Exploring Interactions with Printed Data Visualizations in Augmented Reality

Wai Tong, Chen Zhu-Tian, Meng Xia, Leo Yu-Ho Lo, Linping Yuan, Benjamin Bach, and Huamin Qu

(a) Printed visualizations (b) Augmented digital layers (c) Tilt to pan

(d) Translate to zoom (e) Collocate to link (f) Point to select

Fig. 1. We investigate the possibility of interacting with printed visualizations in Augmented Reality. Suppose a student receives (a) a leaflet about university ranking and wants to analyze three universities' ranking history of interest. Examples of interactions with (b) digital content overlaid: (c) tilt the paper to rescale the y-axis, (d) move (translate) to zoom (e) unfold to show two charts side by side and link them, (f) point to select elements and highlight them in the other chart.

Abstract—This paper presents a design space of interaction techniques to engage with visualizations that are printed on paper and augmented through Augmented Reality. Paper sheets are widely used to deploy visualizations and provide a rich set of tangible affordances for interactions, such as touch, folding, tilting, or stacking. At the same time, augmented reality can dynamically update visualization content to provide *commands* such as pan, zoom, filter, or detail on demand. This paper is the first to provide a structured approach to mapping possible actions with the paper to interaction commands. This design space and the findings of a controlled user study have implications for future designs of augmented reality systems involving paper sheets and visualizations. Through workshops (N=20) and ideation, we identified 81 interactions that we classify in three dimensions: 1) *commands* that can be supported by an interaction, 2) the specific *parameters* provided by an (inter)*action* with paper, and 3) the *number* of paper sheets involved in an interaction. We tested user preference and viability of 11 of these interactions with a prototype implementation in a controlled study (N=12, HoloLens 2) and found that most of the interactions are intuitive and engaging to use. We summarized interactions (e.g., tilt to pan) that have strong affordance to complement "point" for data exploration, physical limitations and properties of paper as a medium, cases requiring redundancy and shortcuts, and other implications for design.

Index Terms—Interaction design, augmented reality, paper interaction, tangible user interface, printed data visualization

◆

1 INTRODUCTION

Interaction with visualizations is necessary for exploration, personalization, and wider engagement with data visualizations. Nevertheless, the specific means and their effectiveness for visualizations still cause considerable controversies and open research questions [27, 38]. For example, visualizations may support direct manipulation [61] through pan&zoom, interactive lenses [75], or brushing&linking. However, these interactions are limited by the current interaction modalities (*e.g.*, mouse, keyboard, touchscreen) and by the visibility and understandability of their interaction affordances [13]. Unnatural interaction, unnoticed affordances, repetitive interactions, ambiguous interaction goals [2], or missing general interaction literacy [3] pose serious obstacles to people engaging with data through visualizations.

To improve affordances and provide for effective interaction, different interaction modalities have been explored [50]. For example, natural language interaction uses voice as a medium for interaction to support querying and creation of visualizations [26]; data physicalizations provide affordances through three-dimensionality, situatedness, tangibility [43], and even dynamicity [72]. Recently, virtual and augmented reality further provide the potential for display and direct interaction [25, 31] as well as offer combinations and hybridizations with tangible means for visualization and interaction [6, 22, 23].

- *Wai Tong, Leo Yu-Ho Lo, Linping Yuan, and Huamin Qu are with the Hong Kong University of Science and Technology. E-mail: {wtong,leoyuho.lo,lyuanaa}@connect.ust.hk, huamin@ust.hk*
- *Chen Zhu-Tian is with Harvard University. E-mail: ztchen@seas.harvard.edu*
- *Meng Xia is with Carnegie Mellon University. E-mail: mengxia@andrew.cmu.edu.*
- *Benjamin Bach is with the University of Edinburgh. E-mail: bbach@exseed.ed.ac.uk.*

Manuscript received 31 March 2022; revised 1 July 2022; accepted 8 August 2022. Date of publication 27 September 2022; date of current version 2 December 2022. This article has supplementary downloadable material available at https://doi.org/10.1109/TVCG.2022.3209386, provided by the authors. Digital Object Identifier no. 10.1109/TVCG.2022.3209386

Complementing this line of research, we explore *paper sheets* as tangible means to interact with visualizations printed onto these paper sheets under augmented reality. We are interested in how far papers can provide affordances and means for direct manipulation with visualizations, and how to inform building systems that use these interactions. This research is motivated by paper being a cheap means to distribute and access information through, *e.g.*, infographics [83], newspapers, posters, books, data comics [5, 81], or zines [54]. Augmented reality (AR)—supported through camera-bearing mobile devices or Head-mounted displays (HMD)—can update such static visualizations [83] and bring interactivity to them by overlaying digital layers. While previous work has demonstrated interaction techniques to interact with data visualization in AR [17, 23], *paper* upon which the visualization is printed provides its very own affordances for interaction. These techniques are only marginally explored yet [6, 44, 65]. For example, paper can serve as a touch surface, which can be moved and rotated in space, bent, folded, moved, tilted, or stacked onto other paper sheets and even torn apart and crumpled. We argue that these interactions might occur naturally and require less training and practice to perform than customized physical tangible devices, such as [22, 23, 32, 40, 63, 67], due to familiarity, yet provide an effective means to interact with the data. Moreover, paper sheets provide tangible surfaces which ease arm fatigue compared to mid-air gestures [23]. With paper interactions, people could easily interact with printed visualization distributed in exhibitions and presentations. Besides, interacting with printed visualizations could be helpful in visualization education [7] and brainstorming [70, 71]. Moreover, paper interactions could possibly facilitate casual collaborative visual analytics [32] due to its low technical barrier and enhance existing authoring tools [24, 62, 82, 83] to support interactive visualizations in AR using paper interactions.

In this work, we present a design space of possible interactions with paper sheets and visualizations enhanced through AR. This design space helps us analyze interactions, inform future interfaces, and point to open research questions. For the purpose of this paper, we define **an *interaction* as the mapping of an *action* onto a *command*, which we denote as a function *action*⇒*command*.** Using terminology from the instrumental interaction framework [9], an *action* is any manipulation applied to an instrument (*e.g.*, a *point*, *rub*, or *fold* to a paper sheet) while a *command* is an interaction task applied to a domain object (*e.g.*, *pan*, *zoom*, and *filter* a data visualization). For example, we can map the action *fold* to the command *filter*. Consequently, we denote this interaction as *fold*⇒*filter* (speak: *"fold-to-filter"* or *"filter-by-fold"*) The parameters that the *fold* action provides, *e.g.*, the degree of folding, can be used to parameterize the *filter* command, *e.g.*, define a threshold for filtering a set of elements from the visualization. Figure 1 shows an example of how a student performs a set of interactions onto a leaflet provided on an university open day.

We collected 146 ideas, 10 commands, and 18 actions from both an extensive literature survey and an ideation workshop with 20 participants (graduate students and researchers in visualization and HCI) (Section 3). Then, we extracted 81 interactions from these ideas and constructed a three-dimensional design space to classify interactions and guide the design of future interactions. The dimensions include 1) the commands supported by an interaction (*e.g.*, *zoom*, *pan*, *filter*, etc.), 2) the specific parameters provided by an interaction (*boolean*, *position/area*, *direction+value*, and *free expression*), and 3) the number of paper sheets involved in an interaction (*1* or *many*). Each interaction, being a combination of an action and a command, can be classified along these three dimensions.

Selecting 11 interactions by focusing on those commands used for view manipulation as described by Heer and Shneiderman [35], we then built an experimental prototype using HoloLens 2 and ran a user study (Section 5) with 12 participants. We were interested in participants' subjective considerations (preference, comfort, intuitiveness, and engagement) as well as interactions' practical viability by observing possible combinations and confounds when using multiple interactions in the same system. Our selected interactions involve eight actions and four commands (*i.e.*, select an interval, zoom, pan, and link&select).

Participants were highly positive towards paper interactions and engaged with the techniques, seamlessly using paper actions to explore static visualizations. We summarize our main findings into six design implications that can inform future designs for interacting with visualizations on paper in AR (Section 6). For example, designers can consider alternative interactions (*e.g.*, *tilt*⇒*pan*) when one interaction (*e.g.*, *point*⇒*pan*) faces technical barriers or is not optimal for different data exploration purposes (*e.g.*, casual exploration). All materials from the user studies and ideation workshops, and a demo video of the experiment prototype can be found at `https://paperinteraction.github.io`.

2 RELATED WORK

Interaction Techniques for AR Visualization. Existing interaction techniques for data visualizations under AR can be roughly classified into five categories based on their modalities: *mid-air*, *tangible*, *touch*, *gaze and speech*, and *spatial* based interfaces. Given that mid-air hand gestures are natural and intuitive for general users, several works have adopted mid-air hand gestures to help users navigate maps [59] and static visualizations projected on projector screens [45]. However, mid-air hand gestures can cause arm fatigue [36] and thus are not suitable for long-term usage. To ease the arm fatigue issue [23], researchers have built AR tangible interfaces [12]. For example, tangible objects, such as paper cards [6, 44, 65], paper spheres [30], embodied axes [22, 63], and custom widgets [32, 40, 67] have been utilized as controllers for users to interact with AR visualizations. As an alternative to the tangible interface, the touch interface on physical objects could be used to manipulate digital information precisely [10, 19, 23, 29]. For example, the touch interface on the tabletop has been utilized for 3D parallel coordinate plot specification and manipulation [19] and 3D selection [10] in AR. Xiao *et al.* [78] further proposed turning every flat surface into a touch screen for head-mounted mixed reality systems. Furthermore, researchers started to utilize gaze and speech interfaces for data visualization interaction [39,45,53] because AR HMDs natively support these interactions (*e.g.*, Microsoft HoloLens 2 and Magic Leap 1). Lastly, spatial user interfaces are also increasingly used for data navigation and placement. For example, researchers extended mobile devices as spatial devices to facilitate 3D data navigation [18] and visualization placement in the 3D environment [39].

