to what degree AI (in our case a LLM) should be used, i.e., would users prefer to speak natural language queries or rather dictate code in the syntax of the visualization grammar to reduce misunderstandings.

### 6.2 Organizational Strategies

Participants exhibited a range of organizational strategies. Some created multiple copies of visualizations for each question, while others manipulated a single visualization, adding or removing elements as needed. Most participants tended to work within a central area, showing reluctance to extensively rotate or move around the space. This behavior challenges the assumed advantage of having access to a large virtual workspace in immersive AR. However, it is possible that a more collaborative sensemaking task would yield a better use of the physical space.

Nevertheless, our overall observations align with previous research on spatial organization in data analysis, such as Andrews et al.’s [25] “Space to Think” and Batch et al.’s [10] work on immersive analysis. In general, people utilize physical space to simplify choice, perception, and internal computation [24]. We saw evidence of all three: placing all representations of a certain dataset on one side (choice), separating different representations (perception), and creating a central workspace at easily accessible heights to facilitate comparison and analysis (internal computation).

While the transition to immersive 3D and AR environments may introduce new considerations for spatial organization strategies, this did not hold up in our results. Contrary to our expectations, participants did not extensively utilize the physical space—such as physical landmarks and environmental features—when arranging virtual content. The lack of awareness of the physical surroundings may be a result of the “attentional tunneling” effect [67] in AR, where digital content takes attention away from physical surroundings. Syiem et al. [68] show that for mobile AR, this effect is task-dependent; the cognitively demanding task in our study may thus have amplified this effect. This raises questions about user proficiency and familiarity with immersive AR. Future research could explore whether more experienced users leverage space more effectively, or if this tendency persists regardless of expertise.

### 6.3 Limitations and Future Work

Our participant pool was predominantly male, which may limit the generalizability of our findings. While recruiting only participants with Vega-Lite experience, we did not specifically ask for their data visualization experience level, which might have affected how they used our system. Within this sample, qualitative themes became repetitive, particularly in the later sessions (P12 and P13), suggesting that thematic saturation had been reached. Performance issues with the AR system may have influenced participant behavior and preferences in, e.g., placing objects in ways to minimize lag induced simulator sickness. Furthermore, due to technical issues, some log data was lost, potentially affecting the comprehensiveness of our analysis. Future studies with more robust hardware and a more dependable logging mechanism could provide further insight.

In our study, the data was not designed to be situated, as there are no particular referents with which visualizations are to be associated. This was of course intentional to assess the authoring systems versatility as well as the spatial organization strategies the

analysts would follow. However, that may also be a factor of why our participating analysts did not employ any physical landmarks in our study when organizing their immersive space. Future studies could explore whether referent-rich settings would yield different spatial organization strategies.

Our tasks and dataset were not complex enough for some participants to make full use of the expressiveness of the system. In fact, participants learned to use our system faster than we expected. Participants did not have to create sophisticated visualizations, nor did they have to correlate findings from multiple displays concurrently. More complex sensemaking tasks may have further challenged the spatial organization of pieces and the use of the voice-based interaction. In other words, we underestimated our participants and their use of our system. Unfortunately, this limited our insights for the sensemaking research question and predictions. Nevertheless, we also did not observe any significant hindrances in our authoring technique that would affect sensemaking in a negative way. Furthermore, while DashSpace [11] does support 3D visualizations, we did not explore authoring of 3D visualizations in this paper. Nevertheless, we believe that many of our findings translate to 3D visualizations as well. The authoring paradigm of DashSpace is the same for 3D as 2D visualizations. The main difference might be the placement of groups in the scene—participants mostly placed them at head-height as the 2D visualizations could be viewed frontally, for the 3D visualization they might have placed them lower so they can view them from an downward angle.

Finally, while our immersive authoring approach aligns with our design goals, it represents just one possible implementation of such an authoring system. Declarative grammars such as Vega-Lite are indeed common in visualization and adopting them means leveraging prior research in the field, but it is not clear how our authoring approach compares to more traditional visual programming paradigms such as block-based, flow-based, or component-based programming.

