

This is the Table I Want! Interactive Data Transformation on Desktop and in Virtual Reality

Sungwon In, Tica Lin, Chris North, Hanspeter Pfister, *Fellow, IEEE*, and Yalong Yang

Abstract—Data transformation is an essential step in data science. While experts primarily use programming to transform their data, there is an increasing need to support non-programmers with user interface-based tools. With the rapid development in interaction techniques and computing environments, we report our empirical findings about the effects of interaction techniques and environments on performing data transformation tasks. Specifically, we studied the potential benefits of direct interaction and virtual reality (VR) for data transformation. We compared gesture interaction versus a standard WIMP user interface, each on the desktop and in VR. With the tested data and tasks, we found time performance was similar between desktop and VR. Meanwhile, VR demonstrates preliminary evidence to better support provenance and sense-making throughout the data transformation process. Our exploration of performing data transformation in VR also provides initial affirmation for enabling an iterative and fully immersive data science workflow.

Index Terms—Immersive Analytics, Data Transformation, Data Science, Interaction, Empirical Study, Virtual/Augmented/Mixed Reality



1 INTRODUCTION

Data transformation is a data science process that converts a data set into the desired format to enable subsequent data science tasks, like visualization and modeling [1]. It is well recognized that data scientists need to spend an excessive amount of time doing data transformation, making it essential but also the most tedious and time-consuming aspect of a data science project [2], [3]. Using a programming language, like SAS, R, or Python, is the standard way of performing data transformation. However, as data science becomes ubiquitous and exposed to people with limited programming knowledge, the prerequisite of knowing to program makes data science inaccessible to a large group of professionals whose workflows involve data [4], which we called *non-technical data workers*. As a result, the back-and-forth communications caused by data science’s iterative and open-ended nature can heavily inhibit insight discovery and decision-making.

In response, like in many other data science processes, there is an increasing trend of providing user interface based (UI-based for short) tools for data transformation (e.g., Tableau [5] for visualization and AutoML [6], [7] for modeling). These UI-based tools lower the entry barrier for data science and also help reduce errors [8]. However, even though there exist several commercial UI-based data transformation tools (e.g., Tableau Prep Builder [9], Trifacta [10], and Alteryx [11]), the field lacks an empirical understanding of people’s experiences in using these tools and the con-

siderations in designing them. Moreover, while the WIMP (windows, icons, menus, pointer) metaphor is typically used for constructing UI-based tools, modern interaction techniques that allow direct manipulation of the visual elements in the same space (named *embedded interaction*) [12], [13] were found to be more time-efficient in specific scenarios, like manipulating visualizations [14], [15]. Most notably, Kandel et al. [2] and Nandi et al. [16], [17], [18] found their UI-based tools to be more efficient in low-level data transformation tasks. However, it is unclear if their findings can be generalized to more realistic and complicated scenarios. **Our first goal** is to investigate the potential benefits of embedded interaction techniques over traditional WIMP interfaces in more realistic data transformation tasks.

On the other hand, in addition to interaction techniques, the rapidly evolved display and interaction environments (e.g., virtual and augmented reality or VR/AR) offer tremendous opportunities for creating innovative human-computer interaction experiences. Specifically, there is a growing interest in using VR/AR for data analysis, bringing an emerging research topic — Immersive Analytics [19], [20]. From recent studies, there are two most frequently reported motivations for using VR/AR in analytics: *large display space* [21], [22], [23], [24] and *embodied interaction* [25], [26], [27]. We believe there is great potential to explore whether those identified benefits can be generalized in improving the data transformation workflow. On the other hand, standard mid-air methods in VR/AR are not suitable for tasks requiring high-precision interactions [28], [29], which could be inevitable in some data transformation tasks. Therefore, **our second goal** is to investigate how these identified pros and cons can affect data transformation tasks in immersive environments.

Data science is iterative by its nature and does not follow a sequential pipeline. Consequently, alternating between different steps is inevitable [1]. For example, after observing some visualizations, analysts may need to perform extra data transformations for the next analysis iteration.

- S. In and C. North are with the Department of Computer Science, Virginia Tech, Blacksburg, VA, 24060. E-mail: {sungwoni, chnorth1}@vt.edu
- Y. Yang is with the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, 30309, and was with the Department of Computer Science, Virginia Tech. E-mail: yalong.yang@gatech.edu
- T. Lin and H. Pfister are with John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, 02138. E-mail: {tlin, pfister}@g.harvard.edu
- Y. Yang is the corresponding author.

Manuscript received April 19, 2005; revised August 26, 2015.

Although visual exploration (e.g., ImAxes [26] and DataHop [30]) is feasible in a fully immersed manner, there is no immersive data transformation tool (i.e., tools for explicitly changing data table formats). When analysts want to use immersive visualization, they have to switch between VR and desktop to complete the iterative data science tasks, causing high overhead for context-switching. To this end, we study immersive data transformation tools to progress to a future where analysts can be fully immersed in VR for the entire data science workflow and maximize the benefits of the next generation of display and interaction environment.

To close these gaps, we developed prototypes with embedded interactions on the desktop and embodied interactions in VR for non-technical data workers to support essential data transformation operations. Compared to a standard WIMP user interface, users can directly manipulate data tables through mouse or physical gestures (e.g., overlay one table on top of another to *merge* them, see Fig. 1). We compared our interaction designs to WIMP for desktop and VR. To best simulate real-world scenarios, instead of testing low-level tasks, we asked participants to transform a set of data tables into a target format. We found that participants required a similar amount of time to complete data transformation in VR and on a desktop. Meantime, VR demonstrated the potential to facilitate strategic thinking and support provenance better. Subjectively, participants found the WIMP user interface on a desktop most familiar, and using VR was more physically demanding. On the positive side, VR was perceived as more engaging, and participants overall preferred the gesture-based experience in VR. **The contributions of this paper are twofold:** *first*, the designs of gesture-based interactions for essential data transformation operations on desktop and in VR; *and second*, a user study systematically investigating the effect of interaction methods (WIMP vs. gesture) and computing environments (desktop vs. VR) on performing essential data transformation operations.

2 RELATED WORK

Our work is built upon two lines of prior work: data transformation tools and interaction methods. We also extend immersive analytics by enabling data transformation, an essential data science workflow, in immersive environments.

2.1 Data Transformation Tools

Programming-based tools are widely used by people with expertise and experience in programming. A wide range of libraries has been developed to support data transformation, like Pandas [31], dplyr [32], tidyr [33], and plyr [33]. Programming-based tools are *expressive*, and users can use their almost “exhaustive” APIs and parameters to complete various data transformation tasks. However, mastering these tools requires extensive training and leads to a steep learning curve. Debugging complicated scripts is also oftentimes challenging [34]. There are a series of attempts to address these issues. As representative examples, DataLore [35] provides code suggestions to speed up the data transformation process. Along a similar line, Wrex [36] uses the notion of programming-by-example to generate data transformation code. On the other hand, Somnus [34] visualize data transformation scripts to help debug and gain

a better overview of the process. Yet, users of these tools are still expected to be experienced in programming, which excludes a large group of non-technical data workers.

UI-based tools allow people to manipulate their data without programming knowledge. Microsoft Excel is undoubtedly the most popular UI-based data transformation tool. Performing simple operations (e.g., editing values, sorting and filtering a column) is straightforward with its WIMP user interface. However, more complicated operations (e.g., merging two tables) requires knowing the specific “secret” menu item or writing code. A few commercial UI-based tools aim to allow the users to quickly find the needed menu items, like Tableau Prep Builder [9], Trifacta [10], and Alteryx [11]. Wrangler [2], meanwhile, provides natural language as ways users can specify the intended operations. However, users could get into trouble with discoverability (i.e., the ability to find and execute features) as a common challenge faced by a conversational user interface [37]. GestureDB project enables gestures to describe the intended database queries [16], [18]. Our gesture design on the desktop environment shares many similar characteristics with their system, and we adapted and extended the gesture-based data transformation method to VR.

Empirical results. Some of the proposed tools have been evaluated, for example, Wrex was found more beneficial than a standard programming interface [36], and Wrangler was found to outperform Excel [2]. Additionally, the GestureDB system was found to be more effective than the programming interface and non-gesture UI-based interface in performing single operations [16], [18]. We focus on UI-based data transformation tools as they lower the barrier for non-technical data workers. Our study aims to enrich the empirical understanding of using UI-based data transformation tools in tasks that require a series of operations.

2.2 Embedded and Embodied Interactions

We consider both embedded and embodied interaction under the same notion of *direct manipulation* of visual representations. Direct manipulation contrasts with the standard WIMP UI design, which requires users to trigger operations on a space-separated area different from the area with visual representations.

Performing direct manipulation on a flat screen is considered **embedded interaction**. It has been widely used in many UIs. For example, when uploading an email attachment, instead of clicking a button to select a file from a newly opened file browser, people can drag&drop the file into the window. In addition to this simple example, it has been used in many other applications, like annotation [38], image editing [39], and content organization [40]. For data science, some work explored embedded interaction in manipulating data visualization [12], [14], [15], [41], [42], analysis [43], [44] and modeling [45]. Most relevant to our work, GestureDB demonstrated some benefits of embedded interaction for elementary operations [16], [18], and we aim to study its potential benefits in more complex data transformation tasks.

Performing direct manipulation using body movement is considered **embodied interaction**. The ability to track physical movement is essential to enable embodied interaction, which is an intrinsic characteristic of VR. As a result,

