

Pen-based Interaction with Spreadsheets in Mobile Virtual Reality

Travis Gesslein¹ Verena Biener¹ Philipp Gagel¹ Daniel Schneider¹ Per Ola Kristensson³
 Eyal Ofek² Michel Pahud² Jens Grubert^{1*}

¹Coburg University of Applied Sciences and Arts ²Microsoft Research
³University of Cambridge




Figure 1: This paper investigates the joint interaction space between an immersive virtual reality headset, and a spatially tracked pen for interacting with spreadsheets on a mobile tablet (a). This approach allows spreadsheet data to be extended around the screen and, for example, enables the user to interact with adjacent spreadsheets (b). It also allows additional information to be provided above the screen and, for example, enables the system to unravel hidden dependencies between cells by directly visualizing the cell relationships in 3D space (c).

ABSTRACT

Virtual Reality (VR) can enhance the display and interaction of mobile knowledge work and in particular, spreadsheet applications. While spreadsheets are widely used yet are challenging to interact with, especially on mobile devices, using them in VR has not been explored in depth. A special uniqueness of the domain is the contrast between the immersive and large display space afforded by VR, contrasted by the very limited interaction space that may be afforded for the information worker on the go, such as an airplane seat or a small work-space. To close this gap, we present a tool-set for enhancing spreadsheet interaction on tablets using immersive VR headsets and pen-based input. This combination opens up many possibilities for enhancing the productivity for spreadsheet interaction. We propose to use the space around and in front of the tablet for enhanced visualization of spreadsheet data and meta-data. For example, extending sheet display beyond the bounds of the physical screen, or easier debugging by uncovering hidden dependencies between sheet's cells. Combining the precise on-screen input of a pen with spatial sensing around the tablet, we propose tools for the efficient creation and editing of spreadsheets functions such as off-the-screen layered menus, visualization of sheets dependencies, and gaze-and-touch-based switching between spreadsheet tabs. We study the feasibility of the proposed tool-set using a video-based online survey and an expert-based assessment of indicative human performance potential.

1 INTRODUCTION

Spreadsheets are widely used data modeling, manipulation, and storage tools that are used in many application areas [5]. Interestingly, the basics of spreadsheet interaction have largely remained unchanged over the last 30 years. Spreadsheet popularity stems from their ease-to-learn, simple principles, and their flexibility [5]. Yet these characteristics are also the spreadsheets' limitations.

Unlike many other applications, spreadsheets are inherently open, and, due to their uniform grid nature, simple to learn for beginners.

*contact author: jens.grubert@hs-coburg.de

For example, navigating the sheet and copying cells are the same actions regardless of the sheet's content. However, the same uniformity also makes it hard to identify the structure of the spreadsheet or to debug complex dependencies in the spreadsheet since it does not allow a display of any type of information that does not conform to the grid. Previous research efforts have tried to visualize the hierarchy of the connection structure using 3D renderings [32, 74], yet any stray from the regular grid display—the user's main workplace—has not managed to achieve wide user adoption.

The spreadsheet also acts as an unconstrained canvas for bringing diverse ideas to life. It might become a simple to-do list, a calendar, a financial model, or a 500,000 cell encompassing inter-disciplinary project. The unconstrained canvas means that no matter how big the project, there is enough sheet space to accommodate it. This flexibility, while it is the spreadsheet's biggest asset, frequently proves to be it's Achilles' heel as well. Users often find themselves in need to scroll again and again along with spreadsheets much larger than the size of their screens. Large chunks of data may be needed to be selected, and, to be assigned as an input to functions multiple times. A possible mitigation strategy is using a larger screen that can show more data as the size of a mobile display proves to be too small. An overview of the spreadsheet may also help to access data outside the screen. However, currently, all available screen space is used for displaying as much of the spreadsheet as possible.

Further, this interaction paradigm—selecting cells, entering data by symbolic input, navigating around an unconstrained canvas—has been carried over from traditional desktop-based environments, which use a physical keyboard and mouse as standard input devices, to mobile settings, where users operate smaller-sized touchscreen devices, such as tablets or smartphones via touch and pen-based input [24]. Both the small screen space and the data entry methods using touch lead to a number of usability problems, first and foremost increased error rates when interacting with spreadsheet data [23]. Still, there is a strong need for accessing spreadsheet data when away from traditional desktop-based computing devices [23].

We see potential in facilitating interaction with spreadsheets in mobile settings using Virtual Reality (VR). First, VR head-mounted displays (HMDs) can help users visualize a much larger display area around them. As a handheld tablet may have the angular field of view of approximately 30 degrees, the accumulative horizontal field of view of VR HMDs can span 100-180 degrees and more.




Figure 2: Concept for utilizing the space around the main spreadsheet (light blue) which spatially coincides with the tablet: The sheet area can be extended (dark blue, dotted lines) beyond the tablet bounds. Additional sheets can be visualized on both sides of the active sheet (purple). Additional data layers can be placed at corresponding (x,y)-positions facilitating faster data association. Behind the sheet additional windows (yellow) can show an overview visualization.

While users may still use the familiar regular grid as their main tool of interaction, this additional display volume, in contrast to the mobile’s small and flat screen, can also reveal information outside the screen. It allows the system to visualize additional data the user can interact with, around and above the main spreadsheet, see Figure 1. Underpinned by the spatial sensing capabilities of modern pen-operated tablets and hover sensing of the pen [20,27], we set out to explore the joint interaction space of VR HMDs and pen-based inputs for interacting with spreadsheets on mobile devices, such as tablets.

In this work, we are not trying to inherently change the nature of spreadsheets. Instead, our objective is to show that embedding spreadsheet interaction within a 3D space can help in exposing the internal structure and allow better interaction possibilities. While there are many aspects and applications that can be analyzed, we chose a number of selected techniques that demonstrate the advantages of using VR for the mobile knowledge worker. While we extend the display of information in front, behind, and around the tablet screen, we do keep the interaction mostly in the limited area of the tablet. The 2D nature of the input device is complemented by pointing and editing information on or near the screen surface by hands that are rested on the tablet or an underlying table, simply by tilting the wrist to raise the pen.

In this paper, we make two contributions: 1) we design and implement techniques that combine VR and pen-based input on and above tablets for efficient interaction with spreadsheets, and 2) we validate those techniques in two indicative studies.

2 RELATED WORK

Our work draws upon the rich resources in spreadsheet interaction, pen-based and in-air interaction, context-menus and VR for knowledge workers.

2.1 Spreadsheet Interaction

According to Burnett et al. [8], the spreadsheet is probably the most popular programming paradigm in use, even though it presents several limitations and challenges to users. Prior work has identified challenges when interacting with spreadsheet software. For example, Mack et al. [55] collected complaints from Reddit and characterized issues users were facing. They identified challenges about tasks, such as importing, managing, querying, and presenting data. Smith et al. [75] investigated both individual and organizational challenges in using spreadsheet software, and, amongst others, identified data pipeline challenges related to importing data. Chambers et al. [11] describes challenges identified by spreadsheet users in a field study and proposed new functions including developing different modes for spreadsheet creation, improving support for spreadsheet reuse, and helping users to find and use features. Flood et al. [21] identified navigation as an issue that affects the performance of users debugging spreadsheets. Birch et al. [5] described success factors of current spreadsheet technologies but also challenges such as hidden errors, comprehensibility, and complexity. Specifically, regarding

comprehensibility Birch et al. [5] note that “One underlying reason for this high error rate is known to be the user’s difficulty in understanding the spreadsheet they are interacting with [46, 64]”, that is, that formulas are hidden by default and that the wider hidden structure of a spreadsheet model is not visible, even if formulas (and their first-order dependent cells) are shown.

Further, several researchers have investigated novel interaction methods and models for spreadsheet use. For example, Miller et al. [62] demonstrate novel features that enable the gradual structuring of spreadsheets based on design patterns of expert spreadsheet modelers. Jones et al. [42] applied Cognitive Dimensions of Notations [6] to their proposal for adding user-defined functions, also called sheet-defined functions, to spreadsheets. Janach et al. [40] adopt techniques from model-based diagnosis to spreadsheet debugging. Kandogan et al. [43] present a spreadsheet-based environment with a task-specific system-administration language for quickly creating small tools or migrating existing scripts to run as web portlets.

Regarding the use of mobile spreadsheet applications, Flood et al. [23] identified a strong need for accessing spreadsheet data when away from traditional desktop-based computing devices. Chintapalli et al. [14] compared four mobile spreadsheet applications according to the following usability criteria: visibility, navigation, scrolling, feedback, interaction, satisfaction, simplicity, and convenience. They found few differences among those applications but identified visibility, navigation, and feedback challenges. In contrast, a systematic review of mobile spreadsheet applications [22] revealed substantial differences in available functions (such as the ability to sort data or to keep headings visible while scrolling) between mobile spreadsheet applications. Flood et al. [24] also identified further challenges mobile users faced when using spreadsheet applications on smartphones, such as inaccurate cell and character selection, unintended actions, and unexpected behaviors.

Based on the insights for both the need for interacting with spreadsheets in mobile settings but also its challenges due to limited input and output capabilities [23], we investigate how the joint interaction space between HMDs and pen-based input can support interaction with spreadsheets on tablets.