Nevertheless, *paper*, leveraging the benefit of touch ability, unique tangibility, and spatial interface, is only marginally explored [6, 44, 65] for data interaction, especially when visualizations can be easily printed on paper. As such, we aim to explore how to utilize the "already there" paper to manipulate data visualizations printed on paper directly.

Paper Interactions in HCI and Visualization. Previous work in the Human-Computer Interaction (HCI) literature investigated the use of paper and its metaphors to achieve better interaction design for tangible interfaces [34, 37], desktops [1], and touchscreens [33, 46, 69, 73]. For example, Holman *et al.* [37] proposed eight paper gestures for interacting with the digital information projected on paper. They are *hold*, *collocate*, *collate*, *flip*, *rub*, *staple*, *point*, and *two-handed pointing*. Utilizing paper interaction has been found to make interaction design more playful and enjoyable, and further help users leverage real-world knowledge in performing the proposed interactions [1].

At the same time, researchers in the visualization community have also used natural interactions with the paper to create more effective ways to interact with visualizations. For example, papers can be utilized as an extra layer on top of a tabletop to interactively display more information [44, 65], while Bach *et al.* [6] used paper cardboard to interact with three-dimensional holograms. Spindler *et al.* [65] further summarized a set of interaction vocabularies for tangible views, such as *translation* and *rotation*. Besides using the paper as a planar, paper could also be used as a prop to interact with 3D visualization of thin fiber structures [41] and a printed wheel chart to interact with volume visualization [68]. Moreover, these paper interactions and their metaphors (*e.g.*, *piling* and *folding*) are heavily used in the traditional desktop and mobile environment for data visualization tasks, such as comparison [74], navigation [28], organization [4, 51], coordination [48], and set operations [57]. Different from utilizing paper

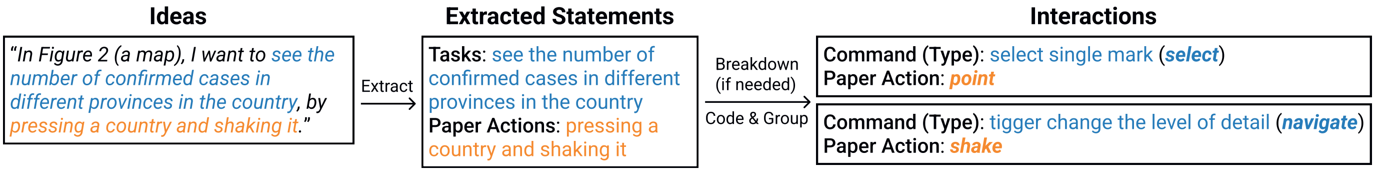

Fig. 2. This figure illustrates the analysis procedure of the ideas collected from the workshop.

interactions as metaphors in the desktop environment, we explore the possibility of using paper as a touch, tangible, and spatial interface for people to interact with digital data intuitively and engagingly in the physical world through AR. We construct a design space to provide designers with a structured way to design systems using paper interactions as well as designing further paper interactions.

Visualization Task Taxonomies. Many works have summarized data exploration tasks as high-level tasks [15, 47, 52] and low-level tasks [15, 35, 80]. High-level tasks, such as *identify*, *compare*, or *summarize*, describe *why* users interact with a visualization [15]. Since low-level tasks are building blocks for high-level tasks [15, 35, 80], we focus on how paper interactions support low-level tasks. Yi *et al.* [80] proposed a set of seven low-level tasks, for instance, *select*, *filter*, and *connect*. Later, Heer and Shneiderman [35] further suggested twelve low-level tasks for data & view specification (*i.e.*, *visualize*, *filter*, *sort*, and *derive*), view manipulation (*i.e.*, *select*, *navigate*, *coordinate*, and *organize*), and analysis process (*record*, *annotate*, *share*, and *guide*). Besides, Brehmer and Munzner [15] added *change* to the low-level tasks.

However, these tasks are mainly explored and summarized in the desktop environment and thus actions beyond the use of the mouse and keyboard are seldom discussed [27, 43, 49]. Our work utilizes the existing low-level tasks as an initial set of *commands* to explore how *actions* on paper sheets can be used to execute these *commands*.

3 SOLICITING INTERACTIONS

To understand the potential of using paper as an interaction medium for data exploration, we conducted an ideation study. Based on the existing literature survey as mentioned in Section 2, we start exploring possible interactions with paper actions: *hold*, *collocate*, *collate*, *flip*, *rub*, *staple*, *point*, and *two-handed pointing* [37], and data visualization commands: *visualize*, *filter*, *sort*, *derive*, *select*, *navigate*, *coordinate*, *organize*, and *change* [15, 35].

3.1 Ideation Workshop

Participants: We invited 20 researchers (one Associate Professor, two Postdoctoral fellows, and 17 Ph.D. students, aged between 22 and 30; 15 males and 5 females) 14 participants came from the same research lab as the first author, four participants joined from other research labs in the same university, and two from other universities. We selected participants with VIS or HCI backgrounds to provide more detailed ideas and start brainstorming in a shorter time due to familiarity with visualization and interaction design [16].

Setup and Materials: To encourage participants to brainstorm more creative ideas (other than familiar point-related gestures), we divided participants into groups of four and across five sessions, inspired by the *partners* technique [55]. We constructed five basic charts (bar chart, pie chart, line chart, scatter chart, and choropleth map) with a Covid-19 dataset (ending on January 17th of 2021) from data repository[1] (containing two-dimensional data, temporal data, and spatial data) to cover the common visualizations and data types encountered in daily life. Since paper action techniques can involve multiple papers/visualizations with the same chart type (*e.g.*, stacking one bar chart on another bar chart), we provided participants two sets of five charts with the confirmed and recovery datasets of Covid-19 cases. Moreover, since the size of the visualization may affect the paper action, we printed two versions of each chart: A4 width and half-A4 width. In total, each participant received 20 ($5 \times 2 \times 2$) charts. Due to the Covid-19 pandemic, all sessions were hosted on Zoom. Participants were asked to print out these materials before the sessions. Each participant received \$10 for compensation.

Procedure: Each session lasted about 90 minutes, consisting of three parts: introduction (15 mins), individual brainstorming (20 mins), and group discussion (55 mins). In the introduction, we first briefly introduced the project background and the workshop's goal. To encourage participants to produce diversified ideas on interactions, we provided two demonstrations created by the authors (flipping a sheet of paper to trigger filtering, and collating two papers to combine two bar charts into one grouped bar chart), inspired by the *priming* technique [55]. Then, we gave participants a task and asked them to spend 20 minutes brainstorming the commands they would like to perform on the printed visualizations and how they would achieve these commands by interacting with the paper. The task description was *Given static Covid-19 figures from a report, what do you want to know more about from the visualizations printed on paper and how will you interact with them?*. To accelerate and encourage participants to brainstorm novel ideas, we provided the initial set of commands and paper actions (we excluded point-related gestures like pointing and two-handed pointing for more diverse ideas) as prompts, adapted from [16]. We also encouraged participants to generate ideas beyond the actions mentioned in the list. Participants were then asked to write down their thoughts without considering any technological restrictions and send them to the host. The host then organized all ideas in a Google Doc for later group discussion. After the individual brainstorming step, each participant shared and demonstrated their ideas to the group. The group then discussed the ideas and brainstormed more ideas (*i.e.*, build upon each other's ideas) and usage scenarios based on these individual ideas in the Google Doc for 55 minutes. As we were interested in collecting a wide variety of ideas for our design space, we did not seek a consensus for a single "ideal" mapping between action and command at the end of each session. All sessions were recorded.

3.2 Data Analysis Procedure

In total, we have gathered 146 ideas from both individual brainstorming and group discussions in all sessions. To extract interactions from these ideas, we performed the following analysis procedures (illustrated in Fig. 2). First, the lead author extracted statements involving tasks and paper actions from the ideas and broke down statements into multiple single action and task mappings if necessary. Next, two authors independently coded the type of commands and actions for the mappings in the first two sessions according to the initial set of data visualization commands [15, 35] and paper actions [37, 65]. For the commands and paper actions that did not fit into the existing taxonomy, the same two authors independently open-coded them and discussed their definitions. For example, participants offered the ideas of *folding* and *tearing* the paper, which were not in the initial set of paper actions. The same two authors iteratively discussed and refined the coding scheme until reaching a Cohen's κ [21] above 0.7 for all classes of actions and commands. The lead author then coded the rest of the sessions. For each command category, the lead author further grouped commands with similar meanings (*e.g.*, "*select a country (from a map)*" and "*select a timestamp (from a line chart)*" are grouped to "*select single mark*"). Finally, we had summarized 81 unique interactions.