Our study opens up several promising avenues for future research. For instance, the potential of immersive visualization authoring in teaching scenarios should be explored further. Exploring block-based drag-and-drop interactions in 2D interfaces alongside immersive 3D environments could provide valuable insights into the strengths and weaknesses of each approach. Additionally, it is unclear which of these potential strengths could also be achieved by large displays.

Further research could explore the integration of more concepts from visual programming or end-user programming. Exploring how multiple analysts might collaborate in shared immersive spaces could uncover new possibilities for team-based data analysis and presentation. Conducting longitudinal studies could help determine if and how spatial organization strategies evolve as analysts become more proficient with immersive AR analytics tools. Feedback on usability will help to improve the authoring system. This includes performance optimizations, text legibility, the size and ease of use of interactive objects, visual feedback for long-use components such as the microphone, and the mapping of features to physical controller buttons.

## 7 CONCLUSION

We have presented a generalizable visual programming technique for immersive authoring, exploration, and presentation of data visualizations, implemented through a spatial system based on direct manipulation. Our user study revealed insights into how analysts

author and use data visualizations in immersive AR environments, focusing on authoring mechanisms, spatial organization strategies, and sensemaking practices. Our findings revealed that participants particularly appreciated the proximity-based direct manipulation, suggesting its potential for intuitive interaction in immersive environments. Contrary to expectations, we observed that spatial organization was not primarily based on physical landmarks in the room, raising questions about the perceived benefits of expansive virtual workspaces in immersive AR—or even AR over VR in itself. The diversity of organizational strategies employed by participants highlights the flexibility of our tool in accommodating various individual preferences and working styles. This work contributes to the growing body of knowledge on immersive analytics and provides a foundation for future investigations into the design and implementation of truly immersive data visualization tools.

## ACKNOWLEDGMENTS

This work was partly supported by the Villum Investigator grant VL-54492 by Villum Fonden, the Aarhus University Research Foundation grant AUFF-E-2022-9-33, and the UKRI Centre for Doctoral Training in Artificial Intelligence, Machine Learning and Advanced Computing (AIMLAC) grant EP/S023992/1. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the funding agency.

## REFERENCES

- [1] N. Elmquist and P. Irani, “Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime,” *Computer*, vol. 46, no. 4, pp. 86–89, 2013.
- [2] T. Chandler, M. Cordeil, T. Czuderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, K. Marriott, F. Schreiber, and E. Wilson, “Immersive analytics,” in *Proceedings of the International Symposium on Big Data Visual Analytics*. Piscataway, NJ, USA: IEEE, 2015, pp. 1–8.
- [3] N. A. M. ElSayed, B. H. Thomas, K. Marriott, J. Piantadosi, and R. T. Smith, “Situating analytics: Demonstrating immersive analytical tools with augmented reality,” *Journal of Visual Languages & Computing*, vol. 36, pp. 13–23, 2016.
- [4] B. Lee, R. Dachselt, P. Isenberg, and E. K. Choe, *Mobile Data Visualization*, ser. AK Peters Visualization Series. Boca Raton, FL, USA: CRC Press, 2022.
- [5] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, Eds., *Immersive Analytics*, ser. Lecture Notes in Computer Science. New York, NY, USA: Springer Publishing, 2018, vol. 11190.
- [6] W. Willett, Y. Jansen, and P. Dragicevic, “Embedded data representations,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 3, pp. 2110–2128, 2017.
- [7] N. Elmquist, “Data analytics anywhere and everywhere,” *Communications of the ACM*, vol. 66, no. 12, pp. 52–63, 2023.
- [8] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott, “ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation,” in *Proceedings of the ACM Symposium on User Interface Software and Technology*. New York, NY, USA: ACM, 2017, pp. 71–83.
- [9] A. Batch, S. Shin, J. Liu, P. W. S. Butcher, P. D. Ritsos, and N. Elmquist, “Evaluating view management for situated visualization in web-based handheld AR,” *Computer Graphics Forum*, vol. 42, no. 3, pp. 349–360, 2023.
- [10] A. Batch, A. Cunningham, M. Cordeil, N. Elmquist, T. Dwyer, B. H. Thomas, and K. Marriott, “There is no spoon: Evaluating performance, space use, and presence with expert domain users in immersive analytics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 1, pp. 536–546, 2020.
- [11] M. Borowski, P. W. S. Butcher, J. B. Kristensen, J. O. Petersen, P. D. Ritsos, C. N. Klokose, and N. Elmquist, “DashSpace: A live collaborative platform for immersive and ubiquitous analytics,” *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–13, 2025.