2.2 Mixed Reality for Knowledge Worker Tasks

The use of Mixed Reality (MR) for supporting knowledge work has attracted recent research interest [28,30,71]. While early work investigated projection systems to extend physical office environments (e.g., [45, 67, 70, 89]), more recently, VR and AR HMDs have been investigated as tools for assisting users with interacting with physical documents (e.g., [26, 52]), focusing on annotating documents displayed on a 2D surfaces. Grubert et al. [28] and McGill et al. [61] explored the positive and negative qualities that VR introduces in mobile scenarios on the go. Other works has investigated VR use in desktop-based environments for tasks such as text entry (e.g., [29, 44, 60]), system control [93, 94] and visual analytics [9, 87]. Research on productivity-oriented desktop-based VR has concentrated on the use of physical keyboards [72], controllers and hands [48, 94], and, recently, tablets [78]. Concurrently with our work, Biener et al. [3] investigated the joint interaction space of VR HMDs and tablets for a variety of knowledge worker tasks.

We complement this prior work by investigating a specific information worker application – spreadsheets using tablets – in detail.

2.3 Pen-based and In-Air Interaction

Besides the commonly used single-point input with pens, enhanced interaction techniques have been explored. Examples include using touch input on the non-dominant hand, supporting pen input in bimanual interaction (e.g., [7, 37, 59, 65]), unimodal surface-based pen-postures [10], bending [19] or using sensors in or around the pen [31, 34, 39, 54, 58, 83] for gestures and postures, and examining pen-grips (e.g., [36, 76, 79]). Our work was inspired by tilting [84]



Figure 3: Interacting with multiple sheets using: initially, solely icons indicating the existent of additionally accessible sheets are visible (a). Neighboring sheets are expanded and each sheet the user gazes at is highlighted with a red frame (b). The user taps with his non-dominant hand on the tablet bezel, causing the selected sheet to slide towards the tablet (c), where the user can edit it using the tablet's touchscreen.




Figure 4: Neighboring sheets can also be accessed via in-air pointing. Here, the user extends the contribution to a function he defines, adding a cell of a neighboring sheet.

and hovering [20, 27] the pen above interactive surfaces, which we use in a VR context.

The use of pens in AR and VR has also been investigated as a standard input device on physical props [81, 82], as well as using grip-specific gestures for mid-air interaction [51]. The accuracy of pen-based mid-air pointing has also been studied [2, 66].

Regarding prior work on combining in-air with touch interaction, Marquardt et al. [57] investigated the use of on and above surface input on a tabletop. Chen et al. [13] explored in-air use of on and above surface input on a tabletop. They propose that interactions can be composed by interweaving in-air gestures before, between, and after touch on a prototype smartphone augmented with hover sensing. Hilliges et al. [33] have been using hover to allow more intuitive interaction with virtual objects that represent physical objects. More recently, Hinckley et al. [35] have been exploring a pre-touch modality on a smartphone including the approach trajectories of fingers to distinguish between different operations. Such technology can be used to connect 3D tracking and touchscreen digitizer for better accuracy of tracking.

Most VR in-air interaction typically aims at using unsupported hands. To enable reliable selection, targets are designed to be sufficiently large and spaced apart [77]. Our focus on mobile knowledge workers on the move dictates small gestures to reduce working fatigue and to retain operationalizability in potentially cramped environments, such as airplane seats. We design gestures to be used by a hand, resting on the screen of a tablet and holding a pen. The pen, or stylus, is a tool that is designed for writing, but also allow precise operation and selection [38] and has buttons to trigger actions, in addition to enabling using handwriting recognition in cells in future extensions of this work. Pen gestures use the fine finger motions to enable fine selection among selections on the screen-surface or above it, as the pen can be tilted up and down to select between layers of menus in the vertical direction to enable new interactions. For example, when using in-place 2D menus, e.g. pie menus, where each selection opens another sub-menu, returning to a parent menu may require a designated gesture. In contrast, in a 3D space above the tablet screen, the user can simply tilt the pen toward a lower menu to re-select it.

2.4 Context Menus, In-place Commands and Bimanual Interaction

In-place or at-hand commands appearing next to the user on demand have the benefits of avoiding a trip to a fixed menu on the screen. Bier et al. [4] explored the benefits of toolglasses and magic lenses manipulated indirectly using a mouse on the preferred-hand, and a trackball/thumb-wheel on the non-preferred-hand. In-place commands also work with direct manipulations on modern pen-and-touch displays by placing the menu near non-preferred-hand fingers and using the pen in the preferred-hand to consume the menu. These types of menus can be very useful on large displays where menus may be out of reach. Xia et al. [91] used an 84 inches Microsoft Surface Hub to explore, seamlessly select or frame content with the non-preferred-hand and manipulate it with the preferred-hand using marking menus [50] appearing in-context at the border of the selection area. Further, large displays' in-place commands can allow multiple users to work together side-by-side [88]. On-demand in-place commands are also useful for small displays, where a screen might have insufficient real-estate to display a fixed menu. A thumb menu may have a special on-demand menu for a hand holding a tablet or a smartphone [65]. Different menu types, such as floating palettes, marking menus, and toolglasses have been investigated for their strengths and weaknesses [56] and lately have been explored in VR [53]. In addition, Tian et al. [84] have been looking into using the orientation (tilt) of a pen to select menus while it's tip is positioned statically on the screen.

Our work complements such prior work by using 3D layered hierarchical pie-menu in close vicinity to a resting touch surface. It allows small hand gestures to easily select among multiple menus.

3 SPREADSHEET INTERACTION

Spreadsheet interaction has been mostly unchanged for decades: scrolling around a 2D grid using mouse or keyboard commands, interaction with cells via the keyboard, mostly on an edit-line outside the grid, and selecting and copying cells using the mouse.

In contrast, working with a pen on a tablet or on a horizontal interactive surface is an interaction style of a different nature (e.g., [7, 37, 59, 65]). To finely control the pen, the palm rests on the tablet screen. Moving the hand around the screen incurs a higher cost than sliding a mouse, resulting in a high incentive to bring the interaction and menus to the vicinity of the pen's reach (e.g., [92]). Using HMDs allows for positioning additional display real-estate near the hand without interfering with the 2D grid content by exploiting the 3D space around the hand.

In this work, we propose a set of tools for (1) enhancing the visibility of spreadsheet elements and meta data around the physical screen using VR; and (2) streamlining workflows for creating and editing spreadsheet functions using pen-based, multimodal interaction techniques. We designed the techniques using an iterative approach with multiple design iterations consisting of conceptualization, implementation and informal user tests ('eat your own dogfood' internal testing) [18, 86].




Figure 5: Extended view of a single sheet beyond the screen area (the semi-transparent green rectangular viewport) (a) By selecting a cell outside the viewport, the spreadsheet data may slide to be aligned with the touchscreen (b). Another option is to keep the viewport fixed in order to be able to reach to all sides of the touchscreen (c). Alternatively, the coordinate system of the interactive space can be tilted to be vertical, allowing a more comfortable display of a large sheet area in front of the user, while her physical hands are moving horizontally over the tablet (d).

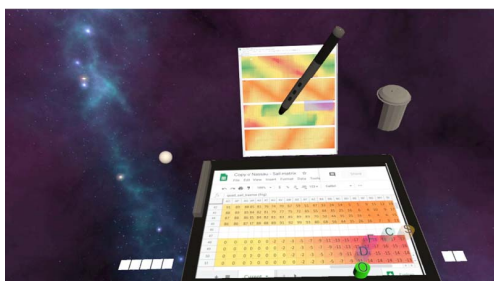


Figure 6: Overview visualization of the current sheet placed behind the active sheet.

3.1 Around Tablet Visualization

Our concept uses the 3D space, around, above, and behind the tablet, to display additional information, both extended views of the spreadsheet as well as in-place contextual information (Figure 2). While the area of the spreadsheet used for 2D interaction lies mostly on the tablet screen area, the extended accumulated field-of-view of the HMD enables extending the visible area of the grid (dark blue area in Figure 2), as well as displaying additional tabs (purple sheets floating next to the tablet in Figure 2). Such tabs can easily become replaced to be the edited tab on the tablet screen using a selection technique that combines eye-gaze and touch gestures. The space beyond the tablet can display additional information, such as a zoom-out overview of the spreadsheet, thereby enabling better discoverability of dependent elements outside the current field of view, and in addition provide fast navigation. Finally, the 3D space above the tablet's screen area is used to display multiple layers of contextual information, in-place, relative to the 2D spreadsheet.

3.1.1 Interaction with Multiple Tabs

Widely used spreadsheet software enables an arrangement of a large spreadsheet into separate tabs, accessible at the bottom of the spreadsheet window. However, as only one tab may be visible at a time, it may hamper visual reference and linking of data across tabs. Using the large field of view of immersive VR HMDs, we propose to display multiple tabs, extending them to the side of the tablet (Figure 3, b), allowing for an easy association of data between neighboring tabs.

As only the current tab aligns with the physical tablet screen, we support two techniques for interacting with tabs. A combined gaze-and-touch interaction, using head-gaze, lets the user look at any tab, which is highlighted when the ray representing the head direction hits this tab, (the red frame shown in Figure 3, b) and tapping a non-dominant-hand finger on the touchscreen bezel will slide the tabs until the selected tab is aligned with the touchscreen

(Figure 3, c, and d). Alternatively, neighboring sheets can directly be accessed by *in-air* pointing with the pen (Figure 4).

Both techniques can be seamlessly combined. While gaze-and-touch enables access of out-of-reach tabs and support of the tablet when editing, it requires time for tab selection. Furthermore, frequent switching of sheets may increase simulator-sickness risk due to frequent head rotations. In-air interaction provides faster access to neighboring tabs within the user's arm's reach, but has less input fidelity due to the lack of screen support and sensing (which can be somewhat mitigated using a table as a resting surface.)