---

[1] `https://github.com/CSSEGISandData/COVID-19/`

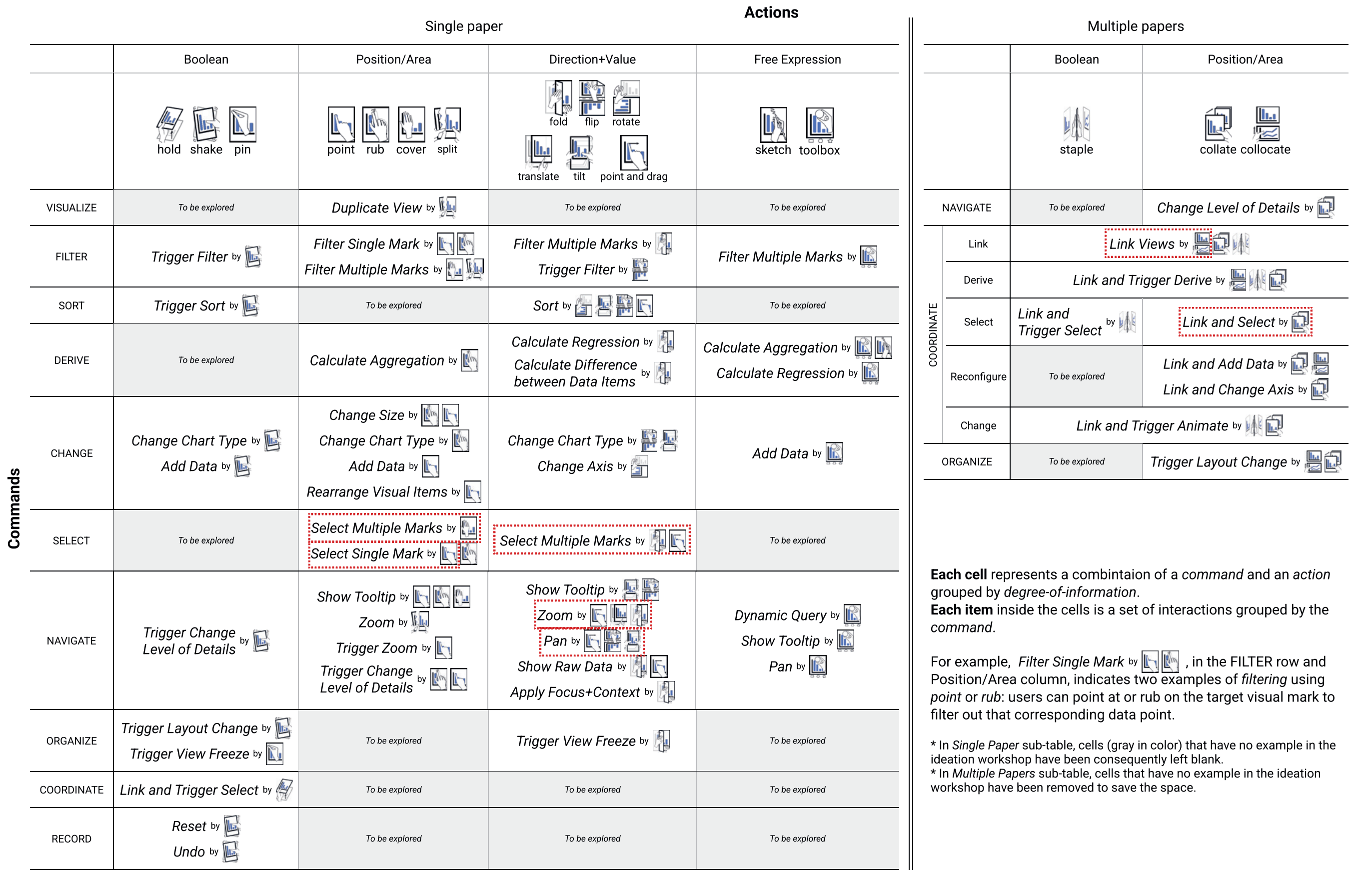

Commands	Boolean	Position/Area	Direction+Value	Free Expression
	hold shake pin	point rub cover split	fold flip rotate translate tilt point and drag	sketch toolbox
VISUALIZE	To be explored	Duplicate View by	To be explored	To be explored
FILTER	Trigger Filter by	Filter Single Mark by Filter Multiple Marks by	Filter Multiple Marks by Trigger Filter by	Filter Multiple Marks by
SORT	Trigger Sort by	To be explored	Sort by	To be explored
DERIVE	To be explored	Calculate Aggregation by	Calculate Regression by Calculate Difference between Data Items by	Calculate Aggregation by Calculate Regression by
CHANGE	Change Chart Type by Add Data by	Change Size by Change Chart Type by Add Data by Rearrange Visual Items by	Change Chart Type by Change Axis by	Add Data by
SELECT	To be explored	Select Multiple Marks by Select Single Mark by	Select Multiple Marks by	To be explored
NAVIGATE	Trigger Change Level of Details by	Show Tooltip by Zoom by Trigger Zoom by Trigger Change Level of Details by	Show Tooltip by Zoom by Pan by Show Raw Data by Apply Focus+Context by	Dynamic Query by Show Tooltip by Pan by
ORGANIZE	Trigger Layout Change by Trigger View Freeze by	To be explored	Trigger View Freeze by	To be explored
COORDINATE	Link and Trigger Select by	To be explored	To be explored	To be explored
RECORD	Reset by Undo by	To be explored	To be explored	To be explored

Multiple papers		Boolean	Position/Area
		staple	collate collocate
NAVIGATE		To be explored	Change Level of Details by
COORDINATE	Link	Link Views by	
	Derive	Link and Trigger Derive by	
	Select	Link and Trigger Select by	Link and Select by
	Reconfigure	To be explored	Link and Add Data by Link and Change Axis by
	Change	Link and Trigger Animate by	
ORGANIZE		To be explored	Trigger Layout Change by

Each cell represents a combintaion of a *command* and an *action* grouped by *degree-of-information*.
Each item inside the cells is a set of interactions grouped by the *command*.

For example, *Filter Single Mark* by , in the FILTER row and Position/Area column, indicates two examples of *filtering* using *point* or *rub*: users can point at or rub on the target visual mark to filter out that corresponding data point.

* In *Single Paper* sub-table, cells (gray in color) that have no example in the ideation workshop have been consequently left blank.
* In *Multiple Papers* sub-table, cells that have no example in the ideation workshop have been removed to save the space.

Fig. 3. Interactions are summarized using the proposed design space. Sub-tables show interactions involving a single paper (left) and multiple papers (right): paper actions grouped by information provided (horizontally) and data visualization commands (vertically). In each cell, we presented the interactions found in the workshop. Interactions highlighted using red dashed rectangles are what we have implemented for the user study.

4 DESIGN SPACE

By analyzing the interactions resulting from the ideation workshop, we constructed a design space to facilitate the organization and creation of paper actions for data exploration in the future. The design space contains three dimensions: *Commands*, *Degree of Information*, and *Number of Paper Sheets Involved*.

4.1 Dimension I: Commands

Dimension I, *Commands*, describes the low-level tasks [15, 35, 66] on data visualizations. We listed out all 10 commands found in the workshop as follows. First, participants wanted to *filter* and *select* data points in the visualizations. To access more details, participants intended to *navigate* (*e.g.*, *zoom*, *pan*, and *show tooltip*) into different charts and *derive* statistical calculations, like `mean`, `min`, and `max`. They might also *change* the chart type for different insights or update the dataset for the latest information. For visual comparison, participants wished to *sort* the data to rearrange the visual marks and *organize* the visualizations in juxtaposition or superimposition. Furthermore, participants wanted to *coordinate* different charts to expand their exploration. For example, one participant wanted *"the information related to this country to be highlighted in another paper (visualization) when one of them is selected."* Lastly, participants proposed to reset the charts or undo some comments (by traversing *recorded* states) to prepare another round of data exploration.

4.2 Dimension II: Degree of Information

Dimension II, *Degree of Information* (DoI), is inspired by the notion of *degree of freedom* in HCI. This dimension describes the number of parameters an action can provide as well as the possible information. Only paper actions that provide a matching DoI can support a given target data visualization command. For example, we can point at the visualization to select a data point because the point action provides the positional information for the system to select the data point in the specified location. However, we cannot shake the paper to select a data point because shaking cannot provide the positional information. Shaking the paper can only trigger a predefined selection. We analyzed the DoI of each action found in the workshop and identified four kinds of DoI, namely, *boolean*, *position/area*, *direction+value*, and *free expression*. We then used these four kinds of DoI to organize the 18 actions found in the workshop (paper actions with an asterisk indicate actions not presented in previous works [37, 65]):

Boolean actions provide a `yes/no` state.

Shake: Move the paper up and down or from side to side forcefully, jerkily, and rapidly. *Shake* provides a boolean information—whether the paper is being shook or not. For simplicity, this will involve some sort of threshold.

Hold: Pick up a piece of paper to the mid-air. *Hold* provides a boolean information—whether the paper is being held or lies flat on a surface.

*Pin**: Anchor the visualization to its current position with the pin hand gesture, similar to fixing a paper on a board using a push pin. *Pin* provides a boolean information–the paper is being pinned, or not.

Staple: Place the papers face to face to mimic the metallic staple effect. *Staple* provides a boolean information–whether papers are being stapled together or not.

Position/Area actions provide the `x`, `y`, `z` value and possibly an `area`.

*Cover**: Put a hand on the paper to block part of the view. Cover provides the position and the area covered by the hand.

Point: Use a finger to point on the paper. *Point* provides the *x*, *y* coordinate of the intended position of the visualization.

Rub: Point on the paper and move the finger back and forth on the paper quickly and repeatedly. Similar to point, *rub* provides the *x*, *y* coordinate of the intended position of the visualization.

Collate: Stack multiple papers together. *Collate* provides the relative position of the upper visualization to the bottom visualization.

Collocate: Organize multiple pieces of paper side-by-side. *Collocate* provides the relative positions of other papers.

Direction+value actions provide a direction and a value.

Flip: Turn the paper's front side back or the back side front. Flipping along different edges of the paper provides the direction information and the current state of the paper—facing up or down.

Tilt: Slant the view plane to a different angle than its normal viewing position. Tilting vertically and horizontally provides different direction information, and the tilt angle provides the value.

Rotate: Reorient the paper to a different angle. Similar to tilt, *rotate* provides the direction of rotation and the rotation degree.