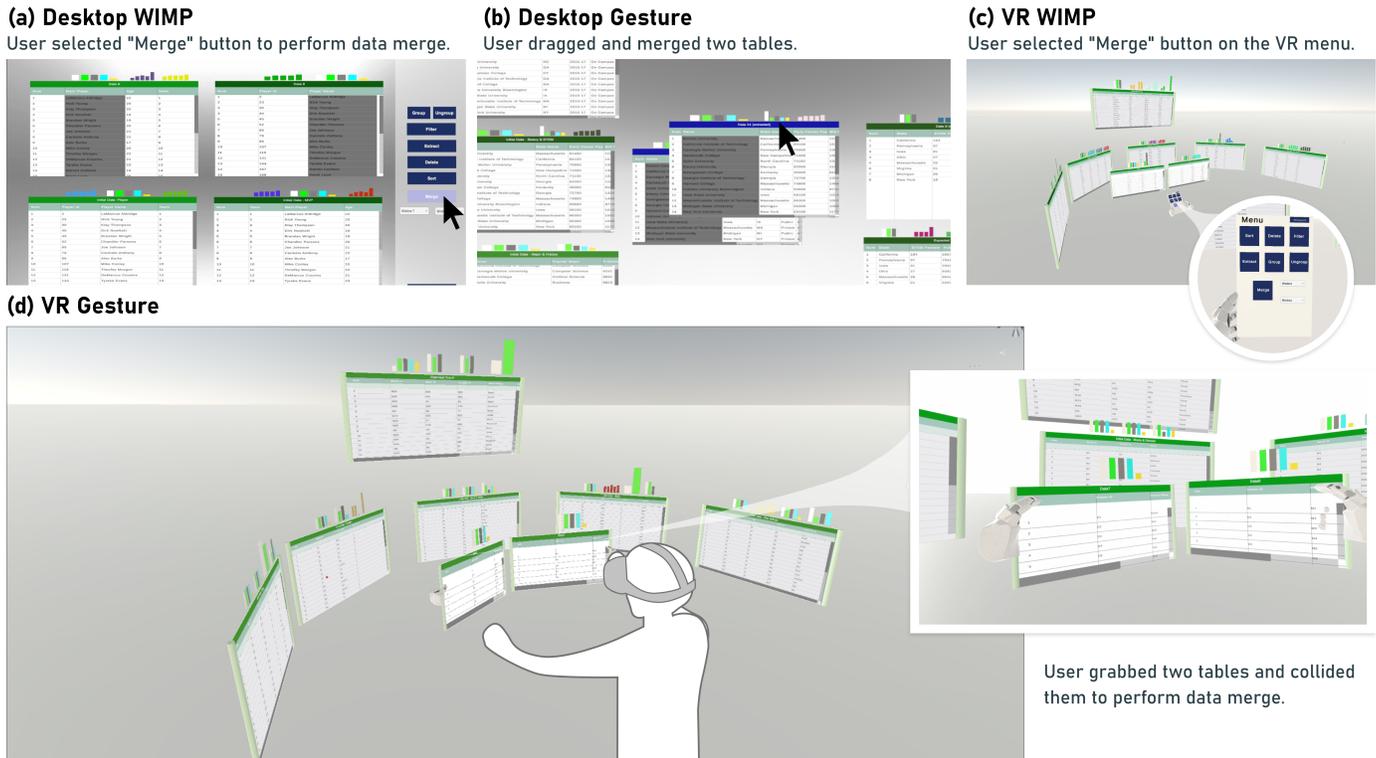


Figure 1. Four conditions designed for performing data transformation in the user study, including a combination of desktop or VR environments, and WIMP or gesture interactions.

many basic VR interactions are embodied. For example, grab&move virtual objects, and rotate the head to change the viewpoint. Embodied interaction has been explored for authoring visualizations [26], navigating in space [23], [46], and switching between different views [25]. We are interested in how we can adapt and extend the gesture designs from desktop to VR and whether the benefits of embodied interaction can be generalized to data transformation tasks.

As evaluated in the aforementioned works, one of the motivations of *direct manipulation* design is to reduce the number of context-switching needed by having the interaction and visual representation in the same display area. We are interested in if this identified benefit can facilitate data transformation tasks.

2.3 Immersive Analytics Toolkits

Immersive Analytics has exploded into a fast-growing body of research on techniques and toolkits [19]. Existing Immersive Analytics research strongly focuses on data visualization [20], [47]. Specifically, a few toolkits enable data scientists to create immersive data visualizations, including DXR [48], VRIA [49], IATK [50], DataHop [30] and ImAxes [26]. While ImAxes and DataHop provide a fully immersive visualization authoring experience, the others require users to create and configure visualizations on the desktop and view visualizations in the immersive environment. More importantly, a data science project is oftentimes iterative and includes more than just data visualization. Immersive data visualization alone cannot fully leverage immersive environments for data analysis. In this study, we explore how we can enable immersive data transformation to make progress toward a fully immersive data science workflow.

3 EMBEDDED AND EMBODIED GESTURE DESIGN

Our study is intended to investigate the opportunities in using novel interaction methods (i.e., embedded and embodied Gesture) and emerging computing environments (i.e., VR) to better support data transformation. We first reviewed the literature to identify the necessary operations and then designed gestures for both Desktop and VR.

3.1 Selecting Data Transformation Operations

Kasica et al. [51] summarized 21 fine-grained data transformation operations across five categories (i.e., create, delete, transform, separate, and combine) on three targets (tables, columns, and rows), see Table 1. As the first attempt to compare data transformation experience on desktop and in VR, we focus on basic operations that are commonly used by non-technical data workers without the need for programming. To this end, we excluded operations that require programming-like input, including *create*, *transform*, *separate*, and *combine* operations on rows and columns. For example, one typical *combine columns* operation can be inputting a formula to calculate the weighted average of selected columns. One exception was the *summarize rows* operation, as it does not require inputting a formula.

In summary, we support 12 operations, including all nine table, one-column, and two-row operations. See Table 1.

3.2 Designing Gestures

After analyzing the selected operations, we found that some operations have differences at the semantic level but imply the same interaction analogy. Specifically, Kasica et al. described *extend*, *supplement*, and *match* as operations to

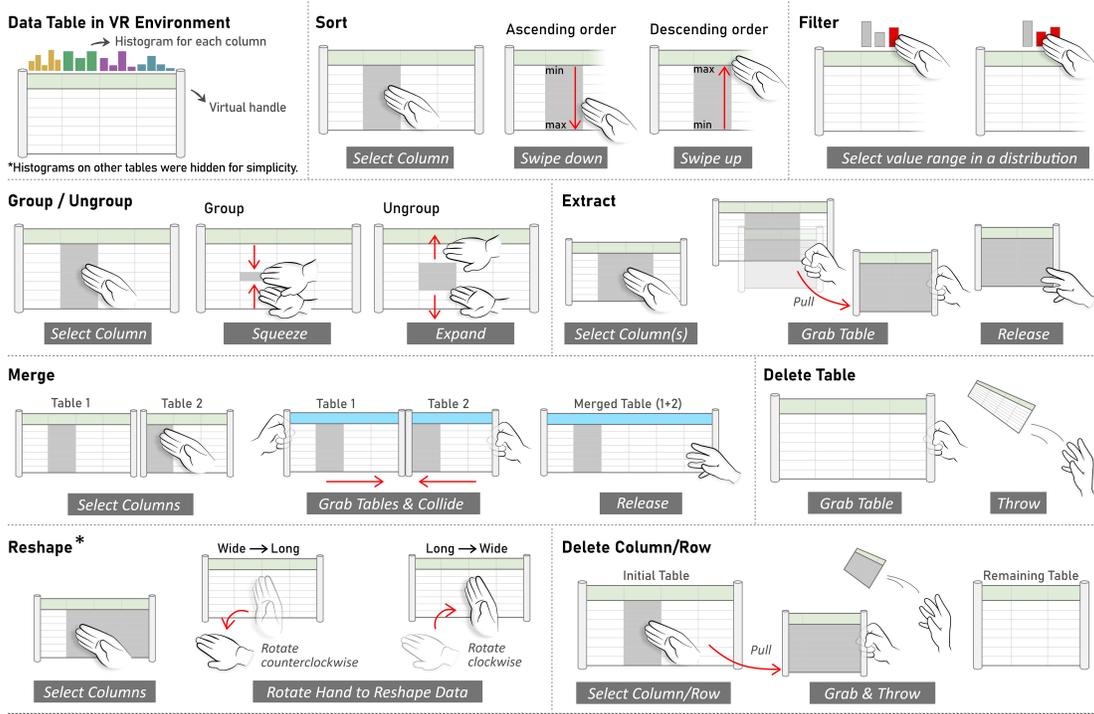


Figure 2. VR data table and gestures for data transformation tasks in the VR+Gesture condition. VR controllers, represented as hand models in the VR prototype and in the figure above, are used to perform these operations. The operation that was excluded from the study is marked with *.

Table 1

A summary of our supported data transformation operations and their matching gestures.

	Operation	Gesture
Create	create tables	✗
	create columns	✗
	create rows	✗
Delete	delete tables	delete tables
	delete columns	delete columns
	delete rows	delete rows
Transform	rearrange tables	sort
	reshape tables	reshape
	transform columns	✗
	transform rows	✗
Separate	subset tables	filter + extract
	decompose tables	filter + extract
	split tables	extract
	separate columns	✗
	separate rows	✗
Combine	extend tables	merge
	supplement tables	merge
	match tables	merge
	summarize rows	group / ungroup
	combine columns	✗
	interpolate rows	✗

combine two data tables, with their difference being row-wise vs. column-wise combination or inner-join vs. outer-join [51]. All three operations suggest an interaction on two tables that results in one table. Thereby, we found one gesture (i.e., *merge*, see Fig. 1) can meet the semantic requirements of all three operations. The same applies to *subset* and *decompose*, and we used a combination of *filter+extract* gestures for them. Deleting tables, rows, and columns shares a similar gesture as the metaphor of throwing things away but differs in the target selection.

We iterate the gesture designs among the team members.

The design objective is to ensure the gestures can intuitively reflect their semantic meanings of data transformation and that no conflicts exist between different gestures. Some initial gesture designs were inspired by the teaching materials of an undergraduate data science course taught by two co-authors, where they frequently used gestures as metaphors for data transformation operations. Specifically, within a data table, horizontal movements were naturally linked to column operations, while vertical movements were considered row operations. Meanwhile, operations involving multiple data tables were intrinsically demonstrated as two-handed gestures. This initial design covered a good range of operations summarized by Kasica et al. [51]. We further introduced more gestures to increase the coverage of data transformation operations, Table 1. In summary, we have designed the following gestures:

- **Extract:** after selecting the target row(s) or column(s), *pull* them out from the original data table to create a new table with the selected content. The original table will be kept, and a new table will be created.
- **Merge:** *move* one table to *collide* with another table and *release* to combine two tables into one. If the key columns [2] are selected and match the criteria, an operation similar to *JOIN* in SQL will be performed. If no key columns are selected, and two tables share the same structure, an operation similar to *UNION* in SQL will be performed. Otherwise, the two tables cannot be combined. After merging, the original tables will be kept, and a new merged table will be created.
- **Filter:** for each column, a histogram is created to show its distribution. *Brush* the histograms to select the range of values to keep. The original table will be updated.
- **Sort:** *swipe* at a column from top to bottom or bottom to top to rearrange data in ascending or descending order. The original table will be updated.

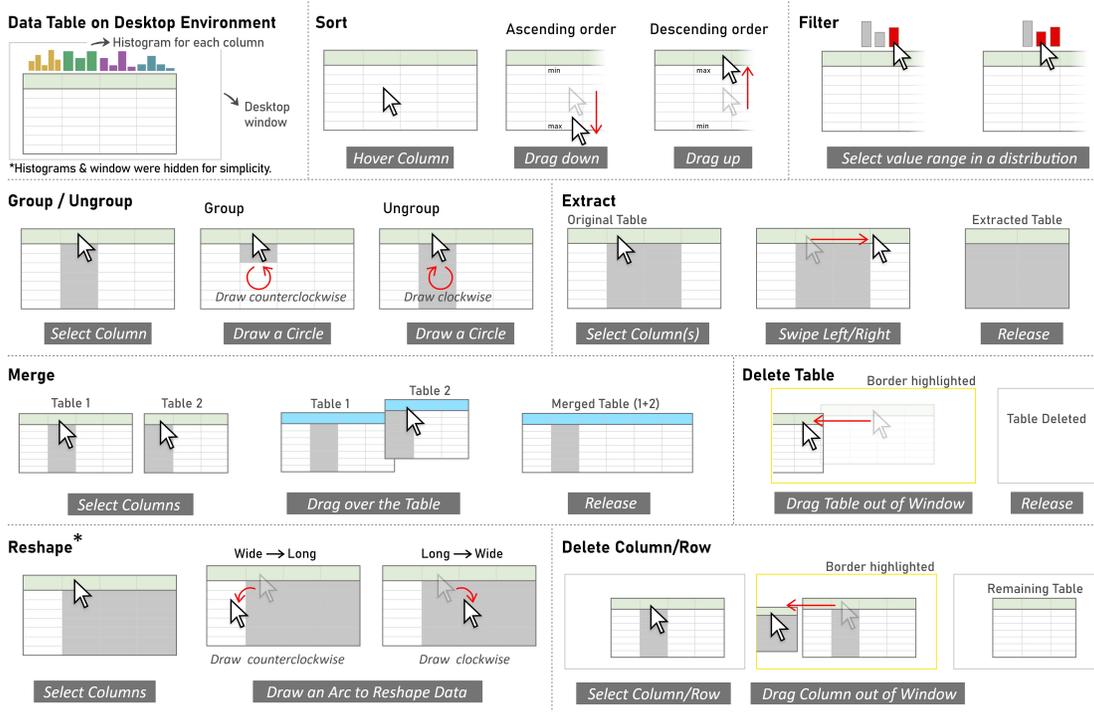


Figure 3. Desktop data table and gestures for data transformation tasks in the Desktop+Gesture condition. The operation that was excluded from the study is marked with *.