- [12] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer, "Vega-Lite: A grammar of interactive graphics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 341–350, 2017.
- [13] D. Saffo, S. Di Bartolomeo, T. Crnovrsanin, L. South, J. Raynor, C. Yildirim, and C. Dunne, "Unraveling the design space of immersive analytics: A systematic review," *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 1, pp. 495–506, 2024.
- [14] S. Shin, A. Batch, P. W. S. Butcher, P. D. Ritsos, and N. Elmqvist, "The reality of the situation: A survey of situated analytics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 8, pp. 5147–5164, 2024.
- [15] M. Belcaid, J. Leigh, R. Theriot, N. Kirshenbaum, R. S. Tabalba, M. L. Rogers, A. E. Johnson, M. D. Brown, L. Renambot, L. Long, A. Nishimoto, C. North, and J. Harden, "Reflecting on the scalable adaptive graphics environment team's 20-year translational research endeavor in digital collaboration tools," *Computing in Science & Engineering*, vol. 25, no. 2, pp. 50–56, 2023.
- [16] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne, T. Dwyer, J. Grubert, J. H. Haga, N. Kirshenbaum, D. Kobayashi, T. Lin, M. Olaosebikan, F. Pointecker, D. Saffo, N. Saquib, D. Schmalstieg, D. A. Szafir, M. Whitlock, and Y. Yang, "Grand challenges in immersive analytics," in *Proceedings of the ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2021, pp. 459:1–459:17.
- [17] M. Cordeil, T. Dwyer, K. Klein, B. Laha, K. Marriott, and B. H. Thomas, "Immersive collaborative analysis of network connectivity: CAVE-style or head-mounted display?" *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 441–450, 2017.
- [18] B. Lee, X. Hu, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer, "Shared surfaces and spaces: Collaborative data visualisation in a co-located immersive environment," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1171–1181, 2021.
- [19] T. Gesslein, V. Biener, P. Gagel, D. Schneider, P. O. Kristensson, E. Ofek, M. Pahud, and J. Grubert, "Pen-based interaction with spreadsheets in mobile virtual reality," in *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*. Los Alamitos, CA, USA: IEEE Computer Society, 2020, pp. 361–373.
- [20] V. Biener, D. Schneider, T. Gesslein, A. Otte, B. Kuth, P. O. Kristensson, E. Ofek, M. Pahud, and J. Grubert, "Breaking the screen: Interaction across touchscreen boundaries in virtual reality for mobile knowledge workers," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 12, pp. 3490–3502, 2020.
- [21] B. Bach, R. Sicat, J. Beyer, M. Cordeil, and H. Pfister, "The Hologram in My Hand: How Effective is Interactive Exploration of 3D Visualizations in Immersive Tangible Augmented Reality?" *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 1, pp. 457–467, 2018.
- [22] Z. Chen, Y. Su, Y. Wang, Q. Wang, H. Qu, and Y. Wu, "MARVisT: Authoring glyph-based visualization in virtual reality for mobile augmented reality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 8, pp. 2645–2658, 2020.
- [23] W. Büschel, A. Mitschick, T. Meyer, and R. Dachsel, "Investigating smartphone-based pan and zoom in 3d data spaces in augmented reality," in *Proceedings of the ACM Conference on Human-Computer Interaction with Mobile Devices and Services*. New York, NY, USA: ACM, 2019.
- [24] D. Kirsh, "The intelligent use of space," *Artificial Intelligence*, vol. 73, no. 1–2, pp. 31–68, 1995.
- [25] C. Andrews, A. Endert, and C. North, "Space to think: large high-resolution displays for sensemaking," in *Proceedings of the ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2010, pp. 55–64.
- [26] G. G. Robertson, M. Czerwinski, K. Larson, D. C. Robbins, D. Thiel, and M. van Dantzich, "Data mountain: Using spatial memory for document management," in *Proceedings of the ACM Symposium on User Interface Software and Technology*. New York, NY, USA: ACM, 1998, pp. 153–162.
- [27] L. Lisle, K. Davidson, J. K. E. Gitre, C. North, and D. A. Bowman, "Sensemaking strategies with immersive space to think," in *Proceedings of the IEEE Conference on Virtual Reality*. Los Alamitos, CA, USA: IEEE Computer Society, 2021, pp. 529–537.
- [28] K. Davidson, L. Lisle, K. Whitley, D. A. Bowman, and C. North, "Exploring the evolution of sensemaking strategies in immersive space to think," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 12, pp. 5294–5307, 2023.
- [29] M. Whitlock, D. A. Szafir, and K. Gruchalla, "HydrogenAR: Interactive data-driven presentation of dispenser reliability," in *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*, 2020, pp. 704–712.