*Fold (bend)**: Fold or bend a piece of paper over to cover other parts of itself. *Fold/bend* provides the folding/bending direction and the portion of the cover.

Translate: Move the paper up and down, left and right, also close or far from the eyes. *Translate* provides the direction and magnitude of the movement.

*Split (Tear/Cut)**: Tear or cut the paper into two parts, splitting up the content. *Tear/Cut* provides the tearing/cutting direction and the size of the resulting parts.

Point&drag: Point and drag one or multiple fingers on the paper, such as drag and pinch gestures. *Point and drag* utilizes time to create the direction and moving distance.

Free expression paper actions can provide an expression beyond numerical values.

Toolbox: Utilize other papers with different shapes, colors, and text annotations as interactive widgets (*e.g.*, buttons, menus, and sliders) for user input. Depending on the design of the paper widget, a toolbox can provide any expression to manipulate the visualization.

Sketch: Use a pen or digital pen to write or draw on the paper. Depending on the predefined commands, free-form sketching or writing can provide any expression to interact with the visualization.

Note that this analysis is capturing only those mappings discussed in the workshop. For example, there could potentially be a multitude of ways to shake a paper, *e.g.*, shake vertically, shake horizontally, shake multiple times. Our design space aims at a first overview of possible and feasible interactions and thus these variations are not considered.

4.3 Dimension III: Number of Paper Sheets Involved

Dimension III, *Number of Paper Sheets Involved*, describes the number of papers involved in the paper action. Paper actions involving one paper target at single view manipulation, while paper actions involving multiple papers supports multiview manipulation and analysis.

Single paper. There are 15 actions (as shown in Fig. 3) found to involve one piece of paper in the workshop. For example, participants pointed at one paper and folded one paper.

Multiple papers. Three actions (*i.e.*, collate, collocate, and staple) were found to involve two or more pieces of paper. These actions allow users to organize multiple sheets of paper into different layouts or use visualization as an object to interact with other visualizations.

4.4 Supporting Commands through Paper Actions

Based on the design space, we describe the collected interactions from the workshop as shown in Fig. 3. The figure shows interactions involving one paper sheet (left) and multiple paper sheets (right). In each sub-table, we have DoI (with the corresponding paper actions) listed horizontally and data visualization commands listed vertically. Each item in a cell represents one or more interactions grouped by the commands. For example, Filter Single Mark by indicates two interactions: *point⇒filter-single-mark*, and *rub⇒filter-single-mark*. Other combinations (*e.g.*, *translate⇒select-single-mark*) that had no practical solutions proposed in the workshop, we have left blank. Below, we explain the details of the design space organized by DoI.

(a) shake to reset (b) rub to filter (c) collate to select and link (d) rotate to sort (e) toolbox to select (f) collocate to link and point to select

Fig. 4. Illustrations of interaction examples. (a) - (e) cover examples in each DoI category. (a) shows a boolean interaction, (b) and (c) show a position interaction and a multiview position interaction, (d) shows a direction+value interaction , and (e) shows a free expression interaction. (f) shows a combination of two interactions to link and select.

Boolean interactions (14/81) are mainly used as a trigger to activate or deactivate commands. For example, it could be used as *shake⇒trigger-filter* (Trigger Filter by) or *shake⇒reset* (Fig. 4(a)). Moreover, we can trigger commands (*e.g.*, derive) with multiview by stapling paper sheets. The number of Boolean interactions is low, probably because the expressiveness (ability to convey users' intentions to the visualization) of these interactions is low.

Position/Area interactions (33/81) allow users to directly communicate with specific visualization components of the visualization since it provides position data for the system to locate visual elements, *i.e.*, *x* and *y* location. In addition to triggering the filter command in the Boolean interaction, users can now specify the visual mark to be filtered out by pointing or rubbing (Filter Single Mark by). Figure 4(b) illustrates that users filter a bar on a printed visualization using the rubbing gesture. Moreover, by involving multiple papers, multiple visualizations can be coordinated using their spatial relationship for more complex multivariate data exploration. There are 12/33 interactions involving multiple papers to perform navigation, coordination, and organization. For example, as shown in Fig. 4(c), users can overlay one paper over the other one to pick a specific timestamp from the bottom visualization and update both visualizations (Link and Select by).

Direction+Value interactions (25/81) provide more information than the position. They allow users to perform commands that require directional information such as sorting in ascending or descending order by rotating, tilting, flipping, or dragging (Sort by) as shown in Fig. 4(d). Moreover, direction+value interactions can also be transformed into area information. For example, users can fold the paper in *x* or *y* direction to select or filter an area of visual marks covered by the folded part of the paper. Some of the interactions, such as *tilt⇒pan* and *translate⇒zoom*, have previously been proposed for mobile devices [64] or tabletop paper lens [65].

Free expression interactions (8/81) can deal with more complex command that are derived and encoded by utilizing extra objects, *i.e.*, pen or customized toolbox. With their expressive power, these interactions can support advanced filtering, querying, or calculation, such as directly picking the elements with the minimum value using a circle shaped representing a `min` command (Calculate Aggregation by , Fig. 4(e)).

Overall, our design space provides designers with a structured way to design paper interactions on printed visualizations. It provides an overview over the feasibility of the interactions. For example, we cannot have *shake⇒select-single-mark* as the *shake* action only provides a Boolean input to the selection command. Designers can look up Position/Area interactions to choose an action for the selection command.

5 USER STUDY

Our design space helps designers to design feasible paper interactions for data visualization. Apart from the design insights, we want to further investigate paper interactions' functionality in real practice.

Table 1. Interactions implemented in this study with their names and a brief description.

Commands	Actions	Description
Select-an-interval	Point&Drag	Swipe the finger on the paper sheet ⟹ select the data points in the range of the axis brushed
	Cover	Cover part of the visualization with flat hand ⟹ select the data points NOT in the range of the axis covered
	Fold	Fold the paper to cover a large (or small) portion of the axis ⟹ select the data points NOT in the range of the axis covered
Zoom	Pinch	Two-finger pinch outward (or inward) on the visualization ⟹ zoom-in (or zoom-out) of the visualization
	Translate	Move the paper closer to (or farther from) the camera ⟹ zoom-in (or zoom-out) of the visualization
	Fold	Fold the paper to cover a large (or small) portion of the axis ⟹ zoom-in (or zoom-out) to the portion of visualization NOT covered
Pan	Point&Drag	Finger scroll left (right, up, or down) on the visualization ⟹ pan the visualization rightward (leftward, downward, or upward)
	Tilt	Slant the paper to the left (right, up, or down) relative to the ground ⟹ pan the visualization rightward (leftward, downward, or upward)
	Flip	Flip the paper from left to right (right to left, up to down, or down to up) ⟹ pan the visualization leftward (rightward, upward, or downward)
Link&Select	Collate	Put one visualization on top of another and center it to specific position relative to the bottom one ⟹ connect the two visualizations, select the data point from the bottom one, and update both visualizations
	Collocate&Point	Put two visualizations side by side and point at the visualization ⟹ connect the two visualizations, select the data point, and update both visualizations

Fig. 5. This figure shows the framework of the prototype and relationship between client, server, and configuration website.

Thus, we conducted a controlled user study with a proof-of-concept prototype to investigate user preferences *(G-Preferences)* and practical viability *(G-Viability)* of the interactions.

5.1 Prototype

We simplify our study by focusing on those commands used for view manipulation as described by Heer and Shneiderman [35]. We implemented the 11 interactions listed in Table 1. An example of using *fold⇒zoom* can be found in Fig. 7. We built the experimental prototype using a client-server model with a WebSocket for the network communication as shown in Fig. 5.

- The **client** app runs on the Hololens 2 to detect users' hand gestures and papers' status. The newer HoloLens 2 provides a diagonal 52° field of view and two-handed fully articulated hand tracking. Hand-pose data was collected through Microsoft's Mixed Reality Toolkit on HoloLens 2. For pose tracking of printed visualizations, we used Vuforia image tracking. Rich feature patterns were added on all sides of the paper for more accurate tracking. Depending on the relative position and direction of fingers and paper sheets, different paper actions and gestures are recognized as action events (Table 1). For example, when two sheets were placed close together, a collocate action event was created (Fig. 1(e)). After observing the events, the client app sends these events to the server.
- The **server** receives the paper action events with their parameters (*e.g.*, positions in the case of a Position/Area action) and updates the visualization through a corresponding command. We maintain a Vega [60] specification for each visualization on the server. The updated visualization is returned to the client.
- The **configuration website** with dropdown menus allows the conductor to change the selected interactions based on the task and participants' needs (Fig. 6(c)). The configuration is updated to the server and the effect is immediately reflected in the client app.

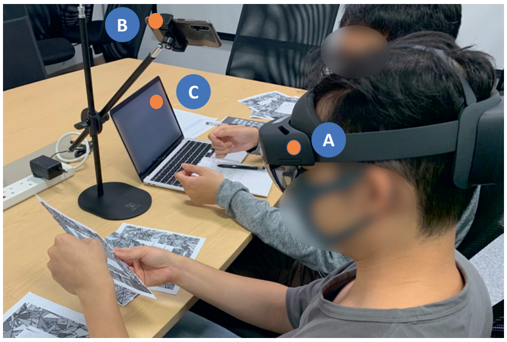

Fig. 6. This figure shows the experiment setting: (a) the participant were wearing the HoloLens 2 to view the digital visualization in AR; (b) both of the participant's hands were recorded by both a smartphone on a phone stand; (c) a laptop was used by one of the authors to change the settings of the current activated interactions.