- **Group/Ungroup:** after selecting the target column(s), *squeeze* to aggregate the values, while *expand* to restore the aggregated values to their original values. The original table will be updated.
- **Reshape:** after selecting the target column(s), *rotate clockwise* to transform LONG data shape to WIDE shape, while *rotate counterclockwise* to transform WIDE data shape to LONG shape. For LONG to WIDE, the values in target column(s) will be grouped into key-value sets, while for WIDE to LONG, the values in target column(s) will be categorized. The original table will be updated.
- **Delete:** after selecting target row(s), column(s) or table(s), *throw* them away to remove the selected content. The original table will be replaced.

Our gesture design on Desktop shares some similarities with the GestureDB system [16], [18]. We further adapted and extended those gestures into embodied interactions in VR. Most gestures require only one input device (i.e., for *extract*, *merge*, *filter*, *sort*, *reshape*, and *delete*). Those gestures are almost identical in interaction behaviors on Desktop and VR (i.e., consist of actions like click, drag, and drop). In VR, the ability to use two input devices (i.e., left and right hand-held controllers, which are visually represented as hands in our prototype) provides an alternative way to *merge* data tables: people can manipulate two data tables simultaneously by grabbing one in each hand and moving them close to *merge* them. The *group/ungroup* gestures also leverage the two input devices in VR to *squeeze* and *expand*, while we use the *draw a circle* counterclockwise and clockwise to imitate the same semantic meaning on the Desktop due to only one input device is available.

The designed gestures for VR are illustrated in Fig. 2, and the desktop gestures are presented in Fig. 3. We conducted a pilot study with three computer science graduate students

who have data science experience to verify the usability of our designed gestures. Throughout this pilot, we confirmed the feasibility of our chosen gestures to complete data transformation tasks. The pilot study revealed no major operational issues and confirmed the intuitiveness of our gestures.

4 USER STUDY

This study involves two primary experimental factors: the computing environment (or ENVIRONMENT, i.e., Desktop vs. VR) and the interaction method (or INTERACTION, i.e., WIMP vs. Gesture), see Fig. 4.

4.1 Study Conditions

To systematically investigate the two primary experimental factors, we included four conditions that cover all their interactions (Fig. 4), namely, Desktop+WIMP, Desktop+Gesture, VR+WIMP, and VR+Gesture (Fig. 1). The conditions are also demonstrated in the supplemental video.

Desktop Conditions: On the Desktop, we provide an *infinite canvas* as the primary working space. The infinite canvas is a *zoomable* and *pannable* canvas with no boundary where the user can place and move digital content (data tables, in our case). It provides extra freedom to content organization and overcomes the size limitation of a physical screen. Various commercial tools (e.g., Miro [52], Google Jamboard [53], Microsoft Whiteboard [54], and SAGE [55]) use it as their working space, and there are also a series of attempts of using it in data science [55], [56], [57], [58]. Following their success and design, we allow the user to zoom in and out of the workspace by scrolling the mouse scroll wheel. The same interaction has been implemented

	WIMP	Gesture
Desktop	DW, see Fig.1 (a)	DG, see Fig.1 (b)
VR	VW, see Fig.1 (c)	VG, see Fig.1 (d)

Figure 4. Our study compares two primary factors, leading to four conditions.

in many widely used Zoomable User Interfaces (ZUI), like Google Maps. The user can also move data tables to any desired location using drag&drop. A mouse is the input device for the two Desktop conditions. In *Desktop+WIMP*, the user clicks buttons to trigger operations (Fig. 1(a)). We provide one button for each operation and place all buttons on a panel that is fixed on the right side of the screen. It is visible to the user all the time. In *Desktop+Gesture*, the user uses our implemented gestures (Fig. 3) to perform operations (Fig. 1(b)).

VR Conditions: In VR, we allow the user to physically move in the space and freely place and move data tables to any location around them. In *VR+WIMP*, operation buttons were placed on a panel and interacted in the same way as in *Desktop+WIMP*. Due to the different display spaces between Desktop and VR, it is unclear where to place this panel. To make the panel placement in VR as close as the *Desktop+WIMP* condition, we initially placed it using a head-reference approach [59], i.e., the panel will move as the user rotate their head, always visible to the user. However, our preliminary test indicates that such a design is distracting and annoying. We then changed the design to attach the panel to a left-hand-held controller according to [23], where the user can easily access or hide it with arm movements. The latter design was clearly preferred by the users we tested with, and was used in the user study (Fig. 1(c)). We also noticed participants struggled to select rows and columns precisely in our pilot study. Participants' comments revealed the need for a real-time visual indicator for selections. Thus, we rendered a "red dot" to indicate the pointer position on the data table. This mitigated the difficulties in selecting rows and columns based on another round of pilot tests. In *VR+Gesture*, the user uses our implemented gestures (Fig. 2) to perform operations (Fig. 1(d)).

Summary: WIMP and Gesture differ in the way they trigger the operations. Additionally, performing operations in WIMP requires the user to move the cursor or pointer back and forth between the data tables and menu panel, which is likely to introduce a context-switching cost, while the gestures are directly operated on the data tables. Desktop and VR both have an "infinite" display space and let the user reposition the table at any location. Regarding the navigation method, the Desktop provides pan&zoom, while VR enables physical navigation.

4.2 Participants

We recruited 20 participants (Male=16, Female=4; Age from 18 to 35) from the university mailing list after screening for their data transformation experiences with a five-question quiz. 20 out of 25 respondents answered at least four questions correctly and were invited to participate in the study. Nine participants indicated they use VR regularly on a weekly basis, another nine only used VR occasionally, and the rest two had no VR experience. All participants had

normal or corrected-to-normal vision. We provided a \$20 Amazon Gift Card as compensation for each participant.

4.3 Experimental Setup

For VR conditions, we used a Meta Quest 2 virtual reality headset with 1920×1832 resolution per eye and a 90 Hz refresh rate. For Desktop conditions, we used a 27" monitor with a 2560×1440 resolution and 75 Hz refresh rate, which is a standard office setup. Both conditions use a PC with an Intel i7-11800H 2.30 GHz processor and NVIDIA GeForce RTX 3070 graphics card. We used the Air Link feature from Meta, which uses the PC for computation and the headset for rendering. Air Link enables a wireless experience while still leveraging the stronger computing power from the PC. Meanwhile, participants could move around more confidently without worrying about being tripped by cables. The study took place in the space of 3.5×3.5 meters ($12.25m^2$), and we let participants freely walk around the given area in VR conditions. The participants were asked to place themselves in the center of the actual space at the beginning of every VR condition. In Desktop conditions, participants sat on a comfortable office chair with the monitor placed in front of them on an office desk.

On the Desktop, every initial data table, including the target data table, was the same size as 530×400 pixels. The initial tables form a grid layout to maximize the use of display space, and the target table was placed in the middle for participants to reference and remember easily. In VR, five initial data tables, each with the size of 1.15×0.65 meters, were placed 1.65 meters in front of the participant's initial position in a semi-circular curved layout, which was identified as an effective space use strategy in VR for multi-window applications [21], [60], which allows all data tables to be placed at the same distance within participant's reach. The target table in VR was placed 35 cm higher than the initial table for the same purpose as in the Desktop conditions, i.e., to be easily distinguishable from other tables. The initial data tables and the target data table cannot be deleted. This is to avoid accidentally deleting those data tables and being unable to complete the task.

4.4 Task and Data

We asked our participants to *perform data transformations with five given data tables to produce a table in the target format*. Participants could move and resize all the given data tables and the target data table. To decide the number of given data tables, we piloted three different options (i.e., three, five, and eight). We found the task was too obvious with three tables and too difficult with eight tables, so we decided to use five tables for the study.

After training, the participants were asked to use data transformation operations to complete the task with each condition. Initially, we included all operations from Sec. 3.2. However, in our pilot tests, we found participants had a hard time conceptually understanding and applying the *reshape* operation, resulting in an unexpectedly long completion time for each trial (>30 minutes). This was aligned with some previous works [2], [51], [61], pointing out the *reshape* operation can be too complicated for novice users. Thus, we decided to remove the *reshape* operation from the

task to target non-technical data workers and control the study duration. We further conducted another test without using *reshape*. Participants could complete the task in a reasonable amount of time (around 15 minutes per trial) without struggling. To ensure the difficulty of each trial was similar, all trials required a minimum of twelve operations. The sequence of the required minimum operations in each trial was different to reduce the learning effect.

We used tabular data sets collected online for training and study tasks. To eliminate the effect of participants' previous experience, we explicitly told them the data was not reflective of real-world information. We also controlled the data size of initial data tables (row count for each: 30, total column count: ~ 20) and target table (row count: 10, column count: ~ 6). We have included all our study stimuli in the supplementary material.

4.5 Design and Procedures

The conducted user study followed a full-factorial within-subjects design. We used a Latin square (4 groups) to balance the study conditions. Each participant completed four study trials, i.e., one for each condition. The user study lasted two hours on average. The participants were first welcomed and reviewed to sign a consent form. Participants were then instructed about the purpose and steps of the study. They will then complete the following components of the study:

Adjustment: We asked participants to adjust the Quest 2 headset (e.g., the IPD) to a comfortable setting for the VR conditions. Similarly, for the Desktop conditions, participants were instructed to adjust the chair height to their preference before starting the tasks. We confirmed that all participants could see the sample text clearly in all conditions before proceeding.

Training: We first introduce the data transformation terms, considering people might use different terminologies for the same data transformation. We then introduced the computing environment (i.e., Desktop or VR) when a participant first encountered it. Sufficient time was provided for them to get familiar with the hardware until the participant asked to continue (usually around five minutes). For each study condition, when it first appeared to the participant, we first asked them to watch a video demonstrating each operation in that study condition. We confirmed that the participant fully understood how to perform each operation before moving to the next one. After that, we asked participants to perform the same task as in the user study but with only three initial data tables and one target data table. In this phase, we encouraged participants to ask questions about interactions and study tasks. All participants completed the training by finishing the task and confirmed familiarity with the study conditions (the training task took around five to seven minutes).