3.1.2 Extending a Single Sheet

In typical spreadsheet applications, accessing cells outside of the view-area of the display's viewport requires scrolling the sheet to reveal new unseen areas. A VR extended display space enables us to show a larger virtual screen (Figure 5), however, only part of the display is supported by the tablet's screen sensing. To enable editing of the entire sheet, it is possible to select a different region of the sheet using gaze-and-touch to slide the region until it is aligned with the tablet area. Another option is to re-target the viewport, as well as the view of the virtual pen from their location aligned with the physical screen toward the area to be edited [29]. Visually, it appears to the user as if the tablet's screen moves from its original location. When using a very large displayed spreadsheet, it may be better to re-target it into a vertical display to avoid the sensation of the plane penetrating the user's body.

A complementary approach may use an overview visualization of the spreadsheet placed behind the tablet (Figure 6). Users can select a region of interest with their pen, sliding the sheet to align the selected location with the physical touchscreen.

3.2 On and Above Tablet Visualization

Enabling the display of additional information beyond the regular grid format can help to show the user important information, such as the flow of data along with the sheet, intermediate debugging results, and other meta information. We propose to use the space above the tablet surface, which is still reachable by the user's hand tilting the pen, to display different semantic information in proximity to the grid cell's content.

3.2.1 Visualizing Hidden Dependencies

Widely user adopted spreadsheet applications display functions' outcomes in a single cell, with no visualization of the cells that contributed to that calculation. While first-order dependencies (i.e., function parameters) may be visualized by coloring the cells when the user clicks on the result (function) cell, higher-level dependencies may not. Inspired by Shiozawa et al. [74]'s visualization, the VR user can lift up a cell to see visual links to the dependent cells (Figure 7, a), as well as higher-order dependencies. Complementary, composed nested functions can also be shown using a stack visualization




Figure 7: Cells are lifted up above their grid location and links visualize their dependency across two levels of computation (the black frame on the bottom show an option to cut the cells from the base layers) (a). While a cell of the grid displays the final result of the composed function, each layer may highlight the intermediate result of each added function (b). A cluster cell (bright green) is selected in a layer hovering above the main grid. The connected cells in the base layer are highlighted in red and also connected via red lines to the cluster cell (c) (shadows are used to better visualize the hovering layer in a monoscopic figure). A quick overview of used cells by masking the unused cells (d) (colored dark grey).

(Figure 7, b). Each layer represents one nesting function until the inner function is visible.

3.2.2 Cluster-cells

Applying a function to a set of values, such as calculating a sum of cells, the user has to select each of the input cells while defining the function. This operation can be carried out by dragging the mouse over each range of cells or by typing their addresses. When dealing with a large amount of data, this selection process can quickly become exhausting, in particular when there is a need to re-select ranges for additional functions, or update ranges (for example, to remove outlines).

Our tool set introduces a concept of a *cluster-cell* that represents groups of cells as a unique entity. While each cluster-node has a unique cell location in the original grid for compatibility with 2D interactions, whenever the user enables a 3D display by tilting the pen from the screen surface toward the approximate height of a layer, or by selecting a button, the cluster-cells are visualized in a layer hovering above the main grid, corresponding geographically to the regular grid (Figure 7, c). Each cluster cell may correspond to one or more cells of similar semantics, originating in lower levels, generating a hierarchy of cluster-cells displayed in multiple overlay levels. Cluster-cells enable easy re-use of selections for different uses (such as different functions, or reuse as different axes in a graph) as well as a means to simplify the visualization of the structure of a spreadsheet by displaying large ranges using a single node. As cluster cells are elevated to their levels, the links to all the original cells they represent are displayed to the user.

There are several ways to define a cluster-cell. A data-to-cell approach requires the user to select input cells from lower levels, either regular cells or preceding cluster cells. They may lie in a continuous range, selected with a single swap of the pen, or be a set of disconnected cells, selected while continuously pressing a pen button. Lifting the pen tip to a higher layer, or selecting from an in-place menu, will generate a cluster-cell and allow the user to label it. Additionally, the user may use a cell-to-data approach and select an existing cluster-cell and add additional input cells by dragging a link (while pressing the pen's button) down to a lower level and selecting one or more cells. Releasing the button attaches the selected cells to the cluster-cell. Input cells of choice may be removed from the cluster by dragging them to the trash bin widget (seen in Figure 8, h).

While displaying overlay layers, unused cells (cells in an overlay grid that have no assigned value) are displayed as transparent to increase visibility of lower layers. It is also possible to color unused cells of the original layer, to allow for a quick overview of the sheet content (Figure 7, d).

3.3 Pen-based on and above Surface Interaction

Prior to presenting our combined in-air and touch techniques, let us consider function creation and editing using standard 2D spreadsheet software. Defining a function in applications such as Google Sheets (Figure 8, bottom row) requires selecting a target cell for the function result (A4) and then specifying a function from either a menu or by typing the function name (Figure 8 i). The source cells are selected using a combination of pointing with the dominant-hand and simultaneously pressing a modifier-key on a physical keyboard (typically CTRL) with the non-dominant hand or by typing a list of cell ranges. When using finger touch or pen only, for example, when there is no keyboard or while holding the tablet, the duration for conducting this procedure becomes substantially longer (see our performance indication in Section 5.2). Specifically, selecting disjoint ranges of cells requires (potentially multiple) switches between selecting source cells, and entering delimiter signs (commas) on a virtual keyboard. Any editing of parameters of a function requires text editing, which is sometimes nontrivial (Figure 8 h). For example, removing a selected cell from a range of cells, (B2 from the range B1:B3), the range needs to be removed and individual cells need to be added again. Furthermore, re-using ranges for multiple uses, such as different functions, usage in graphs, and so on, requires re-selecting, or retyping the ranges.

Based on these observations, we designed a modified workflow combining touch and in-air interaction, which makes use of contextual hierarchical pie menus, visualizing and editing dependency links between cells (using a *telephone-operator metaphor*), and the use of cluster-cells.

3.3.1 Menu Interaction

There has been considerable effort in designing graphical menus [1, 15, 69]. Prior work has investigated the design and use of marking menus in VR using various input modalities such as a phantom [47], fingers [49, 53], hands [16], gaze [68] or controllers with six degrees of freedom input [25, 63] and techniques such as selection through ray-casting [41], crossing [85] or other gestures [90]. This prior body of work envisions the user pointing toward menus that are floating in front of the user in mid-air, spreading options apart, avoiding occlusions, such as overlapping menus, in order to enable easy mid-air selections. This paper, on the other hand, envisions the user working with hands supported by the tablet screen in a very similar fashion to 2D usage.

Inspired by Gebhardt et al. [25], we designed an in-place hierarchical pie-menu that is operated using pen-based interaction while the hand is supported by the surface of the tablet and tilts the pen, lifting its tip to reach menus hovering above the tablet (Figure 9 a-c). While most of the motions are done by the fingers holding the pen, the hand remains static, reducing fatigue. A menu is invoked in-place using button press on the pen, enabling the user to create a

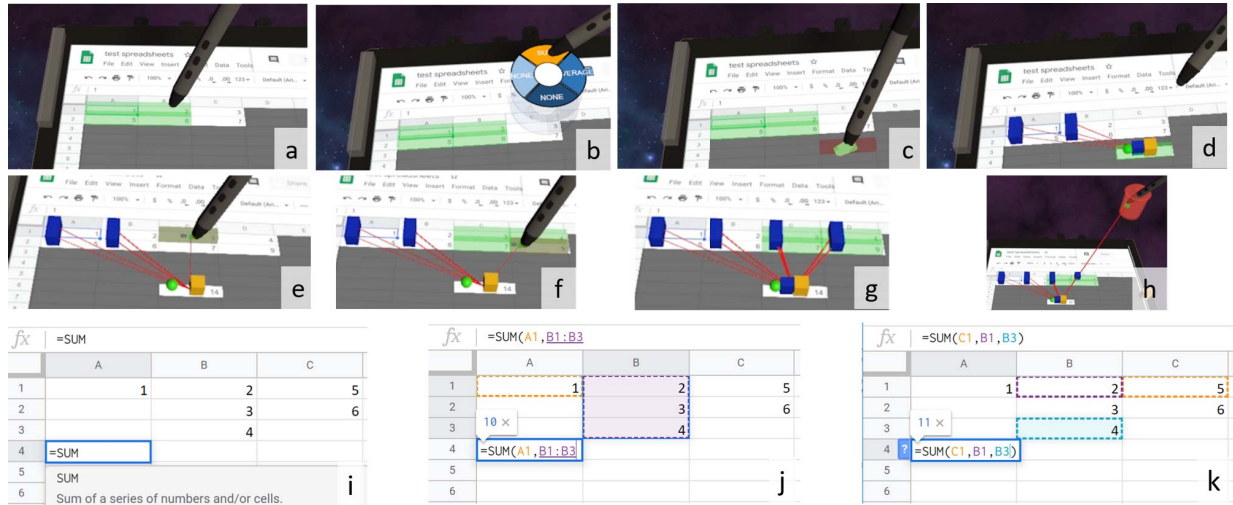


Figure 8: Top row: Adding a function workflow: selecting source cells (a), choosing a function from an in-place pie menu (b), selecting a cell for the function (c). Visualizing a function with links to its source cells (d). The function may be re-positioned by moving the yellow button. Middle row: Adding and removing a function's source cells. Adding cells can be done by selecting with the pen (e), dragging to select a range of cells (f) and lifting the pen and drawing a red link to the function node (g) (inspired by the operation of a legacy telephone operator). Each source node displays a blue box, which can be dragged toward the garbage bin widget to dissociate them from the function (h). Bottom row: Example of creating a function workflow in common spreadsheet application. A target cell (A4) is selected (i), and a function is specified. Source cells are specified as function parameters individually (A1) or as a range (B1:B3) (j). Updating source cells using text entry (k).

function or chart and place it in a target cell in a continuous motion (Figure 9 a–c). The user can confirm the final menu entry by button press. Also, the user may retract a menu by simply lowering the pen tip.