5.2 Setup and Participants

We recruited 12 university students (P1-P12; aged between 22 and 30; 6 males and 6 females). None of them had participated in the ideation workshop. The distribution of their visual analysis experience was "none" (3), "novice" (5), "knowledgeable" (4), and "expert" (0). The distribution of their AR experience was "none" (2), "novice" (9), "knowledgeable" (1), and "expert" (0). The distribution of their daily paper usage was "0 day per week" (1), "1-2 days per week" (2), "3-6 days per week" (1), and "every day" (8). Overall, participants are mainly novices in both visual analysis and AR, and daily paper users. All sessions were recorded using a mobile phone, as shown in Fig. 6(b).

5.3 Procedure

The study consisted of an introduction, two tasks, and a semi-structured interview. Each participant received $13-$17 as compensation according to a 90-120 minute study time.

Introduction (∼10 mins). We first introduced the study background and procedure, and then asked participants to sign the consent form. After that, participants were asked to put on the AR HMD and adjust the device until they felt comfortable and could see the AR content attached to the printed visualization clearly.

Task 1: Unit Evaluation (∼60 mins). To assess *G-Preferences* for each interaction, we asked participants to perform all 11 interactions on a set of printed visualizations showing the latest Covid-19 dataset (*e.g.*, a scatter plot with total confirmed cases against total recovered cases, Fig. 7) , the source of which was the same as the workshop. We counterbalanced the sequence of data visualization commands and also

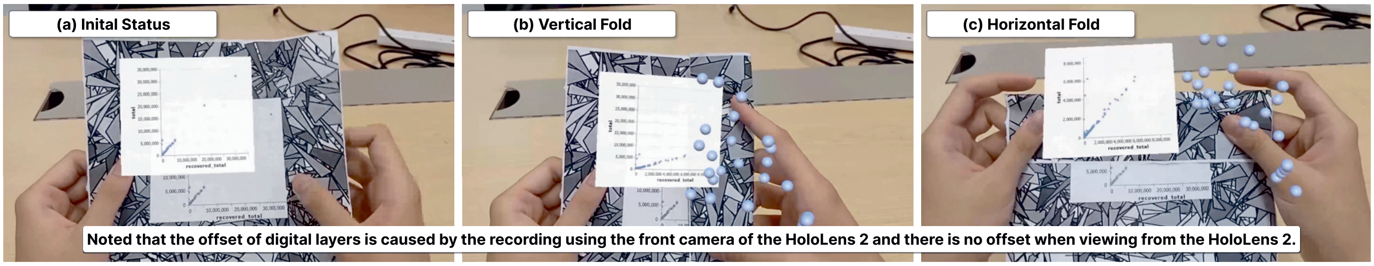

Noted that the offset of digital layers is caused by the recording using the front camera of the HoloLens 2 and there is no offset when viewing from the HoloLens 2.

Fig. 7. This figure demonstrates an example of *fold⇒zoom*: (a) the initial scatter plot and its changes after (b) vertically and (c) horizontally folding the paper. The right hand joints are visualized using the white spheres.

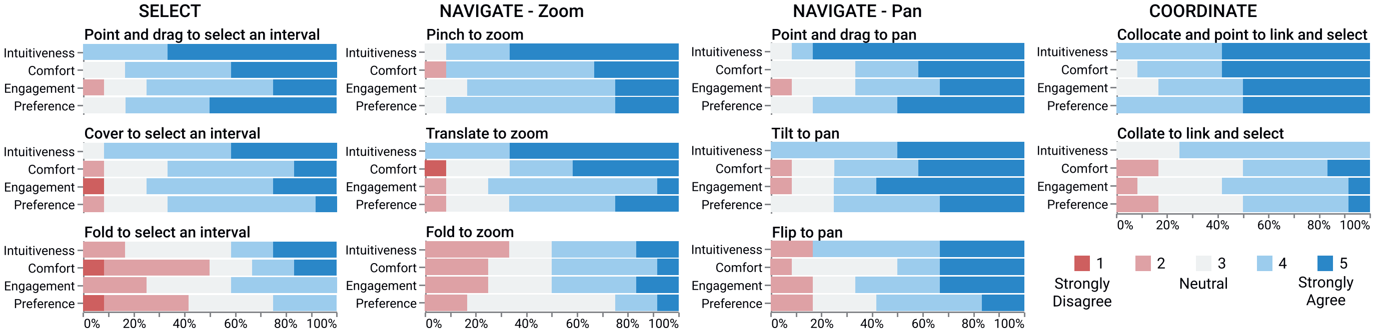

Fig. 8. This figure shows participants' ratings on "intuitiveness", "comfort", "engagement", and "overall preference" for each interactions in task 1.

the sequence of actions per command using the balanced Latin Square method [14]. For each command, we first introduced the task to the participants and then the participants started to try each interaction with the procedure below:

1. The study conductor demonstrated the interaction to the participant.
2. The participant performed the interaction five or more times successfully.
3. The participant rated the interaction on metrics widely adopted in previous research on interactions in AR, *i.e.*, intuitiveness, comfort, engagement, and overall preference [58], by filling in a questionnaire. (To get ratings independently and not implementation specific, participants were told that it was not necessary to compare it with other presented interactions, and it was assumed that the HoloLens 2 worked without technical issues [34].)
4. A series of follow-up questions were asked to obtain further comments and in-depth rationales for the ratings.

After Task 1, participants were asked if they had any discomfort. A five-minute break was given based on the participants' needs.

Task 2: Free-Form Exploration (∼15 mins). For *G-Viability* (whether people can use and how they use interactions), we asked participants to use the above-mentioned interactions to answer a question and explore the data freely within the given 15 minutes time frame. We introduced a set of five visualizations (two visualizations are shown in Fig. 1) on a different dataset from Task 1 (*i.e.*, worldwide university rankings in 2016) to reduce the effect of memorizing the dataset from Task 1. To initialize a set of interactions for participants to perform the Task 2 as the initial setting for free exploration, we picked those interactions with the highest preference for each command from Task 1. For interactions with the same preference rating, we let the participants choose the interaction. In addition, participants were encouraged to tell us when they want to change and explore different interaction mappings, *e.g.*, changing *pinch⇒zoom* to *translate⇒zoom*), and then we change the setting for them. To kick-start exploration, we asked participants to answer the following question: "*Given the line chart, what is the trend of MIT's total score from 2011 to 2016?*". Participants were required to zoom (select, and pan if necessary) since the line chart was complex and cluttered at times, as shown Fig. 1(b). After answering this question, participants could use the remaining time to explore on the five visualizations freely. During free exploration, participants were asked to think aloud about what they were doing and how they planned to perform an interaction.

Post-Study Interview (∼15 mins). We conducted a semi-structured interview with nine questions in three topics: preference, usefulness, and possible new interaction, at the end of the study.

5.4 Results

We report on participants' quantitative ratings and verbal feedback for the interactions from both tasks as well as the semi-structured interview in the user study. Overally, with respect to *G-Preferences*, we found that participants rated the proposed interactions intuitive and engaging to use. With respect to *G-Viability*, they enjoyed interacting with the printed visualizations using paper actions with different affordances.

Preference and feedback for interactions

Figure 8 shows the ratings of 12 participants on *intuitiveness*, *comfort*, *engagement*, and *overall preference* across all interactions for each data visualization command. Overall, participants were positive for most of the interactions from each dimension. Below we present the users' preferences and feedback grouped by the interaction commands. The number inside the brackets indicates the median of the ratings.

Select-an-interval. In this task, participants were asked to select a range of bars in a given bar chart. Participants preferred the *point&drag* action the most (4.5). They explained that *point&drag* was very intuitive (5) as it was *"similar to their current practise [, i.e.,] clicking on desktops"* and *"touching on tablets"*. Participants also preferred *cover* (4), reporting this action to be natural and easy to perform (intuitiveness: (4)). P1 and P9 found themselves engaged and immersed in the data when performing *cover*. P1 said *"it [cover] is fun because I feel involved in the virtual world."* P9 commented that *"it [cover] is just like communicating with data using body language."*

Participants rated *fold* rather neutral (3) because of less comfort (2.5). Reasons reported included that folding damages the paper (P4, P5, P8) and requires extra effort (P2, P3, P4, P12). P4 further commented that both hands would be required to perform the action: *"it is troublesome for me to use both hands to interact with the paper."* Participants added that they were *"sometimes lazy"*.

Zoom. Participants were asked to zoom in and out the given scatterplot to get an overview or obtain detailed information about the data. The two most preferred actions were *pinch* (4) and *translate* (4). One reason was that both actions are intuitive (*pinch*: (5), and *translate*: (5)). Participants reported again that these actions were natural and similar to daily practice. *"I think everyone [who uses smartphones or tablets] is used to pinch,"* P8 said. For *translating* papers, P1 said *"I do this in the real world when I want to see something larger on paper."*

P7 found that *fold⇒zoom* had a unique advantage in terms of preciseness and preferred using it because it allowed for *"controlling the exact amount of zoom in."*

Pan. Participants were asked to pan a given scatterplot in different directions. Overall, *point&drag* was again strongly preferred (4.5) because of participants' familiarity with actions on touchscreen (intuitiveness: (5)). It is exciting that both *tilt* (4.5) and *flip* (4) were also highly ranked. P1 emphasized that *"flip [to pan] is surprisingly easy to understand. It's like there is a bigger visualization behind." Tilt* was ranked intuitive (4.5) and engaging (5), reporting this action to be natural and playful to perform the *pan* command, as well as *"similar to a waterfall"* (P2, P8), *"playing games"* (P2, P3), and *"driving"* (P1).