Study Task: The study task with each study condition started after participants completed the training session for that condition. Before we started the study, we ensured participants were well-informed by providing sufficient context. This included a brief overview of the datasets and the high-level semantic meaning of the target data table. Participants had no time limit for task completion, but we

instructed participants to complete the task as accurately and as fast as they could. For the VR environment, we repositioned participants to the room's center and let them face the same direction before each study task. All participants were able to complete the study task.

Questionnaires. *Post-condition questionnaires:* after completing the study task with each condition, participants were first asked to recall their performed operations in sequence. They were informed about this question at the beginning of the user study. They then filled out a Likert-scale survey adapted from SUS and NASA TLX to rate their subjective experience and provide qualitative feedback about the pros and cons of that condition. *Post-study questionnaires:* after completing all study tasks and post-condition questionnaires, participants were asked to rank all study conditions based on their overall experience. We asked them to provide demographic information at the end.

4.6 Measures

We collected quantitative data and interaction records for each study condition to capture their task performance and sense-making process. Specifically, we used the following measures. **Error score:** we compared the difference between the target table and the participant's result table by rows and columns. Each difference contributed to one error score. The order of rows was considered (as *sorting* operation was included in the task), while the column order did not affect the error score. **Time:** we measured the time from the initial data tables that were first rendered to the participant's task completion. **Number of operations:** we recorded the total number of operations performed by participants to complete the task. **Recall score:** we calculated the *Levenshtein distance* between the participant's actual performed operations sequence and recalled operations sequence that were collected in the questionnaire. A lower value means a closer match and suggests participants can remember their action history more accurately. **Number of performed delete operations:** we recorded the total number of performed table deleting operations in each study trial. **Number of data tables left:** we documented the number of data tables when the participant completed each study trial.

We also collected subjective ratings on a seven-point Likert scale for **mental demand**, **physical demand**, **learnability**, **engagement**, and **usability**. Lower mental and physical demands were considered positive, while higher learnability, engagement, and usability were treated as beneficial. Participants also **ranked** their overall experience. **Qualitative feedback** about each condition's pros and cons were collected from participants. Two authors derived a set of codes from the responses of the first five participants and applied the codes to the remaining responses.

4.7 Hypotheses

We developed our hypotheses based on previous empirical results and our analysis of study conditions, see Sec. 4.1.

Error score. We did not expect any difference in the error score as all required operations were provided consistently across all conditions. We believed that participants could complete the study task for all conditions successfully.

Time. We expected Desktop to outperform VR ($H_{time-env}$) based on previous studies, which found desktop interactions faster than VR interactions due to less required movement [27], [62], [63], [64]. Meanwhile, we anticipated Gesture to be faster than WIMP ($H_{time-interaction}$) considering better completion time of embedded/embedded interaction over WIMP found in earlier research [14], [15].

Number of operations. Our tasks require a sense-making process of foraging and structuring information to solve the problem. Under such a context, we considered VR requires a fewer number of operations than Desktop ($H_{ops-env}$) based on the previous investigation of sense-making in immersive space [24], [65]. We also expected Gesture requires a fewer number of operations than WIMP ($H_{ops-interaction}$). With a lower context-switching cost, Gesture would have fewer disruptions from navigation [66] and require less “mental map” rebuilding for the users [67], [68].

Recall score. We expected participants to have a better recall performance in VR than Desktop ($H_{recall-env}$), as VR with a 3D spatial environment was found to be more effective than Desktop for memorizing and retrieving information [69]. Meanwhile, we foresaw Gesture outperforms WIMP ($H_{recall-interaction}$) since performing body motions with the Gesture has a positive effect on memorability [70].

Number of performed delete operations and data tables left. We predicted a fewer number of delete operations ($H_{del-env}$) and a larger number of data tables left ($H_{left-env}$) in VR than in Desktop. We argue that the larger display space in VR allows participants to keep more intermediate results and reduce the need to delete them. We also anticipated Gesture requires a fewer number of delete operations ($H_{del-interaction}$) and has a larger number of data tables left ($H_{left-interaction}$) than WIMP. As discussed, Gesture has a lower context-switching cost than WIMP, which can reduce the content organization workload and increase the maximum number of data tables participants can handle.

5 RESULTS

We present our statistical results regarding our hypotheses, outline participants’ strategies for using the display space and summarize qualitative feedback for each condition. For dependent variables or their transformed values that met the normality assumption, we used *linear mixed modeling* to evaluate the effect of independent variables on the dependent variables [71]. Compared to repeated measure ANOVA, linear mixed modeling does not have the constraint of sphericity [72, Ch. 13]. We modeled all independent variables (ENVIRONMENT and INTERACTION), and their interactions as fixed effects. A within-subject design with random intercepts was used for all models. We evaluated the significance of the inclusion of an independent variable or interaction terms using a log-likelihood ratio. We then performed Tukey’s HSD posthoc tests for pairwise comparisons using the least square means [73]. We used predicted vs. residual and Q—Q plots to graphically evaluate the homoscedasticity and normality of the Pearson residuals respectively. For other dependent variables that cannot meet the normality assumption, we used the *Friedman* test to evaluate the effect of the independent variable,

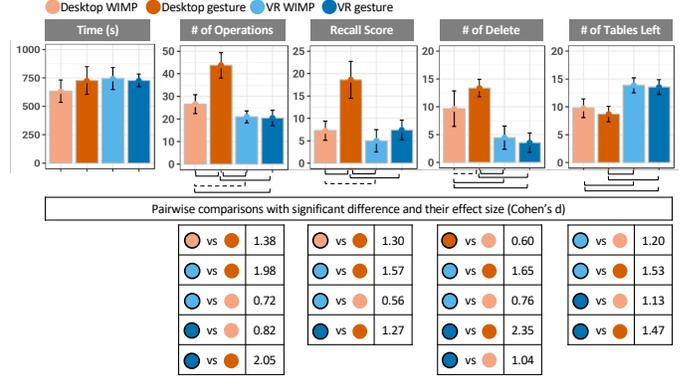


Figure 5. Measurement of time, the total number of operations, recall score, the number of delete operations, and the number of tables left by task. Solid lines indicate statistical significance with $p < 0.05$, and dashed lines indicate $p < 0.1$. The tables below show the effect sizes for pairwise comparison. Circles with black borders indicate the winning conditions.

as well as a Wilcoxon-Nemenyi-McDonald-Thompson test for pairwise comparisons. Significance values are reported for $p < .05(*)$, $p < .01(**)$, and $p < .001(***)$. We also report mean values, 95% confidence intervals (CI), as well as Cohen’s d as an effect size indicator for significant comparisons.

5.1 Quantitative Results

Results are illustrated in Fig. 5, Fig. 6, and Fig. 7. All statistical analyses and anonymized data were included in the supplementary material.

Error score. As expected, all participants could complete the study task correctly under all conditions.

Time. Surprisingly, we found ENVIRONMENT ($p = 0.14$), INTERACTION ($p = 0.51$) and their interaction ($p = 0.42$) did not have a significant effect on time. All conditions took similar amount of time: Desktop+WIMP (634s, CI=73s), Desktop+Gesture (727s, CI=134s), VR+WIMP (744s, CI=116s), and VR+Gesture (728s, CI=96s). Desktop+WIMP tended to be slightly faster but without statistical significance. Thus, we reject $H_{time-env}$ and $H_{time-interaction}$.

Number of operations. We found significant effects of ENVIRONMENT (***) , INTERACTION (***) and their interaction (***) on the total number of performed operations. VR conditions (WIMP with 21.0, CI=2.90 and Gesture with 20.5, CI=2.54) required less number of operations to complete the study task than Desktop conditions (WIMP with 26.6, CI=4.22 and Gesture with 43.8, CI=7.06). All comparisons were statistically significant (*) except for the comparison between VR+WIMP and Desktop+WIMP ($p = 0.057$). On Desktop, Gesture also required more operations than WIMP (***) . In summary, we accept $H_{ops-env}$, and reject $H_{ops-interaction}$.

Recall score. We found significant effects of ENVIRONMENT (***) , INTERACTION (***) on the recall score, with a marginally significant effect from their interaction ($p = 0.081$). Condition-wise, participants can remember their action history better in VR+WIMP (5.0, CI=1.83), VR+Gesture (7.4, CI=2.20), and Desktop+WIMP (7.3, CI=1.94) than Desktop+Gesture (18.6, CI=5.42). VR+WIMP was also marginally more memorable than Desktop+WIMP ($p = 0.078$). For both the WIMP and Gesture, VR could better support the recall

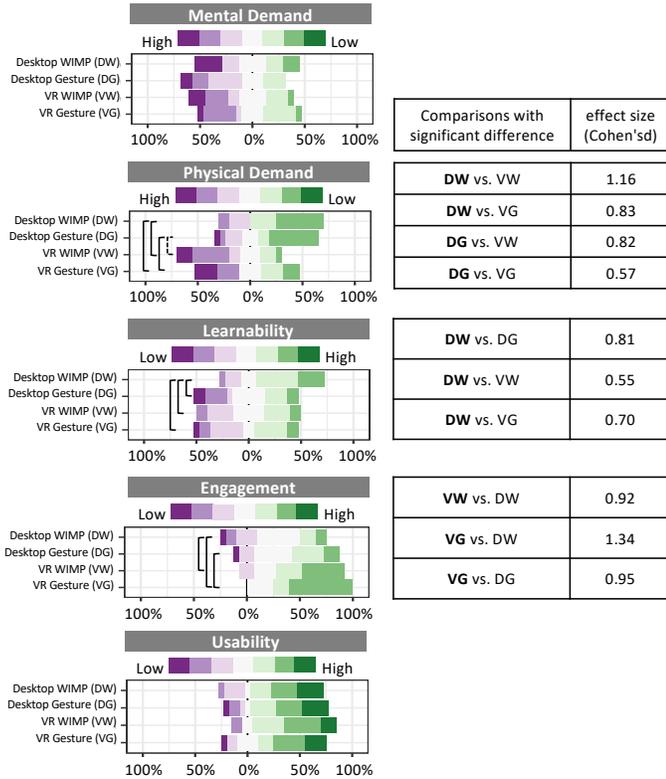


Figure 6. Subjective ratings on mental demand, physical demand, learnability, engagement, and usability by task. Solid lines indicate statistical significance with $p < 0.05$, and dashed lines indicate $p < 0.1$.

process than a Desktop. Desktop+WIMP was also found more memorable than Desktop+Gesture (**). In conclusion, we accept $H_{recall-env}$, and reject $H_{recall-interaction}$.

Number of performed delete operations and number of data tables left. For the number of performed deleting operations and the number of left tables, there was a significant effect from ENVIRONMENT (all ***). INTERACTION only had a significant effect on the number of performed deleting operations (*). Participants deleted fewer tables and kept more data tables in VR than in Desktop (all ***). In summary, along with accepting $H_{del-env}$ and $H_{left-env}$, we reject $H_{del-interaction}$ and $H_{left-interaction}$.

Ratings. We found a significant effect of the study condition on *physical demand* (***), *learnability* (**), and *engagement* (***). Participants found VR conditions (WIMP with 4.0, CI=0.71 and Gesture with 3.7, CI=0.92) more physically demanding than Desktop conditions (WIMP with 2.25, CI=0.69 and Gesture with 2.6, CI=0.88). Participants considered Desktop+WIMP (4.7, CI=0.55) most easy to learn, with Desktop+Gesture (3.45, CI=0.81), VR+WIMP (4.0, CI=0.55), and VR+Gesture (3.7, CI=0.76). VR+Gesture (6.4, CI=0.41) was found more engaging to use than Desktop conditions (WIMP with 4.7, CI=0.73 and Gesture with 5.3, CI=0.64).