An additional arc-menu, on the lower right corner of the tablet, controls the display of meta-features. The arc shape of the menu allows for easy access to entries while the wrist is supported by the tablet, or by a thumb of a hand holding the tablet.

3.3.2 Adding and updating functions

While our prototype supports the 2D adding functions workflow, it also enables a new workflow that uses contextual menus: 1) select a target cell for the function; 2) raise the pen and select a function in the hovering hierarchical pie-menu; and 3) add source cells. Besides the above *function-to-source* workflow, the system also supports the inverse order *source-to-function*: First selecting source-cells, then selecting a function from a menu, and finally storing the function to a target cell in the grid (Figure 8, a–d). Source-to-function enables visualizing results of the function or, for example, a chart content while positioning it in its grid location. Adding source cells to existing functions can be done in a similar way, by first selecting them and then drawing a link to the function node. (Figure 8, e).

To avoid visual clutter, visualization can be toggled individually for each function, see Figure 9, bottom row.

3.3.3 Interacting with Charts

The workflow for the creation of a chart (e.g., a bar chart) within a spreadsheet is very similar to creating a function cell. The user uses a pen to select source cells, then tilt the pen up to select a CHART option from the in-place menu, and completes the motion by selecting the location of the top left corner of the chart on the grid, dragging the pen to define the size of the chart. In this *source-to-chart* flow it is possible to view the chart being rendered while the pen defines this area. The entire workflow can be carried out

with a single fluid motion of the pen. The created chart can be easily modified or resized afterward.

Furthermore, the user can modify the chart with in-air gestures such as adding a trend-line, by an in-air stroke of the pen tip over the chart (Figure 9, e). Other regression models could be added, for example, a polynomial fit can be coupled with a curved in-air gesture.

4 IMPLEMENTATION

The described techniques were implemented in a VE using a Unity game engine¹, commodity PC hardware, and an HTC Vive Pro HMD. The system includes a fully interactive spreadsheet implementation based on Google Sheets². A separate Microsoft Surface Pro 4 tablet was used to sense pen input data including touches and button press events and sent them to the Unity application via UDP network communication. Both the tablet and the pen were spatially tracked via an OptiTrack motion-tracking system, and were rendered in the Unity virtual environment by similar sized 3D models.

Google Sheets website pages are rendered inside Unity using embedded instances of a Chromium browser engine, provided by the ZFBrowser Embedded Browser Unity plugin³. ZFBrowser renders the contents of opened web pages into a texture with the same resolution as the physical Surface Pro 4 to achieve equivalent scaling of rendered HTML elements between the physical and virtual notebook. This texture is then mapped onto the screen of the virtual notebook, and touch interaction coordinates received via UDP from the notebook can be mapped onto this texture in a one-to-one fashion. Optionally, input coordinates can be normalized to a [0,1] range in x and y directions by dividing by the physical screen width and height to remap the input onto output surfaces of arbitrary dimensions.

¹<https://unity.com/> Last access May 18, 2020

²<https://www.google.com/sheets/about/> Last access May 18, 2020

³<https://assetstore.unity.com/packages/tools/gui/embedded-browser-55459> Last access May 18, 2020

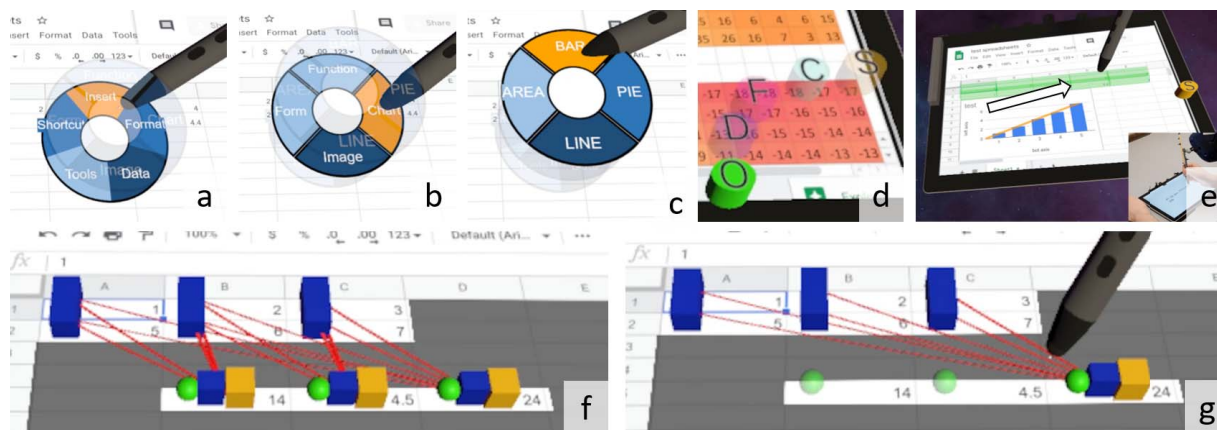


Figure 9: Top row: the hierarchical pie menu is invoked in-place and navigated from bottom to top through in-air movements of the pen (a-c). An arc menu in the corner of the tablet controls visualization options (activated - solid color, non activated - semi-transparent) (d). Adding a trend-line to a chart through an in-air stroke gesture (e). Bottom row: Three functions with their respective links to their source cells highlighted (f). Green spherical buttons toggle displays of links per function. Here only the center function links display is enabled (g). Abbreviations: O: overview window, D: dependency links, F: nested functions, C: cluster-nodes. S: additional sheets (tabs).

While interacting with the sheets web interface, there is no way to define or change functions and charts in one step, so we used the Google Sheets cloud API to apply cell transformations. Since the Google Sheets API exposes no functionality to track client-side interactions, operations of users with the web page were tracked inside Unity. In particular, tracking of cell selection was implemented by constructing virtual-cells in the Unity space using oriented bounding boxes, and spatially position them in their corresponding places to fit the spreadsheet texture. Tracking of the pen and the Unity collision detection mechanism is used to detect whether the pen tip lies inside a certain cell⁴.

Enabling a display of only used cells in overlay layers was implemented using a custom alpha masking technique. First, to identify empty cells we use the Google Sheets API to retrieve the cell values for each cell in a displayed spreadsheet. Then, a virtual camera renders cells that should appear transparent into a separate monochrome texture, called an alpha mask. The cells are rendered into this texture such that the alpha mask is registered with the rendered browser texture. A custom shader uses the corresponding alpha mask values and rendered browser pixels transparent accordingly.

To facilitate further research and development of spreadsheet interaction within VR, our code is available under <https://gitlab.com/mixedrealitylab/spreadsheetr>.

5 EVALUATION

We validated our prototype through an online survey and by gathering performance data from expert users. Both evaluations are described next.

5.1 Video-based Survey

While we used informal user tests throughout the iterative design process of the described techniques, we gathered further user feedback on the potential usefulness and attractiveness of the techniques through a video-based online survey.

5.1.1 Design and Procedure

We ran an online experiment using a within-subjects design, in which the individual techniques were presented to users as video prototypes.

⁴<https://docs.unity3d.com/ScriptReference/Collider.OnTriggerEnter.html>
Last access May 18, 2020

Please note that while users were not able to try the techniques out for themselves as interactive prototypes, collecting user feedback based on video prototypes is an established practice (e.g., [12, 17, 80]). The videos showed the following twelve techniques: creating functions (marked as CF), manipulating functions (adding, removing data cells) (MF), creating cluster-nodes (CC), selection across multiple tabs using gaze-and-touch interaction (SE), selection across multiple tabs using mid-air interaction (SA), Display of an overview window of the spreadsheet (OV), extended view of a spreadsheet beyond the tablet's view-port (EV), visualization of cell dependencies (DV), visualization of nested functions (NF), masking unused cells, (MC) chart interaction (CT) as well as a close-up of the hierarchical pie-menu interaction (PM) (which was also used as integral part of other techniques).

For each presented technique, users were asked to rate its usefulness, how easy-to-use it looks, if they would recommend it to a friend and how they would rate the technique overall. Further, participants were free to add further comments on each technique.

Please note, that the presentation of the techniques was not counterbalanced. While this could lead to ordering effects, some techniques were based on using prior presented techniques (e.g. adding or manipulating functions). With counterbalancing in place, this could have limited participants' comprehension of the respective techniques.

After all techniques were presented, the participants were asked to select one technique that they liked best overall, as well as one technique, that they liked least overall. They were then asked to state reasons for their respective choices.

Overall, the evaluation took about 30 minutes per participant. No compensation was paid to participants.

5.1.2 Participants

We recruited 18 participants (3 female, 15 male, mean age 31 years, $sd = 8.39$). Four participants reported a very high level of experience with virtual reality, eight stated a high experience level, four said that they were moderately experienced and two indicated little experience. Three participants indicated a very high level of experience with pen-based systems, five stated a high experience level, four reported moderate experience, three indicated little experience and three had no experience at all. Six participants reported a very high

level of experience with spreadsheet applications, six reported a high level of experience, five stated they were moderately experienced and one reported little experience. Three participants reported that they very frequently use spreadsheet applications on mobile devices. Two stated that they used them frequently, three participants use them sometimes, seven rarely and three never.

5.1.3 Results

The results for the user ratings are depicted in Figure 10. Friedman tests revealed significant differences, between conditions, see Table 1. However, post-hoc tests using Wilcoxon signed rank tests with Bonferroni correction, did not reveal significant pairwise differences.