Link&Select. Participants were asked to first select countries from a bar chart by selecting a continent in a pie chart (Fig. 4(f)) and then select a time in a timeline to show different data in the pie chart about that specified time (Fig. 4(c)). *Collocate&point* was strongly preferred (4.5) due to its naturalness and strong familiarity of selection by "touching" (5). *Collate* was less preferred (3.5) because of occlusion (P1-5, P7, P8, and P11). However, half of the participants (P1-3, P5, P8, and P12) appreciated *collate* for revealing temporal changes. They described *collate* as novel and engaging and that they could easily focus the changes at the top visualization.

Practical Viability and Observation

All participants could complete the question and explore the data in Task 2. During exploration, there are four participants (P3, P5, P6 and P12) changed the interactions. Two participants (P5 and P12) had switched from *pinch⇒zoom* to *fold⇒zoom* when answering the question. While they first used *pinch⇒zoom* and further *point&drag⇒pan* to the cluttered lines, it required several trials of zoom and pan to observe the trend, which was tedious. Thus, they tried using *fold⇒zoom* because they noticed that folding the paper might possibly zoom in to that specific area easier. Moreover, P6 and P3 found alternatives to *point* actions. P6 has switched *collocate&point⇒link&select* to *collate⇒link&select* for more accurate selection and P3 had changed *point&drag⇒pan* to *tilt⇒pan* for the free exploration in the remaining time. P3 commented that *"it [tilt] is much easier to perform than dragging when the hand tracking is not working well."*

Users' attitudes and reactions

All participants enjoyed interacting with the printed visualizations using paper actions and looked forward to the complete prototype system. As key strengths for data exploration participants reported that paper interactions *"increase the capability of static visualizations"* (5/12 participants) and were generally *"convenient"* (4/12 participants). For example, P3 commented that *"direct interacting with papers is more convenient than using PC for data exploration."* Yet, the key weaknesses were reported to be *"the durability of the paper"* (4/12 participants) and *"ergonomic issues brought by the HMD"* (4/12 participants). Overall, all participants stated that they would use these interactions for data analysis in other contexts. Five participants would have liked to perform multiview analyses on experiment reports and academic papers. Seven participants envisioned that paper interactions could be used in presentations to interact with data directly on the printed reports. Moreover, four participants can see paper interactions being used in education due to the interactions engagingness. For instance, P10 stated *"It would be great if students could interact with the map directly to learn about geography."* Furthermore, due to the ubiquity of paper and the ease of deploying interactions, three participants imagined using it inside shopping malls and exhibitions to interact with the materials (*e.g.*, leaflets) received.

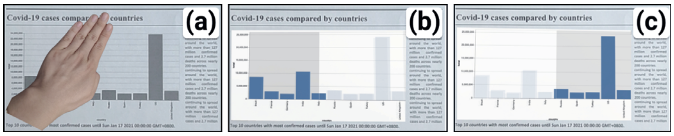

Fig. 9. When a user (a) covers five bars in a bar chart, (b) the covered bars or (c) the uncovered bars can be selected.

6 DESIGN IMPLICATIONS

Based on the results and observations in the study, we derived six design implications (*i.e.*, I1-I6) for future designs and studies.

I1. Provide Redundant Actions to *Point*. Due to the limited accuracy of fingertip detection using only computer vision algorithms, we suggest designers provide redundant actions, such as cover, tilt, and translate, to *point* for selecting an interval, panning, and zooming, respectively. Although point-related paper actions are the most preferred paper actions for all commands presented in the study, it is still challenging to provide a good and precise pointing experience similar to touchscreen due to the limitations of fingertip tracking with occluded and fast-moving fingers [8] and depth estimation with deformable papers. In addition, imprecise pointing required participants to increase their fingers' movement when they conducted data exploration using pointing gestures. Moreover, the study from Spindler *et al.* [64] showed that the task completion time of spatial input (i.e., 3D translation) for 2D document navigation on mobile phones was faster than conventional point-based input. As a result, we suggest that designers use cover, translate, and tilt to complement pointing in range selection, zoom, and pan, respectively, with more accurate detection and similar high scores for intuitiveness, comfort, and engagement.

I2. Make Use of Different Actions for Command Shortcuts. Paper actions can be utilized as shortcuts to save users' efforts. Paper sheets can provide additional actions (*e.g.*, flip and fold) compared to the mouse, touchscreen, and keyboard (*i.e.*, click, touch, and keypress). We can utilize some of the proposed interactions (*i.e.*, *fold⇒zoom*, *flip⇒pan*, and *collate⇒link&select*) that provide unique advantages as shortcuts. For *flip⇒pan*, participants agreed on its intuitiveness and showed its strength in panning a long distance, which can relieve users from pointing and dragging multiple times and tilting for a long time. Furthermore, *fold⇒zoom* is beneficial when dealing with skewed data distribution (*i.e.*, dense points in the corner in the visualization). Lastly, users are engaged to use *collate⇒link&select* to quickly link two visualizations and make a selection simultaneously to focus on the temporal changes of the visualization on the top without context switching compared with *collocate&point⇒link&select*.

I3. Support both Selection and Inverse Selection for *Cover*. Our study suggests supporting *cover⇒select-an-interval* for both selection and inverse selection. *Cover* an area can be treated as selecting wanted data or excluding unwanted data, as shown in Fig. 9. In the user study, while it is easy to use the cover gesture to hide a small set of outliers and focus on the main area of the data, it becomes difficult to cover a large portion of the visualization to select a small amount of uncovered data. This trade-off has also been stated in [72]. As a result, designers can provide both selection and inverse selection by using different gestures, such as palm up and palm down.

I4. Utilize the Semantic Meaning of Paper Actions. Designers should consider the semantic meaning of the paper action to increase the intuitiveness when designing new interaction. Although paper actions within the same DoI could be used for a command, they provide different semantic meanings related to day-to-day usages of papers. For example, participants in the workshop preferred rubbing as a filtering action (as a rubber) or a revealing action (as a cleaner). Rotation-based actions, *i.e.*, tilt and rotate, correlate physics-based metaphor, such as gravity. Moreover, moving a paper sheet back and forth has implicitly provided a zooming metaphor. As such, *translate⇒zoom* and *tilt⇒pan* provide strong semantic meaning and support a strong mental model [42, 56], thus getting high ratings in intuitiveness.

I5. Intuitiveness of Actions vs. Readability of Text. Designers need to consider the trade-off between the intuitiveness of the paper action and the readability of the digital visualization. Users need to read both the text and the visual marks on the visualization for value retrieval and pattern recognition. Therefore, ensuring text readability is essential for designing interactions. In the study, some interactions are intuitive and engaging. However, they may not be optimal for readability. For example, *tilt⇒pan* causes the text to be hard to read because the text is tilted. Thus, designers might consider ensuring readability when using actions that involve movements on paper, such as tilting and translating. One possible solution is to fix some visualization components in place, such as the title, axis, and legend, while translating and tilting. While the text is fixed at a certain distance for reading, the visual pattern could follow the movement for intuitiveness and engagement. This implication could value beyond the domain of visualization to general physical documents.

I6. Effects of Paper's Physical Properties on Actions. Visualizations can be printed on, *e.g.*, books, A4-sized leaflets, and small paper cards. During the investigation of paper actions for data exploration, we found that the physical properties (*e.g.*, size, weight, thickness, and physical constraints) could affect the usage of interactions. First, people prefer large-sized papers to perform actions with large movements or high precision. Participants in the ideation workshop preferred to interact with visualizations on a larger paper size because they can perform actions requiring large movements easier, such as folding. On the other hand, it becomes hard to select a large area with *cover⇒select-an-interval* if the size of the paper is too large. Second, the weight of the paper used should be light, so that none of the participants reported that interacting with the paper for about an hour was tiring. Furthermore, participants can easily pick up the paper sheet for a better angle to view the visualization. Third, the thickness affects the use of paper action. In the user study, we used standard office paper, which is thin. While participants can easily fold the papers, it is hard to perform point-related gestures, as the thin paper cannot support the force given by the participants' fingers. Media like paper cards may provide an ideal experience for these actions. Fourth, designers should consider the paper format. Paper actions (*i.e.*, *fold*, *collocate*, *collate*, *flip*, and *staple*) are constrained if papers are bounded together. For example, participants can only fold one side of the paper and cannot perform multi-paper actions if papers are bounded as books and magazines. These findings might still be valid outside the field of visualization.

7 DISCUSSION

Paper Interaction Design Space. Our design space captures actions and their mappings to commands. It shows the mappings we have investigated in this study while leaving the possibility and feasibility of other mappings to future work. The design space and our findings can help choosing mappings for real-life systems. It shows that for some combinations of actions and commands, multiple options exist.

Specifically, our design space and study help making more informed decisions for creating systems based on paper interactions. For example, participants in our study preferred point-based gestures or spatial paper actions (*e.g.*, *tilting*, *translating*) over *folding*. Designers should also consider redundant actions for the same command, especially when a command requires different levels of granularity. For example, tilt could be used for fine-grained panning, while flip could be used to quickly pan over large distances.

Possible usage scenarios. Our study demonstrates that interacting with printed visualizations is fun and practically viable. We envision five application opportunities for applying paper interactions with data analysis: (1) **education**: we can add interactivity to paper sheets that could benefit classroom teaching that are still common to use paper and data visualization, such as teaching data visualization [7], geography (*e.g.*, printed maps), chemistry (*e.g.*, printed experiment results), and math (printed or hand drawn plots); (2) **brainstorming sessions**: UX designers may consider making use of different types of papers for different tasks. In addition to qualitative data analysis with printed reports, sticky note is a common tool for supporting brainstorming [71]. Supporting interactions directly on papers can reduce the context switching between the desktop visual analytics tools and sticky notes [70]; (3) **exhibitions/presentations**: audiences could directly interact with paper handouts (*e.g.*, worksheets, leaflets, and pamphlets; as shown in Fig. 1) provided without switching back and forth between mobile phones and handouts to seek more information during visiting an exhibition and attending a presentation. (4) **collaboration**: with intuitive paper actions, people could quickly explore the data printed on the paper sheets, which can support short analytical sprints. It helps enhancing collaboration between diverse domains [32]. (5) **AR-based authoring tools for interactive visualization**: current AR-based authoring tools [24,62,82,83] only support minimal or even no interaction configuration, which hinders users to interact with the created visualization. Our paper interaction design space could help developers and researchers to further extend their tools to support feasible paper interactions for interactive AR visualizations.