Ranking. Participants ranked the VR+Gesture as providing the best overall experience, with 65% ranked it as first place, and 25% ranked it as second place (i.e., 90% in total ranked VR+Gesture as the first or second place).

5.2 Layout Strategies

To better understand how participants used the display space, we grouped the final layout of each trial. We found participants had different strategies in Desktop and VR, but

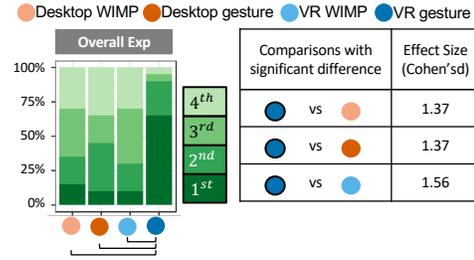


Figure 7. User ranking of overall user experience for each condition. Solid lines indicate significant differences with $p < 0.05$.

used a similar layout between Gesture and WIMP. All final layouts are included in the supplementary material.

On Desktop, we identified four different layout strategies (Fig. 8): **Grid** (five). Participants placed their data tables in a regular layout, closely forming a grid shape. **Piling** (one). The participant created a few piles of data tables. **Grid+Piling** (five). The participants created a regular layout with some piles of data tables. **No obvious pattern** (nine). Roughly half of the participants did not demonstrate a clear layout pattern and created an “organic” layout.

In VR, 19 out of 20 participants almost did not move the five initial tables or the target table. Participants seemed to treat these given tables in as *strong anchors* and were reluctant to manipulate them. This observation might reflect participants’ intention to keep data provenance. Thus, VR might increase the awareness of provenance and allow the track of provenance to be more manageable. Regarding the final layouts, we found three different strategies (Fig. 9). 17 participants performed data transformation behind themselves: **Behind-cluster** (16). Participants created a few clusters at the back of their initial orientation. **Behind-piling** (one). Participants piled data tables in one cluster and formed a roughly vertical line at the back of their initial orientation. **Front** (three). Participants used the space in front of their initial orientation, and to avoid occlusions, they had to delete data tables more frequently. In all strategies, we observed that participants preferred to stay in the center and place data tables close to them to reduce physical movements.

Due to limited data points in some strategies, we could not identify significant performance differences between different strategies, except that the *Front* strategy required more *deletion and total operations* and had fewer *tables left* than other strategies utilizing the entire 360° circular space in VR.

5.3 Qualitative Feedback

We performed qualitative coding to extract common themes from user feedback on each condition. For each condition, we listed the top three mentioned codes and those mentioned more than five times by the participants (frequency shown in parenthesis). We further highlighted the top codes with other frequently associated codes. Finally, we summarized the overall insights across all conditions. The complete coding results can be found in the supplemental materials.

Desktop+WIMP was considered *straightforward* (7), and *familiar* (5). The downsides were *limited space* (11), *hard to interact* (6), and button interaction were *not intuitive* (5). Specifically, limited space was the primary concern,

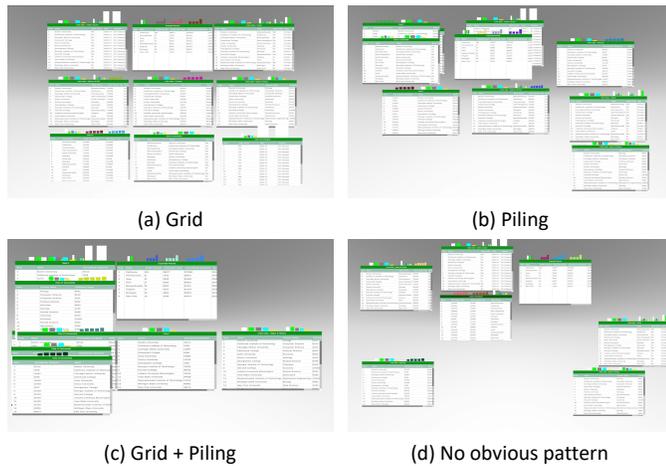


Figure 8. Four layout strategies used by our participants on Desktop.

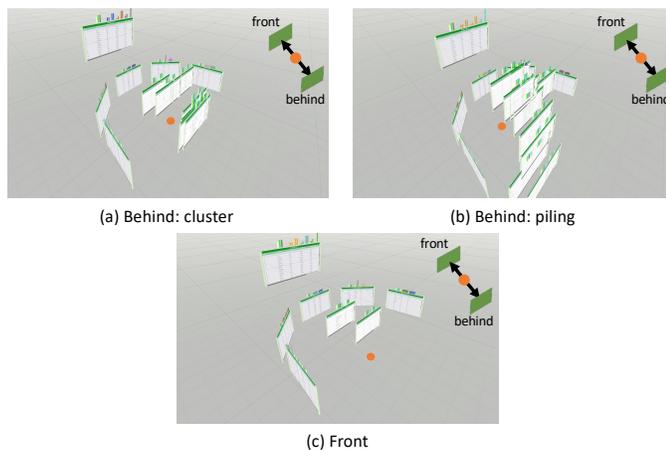


Figure 9. Three layout strategies used by our participants in VR. The red dot indicates the position of the participant.

as shown by its association with several codes, including hard-to-interact (3) and the constant shift in focus from the working window to the buttons created task disruptions when executing operations (2).

Desktop+Gesture was considered *intuitive* (15), *better than button* (8), and *easy to use* (8). Because of the intuitive feeling, it was considered better than buttons and easy to use by five users, good for merging by four users, and good for sorting by two users. The primary issues included *limited space* (9), *hard to interact* (8), *discoverability issues* (6), and *functionality* (5). Similar to the Desktop+WIMP, the limited space was the predominant concern linked to other codes. This included instances where gestures were occasionally challenging to execute due to the limited space, which resulted in creating task disruptions (3), particularly when executing certain functions (3) that demanded more space, such as extraction tasks.

VR+WIMP was praised with *large space* (12), *better than desktop* (9), *grabbed&moved* navigation was *intuitive* (7), *easy to use* (6), and *flexible* (6). Specifically, the major benefit of the large space was associated with flexibility by five users, easy to use by two users, and good understanding by two users. Among comments on better than desktop, people referred to it as easy to use (5), easy to organize (2), and accessible (2). Conversely, the issues related to *functionality* (13). In addition, one participant pointed out that using a

pointer to select two tables from a dropdown menu caused greater disruption when trying to execute merge operations compared to the Desktop+WIMP setup. The concerns arose regarding the constant need to hold buttons on the left hand, causing task disruptions (6), along with another reason similar to those encountered with the Desktop+WIMP. People had diverse opinions on the functionality, such as hard-to-select (3), heavy headset (2), and other technical issues like resolutions.

VR+Gesture was found to be *intuitive* (17), *better than button* (10), *easy to use* (7), *straightforward* (5), and *flexible* (5). The majority (17) found it intuitive and subsequently associated it with better than buttons (7), easy to use (6), and flexible (4). Three also cited more accessible, and the other three cited large spaces. Furthermore, people considered it better than buttons mainly because it is easy to use (4), straightforward (4), accessible (3), and promotes good understanding (2). Participants particularly spoke highly about gesture interactions and grab&move navigation. On the other hand, the major concern was *functionality* (12), mostly with resolution and technical issues (7), but no specific functionality issues were found to perform operations in VR+WIMP.

Summary. All conditions were considered *easy to use*. In Gesture conditions (Desktop+Gesture and VR+Gesture), the majority cited *intuitive* (15 and 17 times, respectively) and considered them *better than button* (8 and 10 times). VR conditions (VR+WIMP and VR+Gesture) were considered *flexible* (6 and 5 times) and *better than desktop* by several users (9 and 4 times). Interestingly, Desktop+WIMP and VR+Gesture each were cited *straightforward* by a handful of users (7 and 5 times), indicating that users have expected that WIMP is natural for desktop and 3D gestures for VR environment. More discussions were presented in Sec. 5.

On the other hand, Desktop conditions hinder interaction due to *limited space* (11 and 9 times) and *hard to interact* (6 and 8 times). With VR conditions, the majority have commented on *functionality* (13 and 12 times) due to the unfamiliarity with data transformation in VR, such as technical issues and feature suggestions. In VR+WIMP, people found it *hard to interact* and *select* (5 times), leading to creating *disturbed* (6 times).

6 KEY FINDINGS AND DISCUSSIONS

Performing data transformation in VR and on a Desktop had similar time performance. Previous studies produced mixed results in comparing VR and Desktop. In immersive analytics, for completion time, VR was found to be primarily beneficial for visualizing spatial data and 3D shape perception [19], [47], and to be slower than Desktop due to more required movements (see our *time hypothesis* in Sec. 4.7). Our tested task did not involve perceiving 3D or spatial visualization, so we expected Desktop to be faster than VR. Surprisingly, our four tested conditions had similar time performance. We believe the empirical results of VR being slower than Desktop in performing interactions still applied to our case, which is partially reflected by the fact that VR conditions were considered significantly more physically demanding than Desktop conditions. On the other hand, unlike the previously tested tasks, our tested task

not only involved low-level interactions but also required participants to actively make sense of the data to come up with a sequence of interactions to complete the task. We could reasonably expect participants to put a significant amount of effort into thinking and planning strategies for a complicated task like ours. Based on this assumption and our results, we suggest that VR allows participants to complete the high-level sense-making components faster than Desktop, which supplements its extra time costs in performing low-level interactions. Below, we elaborate more on the high-level sense-making process from the *provenance* and *strategic thinking* perspectives.

☐ **VR showed the potential to provide improved provenance over ☐ Desktop.** Provenance is about the lineage and processing history of data. It provides a detailed record of the origins of the data, how it has been processed or modified, and where and when it transformed over time [74]. The ability to track provenance can support many applications, like data quality control, data auditing, and replication [75]. In our study, we used the *recall score*, the *number of performed delete operations*, and the *number of data tables left* as indicators of certain aspects of provenance. The *recall score* measures the ability to recall the operation history, reflecting when the data was transformed. The latter two metrics offer objective measures of the amount of kept information.