- [30] L. Grammel, C. Bennett, M. Tory, and M. D. Storey, "A survey of visualization construction user interfaces," in *Short Paper Proceedings of the Eurographics/IEEE VGTC Conference on Visualization*. Eindhoven, The Netherlands: Eurographics Association, 2013, pp. 19–23.
- [31] C. Stolte, D. Tang, and P. Hanrahan, "Polaris: A system for query, analysis, and visualization of multidimensional relational databases," *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 1, pp. 52–65, 2002.
- [32] H. Wickham, "A layered grammar of graphics," *Journal of Computational and Graphical Statistics*, vol. 19, no. 1, pp. 3–28, 2010.
- [33] D. Ren, T. Höllerer, and X. Yuan, "iVisDesigner: Expressive interactive design of information visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2092–2101, 2014.
- [34] A. Satyanarayan and J. Heer, "Lyra: An interactive visualization design environment," *Computer Graphics Forum*, vol. 33, no. 3, pp. 351–360, 2014.
- [35] J. Zong, D. Barnwal, R. Neogy, and A. Satyanarayan, "Lyra 2: Designing interactive visualizations by demonstration," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 304–314, 2021.
- [36] Z. Liu, J. Thompson, A. Wilson, M. Dontcheva, J. Delorey, S. Grigg, B. Kerr, and J. T. Stasko, "Data Illustrator: Augmenting vector design tools with lazy data binding for expressive visualization authoring," in *Proceedings of the ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2018, pp. 123:1–123:13.
- [37] D. Ren, B. Lee, and M. Brehmer, "Charticulator: Interactive construction of bespoke chart layouts," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 789–799, 2019.
- [38] A. Satyanarayan, B. Lee, D. Ren, J. Heer, J. Stasko, J. Thompson, M. Brehmer, and Z. Liu, "Critical reflections on visualization authoring systems," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 1, pp. 461–471, 2020.
- [39] Y. Wang, Z. Hou, L. Shen, T. Wu, J. Wang, H. Huang, H. Zhang, and D. Zhang, "Towards Natural Language-Based Visualization Authoring," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 1, pp. 1222–1232, 2023.
- [40] V. Setlur, S. E. Battersby, M. Tory, R. Gossweiler, and A. X. Chang, "Eviza: A Natural Language Interface for Visual Analysis," in *Proceedings of the ACM Symposium on User Interface Software and Technology*. New York, NY, USA: ACM, 2016, p. 365–377.
- [41] A. Srinivasan, N. Nyapathy, B. Lee, S. M. Drucker, and J. Stasko, "Collecting and Characterizing Natural Language Utterances for Specifying Data Visualizations," in *Proceedings of the ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2021, pp. 464:1–464:10.
- [42] H. K. Bako, A. Bhutani, X. Liu, K. A. Cobbina, and Z. Liu, "Evaluating the Semantic Profiling Abilities of LLMs for Natural Language Utterances in Data Visualization," in *Proceedings of the IEEE Conference on Visualization and Visual Analytics*, 2024, pp. 261–265.
- [43] M. Resnick, J. Maloney, A. Monroy-Hernández, N. Rusk, E. Eastmond, K. Brennan, A. Millner, E. Rosenbaum, J. Silver, B. Silverman, and Y. Kafai, "Scratch: programming for all," *Communications of the ACM*, vol. 52, no. 11, p. 60–67, Nov. 2009.
- [44] S. Cooper, W. Dann, and R. Pausch, "Alice: a 3-D tool for introductory programming concepts," *J. Comput. Sci. Coll.*, vol. 15, no. 5, p. 107–116, Apr. 2000.
- [45] S. A. Panéels, P. D. Ritsos, P. J. Rodgers, and J. C. Roberts, "Prototyping 3D haptic data visualizations," *Computers & Graphics*, vol. 37, no. 3, pp. 179–192, 2013.
- [46] D. Weintrop and U. Wilensky, "To block or not to block, that is the question: students' perceptions of blocks-based programming," in *Proceedings of the ACM Conference on Interaction Design and Children*. New York, NY, USA: ACM, 2015, p. 199–208.
- [47] D. Bau, J. Gray, C. Kelleher, J. Sheldon, and F. Turbak, "Learnable programming: blocks and beyond," *Communications of the ACM*, vol. 60, no. 6, pp. 72–80, 2017.
- [48] M. Hedlund, A. Jonsson, C. Bogdan, G. Meixner, E. Ekblom Bak, and A. Matvienko, "BlocklyVR: Exploring block-based programming in virtual reality," in *Proceedings of the ACM Conference on Mobile and Ubiquitous Multimedia*. New York, NY, USA: ACM, 2023, p. 257–269.
- [49] P. W. S. Butcher, N. W. John, and P. D. Ritsos, "VRIA: A web-based framework for creating immersive analytics experiences," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 7, pp. 3213–3225, 2021.
- [50] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer, "IATK: An immersive analytics toolkit," in *Proceedings of the IEEE Conference on Virtual Reality*. Los Alamitos, CA, USA: IEEE Computer Society, 2019, pp. 200–209.