We also asked users to identify their most preferred and least preferred function, the results are depicted in Figure 11.

Besides, rating the techniques, we asked users to comment on the individual techniques.

creating functions: Participant 7 (P7) stated that "Looks like it would be fun to work with". as well as "working with symbols [boxes and links] is much quicker than working with large numbers or functions." P14 stated that "allows to see where the values are coming from." While users stated that selecting functions from pie-menus are "pretty straightforward" to use, this method "may not work for the many functions not on the menu" (P8). P16 was concerned about "the high degree of motor control that the technique seems to require" and P17 noted that 3D cells "are potentially distracting and may occlude important content (in contrast to the more typical use of border colours to indicate a cell range)"

manipulating functions: P8 mentioned that adding source cells and toggling the links "appears to be easy". P14 stated that buttons on the pen "would be better for activating the different modalities." P4 mentioned that this technique "seems useful and faster than the "normal" approach in excel." P13 stated that the technique is "easy to understand. shows links. could be used in normal excel." Three users stated that deleting is "too complicated" due to added dragging toward the trash can." P14 suggested a possible solution "to show the bin closer to the hand once a cell is activated". Three participants proposed to using additional pen buttons to overload functions (e.g., switching between adding and deleting of cells).

cluster cells: P8 mentioned "the technique is useful. I had not seen aliases for cell ranges before", that the motion "seems easy", but also wondered if it would be "frequently triggered accidentally". P14 said that the technique is "really good" and that "creating on the fly shortcuts or templates would reduce a lot of time" as it's "very common in excel to do the same operations for lots of combinations of columns, where the same function is copied and modified over and over." P9 mentioned "It makes it easy to group data without having to use multiple sheets and thus reduces complexity."

interaction across multiple sheets using gaze-and-touch: P7 stated the technique seems "very easy" as well as "I can get a better overview of my different worksheets and select them very quickly. This technique would be a great relief in my daily work." P10 stated "Great way to visualize the connection to other sheets and I assume it works better with a larger number of sheets than the in-air technique and therefore I prefer it." P18 stated "It made something that is difficult to do today much easier."

On the contrary, three users commented that head motions seem unnecessary large. P8 suggested that a better layout of neighboring tabs may allow for smaller head motions.

interaction across multiple sheets using in-air interactions: P8 suggested to extend the technique to far away tabs, by "magnify[ing]" them once the user reaches out for them. P17 liked the ability to work with two sheets side-by-side. "Main issue here is that you are forced to operate at greater distance on second sheet which will negatively impact legibility and selection accuracy."

overview visualization: P14 thought it was "really good and make[s] more sense for navigating spreadsheets rather than scrolling". P17 called it "a great idea" in particular for data sets. P5 called the technique "very handy to navigate large data sheets" and P16 liked "the separation of views".

extended view beyond the tablet view-port: P16 liked "the spreadsheet breaks out of the boundaries" of the tablet and P17 called "leveraging [the] expanded display region and orientation freedom" a great idea. P16 also mentioned a possible way for improvement: "perhaps using the same approach to maps (where one zooms out when moving across regions) might make this technique better."

dependency visualization: P7 mentioned that the technique seems "very useful to check quickly the correctness" of spreadsheets calculations. P8 liked the idea of using the depth to separate out layers of information. P3 called it "extremely useful" for large sheets. P8 stated that "Tracking down cell dependencies in Excel is one of the banes of my existence. The technique is easy to use and presents information in an intuitive way. Also, I feel that it is suitable to VR and leverages the affordances of VR." P14 mentioned "This can be very useful to check and verify that certain functions are correctly set up."

masking unused cells, P9 thought masking unused cells "makes it easier to find important content".

chart interaction: P12 said this technique seems "easy and natural to use."

hierarchical in-place pie menus: P7 was reminded of "drawing with the essential fine motor skills." P17 thought it was "cool" but warns of a potential problem as "is that place in hierarchy and path to current menu is partially hidden." P12 mentioned the menu seems "quite confusing and hard to manipulate" and P16 was "not a fan of mid-air stacks as a selection mechanism".

5.2 Indicative Human Performance Potential

We had five expert users (four male, one female, mean age 28.4 years $sd = 5.13$), who trained to be proficient in executing common tasks in spreadsheet software.

The lab-based experiment had two independent variables. The first one was the interaction technique which was either VR, TABLET AND PEN or KEYBOARD AND TOUCHPAD. The second one was the task which included creating a function (CF), adding individual cells to a function (AIC), adding a range to a function (AR), remove single cells from a function (RC), reusing cells when creating a function (RE), adding a chart (AC) and adding a trendline (AT) and adding individual cells from another tab (AICS). The participants executed all of these tasks and repeated them five times with each interaction technique (see video in the supplementary material for footage of a selected expert user). We used counterbalancing, for the tasks as well as for the interaction techniques, to mitigate learning effects. Please note that we used simple tasks on purpose, as more complex tasks are typically composed of those atomic parts, even though this might be unfavorable for the VR techniques. For example, for reusing cells when creating a function, participants could reuse cells by copying a cell range initially and pasting it into several functions using TABLET AND PEN as well as KEYBOARD AND TOUCHPAD techniques. This process of copying once and pasting multiple times might not work efficiently in scenarios, in which several cell ranges and multiple functions are used. Here, copy and paste could

Table 1: Results of Friedman tests with regards to the four different categories (usefulness, ease of use, recommendation and overall rating)

	$\chi^2(11)$	p	W
Usefulness	26.0	.006	0.13
Ease of use	31.6	.0008	0.16
Recommendation	25.6	.007	0.13
Overall rating	26.2	.006	0.13

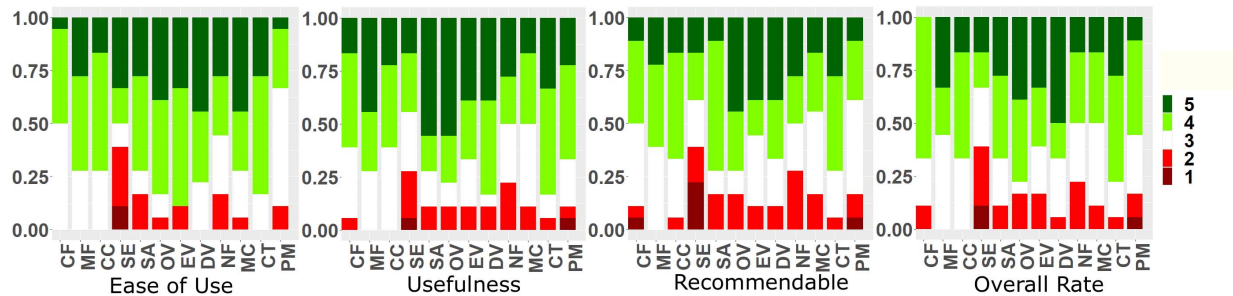


Figure 10: Rating overview on ease of use, usefulness, recommendability, and overall rating on a five-item Likert scale (1: strongly disagree, 5 strongly agree). Techniques: CF: creating functions, MF: manipulating functions, CC: clustering cells, SE: gaze+touch interaction across multiple tabs, SA: in-air interaction across multiple tabs, OV: overview visualization, EV: extended viewport, DV: dependency visualization, NF: visualization of nested functions, MC: masking of unused cells, CT: Adding a trendline to a chart, PM: hierarchical pie menu. The y-axis depicts the number of participants in percent (1.0 = 100% of participants).




Figure 11: Most and least preferred techniques as stated by the participants. Techniques: CF: creating functions, MF: manipulating functions, CC: clustering cells, SE: gaze+touch interaction across multiple sheets, SA: in-air interaction across multiple sheets, OV: overview visualization, EV: extended viewport, DV: dependency visualization, NF: visualization of nested functions, MC: masking of unused cells, CT: Adding a trendline to a chart, PM: hierarchical pie. The x-axis depicts the number of participants in percent (1.0 = 100% of participants) menu.

result in higher coordination efforts (i.e. to avoid copying the same range multiple times, a range first has to be copied in all functions, then the next range should be copied and pasted). Besides TABLET AND PEN we included KEYBOARD AND TOUCHPAD interaction as an additional reference point, when users would have access to physical keyboards, such as on notebooks or on selected slates with detachable keyboard. When using text entry in either TABLET AND PEN or KEYBOARD AND TOUCHPAD, closing brackets could be omitted to speed up text entry. Please also note, that in the VR condition, users did not need to explicitly enter text. We tried to design the best possible techniques in VR and indeed found that they do not require text entry for the tasks at hand.

The task completion times were computed as the duration between the first selection of a relevant cell (e.g., for entering a function or selecting source cells) until the result was displayed (e.g., the function result, a chart). This way, initial travel times from various potential starting points have not been included. Overall, the data collection took around 90 minutes per participant. No compensation was paid.

The results for task completion time are depicted in Table 2. Please note, that these times should solely indicate the performance that can be achieved with sufficient training. While we depict results of repeated measures analysis of variance (RM-ANOVA, data was tested for normality with the Shapiro-Wilk normality test) along with Bonferroni adjusted post-hoc tests (initial significance level $\alpha = .05$), the results are solely indicative, both due to the small sample size and the lack of integrating users with a more representative background.

Table 2: Task completion time in seconds (average, standard deviation in parenthesis below) for the different applications and the three interfaces virtual reality (VR), tablet and pen (TP) as well as keyboard and touchpad (KT). Tasks (T): creating a function (CF), adding individual cells to a function (AIC), adding a range to a function (AR), remove single cells from a function (RC), reusing cells when creating a function (RE), adding a chart (AC) and adding individual cells from another sheet (AIS). For AICS, both gaze+touch (G+T) as well as in-air selection (IA). Pairwise significant differences at an initial significance level of $\alpha = .05$ are indicated in the * column.