In addition to the previously confirmed memorability advantage in VR (see our provenance hypothesis in Sec. 4.7), we believe the large display space and embodied navigation in VR also helped participants keep track of provenance. *First*, participants had more space to place tables in VR, whereas on Desktop, they were more likely to delete tables to free display space to reduce clutter. Our results of fewer table deletions and more kept tables in VR than on Desktop are well aligned with provenance tracking ability. The data tables, kept persistently visible, provided a reliable reference for participants to confirm the success of their current operations. Moreover, they were a handy tool for participants to return to whenever they made mistakes. This was likely due to the large display space provided within the VR setup. Particularly, nine participants explicitly complained about the display space on the Desktop; for example, *“I have a very limited workspace (on the desktop), and cannot see all the tables at once, which really hurts my performance (P15).”* This limited space on the Desktop lets participants continually delete the data tables, which leads to a loss of opportunities to keep track of the processing history. *Second*, we consider physical navigation more efficient than virtual navigation. Previous research identified the benefits of using physical movements to navigate large displays over virtual navigation (i.e., zooming) [76], [77]. Our results partially re-confirmed their findings in VR. 12 participants’ comments reconciled with our assumption, like, *“I have all data tables in front of me without zooming in & out, and I can focus on the task (P12).”* We found that the participants prefer physical navigation, as it offers a more intuitive and faster method for altering the perspective view of their workspace for accessing all kept information. One participant also specifically described the use of large display space to improve spatial memory and embodied navigation in VR: *“I like the point that I put the data table behind me so that I can come back if I make mistakes (P7).”* This feedback aligns well with the improved provenance in

VR, which enhances data quality through auditing without manually documenting every detail of each step. However, provenance is still a relatively abstract concept, and our measurements also only captured certain aspects of it. More effort is needed to formally define and quantify provenance in data transformation.

☐ **VR demonstrated preliminary evidence of promoting strategic thinking.** Participants needed to continuously develop the next steps in our tested task and evaluate their progress. We found VR required fewer operations to complete transformation tasks than Desktop, which points to the potential advantages of VR in supporting strategic planning. With the improved provenance, it is likely that VR users can better track the progress and take more efficient steps. Another potential reason is the flexible arrangement of VR tables that support easier visual comparison. We observed *all 20* participants in VR *grabbed&moved* their working table under the target table for comparison. Meanwhile, participants rarely performed a similar interaction on the Desktop. We believe such frequent comparisons enabled continuous progress evaluation and promoted strategic thinking in VR. We anticipate the embodied table management (i.e., the *grabbed&moved*) and large display space in VR provide a natural mechanism to support such an approach. In contrast, the need for repetitive zooming in/out on Desktop made participants reluctant to perform the same strategy. Additionally, participants also commented on their experiences along this line, for instance, *“I have many more places to put the data table (in VR). Organizing the table was way more manageable (P14).”* *“The interaction in VR was way better than the monitor version; the task became way more accessible, and organizing was easy (P3).”*

☐ **WIMP tended to be more suitable for ☐ Desktop than ☐ Gesture.** Unlike previous studies [14], [15], we did not find positive effects on the collected measures of using Gesture over WIMP on a Desktop. We believe the task and gesture complexities could be the main reasons for our contradictory results. We tested a more complicated high-level task (average completion time was 10+ minutes per trial) than the previous low-level tasks (average completion time was around one minute per trial). Performing a single interaction using Gesture could outperform WIMP in our study, but the complexity of Gesture might introduce other overheads that decreased its performance significantly in our tested task. The subjective ratings partially confirmed our assumption: Desktop+Gesture was considered harder to learn than Desktop+WIMP. Participants had to remember more gestures (eight) than in the previous studies (three in [14] and five in [15]), which might introduce a high working memory and affect their performance, especially in a more complicated task. The fact that participants struggled with recalling their action history in Desktop+Gesture resonated with our conclusion, with some representative comments like: *“The interaction was confusing (in Desktop+Gesture). I think I would like[ly] have more errors than other conditions.”*

No noticeable difference in time, efficiency, and provenance measures between ☐ WIMP and ☐ Gesture in ☐ VR. Similar to Desktop, the identified benefits of body motion [70] might also be affected by the task and gesture complexities in VR. On the other hand, we did not find any significant difference between VR+WIMP and VR+Gesture,

indicating the overhead introduced by learning and remembering Gesture in VR could be negligible due to its intuitiveness. Participants ranked VR+Gesture with the best overall experience, and some commented on the benefits of using Gesture in VR explicitly, like *“Embodied interaction helps me a lot in understanding the data story (P2)”* and *“Using both hands was very helpful in performing the task (P4).”* More importantly, the embodied table management (i.e., the grabbed&moved) and physical navigation were provided in all VR conditions. Compared to interactions for triggering the operations, these two features might contribute a heavier weight in the final performance, making both VR conditions have similar performances. Furthermore, the perceived intuitiveness of the gestures and embodied interactions was reflected in the feedback we collected. Comparably, none of the participants explicitly mentioned Desktop+WIMP to be intuitive.

Various layout strategies exist in both environments. The “infinite” space in both the desktop and VR conditions allowed participants to freely lay out content in their workspaces. We found two distinct dimensions in their strategies: *space usage* and *view management methods*. For space usage, on Desktop, most participants only used space slightly larger than the initial setup (Fig. 8) with 6 to 9 data tables and adjusted the view by zooming. On the other hand, the majority of participants in VR utilized the entire 360° circular space around them to fit more data tables (12 to 15) (Fig. 9). In VR, previous studies observed using semi-circular layouts were more frequent [21] and beneficial over 360° layouts [60], [78] for performing low-level tasks (e.g., search and comparison). However, for high-level tasks, our observations aligned well with some other works, such as sensemaking [79] and visual exploration [80], where people preferred using the entire fully circular space. We anticipated that placing content within the field of view is beneficial for time-constrained low-level tasks, while fully utilizing the display space is essential for the larger amount of content generated in high-level tasks. In terms of view management methods, the layout on the Desktop was considerably more organized than in VR, showing that precise view placement is easier with a mouse than with a VR controller. However, user behaviors were similar across both Desktop and VR. Most participants intended to create occlusion-free clusters (e.g., a grid layout in Fig. 9b), yet some participants created piled clusters (e.g., a piling layout in Fig. 9a), which was also observed in a recent AR study [81]. We also noticed the use of space to record the operation history from some participants (e.g., moving older tables to the top), which has been identified on large 2D displays [46].

7 GENERALIZATIONS, LIMITATIONS AND FUTURE WORK

Applications. Our study provides empirical evidence of the benefits of using VR for data transformation. We believe the large display space, spatial memory, and embodied navigation offered by VR are the primary factors for its improved performance. Since these are general characteristics of VR, we expect our results can be generalized to similar tasks that require organizing a large number of entities in space for sense-making. Some preliminary research explored the

use of VR in such applications, like multiple view visualization [26], [82], multi-scale geographic navigation [21], and large-scale document analysis [24], [65], [81]. However, as highlighted by Ens et al. [19], there is a lack of fundamental study for comparing immersive versus non-immersive platforms for analytics purposes. Our study provides a preliminary assessment for data transformation under this context. Future studies may test other applications and delineate the benefits of immersive environments for a broader range of analytical tasks.

Target users and functionality. We focused on investigating intuitive data transformation tools for non-technical data workers and intentionally excluded the programming requirements in our study. However, programming-like operations are essential for more experienced users and more complicated tasks. The Gesture has limitations in how much information and intention it can express. To extend this work in this direction, future work needs to integrate a programming interface into the prototypes, like many UI-based commercial tools (e.g., Tableau and Excel) and research prototypes [83], [84]. To achieve this in VR, we need to consider the most appropriate text input method [85]. Alternatively, we may also consider using natural language as a user interface with higher expressiveness [86], [87]. On the other hand, we also see opportunities to develop new interaction techniques (e.g., embodied gestures) to lower the learning curve of complicated data transformation operations for non-technical data workers.

Techniques. We believe our designed gestures were intuitive and natural to their represented data transformation operations. However, the Gesture did not contribute to the participants’ performance as we expected. We believe Gesture can still outperform WIMP in low-level tasks when the user knows the precise operation to perform [15], [16], [41]. Specifically, Nandi et al. found gestures to have better performance and discoverability than WIMP for a single data transformation operation [16]. However, discoverability might become a more severe issue in high-level tasks where the user must continuously develop the next steps. Although during the training phase, the participants were able to complete training tasks with a smaller dataset, it was still reasonable for participants to spend extra time and effort to recall the gestures in the longer study tasks. Future work may look at further improving the discoverability and learnability of gestures [88], [89], [90] or conduct a longitudinal study to reduce this effect.

Additionally, there are a few other directions to extending our work. Although we improved the selection experience in VR, precisely selecting and manipulating objects can still be challenging. Along with improving mid-air interactions, one may consider other input devices, like using a mouse in VR [91], [92] and other tangible proxies [93], [94]. Besides, we consider visualizing data processing history could further improve provenance, like a data flow system [58], [95].

Computing environments. We tested Desktop and VR, as Desktop is the most widely used, and VR is emerging. Testing other computing environments with different input modalities could bring more insights into interaction designs, like a tablet [15], [16], [41], [44] and display wall [96]. Tablets offer multitouch capabilities and are considered

more natural for performing gestures than Desktop. However, tablets are usually limited in size, and users may suffer from “fat finger” issues in precise interactions. A larger touch screen may alleviate the issue, but not easily accessible to many people. Nevertheless, testing different touching devices is an exciting future direction. Moreover, specific to our study, our gestures primarily only involve click and drag, which could be easily completed with a mouse, as demonstrated in previous work [13], [14]. Increasing the display space on the Desktop is likely to increase the performance (like using multiple monitors or a display wall), as one of the most notable benefits of VR is its large display space. The effect of display size has been previously explored [97]. Future work can follow the same methodology to study the effect of display size on the data transformation workspace. Furthermore, we found that our participants might not be familiar with our zoomable and pannable interface on the Desktop. However, there’s an increasing trend of no-code data science tools designed with zoomable and pannable interfaces, as we discussed in Section 4.1. It gradually becomes an essential design for non-technical data workers. Despite this, we also see other opportunities for managing multiple data tables, such as using tabs.

Scalability and ecological validity. We tested relatively small data in our study (see Sec. 4.4), as we want to control the study duration. Meanwhile, we had an interesting comment from one participant that they did not actively check the rows of data tables but focused on the columns more. This comment aligns well with the design of the GestureDB system [16], [18], where they present the columns all the time and only present rows on demand or for confirmation purposes. Larger data should be tested in future studies, especially data with more columns. Meanwhile, it is challenging to display all columns for large data sets (e.g., with 20 columns). As such, in our tested data, we intentionally have columns that cannot fit into the table view. The width of each column was determined by its longest data entry, causing the total width of all columns to exceed the table’s width. Hence, even with a limited number of columns, columns were not visible at once, and scrolling was necessary to finish the task. From this perspective, we anticipate that, with increasing data size, participants need to scroll more, which will make the task more challenging for all conditions. Precision is often required in scrolling, so increasing the data size may affect VR more than the Desktop. On the other hand, it is possible to leverage the large display space in VR to alleviate this issue: we can use a large space to ensure all columns are visible without the need to scroll in VR. The user then can physically move in the space to navigate the large data tables. Such physical navigation has been found more effective than virtual navigation like panning or scrolling [77]. However, additional interaction supports are likely to be needed (e.g., shrinking the size of the table when it is grabbed) and should be tested in future work. Our study did not include commercial data transformation tools on the desktop like Tableau Pre Builder and Trifacta. Those tools include many more features besides the basic data transformation operations, which may distract us from studying our intended main effects (e.g., WIMP vs. Gesture and Desktop vs. VR). Adapting necessary features from those commercial desktop tools to VR is a

promising next step, along with an ecological investigation of a complete system. Beyond the size of the datasets used, conducting user studies with a larger number of participants also should be tested. We expect this to lower the variance in several measures and further strengthen our hypothesis.