Study Limitations. Despite our best effort, this study has some limitations to be aware of. First, our investigation was mainly based on the workshop's outcomes and unable to investigate all possible designs conclusively and exhaustively. Second, AR technology is still premature. For example, the inaccurate detection of fingers and paper hindered the user experience; blurry text, due to the fixed focal length of HoloLens 2, caused eye fatigue and strain (reported by half of the participants). Furthermore, as an exploratory study, we have implemented a subset of 11 interactions with two paper-specific actions (*i.e.*, *cover* and *fold*) to complete simple tasks. The sample of the user study is also small and did not consider the analytic benefits of interactions. However, our design space with 81 interactions, our grouping of paper actions and commands, as well as our study results provide a good framework for a more systematic exploration of paper interactions in the future.

Future Work. More studies could be done to expand the design space for the analysis process & provenance [35]. The prototype could also be extended to conduct more studies for assessing other aspects (*e.g.*, task accuracy and completion time for analytic benefits and memory test for intuition), as well as complex tasks for authoring visualizations and immersive collaborative analysis [11, 32, 84]. It could further include more paper-specific actions to support more commands, such as *dogearing⇒pin-view*. Furthermore, artificial intelligence could be introduced to facilitate better interaction support in AR [76, 77], and better paper detection and finger detection with depth cameras and extra sensors. Last but not least, it is interesting to explore the possibility of using paper action as metaphors for intuitive gesture design in the air (without actually interacting with physical papers) to interact with data visualizations in virtual reality (VR) [20, 79]. Although the haptic feedback might be lost, there are more design choices when deploying paper interactions without the physical paper. For instance, undoing a tearing action on a virtual paper sheet is possible in VR.

8 CONCLUSION

This paper explores the use of paper sheets as a new means for interaction with data visualization. We first conducted an ideation workshop with 20 VIS and HCI researchers to solicit 81 interactions. Furthermore, we construct a three-dimensional design space (*i.e.*, *Commands*, *Degree of Information* and *Number of Paper Sheets Involved*) to describe and create possible interactions and verify the feasibility of interactions. Lastly, we built a proof-of-concept prototype and conducted a user study with 12 participants to provide initial insights by evaluating 11 interactions. Our findings show that all participants considered these interactions intuitive and engaging. Based on the findings, we developed six design implications. We found strong affordances for some interactions, physical limitations and properties of paper as a medium, cases requiring redundancy and shortcuts, and other implications for design. We hope that our work can inspire future work on developing interactions for data exploration more intuitively and engagingly.

ACKNOWLEDGMENTS

This work is partially supported by Hong Kong RGC GRF Grant (No. 16210321).

REFERENCES

[1] A. Agarawala and R. Balakrishnan. Keepin'it real: Pushing the Desktop Metaphor with Physics, Piles and the Pen. In *Proceedings of the Conference on Human Factors in Computing Systems*, pp. 1283–1292, 2006.

[2] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.*, pp. 111–117. IEEE, 2005.

[3] B. Bach. Ceci n'est pas la data: Towards a notion of interaction literacy for data visualization., 2018.

[4] B. Bach, N. Henry-Riche, T. Dwyer, T. Madhyastha, J.-D. Fekete, and T. Grabowski. Small MultiPiles: Piling Time to Explore Temporal Patterns in Dynamic Networks. In *Computer Graphics Forum*, vol. 34, pp. 31–40. Wiley Online Library, 2015.

[5] B. Bach, N. H. Riche, S. Carpendale, and H. Pfister. The Emerging Genre of Data Comics. *IEEE Computer Graphics and Applications*.

[6] 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*, 24(1):457–467, 2017.

[7] S. Bae, R. Yang, P. Gyory, J. Uhr, D. A. Szafir, and E. Y.-L. Do. Touching information with diy paper charts & ar markers. In *Interaction Design and Children*, pp. 433–438, 2021.

[8] A. Bandini and J. Zariffa. Analysis of the hands in egocentric vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 2020.

[9] M. Beaudouin-Lafon. Instrumental interaction: An interaction model for designing post-wimp user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '00, p. 446–453. Association for Computing Machinery, New York, NY, USA, 2000. doi: 10.1145/332040.332473

[10] H. Benko and S. Feiner. Balloon selection: A multi-finger technique for accurate low-fatigue 3d selection. In *2007 IEEE Symposium on 3D User Interfaces*. IEEE, 2007.

[11] M. Billinghurst, M. Cordeil, A. Bezerianos, and T. Margolis. Collaborative immersive analytics. In *Immersive Analytics*, pp. 221–257. Springer, 2018.

[12] M. Billinghurst, H. Kato, I. Poupyrev, et al. Tangible augmented reality. *ACM Siggraph Asia*, 7(2):1–10, 2008.

[13] J. Boy, L. Eveillard, F. Detienne, and J.-D. Fekete. Suggested interactivity: Seeking perceived affordances for information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):639–648, 2016. doi: 10.1109/TVCG.2015.2467201

[14] J. V. Bradley. Complete counterbalancing of immediate sequential effects in a latin square design. *Journal of the American Statistical Association*, 53(282):525–528, 1958.

[15] M. Brehmer and T. Munzner. A Multi-level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, 2013.

[16] N. Bressa, K. Wannamaker, H. Korsgaard, W. Willett, and J. Vermeulen. Sketching and ideation activities for situated visualization design. In *Proceedings of the 2019 on Designing Interactive Systems Conference*, pp. 173–185, 2019.

[17] W. Büschel, J. Chen, R. Dachselt, S. Drucker, T. Dwyer, C. Görg, T. Isenberg, A. Kerren, C. North, and W. Stuerzlinger. Interaction for immersive analytics. In *Immersive analytics*, pp. 95–138. Springer, 2018.

[18] W. Büschel, A. Mitschick, T. Meyer, and R. Dachselt. Investigating smartphone-based pan and zoom in 3d data spaces in augmented reality. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*, pp. 1–13, 2019.

[19] S. Butscher, S. Hubenschmid, J. Müller, J. Fuchs, and H. Reiterer. Clusters, trends, and outliers: How immersive technologies can facilitate the collaborative analysis of multidimensional data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2018.

[20] X. Chu, X. Xie, S. Ye, H. Lu, H. Xiao, Z. Yuan, C. Zhu-Tian, H. Zhang, and Y. Wu. TIVEE: visual exploration and explanation of badminton tactics in immersive visualizations. *IEEE Trans. Vis. Comput. Graph.*, 28(1):118–128, 2022. doi: 10.1109/TVCG.2021.3114861

[21] J. Cohen. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46, 1960.

[22] 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*, pp. 1–12, 2020.

[23] M. Cordeil, B. Bach, Y. Li, E. Wilson, and T. Dwyer. Design space for spatio-data coordination: Tangible interaction devices for immersive information visualisation. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 46–50. IEEE, 2017.

[24] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer. Iatk: An immersive analytics toolkit. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 200–209. IEEE, 2019.

[25] 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*, pp. 71–83, 2017.

[26] K. Cox, R. E. Grinter, S. L. Hibino, L. J. Jagadeesan, and D. Mantilla. A multi-modal natural language interface to an information visualization environment. *International Journal of Speech Technology*, 4(3):297–314, 2001.

[27] E. Dimara and C. Perin. What is interaction for data visualization? *IEEE transactions on visualization and computer graphics*, 26(1):119–129, 2019.

[28] N. Elmqvist, N. Henry, Y. Riche, and J.-D. Fekete. Melange: space folding for multi-focus interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1333–1342, 2008.

[29] N. A. ElSayed, R. T. Smith, K. Marriott, and B. H. Thomas. Blended ui controls for situated analytics. In *2016 Big Data Visual Analytics (BDVA)*, pp. 1–8. IEEE, 2016.

[30] D. Englmeier, J. Dörner, A. Butz, and T. Höllerer. A tangible spherical proxy for object manipulation in augmented reality. In *IEEE Conference on Virtual Reality and 3D User Interfaces*, pp. 221–229. IEEE, 2020.

[31] 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. 2020.

[32] B. Ens, S. Goodwin, A. Prouzeau, F. Anderson, F. Y. Wang, S. Gratzl, Z. Lucarelli, B. Moyle, J. Smiley, and T. Dwyer. Uplift: A tangible and immersive tabletop system for casual collaborative visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 2020.

[33] A. Girouard, A. Tarun, and R. Vertegaal. Displaystacks: interaction techniques for stacks of flexible thin-film displays. In *Proceedings of the 2012 CHI Conference on Human Factors in Computing Systems*, pp. 2431–2440, 2012.

[34] A. Gupta, B. R. Lin, S. Ji, A. Patel, and D. Vogel. Replicate and reuse: Tangible interaction design for digitally-augmented physical media objects. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2020.

[35] J. Heer and B. Shneiderman. Interactive Dynamics for Visual Analysis. *Queue*, 10(2):30–55, 2012.

[36] J. D. Hincapié-Ramos, X. Guo, P. Moghadasian, and P. Irani. Consumed endurance: a metric to quantify arm fatigue of mid-air interactions. In *Proceedings of the 2014 CHI Conference on Human Factors in Computing Systems*, pp. 1063–1072, 2014.