- [51] R. Sicat, J. Li, J. Choi, M. Cordeil, W.-K. Jeong, B. Bach, and H. Pfister, "DXR: A toolkit for building immersive data visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 715–725, 2019.
- [52] P. Fleck, A. Sousa Calepso, S. Hubenschmid, M. Sedlmair, and D. Schmalstieg, "RagRug: A toolkit for situated analytics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 7, pp. 3281–3297, 2023.
- [53] A. Batch, P. W. S. Butcher, P. D. Ritsos, and N. Elmqvist, "Wizualization: A "hard magic" visualization system for immersive and ubiquitous analytics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 1, pp. 501–517, 2024.
- [54] S. Jang, N. Elmqvist, and K. Ramani, "MotionFlow: Visual abstraction and aggregation of sequential patterns in human motion tracking data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 21–30, 2016.
- [55] —, "GestureAnalyzer: visual analytics for pattern analysis of mid-air hand gestures," in *Proceedings of the ACM Symposium on Spatial User Interaction*. New York, NY, USA: ACM, 2014, pp. 30–39.
- [56] K. Cheng, J. Pu, Z. Cheng, J. Huang, and X. Cong, "NCARVis: No-code visualization creation system based on free-hand," in *Proceedings of the IEEE Pacific Symposium on Visualization*. Los Alamitos, CA, USA: IEEE Computer Society, 2023, pp. 1–5.
- [57] M. Satkowski, W. Luo, and R. Dachsel, "Towards in-situ authoring of AR visualizations with mobile devices," in *Adjunct Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*. Los Alamitos, CA, USA: IEEE Computer Society, 2021, pp. 324–325.
- [58] M. Bostock, V. Ogievetsky, and J. Heer, "D<sup>3</sup>: Data-Driven Documents," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2301–2309, Dec. 2011.
- [59] D. M. Russell, M. Stefik, P. Pirolli, and S. K. Card, "The cost structure of sensemaking," in *Proceedings of the ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 1993, pp. 269–276.
- [60] N. Elmqvist, A. Vande Moere, H.-C. Jetter, D. Cernea, H. Reiterer, and T. Jankun-Kelly, "Fluid interaction for information visualization," *Information Visualization*, vol. 10, no. 4, pp. 327–340, Oct. 2011.
- [61] A. A. diSessa and H. Abelson, "Boxer: A reconstructible computational medium," *Communications of the ACM*, vol. 29, no. 9, pp. 859–868, 1986.
- [62] C. N. Klokrose, J. R. Eagan, S. Baader, W. Mackay, and M. Beaudouin-Lafon, "Webstrates: Shareable dynamic media," in *Proceedings of the ACM Symposium on User Interface Software and Technology*. New York, NY, USA: ACM, 2015, pp. 280–290.
- [63] K. S. B. Robert S. Kennedy, Norman E. Lane and M. G. Lilienthal, "Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness," *The International Journal of Aviation Psychology*, vol. 3, no. 3, pp. 203–220, 1993.
- [64] J. Brooke, "SUS - a 'quick and dirty' usability scale," in *Usability Evaluation In Industry*. Boca Raton, FL, USA: CRC Press, 1996, pp. 189–194.
- [65] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022. [Online]. Available: <https://arxiv.org/abs/2212.04356>
- [66] A. Bangor, P. Kortum, and J. Miller, "Determining what individual SUS scores mean: Adding an adjective rating scale," *Journal of Usability Studies*, vol. 4, no. 3, p. 114–123, 2009.
- [67] C. D. Wickens and A. L. Alexander, "Attentional tunneling and task management in synthetic vision displays," *International Journal of Aviation Psychology*, vol. 19, no. 2, pp. 182–199, 2009.
- [68] B. V. Syiem, R. M. Kelly, J. Goncalves, E. Velloso, and T. Dingler, "Impact of task on attentional tunneling in handheld augmented reality," in *Proceedings of the ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2021, pp. 193:1–193:14.