VR vs. AR. Our study was conducted in VR, as the VR hardware is more mature than AR headsets (e.g., higher resolution and larger field of view). On the positive side, our designed gestures and implemented prototypes can be easily migrated from VR to AR thanks to the developing ecosystem, i.e., Unity3D. We also believe our identified benefits in VR can be generalized to AR or hybrid systems of AR and Desktop [98], because the same as VR, AR also provides embodied interaction and large display space. With AR hardware improving, it is important to investigate AR’s unique challenges and benefits. For example, observing the physical environment can improve people’s spatial memory [81] and seamlessly switch between different computing environments [98]. We will systematically analyze and evaluate data transformation in AR in a future study once better AR headsets are released.

8 CONCLUSION

In this paper, we presented our prototypes to enable embedded interactions on the Desktop and embodied interactions in VR for data transformation. We conducted a controlled user study to systematically evaluate the effect of the computing environments (Desktop vs. VR) and interaction methods (WIMP vs. Gesture). We found initial evidences showing the benefits of using VR for data transformation: VR had the potential to provide improved provenance over Desktop and demonstrated preliminary evidence of promoting strategic thinking. VR also had a similar time performance compared to Desktop. Additionally, participants found VR to be more engaging, and VR+Gesture provided the best overall experience. Considering the iterative nature of data science, we foresee strong initiatives to combine immersive data transformation and visualization to enable a more complete *Immersive Analytics* workflow, reducing the overhead of switching between different computing environments. On the other hand, we did not find a better performance of Gesture over WIMP. We still believe there is potential to improve Gestures, for example, by providing long-term training and real-time memory assistance. In summary, our results provide preliminary evidence that the large display space, spatial memory, and embodied navigation in VR are beneficial for high-level data transformation tasks.

ACKNOWLEDGMENTS

This work was supported in part by NSF I/UCRC CNS-1822080 via the NSF Center for Space, High-performance, and Resilient Computing (SHREC) and NSF grant III-2107328.

REFERENCES

- [1] H. Wickham and G. Grolemund, *R for data science: import, tidy, transform, visualize, and model data.* " O’Reilly Media, Inc.", 2016.

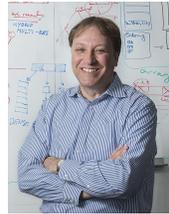
- [2] S. Kandel, A. Paepcke, J. Hellerstein, and J. Heer, "Wrangler: Interactive visual specification of data transformation scripts," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Vancouver BC Canada: ACM, May 2011, pp. 3363–3372.
- [3] P. J. Guo, S. Kandel, J. M. Hellerstein, and J. Heer, "Proactive wrangling: Mixed-initiative end-user programming of data transformation scripts," in *Proceedings of the 24th annual ACM symposium on User interface software and technology*, 2011, pp. 65–74.
- [4] S. Kandel, J. Heer, C. Plaisant, J. Kennedy, F. Van Ham, N. H. Riche, C. Weaver, B. Lee, D. Brodbeck, and P. Buono, "Research directions in data wrangling: Visualizations and transformations for usable and credible data," *Information Visualization*, vol. 10, no. 4, pp. 271–288, 2011.
- [5] T. Desktop, "Tableau desktop," <https://www.tableau.com/products/desktop>, 2022, accessed: 2022-08-28.
- [6] Amazon, "Aws auto ml solutions," <https://aws.amazon.com/machine-learning/automl>, 2022, accessed: 2022-08-28.
- [7] Google. (2022) Automl. <https://cloud.google.com/automl/>. Accessed: 2022-08-28.
- [8] D. Wang, J. Andres, J. D. Weisz, E. Oduor, and C. Dugan, "AutoDS: Towards Human-Centered Automation of Data Science," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Yokohama Japan: ACM, May 2021, pp. 1–12.
- [9] T. P. Builder. (2022) Tableauprep. <https://www.tableau.com/products/prep>. Accessed: 2022-08-28.
- [10] Trifacta. (2022) Trifacta. <https://www.trifacta.com/>. Accessed: 2022-08-28.
- [11] Alteryx. (2022) Alteryx. <https://www.alteryx.com/>. Accessed: 2022-08-28.
- [12] B. Saket, A. Srinivasan, E. D. Ragan, and A. Endert, "Evaluating interactive graphical encodings for data visualization," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 3, pp. 1316–1330, 2017.
- [13] B. Saket, S. Huron, C. Perin, and A. Endert, "Investigating direct manipulation of graphical encodings as a method for user interaction," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 1, pp. 482–491, 2019.
- [14] A. Sarvghad, B. Saket, A. Endert, and N. Weibel, "Embedded Merge & Split: Visual Adjustment of Data Grouping," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 800–809, Jan. 2019.
- [15] S. M. Drucker, D. Fisher, R. Sadana, J. Herron, and m. schraefel, "TouchViz: A case study comparing two interfaces for data analytics on tablets," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Paris France: ACM, Apr. 2013, pp. 2301–2310.
- [16] A. Nandi, L. Jiang, and M. Mandel, "Gestural query specification," *Proceedings of the VLDB Endowment*, vol. 7, no. 4, pp. 289–300, 2013.
- [17] A. Nandi and M. Mandel, "The interactive join: recognizing gestures for database queries," in *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, 2013, pp. 1203–1208.
- [18] L. Jiang, M. Mandel, and A. Nandi, "Gesturequery: A multitouch database query interface," *Proceedings of the VLDB Endowment*, vol. 6, no. 12, pp. 1342–1345, 2013.
- [19] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne *et al.*, "Grand challenges in immersive analytics," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–17.
- [20] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, *Immersive analytics*. Springer, 2018, vol. 11190.
- [21] K. A. Satriadi, B. Ens, M. Cordeil, T. Czauderna, and B. Jenny, "Maps around me: 3d multiview layouts in immersive spaces," *Proceedings of the ACM on Human-Computer Interaction*, vol. 4, no. ISS, pp. 1–20, 2020.
- [22] O.-H. Kwon, C. Muelder, K. Lee, and K.-L. Ma, "A study of layout, rendering, and interaction methods for immersive graph visualization," *IEEE transactions on visualization and computer graphics*, vol. 22, no. 7, pp. 1802–1815, 2016.
- [23] Y. Yang, M. Cordeil, J. Beyer, T. Dwyer, K. Marriott, and H. Pfister, "Embodied navigation in immersive abstract data visualization: Is overview+ detail or zooming better for 3d scatterplots?" *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1214–1224, 2020.
- [24] L. Lisle, X. Chen, J. Edward Gitre, C. North, and D. A. Bowman, "Evaluating the Benefits of the Immersive Space to Think," in *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. Atlanta, GA, USA: IEEE, Mar. 2020, pp. 331–337.
- [25] Y. Yang, T. Dwyer, K. Marriott, B. Jenny, and S. Goodwin, "Tilt map: Interactive transitions between choropleth map, prism map and bar chart in immersive environments," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 12, pp. 4507–4519, 2020.
- [26] 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 30th annual ACM symposium on user interface software and technology*, 2017, pp. 71–83.
- [27] 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, 2017.
- [28] T. Drey, J. Gugenheimer, J. Karlbauer, M. Milo, and E. Rukzio, "Vrsketchin: Exploring the design space of pen and tablet interaction for 3d sketching in virtual reality," in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–14.
- [29] D. Mendes, F. Relvas, A. Ferreira, and J. Jorge, "The benefits of dof separation in mid-air 3d object manipulation," in *Proceedings of the 22nd ACM conference on virtual reality software and technology*, 2016, pp. 261–268.
- [30] D. Hayatpur, H. Xia, and D. Wigdor, "Datahop: Spatial data exploration in virtual reality," in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, 2020, pp. 818–828.
- [31] W. McKinney *et al.*, "Data structures for statistical computing in python," in *Proceedings of the 9th Python in Science Conference*, vol. 445. Austin, TX, 2010, pp. 51–56.
- [32] H. Wickham, R. François, L. Henry, and K. Müller. (2022) dplyr: A grammar of data manipulation. <https://dplyr.tidyverse.org>. Accessed: 2022-08-28.
- [33] H. Wickham and M. Girlich. (2022) tidyr: Tidy messy data. <https://tidyr.tidyverse.org>. Accessed: 2022-08-28.
- [34] K. Xiong, S. Fu, G. Ding, Z. Luo, R. Yu, W. Chen, H. Bao, and Y. Wu, "Visualizing the Scripts of Data Wrangling with SOMNUS," *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2022.
- [35] JetBrains. (2020) Datalore. <https://datalore.jetbrains.com/>. Accessed: 2022-08-28.
- [36] I. Drosos, T. Barik, P. J. Guo, R. DeLine, and S. Gulwani, "Wrex: A unified programming-by-example interaction for synthesizing readable code for data scientists," in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–12.
- [37] P. Kirschthaler, M. Porcheron, and J. E. Fischer, "What can i say? effects of discoverability in vuis on task performance and user experience," in *Proceedings of the 2nd Conference on Conversational User Interfaces*, 2020, pp. 1–9.
- [38] B. Shneiderman and H. Kang, "Direct annotation: A drag-and-drop strategy for labeling photos," in *2000 IEEE Conference on Information Visualization. An International Conference on Computer Visualization and Graphics*. IEEE, 2000, pp. 88–95.
- [39] J. Jia, J. Sun, C.-K. Tang, and H.-Y. Shum, "Drag-and-drop pasting," *ACM Transactions on graphics (TOG)*, vol. 25, no. 3, pp. 631–637, 2006.
- [40] F. Lekschas, X. Zhou, W. Chen, N. Gehlenborg, B. Bach, and H. Pfister, "A generic framework and library for exploration of small multiples through interactive piling," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 358–368, 2020.
- [41] J. M. Rzeszutarski and A. Kittur, "Kinetica: Naturalistic multi-touch data visualization," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014, pp. 897–906.
- [42] P. Pirolli and R. Rao, "Table lens as a tool for making sense of data," in *Proceedings of the workshop on Advanced visual interfaces*, 1996, pp. 67–80.
- [43] T. Horak, P. Berger, H. Schumann, R. Dachsel, and C. Tominski, "Responsive matrix cells: A focus+ context approach for exploring and editing multivariate graphs," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1644–1654, 2020.
- [44] C. J. Burley and A. Nandi, "Arquery: Hallucinating analytics over real-world data using augmented reality," in *CIDR*, 2019.
- [45] A. Crotty, A. Galakatos, E. Zraggen, C. Binnig, and T. Kraska, "Vizdom: interactive analytics through pen and touch," *Proceedings of the VLDB Endowment*, vol. 8, no. 12, pp. 2024–2027, 2015.