T	VR	TP	KT	$F_{(2,8)}$	p	η_p^2	*
CF	6.62 (1.30)	12.7 (3.62)	4.96 (1.49)	15.2	.002	.79	VR-TP KT-TP
AIC	2.70 (0.59)	6.70 (1.11)	3.07 (0.47)	29.5	< .001	.88	VR-TP KT-TP
AR	3.53 (0.79)	10.4 (1.95)	3.93 (0.81)	40.2	< .001	.91	VR-TP KT-TP
RC	2.21 (0.61)	12.3 (4.86)	6.75 (1.65)	18.3	.001	.82	VR-TP KT-TP
RE	5.71 (1.97)	10.6 (3.26)	3.06 (1.53)	11.0	.005	.73	VR-TP KT-TP
AC	9.03 (1.74)	7.17 (2.92)	4.93 (0.74)	3.97	.06	.50	
	IA	G+T		$F_{(3,12)}$			
AIS	3.94 (0.99)	6.77 (2.30)	8.39 (2.56)	5.72 (0.60)	4.22	.03	0.51 IA-TP

6 DISCUSSION

We presented a first step in investigating the use of VR to improve the user experience of spreadsheet work on the go. We use VR to extend the user display space and enable better visualization and interaction of spreadsheets while simultaneously maintaining a small interaction-space around the tablet. We also want to maximize users' familiarity with common 2D spreadsheet workflows and use the support of the tablet for the user's hands to enable long period of work with less fatigue.

This work has raised a set of principles, such as a separation between the input-space and the data/display-space, compatibility of representations between the VR and 2D workflows and limiting the input around familiar touchscreen devices. While we expect more interaction techniques to be suggested in the future, we also expect these principles to be maintained.

Our indicative evaluations revealed that the techniques, were mostly deemed usable and useful by participants in a video-based

survey. No technique was clearly preferred by the majority of users. However, the feedback users provided on individual functions, highlighted both opportunities for further improvement of the techniques as well as potential challenges when running a video-based evaluation. For example, gaze-and-touch may look confusing when observed in a video, as indicated by participant ratings, due to the fast apparent head motions. In contrast, when developing this technique, in our informal tests, we found the technique to be very comfortable when wearing an HMD and using it in-situ. Similarly, several users were concerned with in-air interaction with hierarchical pie-menus as well as with visual clutter. Again, when using 'dogfood' prototype in internal testing, we found the technique to be comfortable and efficient due to the support of a resting wrist on the table or tablet, with little rotations of the pen for menu selection (instead of lifting up the entire hand or arm). The transparency of the non selected layers, which may look cluttered, enables a glance at the next level of menus without covering the underlying opaque layer when seen through a stereo HMD.

The indicative human performance evaluation with expert users further revealed that the proposed techniques can be efficient to use. For several base functions such as creating functions, adding (individual or multiple) cells, removing or reusing cells, VR was significantly faster than tablet and pen interaction. However, please note that these evaluations should be seen as indicative. In future work, walk-up-and-use usability and performance should be tested with non-expert users in interactive sessions. Further, the performance of compound tasks (complementary to the atomic tasks used in the expert evaluation) could be further studied. Nonetheless, the expert evaluation provided evidence that the system can be usable and be used efficiently and the video-based evaluation revealed some preliminary evidence of wider user interest in this new approach to mobile spreadsheet interaction.

The choices of the experimental designs were influenced and limited by the COVID-19 restrictions put in place at our university. Therefore, the results of the studies have to be interpreted with caution and further studies should be conducted.

The current prototype was implemented in a lab, using an external OptiTrack device for tracking the pen and the tablet to achieve the best spatial tracking we can get. However, similar tracking capabilities are becoming accessible for mobile settings, even though the accuracy of current generation mobile tracking systems is substantially lower compared to dedicated outside-in tracking systems like OptiTrack [73]. The use of head mounted cameras on HMDs has become prevalent for both inside-out tracking, and for streaming outside video for video-based augmented reality applications. Using the video stream from those cameras, the tablet and the pen can be tracked (the tablet screen may display fiducials for pose estimation as it is not seen by the user). Finger tracking technologies are appearing in recent HMDs, such as Hololens 2.0, Oculus Quest, Vive Pro and others. This would also enable to study spreadsheet interactions in various mobile real-world scenarios.

Further, the examples in this paper were executed using a pen, as this input device combines several advantages: it enables fine accuracy of touch input, which can be used for text entry with hand writing. This is easy to do in-place and using one hand, while the non-dominant-hand holds the tablet. It is also easy to track the pen by adding a tracker at its back tip, and it has buttons that can be used to control operations, such as selecting multiple cells one after the other. However, our techniques can be extended to be used with other input modalities, such as a finger touch, where the non-dominant-hand touch may simulate a button, a soft keyboard is used for typing, and the raising index finger replaces pen tilting. If space allows, a mouse and a keyboard may also be used, where the mouse pointer is the location of the pen, and the mouse wheel or a touch sensitive area on the mouse may be used to tilt the pen upward.

In most of the figures of this paper, we rendered the display

space parallel to the surface of the tablet. Such a display is easy to understand, as it is clear where the tablet and the user's pen are. However, in many applications, where the users have to work in very limited places, such as an airplane seat, we envision users moving the display space to a vertical plane ahead of them (such as in Figure 5), while the hands are still supported on the tablet. Every movement of the pen or the hands is represented as a corresponding movement of the virtual pen or hands in the display space (as explored by Grubert et. al. [29]). Finally, the proposed interaction techniques could be explored in augmented reality (AR) using optical see-through (OST) HMDs. While the field of view of immersive VR HMDs is typically substantially larger compared to OST HMDs, most VR HMDs do not match the output resolution of a tablet's screen. Hence, in our work, the number of cells being concurrently legible was smaller than on a typical 4K tablet screen, potentially leading to an increased need for navigating between cells that are far apart. OST HMDs could potentially be leveraged to show additional information about spreadsheets while still allowing use of a high resolution screen.

7 CONCLUSIONS AND FUTURE WORK

Within this work, we presented a toolset for enhancing mobile spreadsheet interaction on tablets using immersive VR HMDs and pen-based input. We proposed to use the space around the tablet for enhanced visualization of spreadsheet data and presented a set of interaction techniques for around and above tablet visualization of spreadsheet-related information, for example, to allow previewing cells beyond the bounds of the physical touchscreen from the same or adjacent tabs or to uncover hidden dependencies between cells. Combining the input capabilities of precise on-surface, pen-based interaction along with spatial sensing around the tablet, we demonstrated techniques for efficient creation and editing of spreadsheets functions. Our indicative studies showed the feasibility of our approach. In future work, our initial explorations and evaluations should be extended to cover walk-up-and-use scenarios, more complex compound tasks. Further, we are interested in exploring how VR spreadsheet interactions could work in collaborative scenarios. For instance, when interacting with multiple sheets, being able to visualize the hands of other users working on different sheets for awareness or to facilitate collaboratively working with them. We are also interesting in exploring how users could edit or modify a spreadsheet's content in VR using the pen, such as using handwriting in the cells, or using or adjusting values in cells with real-time updates of the charts by using relative movements, for example, similar to Pfeuffer et. al [65] but in VR. This work has been focusing on pen and touch interactions, but there might be situations where users do not have a pen or prefer using touch only, such as in an airplane where they could be worried to lose the pen. To that end, we are interested in studying a variant of the spreadsheet VR experience with only touch and compare the usability and performance with pen and touch. We are also contemplating the possibility to extend our research to explore heads-up experiences by using pen and touch to manipulate spreadsheets indirectly while having the visuals situated in front of the user. Further, in future tasks that also require text entry in VR could be studied. Also, while for our performance study we reduced the number of concurrently visible cells in favor of a comparable field of view between the virtual and physical screen, one could also change the field of view of the virtual screen to match a higher number of visible cells typically concurrently visible on physical 4K screens. Finally, the use of augmented keyboards [72] for supporting spreadsheet interaction should be further investigated, for example, by re-purposing keys for navigating between data types (e.g., to jump to next numeric or character).

REFERENCES

- [1] G. Bailly, E. Lecolinet, and L. Nigay. Visual menu techniques. *ACM Computing Surveys (CSUR)*, 49(4):1–41, 2016.