[37] D. Holman, R. Vertegaal, M. Altosaar, N. Troje, and D. Johns. Paper windows: interaction techniques for digital paper. In *Proceedings of the 2005 CHI Conference on Human Factors in Computing Systems*, pp. 591–599, 2005.

[38] K. Hornbæk, A. Mottelson, J. Knibbe, and D. Vogel. What do we mean by "interaction"? an analysis of 35 years of chi. *ACM Transactions on Computer-Human Interaction*, 26:1–30, 07 2019. doi: 10.1145/3325285

[39] S. Hubenschmid, J. Zagermann, S. Butscher, and H. Reiterer. Stream: Exploring the combination of spatially-aware tablets with augmented reality head-mounted displays for immersive analytics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–14, 2021.

[40] P. Issartel, F. Guéniat, and M. Ammi. Slicing techniques for handheld augmented reality. In *2014 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 39–42. IEEE, 2014.

[41] B. Jackson, T. Y. Lau, D. Schroeder, K. C. Toussaint, and D. F. Keefe. A lightweight tangible 3d interface for interactive visualization of thin fiber structures. *IEEE transactions on visualization and computer graphics*, 19(12):2802–2809, 2013.

[42] R. J. Jacob, A. Girouard, L. M. Hirshfield, M. S. Horn, O. Shaer, E. T. Solovey, and J. Zigelbaum. Reality-based interaction: a framework for post-wimp interfaces. In *Proceedings of the 2008 CHI Conference on Human Factors in Computing Systems*, pp. 201–210, 2008.

[43] Y. Jansen and P. Dragicevic. An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2396–2405, 2013.

[44] K. Kim and N. Elmqvist. Embodied lenses for collaborative visual queries on tabletop displays. *Information Visualization*, 11(4):319–338, 2012.

[45] T. Kim, B. Saket, A. Endert, and B. MacIntyre. Visar: Bringing interactivity to static data visualizations through augmented reality. *arXiv preprint arXiv:1708.01377*, 2017.

[46] B. Lahey, A. Girouard, W. Burleson, and R. Vertegaal. Paperphone: understanding the use of bend gestures in mobile devices with flexible electronic paper displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1303–1312, 2011.

[47] H. Lam, M. Tory, and T. Munzner. Bridging from Goals to Tasks with Design Study Analysis Reports. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):435–445, 2017.

[48] R. Langner, T. Horak, and R. Dachselt. Vistiles: Coordinating and combining co-located mobile devices for visual data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):626–636, 2017. doi: 10.1109/TVCG.2017.2744019

[49] B. Lee, P. Isenberg, N. H. Riche, and S. Carpendale. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2689–2698, 2012.

[50] B. Lee, A. Srinivasan, J. Stasko, M. Tory, and V. Setlur. Multimodal interaction for data visualization. In *Proceedings of the 2018 International Conference on Advanced Visual Interfaces*, pp. 1–3, 2018.

[51] 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*, 2020.

[52] Z. Liu and J. Stasko. Mental Models, Visual Reasoning and Interaction in Information Visualization: A Top-down Perspective. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):999–1008, 2010.

[53] T. Mahmood, W. Fulmer, N. Mungoli, J. Huang, and A. Lu. Improving information sharing and collaborative analysis for remote geospatial visualization using mixed reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 236–247. IEEE, 2019.

[54] A. McNutt. On the potential of zines as a medium for visualization. 08 2021.

[55] M. R. Morris, A. Danielescu, S. Drucker, D. Fisher, B. Lee, M. Schraefel, and J. O. Wobbrock. Reducing legacy bias in gesture elicitation studies. *interactions*, 21(3):40–45, 2014.

[56] J. M. Rzeszotarski and A. Kittur. Kinetica: Naturalistic multi-touch data visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 897–906, 2014.

[57] R. Sadana, T. Major, A. Dove, and J. Stasko. Onset: A visualization technique for large-scale binary set data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1993–2002, 2014.

[58] A. Samini and K. L. Palmerius. Popular performance metrics for evaluation of interaction in virtual and augmented reality. In *2017 International Conference on Cyberworlds (CW)*, pp. 206–209. IEEE, 2017.

[59] K. A. Satriadi, B. Ens, M. Cordeil, B. Jenny, T. Czauderna, and W. Willett. Augmented reality map navigation with freehand gestures. In *IEEE Conference on Virtual Reality and 3D User Interfaces*, pp. 593–603. IEEE, 2019.

[60] A. Satyanarayan, R. Russell, J. Hoffswell, and J. Heer. Reactive Vega: A Streaming Dataflow Architecture for Declarative Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):659–668, 2015.

[61] B. Shneiderman. Direct manipulation: A step beyond programming languages. In *Proceedings of the Joint Conference on Easier and More Productive Use of Computer Systems.(Part-II): Human Interface and the User Interface-Volume 1981*, p. 143, 1981.

[62] 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*, 25(1):715–725, 2018.

[63] J. Smiley, B. Lee, S. Tandon, M. Cordeil, L. Besançon, J. Knibbe, B. Jenny, and T. Dwyer. The made-axis: A modular actuated device to embody the axis of a data dimension. *Proceedings of the ACM on Human-Computer Interaction*, 5(ISS):1–23, 2021.

[64] M. Spindler, M. Schuessler, M. Martsch, and R. Dachselt. Pinch-drag-flick vs. spatial input: rethinking zoom & pan on mobile displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1113–1122, 2014.

[65] M. Spindler, C. Tominski, H. Schumann, and R. Dachselt. Tangible views for information visualization. In *ACM International Conference on Interactive Tabletops and Surfaces*, pp. 157–166, 2010.

[66] A. Srinivasan, B. Lee, N. Henry Riche, S. M. Drucker, and K. Hinckley. Inchorus: Designing consistent multimodal interactions for data visualization on tablet devices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020.

[67] S. Y. Ssin, J. A. Walsh, R. T. Smith, A. Cunningham, and B. H. Thomas. Geogate: Correlating geo-temporal datasets using an augmented reality space-time cube and tangible interactions. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 210–219. IEEE, 2019.

[68] S. Stoppel and S. Bruckner. Vol 2 velle: Printable interactive volume visualization. *IEEE transactions on visualization and computer graphics*, 23(1):861–870, 2016.

[69] P. Strohmeier, J. Burstyn, J. P. Carrascal, V. Levesque, and R. Vertegaal. Reflex: A flexible smartphone with active haptic feedback for bend input. In *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*, pp. 185–192, 2016.

[70] H. Subramonyam, E. Adar, and S. Drucker. Composites: A tangible interaction paradigm for visual data analysis in design practice. In *Proceedings of the 2022 International Conference on Advanced Visual Interfaces*, pp. 1–9, 2022.

[71] H. Subramonyam, S. M. Drucker, and E. Adar. Affinity lens: data-assisted affinity diagramming with augmented reality. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pp. 1–13, 2019.

[72] F. Taher, J. Hardy, A. Karnik, C. Weichel, Y. Jansen, K. Hornbæk, and J. Alexander. Exploring interactions with physically dynamic bar charts. In *Proceedings of the 33rd annual acm conference on human factors in computing systems*, pp. 3237–3246, 2015.

[73] L. Terrenghi, D. Kirk, A. Sellen, and S. Izadi. Affordances for manipulation of physical versus digital media on interactive surfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 1157–1166, 2007.

[74] C. Tominski, C. Forsell, and J. Johansson. Interaction support for visual comparison inspired by natural behavior. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2719–2728, 2012.

[75] C. Tominski, S. Gladisch, U. Kister, R. Dachselt, and H. Schumann. A survey on interactive lenses in visualization. In *EuroVis*, 2014.

[76] A. Wu, W. Tong, H. Li, D. Moritz, Y. Wang, and H. Qu. Computableviz: Mathematical operators as a formalism for visualisation processing and analysis. In *CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2022.

[77] A. Wu, Y. Wang, X. Shu, D. Moritz, W. Cui, H. Zhang, D. Zhang, and H. Qu. Ai4vis: Survey on artificial intelligence approaches for data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2021.

[78] R. Xiao, J. Schwarz, N. Throm, A. D. Wilson, and H. Benko. MRTouch: Adding Touch Input to Head-mounted Mixed Reality. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1653–1660, 2018.

[79] S. Ye, C. Zhu-Tian, X. Chu, Y. Wang, S. Fu, L. Shen, K. Zhou, and Y. Wu. Shuttlespace: Exploring and analyzing movement trajectory in immersive visualization. *IEEE Trans. Vis. Comput. Graph.*, 27(2):860–869, 2021. doi: 10.1109/TVCG.2020.3030392

[80] J. S. Yi, Y. ah Kang, J. Stasko, and J. A. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, 2007.

[81] Z. Zhao, R. Marr, and N. Elmqvist. Data comics: Sequential art for data-driven storytelling. *tech. report*, 2015.

[82] C. Zhu-Tian, Y. Su, Y. Wang, Q. Wang, H. Qu, and Y. Wu. Marvist: Authoring glyph-based visualization in mobile augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(8):2645–2658, 2020. doi: 10.1109/TVCG.2019.2892415

[83] C. Zhu-Tian, W. Tong, Q. Wang, B. Bach, and H. Qu. Augmenting static visualizations with paparvis designer. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2020.

[84] C. Zhu-Tian, S. Ye, X. Chu, H. Xia, H. Zhang, H. Qu, and Y. Wu. Augmenting sports videos with viscommentator. *IEEE Trans. Vis. Comput. Graph.*, 28(1):824–834, 2022. doi: 10.1109/TVCG.2021.3114806