- [46] C. Andrews, A. Endert, and C. North, "Space to think: large high-resolution displays for sensemaking," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2010, pp. 55–64.
- [47] A. Fonet and Y. Prie, "Survey of immersive analytics," *IEEE transactions on visualization and computer graphics*, vol. 27, no. 3, pp. 2101–2122, 2019.
- [48] 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, 2018.
- [49] P. W. 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, 2020.
- [50] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer, "Iatc: An immersive analytics toolkit," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2019, pp. 200–209.
- [51] S. Kasica, C. Berret, and T. Munzner, "Table Scraps: An Actionable Framework for Multi-Table Data Wrangling From An Artifact Study of Computational Journalism," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 957–966, Feb. 2021.
- [52] Miro. (2020) Miro. <https://miro.com/>. Accessed: 2022-08-28.
- [53] Google. (2020) Jamboard. <https://jamboard.google.com/>. Accessed: 2022-08-28.
- [54] Microsoft. (2020) Whiteboard. <https://whiteboard.office.com>. Accessed: 2022-08-28.
- [55] T. Marrinan, J. Aurisano, A. Nishimoto, K. Bharadwaj, V. Mateevitsi, L. Renambot, L. Long, A. Johnson, and J. Leigh, "Sage2: A new approach for data intensive collaboration using scalable resolution shared displays," in *10th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing*. IEEE, 2014, pp. 177–186.
- [56] N. Kodagoda, S. Attfield, B. L. W. Wong, C. Rooney, and S. Choudhury, "Using Interactive Visual Reasoning to Support Sense-Making: Implications for Design," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2217–2226, Dec. 2013.
- [57] N. Elmqvist, J. Stasko, and P. Tsigas, "DataMeadow: A Visual Canvas for Analysis of Large-Scale Multivariate Data," in *2007 IEEE Symposium on Visual Analytics Science and Technology*. Sacramento, CA, USA: IEEE, Oct. 2007, pp. 187–194.
- [58] W. Javed and N. Elmqvist, "ExPlates: Spatializing Interactive Analysis to Scaffold Visual Exploration," *Computer Graphics Forum*, vol. 32, no. 3pt4, pp. 441–450, Jun. 2013.
- [59] W. S. Lages and D. A. Bowman, "Walking with adaptive augmented reality workspaces: design and usage patterns," in *Proceedings of the 24th International Conference on Intelligent User Interfaces*, 2019, pp. 356–366.
- [60] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer, "Design and evaluation of interactive small multiples data visualisation in immersive spaces," in *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2020, pp. 588–597.
- [61] V. Raman and J. M. Hellerstein, "Potter's wheel: An interactive data cleaning system," in *VLDB*, vol. 1, 2001, pp. 381–390.
- [62] J. Chen, H. Cai, A. P. Auchus, and D. H. Laidlaw, "Effects of stereo and screen size on the legibility of three-dimensional streamtube visualization," *IEEE transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2130–2139, 2012.
- [63] J. A. Wagner Filho, M. F. Rey, C. M. Freitas, and L. Nedel, "Immersive visualization of abstract information: An evaluation on dimensionally-reduced data scatterplots," in *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2018, pp. 483–490.
- [64] L. Arms, D. Cook, and C. Cruz-Neira, "The benefits of statistical visualization in an immersive environment," in *Proceedings IEEE Virtual Reality (Cat. No. 99CB36316)*. IEEE, 1999, pp. 88–95.
- [65] L. Lisle, K. Davidson, E. J. Gitre, C. North, and D. A. Bowman, "Sensemaking Strategies with Immersive Space to Think," in *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. Lisboa, Portugal: IEEE, Mar. 2021, pp. 529–537.
- [66] R. Ball and C. North, "Analysis of user behavior on high-resolution tiled displays," in *IFIP Conference on Human-Computer Interaction*. Springer, 2005, pp. 350–363.
- [67] H. C. Purchase, E. Hoggan, and C. Görg, "How important is the "mental map"?—an empirical investigation of a dynamic graph layout algorithm," in *International Symposium on Graph Drawing*. Springer, 2006, pp. 184–195.
- [68] K. Reda, A. E. Johnson, M. E. Papka, and J. Leigh, "Effects of display size and resolution on user behavior and insight acquisition in visual exploration," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015, pp. 2759–2768.
- [69] F. Yang, J. Qian, J. Novotny, D. Badre, C. D. Jackson, and D. H. Laidlaw, "A virtual reality memory palace variant aids knowledge retrieval from scholarly articles," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 12, pp. 4359–4373, 2020.
- [70] J. Zagermann, U. Pfeil, D. Fink, P. von Bauer, and H. Reiterer, "Memory in motion: The influence of gesture-and touch-based input modalities on spatial memory," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 1899–1910.
- [71] D. Bates, M. Mächler, B. Bolker, and S. Walker, "Fitting Linear Mixed-Effects Models Using **lme4**," *Journal of Statistical Software*, vol. 67, no. 1, 2015.
- [72] A. Field, J. Miles, and Z. Field, *Discovering statistics using R*. Sage publications, 2012.
- [73] R. V. Lenth, "Least-Squares Means: The R Package **lsmeans**," *Journal of Statistical Software*, vol. 69, no. 1, 2016.
- [74] J. Wang, D. Crawl, S. Purawat, M. Nguyen, and I. Altintas, "Big data provenance: Challenges, state of the art and opportunities," in *2015 IEEE international conference on big data (Big Data)*. IEEE, 2015, pp. 2509–2516.
- [75] Y. L. Simmhan, B. Plale, and D. Gannon, "A survey of data provenance in e-science," *ACM Sigmod Record*, vol. 34, no. 3, pp. 31–36, 2005.
- [76] R. Ball and C. North, "Effects of tiled high-resolution display on basic visualization and navigation tasks," in *CHI'05 extended abstracts on Human factors in computing systems*, 2005, pp. 1196–1199.
- [77] R. Ball, C. North, and D. A. Bowman, "Move to improve: promoting physical navigation to increase user performance with large displays," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2007, pp. 191–200.
- [78] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer, "Effects of display layout on spatial memory for immersive environments," *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. ISS, pp. 468–488, 2022.
- [79] 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*, 2022.
- [80] A. Batch, A. Cunningham, M. Cordeil, N. Elmqvist, 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, 2019.
- [81] W. Luo, A. Lehmann, H. Widengren, and R. Dachselt, "Where should we put it? layout and placement strategies of documents in augmented reality for collaborative sensemaking," in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–16.
- [82] 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, 2020.
- [83] T. Kraska, "Northstar: An interactive data science system," *Proc. VLDB Endow.*, vol. 11, no. 12, p. 2150–2164, aug 2018. [Online]. Available: <https://doi.org/10.14778/3229863.3240493>
- [84] Z. Shang, E. Zraggen, B. Buratti, P. Eichmann, N. Karimeddiny, C. Meyer, W. Runnels, and T. Kraska, "Davos: a system for interactive data-driven decision making," *Proceedings of the VLDB Endowment*, vol. 14, no. 12, pp. 2893–2905, 2021.
- [85] T. J. Dube and A. S. Arif, "Text entry in virtual reality: A comprehensive review of the literature," in *International Conference on Human-Computer Interaction*. Springer, 2019, pp. 419–437.
- [86] L. Shen, E. Shen, Y. Luo, X. Yang, X. Hu, X. Zhang, Z. Tai, and J. Wang, "Towards natural language interfaces for data visualization: A survey," *IEEE transactions on visualization and computer graphics*, vol. PP, 2022.
- [87] A. Narechania, A. Srinivasan, and J. Stasko, "NI4dv: A toolkit for generating analytic specifications for data visualization from natural language queries," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 369–379, 2020.

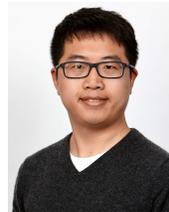
- [88] O. Bau and W. E. Mackay, "Octopocus: a dynamic guide for learning gesture-based command sets," in *Proceedings of the 21st annual ACM symposium on User interface software and technology*, 2008, pp. 37–46.
- [89] A. Cockburn, C. Gutwin, J. Scarr, and S. Malacria, "Supporting novice to expert transitions in user interfaces," *ACM Computing Surveys (CSUR)*, vol. 47, no. 2, pp. 1–36, 2014.
- [90] M. A. Nacenta, Y. Kamber, Y. Qiang, and P. O. Kristensson, "Memorability of pre-designed and user-defined gesture sets," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2013, pp. 1099–1108.
- [91] Q. Zhou, G. Fitzmaurice, and F. Anderson, "In-depth mouse: Integrating desktop mouse into virtual reality," in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–17.
- [92] L. Pavanatto, C. North, D. A. Bowman, C. Badea, and R. Stoakley, "Do we still need physical monitors? an evaluation of the usability of ar virtual monitors for productivity work," in *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2021, pp. 759–767.
- [93] M. Cordeil, B. Bach, A. Cunningham, B. Montoya, R. T. Smith, B. H. Thomas, and T. Dwyer, "Embodied axes: Tangible, actuated interaction for 3d augmented reality data spaces," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–12.
- [94] K. A. Satriadi, J. Smiley, B. Ens, M. Cordeil, T. Czauderna, B. Lee, Y. Yang, T. Dwyer, and B. Jenny, "Tangible globes for data visualisation in augmented reality," in *CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–16.
- [95] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort, "The sandbox for analysis: concepts and methods," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 2006, pp. 801–810.
- [96] S. K. Badam, F. Amini, N. Elmqvist, and P. Irani, "Supporting visual exploration for multiple users in large display environments," in *2016 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2016, pp. 1–10.
- [97] M. Rønne Jakobsen and K. Hornbæk, "Sizing up visualizations: effects of display size in focus+ context, overview+ detail, and zooming interfaces," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, pp. 1451–1460.
- [98] X. Wang, L. Besançon, D. Rousseau, M. Sereno, M. Ammi, and T. Isenberg, "Towards an understanding of augmented reality extensions for existing 3d data analysis tools," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.



Chris North is a professor of computer science with Virginia Tech. He is the associate director of the Sanghani Center for Artificial Intelligence and Data Analytics. His research seeks to create effective methods for human-in-the-loop analytics of big data. His work falls in visual analytics, information visualization, human-computer interaction, large high-resolution display and interaction techniques, and visualization evaluation methods, with applied work in intelligence analysis, cyber security, bioinformatics, and GIS.



Hanspeter Pfister is the Academic Dean of Computational Sciences and Engineering and An Wang Professor of Computer Science in the John A. Paulson School of Engineering and Applied Sciences at Harvard University. He has a Ph.D. in Computer Science from Stony Brook University and an MS degree in electrical engineering from ETH Zurich, Switzerland. His research in visual computing lies at the intersection of visualization, computer graphics, and computer vision and spans a wide range of topics, including biomedical visualization, image, and video analysis, machine learning, and data science.



Yalong Yang is an Assistant Professor in the School of Interactive Computing, Georgia Institute of Technology. He was an Assistant Professor at Virginia Tech, a Postdoctoral Fellow at Harvard University, and a Ph.D. student at Monash University, Australia. His research designs and evaluates interactive human-data interaction systems on both conventional 2D screens and in 3D immersive environments (VR/AR).



Sungwon In is a Ph.D. student at Virginia Tech. Before moving to Virginia Tech, he earned Bachelor's Degree in Computer Science from Indiana University, Bloomington, and he was an A.I. Software Engineer at FiniView. His research interests include the field of data transformation, data visualization, immersive analytics, and human-computer interaction, especially in 3D immersive environments (VR/AR).



Tica Lin is a Ph.D. student at the Visual Computing Group at Harvard University. Prior to joining Harvard, she was a Data Visualization Designer at Visa and a UX Developer at NBA 76ers. Her research interests include data visualization, immersive analytics, and human-computer interaction. In particular, she explores novel visualization and interaction design in Augmented Reality.