- [2] A. U. Batmaz, A. K. Mutasim, and W. Stuerzlinger. Precision vs. power grip: A comparison of pen grip styles for selection in virtual reality.
- [3] V. Biener, D. Schneider, T. Gesslein, A. Otte, S. Kuth, P. O. Kristensson, E. Ofek, M. Pahud, and J. Grubert. Breaking the screen: Interaction across touchscreen boundaries in virtual reality for mobile knowledge workers. In *IEEE transactions on visualization and computer graphics*, 2020.
- [4] E. A. Bier, M. Stone, K. Pier, W. Buxton, and D. T. Toolglass and magic lenses: the see-through interface. pp. 73–80. ACM, 1993.
- [5] D. Birch, D. Lyford-Smith, and Y. Guo. The future of spreadsheets in the big data era. *arXiv preprint arXiv:1801.10231*, 2018.
- [6] A. Blackwell and T. Green. Notational systems—the cognitive dimensions of notations framework. *HCI models, theories, and frameworks: toward an interdisciplinary science*. Morgan Kaufmann, 2003.
- [7] P. Brandl, C. Forlines, D. Wigdor, M. Haller, and C. Shen. Combining and measuring the benefits of bimanual pen and direct-touch interaction on horizontal interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces, AVI '08*, p. 154–161. Association for Computing Machinery, New York, NY, USA, 2008. doi: 10.1145/1385569.1385595
- [8] M. Burnett, J. Atwood, R. W. Djang, J. Reichwein, H. Gottfried, and S. Yang. Forms/3: A first-order visual language to explore the boundaries of the spreadsheet paradigm. *Journal of functional programming*, 11(2):155–206, 2001.
- [9] 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.
- [10] D. Cami, F. Matulic, R. G. Calland, B. Vogel, and D. Vogel. Unimanual pen+ touch input using variations of precision grip postures. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pp. 825–837, 2018.
- [11] C. Chambers and C. Scaffidi. Struggling to excel: A field study of challenges faced by spreadsheet users. In *2010 IEEE Symposium on Visual Languages and Human-Centric Computing*, pp. 187–194. IEEE, 2010.
- [12] X. Chen, T. Grossman, D. J. Wigdor, and G. Fitzmaurice. Duet: exploring joint interactions on a smart phone and a smart watch. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 159–168, 2014.
- [13] X. Chen, J. Schwarz, C. Harrison, J. Mankoff, and S. E. Hudson. Air+ touch: interweaving touch & in-air gestures. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*, pp. 519–525, 2014.
- [14] V. V. Chintapalli, W. Tao, Z. Meng, K. Zhang, J. Kong, and Y. Ge. A comparative study of spreadsheet applications on mobile devices. *Mobile Information Systems*, 2016, 2016.
- [15] R. Dachselt and A. Hübner. Three-dimensional menus: A survey and taxonomy. *Computers & Graphics*, 31(1):53–65, 2007.
- [16] M. M. Davis, J. L. Gabbard, D. A. Bowman, and D. Gracanin. Depth-based 3d gesture multi-level radial menu for virtual object manipulation. In *2016 IEEE Virtual Reality (VR)*, pp. 169–170. IEEE, 2016.
- [17] B. Dhillon, P. Banach, R. Kocielnik, J. P. Empananza, I. Politis, A. Rzewaska, and P. Markopoulos. Visual fidelity of video prototypes and user feedback: a case study. In *Proceedings of HCI 2011 The 25th BCS Conference on Human Computer Interaction 25*, pp. 139–144, 2011.
- [18] A. Drachen, P. Mirza-Babaei, and L. E. Nacke. *Games user research*. Oxford University Press, 2018.
- [19] N. Fellion, T. Pietrzak, and A. Girouard. Flexstylus: Leveraging bend input for pen interaction. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology, UIST '17*, p. 375–385. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3126594.3126597
- [20] G. Fitzmaurice, A. Khan, R. Pieké, B. Buxton, and G. Kurtenbach. Tracking menus. In *UIST*, pp. 71–79. ACM Press, 2003.
- [21] D. Flood, K. M. Daid, F. M. Caffery, and B. Bishop. Evaluation of an intelligent assistive technology for voice navigation of spreadsheets. *arXiv preprint arXiv:0809.3571*, 2008.
- [22] D. Flood, R. Harrison, C. Martin, and K. McDaid. A systematic evaluation of mobile spreadsheet apps. In *IADIS International Conference Interfaces and Human Computer Interaction*, pp. 1–8, 2011.
- [23] D. Flood, R. Harrison, and K. McDaid. Spreadsheets on the move: An evaluation of mobile spreadsheets. *arXiv preprint arXiv:1112.4191*, 2011.
- [24] D. Flood, R. Harrison, and A. Nosseir. Useful but tedious: An evaluation of mobile spreadsheets. In *PPIG*, p. 6. Citeseer, 2011.
- [25] S. Gebhardt, S. Pick, F. Leithold, B. Hentschel, and T. Kuhlen. Extended pie menus for immersive virtual environments. *IEEE transactions on visualization and computer graphics*, 19(4):644–651, 2013.
- [26] R. Grasset, A. Duenser, H. Seichter, and M. Billinghurst. The mixed reality book: a new multimedia reading experience. In *CHI'07 extended abstracts on Human factors in computing systems*, pp. 1953–1958. ACM, 2007.
- [27] T. Grossman, K. Hinckley, P. Baudisch, M. Agrawala, and R. Balakrishnan. Hover widgets: Using the tracking state to extend the capabilities of pen-operated devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06*, p. 861–870. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1124772.1124898
- [28] J. Grubert, E. Ofek, M. Pahud, P. O. Kristensson, F. Steinicke, and C. Sandor. The office of the future: Virtual, portable, and global. *IEEE computer graphics and applications*, 38(6):125–133, 2018.
- [29] J. Grubert, L. Witzani, E. Ofek, M. Pahud, M. Kranz, and P. O. Kristensson. Text entry in immersive head-mounted display-based virtual reality using standard keyboards. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 159–166. IEEE, 2018.
- [30] J. Guo, D. Weng, Z. Zhang, H. Jiang, Y. Liu, Y. Wang, and H. B.-L. Duh. Mixed reality office system based on maslow’s hierarchy of needs: Towards the long-term immersion in virtual environments. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 224–235. IEEE, 2019.
- [31] K. Hasan, X.-D. Yang, A. Bunt, and P. Irani. A-coord input: Coordinating auxiliary input streams for augmenting contextual pen-based interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12*, p. 805–814. Association for Computing Machinery, New York, NY, USA, 2012. doi: 10.1145/2207676.2208519
- [32] F. Hermans, M. Pinzger, and A. van Deursen. Supporting professional spreadsheet users by generating leveled dataflow diagrams. In *Proceedings of the 33rd International Conference on Software Engineering, ICSE '11*, p. 451–460. Association for Computing Machinery, 2011. doi: 10.1145/1985793.1985855
- [33] O. Hilliges, S. Izadi, A. Wilson, S. Hodges, A. Garcia-Mendoza, and A. Butz. Interactions in the air: Adding further depth to interactive tabletops. In *UIST '09 Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pp. 139–148. ACM, October 2009.
- [34] K. Hinckley, X. hen, and H. Benko. Motion and context sensing techniques for pen computing. In *Proceedings of Graphics Interface 2013, GI '13*, p. 71–78. Canadian Information Processing Society, CAN, 2013.
- [35] K. Hinckley, S. Heo, M. Pahud, C. Holz, H. Benko, A. Sellen, R. Banks, K. O’Hara, G. Smyth, and B. Buxton. Pre-touch sensing for mobile interaction. In *CHI '16 Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 2869–2881. ACM, May 2016.
- [36] K. Hinckley, M. Pahud, H. Benko, P. Irani, F. Guimbretière, M. Gavrilu, X. A. Chen, F. Matulic, W. Buxton, and A. Wilson. Sensing techniques for tablet+stylus interaction. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology, UIST '14*, p. 605–614. Association for Computing Machinery, New York, NY, USA, 2014. doi: 10.1145/2642918.2647379
- [37] K. Hinckley, K. Yatani, M. Pahud, N. Coddington, J. Rodenhouse, A. Wilson, H. Benko, and B. Buxton. Pen + touch = new tools. In *UIST '10 Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pp. 27–36. ACM, October 2010.
- [38] K. Hinckley, S. Zhao, R. Sarin, P. Baudisch, E. Cutrell, M. Shilman, and D. Tan. Inkseine: In situ search for active note taking. In *CHI '07 Proceedings of the SIGCHI Conference on Human Factors in Comput-*