**Peter W. S. Butcher** received the Ph.D. degree in 2020 from the University of Chester, UK. He is a Lecturer in the School of Computer Science and Engineering, Bangor University, UK. His research interests include human-computer interaction, information visualization, and immersive analytics.



**Shannon J. Jones** received the Bachelor's degree in computer information systems in 2023 from Bangor University, UK. She is a Ph.D. candidate in the School of Computer Science and Engineering, Bangor University, UK. Her research interests include internet of things, mixed reality, and artificial intelligence.



**Panagiotis D. Ritsos** received the Ph.D. degree in 2006 from the University of Essex, Colchester, United Kingdom. He is a Senior Lecturer in the School of Computer Science and Engineering, Bangor University, United Kingdom. His research interests include mixed and virtual reality, information visualization, and visual analytics. He is a member of the IEEE, the IEEE Computer Society, and the ACM.



**Clemens N. Klokrose** received the Ph.D. degree in computer science in 2009 from Aarhus University in Aarhus, Denmark. He is an Associate Professor in the Department of Computer Science at Aarhus University. His research interests include human-computer interaction, interactive systems, and computational media.



**Niklas Elmqvist** received the Ph.D. degree in 2006 from Chalmers University of Technology in Göteborg, Sweden. He is a Villum Investigator and professor in the Department of Computer Science at Aarhus University in Aarhus, Denmark. He was previously faculty at University of Maryland, College Park from 2014 to 2023, and at Purdue University from 2008 to 2014. His research interests include visualization, HCI, and human-centered AI. He is a Fellow of the IEEE and the ACM.



**Marcel Borowski** received the Ph.D. degree in computer science in 2022 from Aarhus University in Aarhus, Denmark. He is a Postdoctoral Research Fellow in Department of Computer Science at Aarhus University. His research interests include human-computer interaction, interactive systems, and computational media.