- ing Systems, pp. 251–260. ACM, April 2007.
- [39] S. Hwang, A. Bianchi, M. Ahn, and K. Wohn. Magpen: Magnetically driven pen interactions on and around conventional smartphones. In *Proceedings of the 15th International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '13, p. 412–415. Association for Computing Machinery, New York, NY, USA, 2013. doi: 10.1145/2493190.2493194
- [40] D. Jannach and T. Schmitz. Model-based diagnosis of spreadsheet programs: a constraint-based debugging approach. *Automated Software Engineering*, 23(1):105–144, 2016.
- [41] S. Jeong, E. S. Jung, and Y. Im. Ergonomic evaluation of interaction techniques and 3d menus for the practical design of 3d stereoscopic displays. *International Journal of Industrial Ergonomics*, 53:205–218, 2016.
- [42] S. P. Jones, A. Blackwell, and M. Burnett. A user-centred approach to functions in excel. In *Proceedings of the eighth ACM SIGPLAN international conference on Functional programming*, pp. 165–176, 2003.
- [43] E. Kandogan, E. Haber, R. Barrett, A. Cypher, P. Maglio, and H. Zhao. A1: end-user programming for web-based system administration. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pp. 211–220, 2005.
- [44] P. Knierim, V. Schwind, A. M. Feit, F. Nieuwenhuizen, and N. Henze. Physical keyboards in virtual reality: Analysis of typing performance and effects of avatar hands. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, p. 345. ACM, 2018.
- [45] M. Kobayashi and H. Koike. Enhanceddesk: integrating paper documents and digital documents. In *Proceedings. 3rd Asia Pacific Computer Human Interaction (Cat. No. 98EX110)*, pp. 57–62. IEEE, 1998.
- [46] A. Kohlhase, M. Kohlhase, and A. Guseva. Context in spreadsheet comprehension. In *SEMS@ ICSE*, pp. 21–27, 2015.
- [47] R. Komerska and C. Ware. A study of haptic linear and pie menus in a 3d fish tank vr environment. In *12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2004. HAPTICS'04. Proceedings.*, pp. 224–231. IEEE, 2004.
- [48] P. G. Kry, A. Pihuit, A. Bernhardt, and M.-P. Cani. Handnavigator: Hands-on interaction for desktop virtual reality. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology*, pp. 53–60. ACM, 2008.
- [49] A. Kulshreshtha and J. J. LaViola Jr. Exploring the usefulness of finger-based 3d gesture menu selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1093–1102, 2014.
- [50] G. Kurtenbach and W. Buxton. User learning and performance with marking menus. pp. 258–264. ACM Press, 1994.
- [51] N. Li, T. Han, F. Tian, Huang, S. Jin, P. Minhui, Irani, and J. Alexander. Get a grip: Evaluating grip gestures for vr input using a lightweight pen. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 2020.
- [52] Z. Li, M. Annett, K. Hinckley, K. Singh, and D. Wigdor. Holodoc: Enabling mixed reality workspaces that harness physical and digital content. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, p. 687. ACM, 2019.
- [53] Z. H. Lim and P. O. Kristensson. An evaluation of discrete and continuous mid-air loop and marking menu selection in optical see-through hmds. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*, pp. 1–10, 2019.
- [54] S. Liu and F. Guimbretière. Flexaura: A flexible near-surface range sensor. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*, UIST '12, p. 327–330. Association for Computing Machinery, New York, NY, USA, 2012. doi: 10.1145/2380116.2380158
- [55] K. Mack, J. Lee, K. Chang, K. Karahalios, and A. Parameswaran. Characterizing scalability issues in spreadsheet software using online forums. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–9, 2018.
- [56] W. Mackay. Which interaction technique works when?: floating palettes, marking menus and toolglasses support different task strategies. pp. 203–208, 01 2002. doi: 10.1145/1556262.1556294
- [57] N. Marquardt, R. Jota, S. Greenberg, and J. A. Jorge. The continuous interaction space: interaction techniques unifying touch and gesture on and above a digital surface. In *IFIP Conference on Human-Computer Interaction*, pp. 461–476. Springer, 2011.
- [58] F. Matulic, R. Arakawa, B. Vogel, and D. Vogel. Pensight: Enhanced interaction with a pen-top camera. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 2020.
- [59] F. Matulic and M. C. Norrie. Pen and touch gestural environment for document editing on interactive tabletops. In *Proceedings of the 2013 ACM international conference on Interactive tabletops and surfaces*, pp. 41–50, 2013.
- [60] M. McGill, D. Boland, R. Murray-Smith, and S. Brewster. A dose of reality: Overcoming usability challenges in vr head-mounted displays. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 2143–2152. ACM, 2015.
- [61] M. McGill, J. Williamson, A. Ng, F. Pollick, and S. Brewster. Challenges in passenger use of mixed reality headsets in cars and other transportation. *Virtual Reality*, pp. 1–21, 2019.
- [62] G. Miller and F. Hermans. Gradual structuring in the spreadsheet paradigm. In *2016 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, pp. 240–241. IEEE, 2016.
- [63] P. Monteiro, H. Coelho, G. Gonçalves, M. Melo, and M. Bessa. Comparison of radial and panel menus in virtual reality. *IEEE Access*, 7:116370–116379, 2019.
- [64] R. Panko. What we don't know about spreadsheet errors today: The facts, why we don't believe them, and what we need to do. *arXiv preprint arXiv:1602.02601*, 2016.
- [65] K. Pfeuffer, K. Hinckley, M. Pahud, and B. Buxton. Thumb + pen interaction on tablets. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, p. 3254–3266. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3025453.3025567
- [66] D.-M. Pham and W. Stuerzlinger. Is the pen mightier than the controller? a comparison of input devices for selection in virtual and augmented reality. In *25th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–11, 2019.
- [67] C. Pinhanez. The everywhere displays projector: A device to create ubiquitous graphical interfaces. In *International conference on ubiquitous computing*, pp. 315–331. Springer, 2001.
- [68] J. R. F. Pollock. *CountMarks: Multi-Finger Marking Menus for Mobile Interaction with Head-Mounted Displays*. PhD thesis, Carleton University, 2019.
- [69] D. R. Raymond. *A survey of research in computer-based menus*. Cite-seer, 1986.
- [70] J. Rekimoto and M. Saitoh. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pp. 378–385. ACM, 1999.
- [71] A. Ruvimova, J. Kim, T. Fritz, M. Hancock, and D. C. Shepherd. "transport me away": Fostering flow in open offices through virtual reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, 2020.
- [72] D. Schneider, A. Otte, T. Gesslein, P. Gagel, B. Kuth, M. S. Damlakhi, O. Dietz, E. Ofek, M. Pahud, P. O. Kristensson, et al. Reconfiguration: Reconfiguring physical keyboards in virtual reality. *IEEE transactions on visualization and computer graphics*, 2019.
- [73] D. Schneider, A. Otte, A. S. Kublin, A. Martschenko, P. O. Kristensson, E. Ofek, M. Pahud, and J. Grubert. Accuracy of commodity finger tracking systems for virtual reality head-mounted displays. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pp. 805–806. IEEE, 2020.
- [74] H. Shiozawa, K.-i. Okada, and Y. Matsushita. 3d interactive visualization for inter-cell dependencies of spreadsheets. In *Proceedings of the 1999 IEEE Symposium on Information Visualization*, INFOVIS '99, p. 79. IEEE Computer Society, 1999.
- [75] J. Smith, J. A. Middleton, and N. A. Kraft. Spreadsheet practices and challenges in a large multinational conglomerate. In *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, pp. 155–163. IEEE, 2017.
- [76] H. Song, H. Benko, F. Guimbretiere, S. Izadi, X. Cao, and K. Hinckley.

- Grips and gestures on a multi-touch pen. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, p. 1323–1332. Association for Computing Machinery, New York, NY, USA, 2011. doi: 10.1145/1978942.1979138
- [77] M. Speicher, A. M. Feit, P. Ziegler, and A. Krüger. Selection-based text entry in virtual reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, p. 1–13. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3173574.3174221
- [78] H. B. Surale, A. Gupta, M. Hancock, and D. Vogel. Tabletinvr: Exploring the design space for using a multi-touch tablet in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, p. 13. ACM, 2019.
- [79] Y. Suzuki, K. Misue, and J. Tanaka. Interaction technique for a pen-based interface using finger motions. In *International Conference on Human-Computer Interaction*, pp. 503–512. Springer, 2009.
- [80] D. S. Syrdal, N. Otero, and K. Dautenhahn. Video prototyping in human-robot interaction: Results from a qualitative study. In *Proceedings of the 15th European conference on Cognitive ergonomics: the ergonomics of cool interaction*, pp. 1–8, 2008.
- [81] Z. Szalavári and M. Gervautz. The personal interaction panel—a two-handed interface for augmented reality. In *Computer graphics forum*, vol. 16, pp. C335–C346. Wiley Online Library, 1997.
- [82] Z. Szalavári and M. Gervautz. Using the personal interaction panel for 3d interaction. In *proceedings of the Conference on Latest Results in Information Technology*, p. 36. Citeseer, 1997.
- [83] M. Teysier, G. Bailly, and E. Lecolinet. Versapen: An adaptable, modular and multimodal i/o pen. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '17, p. 2155–2163. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3027063.3053159
- [84] F. Tian, L. Xu, H. Wang, X. Zhang, Y. Liu, V. Setlur, and G. Dai. Tilt menu: Using the 3d orientation information of pen devices to extend the selection capability of pen-based user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, p. 1371–1380. Association for Computing Machinery, New York, NY, USA, 2008. doi: 10.1145/1357054.1357269
- [85] H. Tu, S. Huang, J. Yuan, X. Ren, and F. Tian. Crossing-based selection with virtual reality head-mounted displays. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–14, 2019.
- [86] R. Unger and C. Chandler. *A Project Guide to UX Design: For user experience designers in the field or in the making*. New Riders, 2012.
- [87] J. A. Wagner Filho, C. M. D. S. Freitas, and L. Nedel. Virtualdesk: a comfortable and efficient immersive information visualization approach. In *Computer Graphics Forum*, vol. 37, pp. 415–426. Wiley Online Library, 2018.
- [88] P. M. H. K. Webb, Andrew and B. Buxton. Wearables as context for guard-abiding bimanual touch. In *UIST '16 Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pp. 287–300. ACM, October 2016.
- [89] P. D. Wellner. Interacting with paper on the digitaldesk. Technical report, University of Cambridge, Computer Laboratory, 1994.
- [90] S. White, D. Feng, and S. Feiner. Interaction and presentation techniques for shake menus in tangible augmented reality. In *2009 8th IEEE International Symposium on Mixed and Augmented Reality*, pp. 39–48. IEEE, 2009.
- [91] H. Xia, K. Hinckley, M. Pahud, X. Tu, and B. Buxton. Writlarge: Ink unleashed by unified scope, action, and zoom. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, May 2017. Honorable Mention.
- [92] Y. Zhang, M. Pahud, C. Holz, H. Xia, G. Laput, M. McGuffin, X. Tu, A. Mittereder, F. Su, B. Buxton, and K. Hinckley. Sensing posture-aware pen+touch interaction on tablets. In *CHI 2019 Conference on Human Factors in Computing Systems*. ACM, May 2019.
- [93] D. Zielasko, M. Krüger, B. Weyers, and T. W. Kuhlen. Menus on the desk? system control in deskvr. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1287–1288. IEEE, 2019.
- [94] D. Zielasko, M. Krüger, B. Weyers, and T. W. Kuhlen. Passive haptic menus for desk-based and hmd-projected virtual reality. In *2019 IEEE 5th Workshop on Everyday Virtual Reality (WEVR)*, pp. 1–6. IEEE, 2019.