VirtualDesk: A Comfortable and Efficient Immersive Information Visualization Approach

J. A. Wagner Filho, C. M.D.S. Freitas and L. Nedel

Institute of Informatics, Federal University of Rio Grande do Sul, Brazil


Figure 1: *In the VirtualDesk prototype, data is rendered at arm's reach and manipulated only by mid-air natural hand gestures (center). A reproduction of the analyst's real desk is included (left) to enable tangible interaction with coordinated views and controls (right).*

Abstract
3D representations are potentially useful under many circumstances, but suffer from long known perception and interaction challenges. Current immersive technologies, which combine stereoscopic displays and natural interaction, are being progressively seen as an opportunity to tackle this issue, but new guidelines and studies are still needed, especially regarding information visualization. Many proposed approaches are impractical for actual usage, resulting in user discomfort or requiring too much time or space. In this work, we implement and evaluate an alternative data exploration metaphor where the user remains seated and viewpoint change is only realisable through physical movements. All manipulation is done directly by natural mid-air gestures, with the data being rendered at arm's reach. The virtual reproduction of the analyst's desk aims to increase immersion and enable tangible interaction with controls and two dimensional associated information. A comparative user study was carried out against a desktop-based equivalent, exploring a set of 9 perception and interaction tasks based on previous literature and a multidimensional projection use case. We demonstrate that our prototype setup, named VirtualDesk, presents excellent results regarding user comfort and immersion, and performs equally or better in all analytical tasks, while adding minimal or no time overhead and amplifying user subjective perceptions of efficiency and engagement. Results are also contrasted to a previous experiment employing artificial flying navigation, with significant observed improvements.

CCS Concepts
•Human-centered computing → Empirical studies in visualization; Virtual reality;

1. Introduction

Three dimensional representations are known to offer advantages under many circumstances. For inherently spatial data, these representations contribute to a quicker construction of the mental model [Mun14]. For abstract information, they have also been demonstrated to be useful, allowing clearer spatial separation in large graphs [WM08], detection of trivariate patterns in scatter-

plots [SBSSB15] and more accurate projection of multidimensional data [GGR*16, PEP*11]. Nonetheless, their use has also been long controversial, mostly because they are hindered by the occurrence of known perceptual issues: perspective distortion, foreshortening and occlusion make data exploration cumbersome and error-prone. Moreover, there is a relevant mismatch between 3D visualizations and conventional 2D interaction devices such as the mouse.

Immersive approaches, which combine stereoscopic displays with natural interaction, have the potential to change this scenario, with multiple favourable results having been demonstrated [CML*11, DSH*15, UKF*17]. Nonetheless, new evaluations and guidelines are still needed, particularly when dealing with abstract information visualization [GHAWK16]. Navigation, especially, is an open topic. Many works employ flying metaphors, which are time-consuming and often result in simulator sickness [WRFN18]. Other approaches, such as real walking, are also unnecessarily inefficient, both in terms of time and space requirements. Moreover, how to display inherently 2D content and texts in the virtual environment is also a problem.

In this work, we propose, implement and evaluate an alternative data exploration approach where viewpoint change is only realisable through head movements. All data manipulation is done directly by mid-air gestures, with the data being rendered at arm's reach [MBJS97]. To increase immersion and enable the display of two-dimensional associated views and interaction with tangible controls, we also build upon previous work and reproduce in the virtual environment an exact copy of the analyst's desk [ZWB*17]. Important data exploration resources are provided, including coordinated views, combinable filters and annotation tools.

Our main hypothesis is that our prototype setup, named VirtualDesk, will enhance user perception and decrease workload, while remaining time-efficient and not inducing cybersickness. A controlled comparative user study was carried out against a desktop-based equivalent 3D environment, implemented with the same functionalities and following typical mouse and keyboard interaction approaches. Both conditions employ a visualization use case with multidimensional data projected to three dimensions. In order to identify the strengths and weaknesses of each, participants were asked to complete a set of 9 representative perception and interaction tasks, inspired by previous literature.

Our main contribution is the recommendation of a so-far atypical data exploration metaphor, which under controlled evaluation presented excellent results for user comfort and immersion, and performed equally or better than a desktop-based solution for all proposed tasks, while adding minimal time overhead and amplifying subjective user perception of accuracy and engagement. Additionally, we also contribute to future similar studies, identifying a set of relevant tasks, providing specific baseline results to be used and pointing paths for improvement.

The remainder of this work is organized as follows. Firstly, related work is reviewed (§2) and the VirtualDesk introduced (§3). Then, our evaluation methodology is presented in the form of a user study (§4). These results are finally presented (§5) and discussed (§6), leading to the conclusions and future works (§7).

2. Related Work

Recently, multiple technological advancements and the release of consumer-grade Head-Mounted Displays (HMDs) have brought attention to Immersive Analytics approaches as candidates to tackle 3D visualization challenges. Chandler et al. defined this field as the investigation of how novel interaction and display technologies combined may support the analytical reasoning [CCC*15]. Some researchers also explore immersion in CAVE rooms [KH14], spaces surrounded by retro-projected walls, but their application is more limited due to the complexity of the infrastructure required.

For scientific visualization, immersive applications have already achieved a somewhat consolidated usage [GHAWK16]. Interesting results were provided, for example, for analysis of volumes [CML*11], materials [DSH*15] and neuron tracing [UKF*17]. In the case of abstract information visualization, research and guidelines are still limited, but promising results have also been demonstrated. A commonly explored representation is node-link diagrams. Halpin et al. implemented a generic semantic social network visualization software for CAVE-like environments, named *Redgraph*, and, in a user study, observed significant performance improvements for fine-grained questions using the immersive condition [HZBK08]. Kwon et al. explored different techniques in an HMD-based environment, proposing the use of a new spheric layout that offered performance increase especially for more difficult tasks [KMLM16]. Cordeil et al. presented a comparative study between CAVE-style and HMD-based environments for collaborative analysis of graphs, and were able to obtain high accuracy scores in both [CDK*17]. Nonetheless, users in the HMD condition were found to be substantially faster, regardless of the collaborative strategies adopted. Ware and Mitchell observed an order of magnitude increase over 2D displays in a path tracing task, using high resolution displays and a mirror stereoscope [WM08].

Point clouds have also been targeted due to their broad application in various areas, from multidimensional data to spatio-temporal visualization. Donalek et al. implemented *iViz*, an HMD-based platform for visualization of multidimensional data mapped to different attributes of points in a 3D scatterplot [DDC*14]. Bach et al. recently evaluated the effectiveness of Augmented Reality (AR) approaches, using tablets or see-through HMDs combined with tangible markers, in four different tasks [BSB*17]. They observed that the proposed direct hologram interaction was helpful in highly interactive tasks, but the desktop alternative was still the quickest and most accurate in most cases. Meanwhile, Cordeil et al. proposed an interactive Virtual Reality (VR) tool, *ImAxes*, where variable axes can be combined, through embodied interaction, to construct different representations, such as 2D or 3D scatterplots and parallel coordinates plots [CCD*17]. Keiriz et al. also employed 3D point clouds for the immersive exploration of brain connectivity in their *NeuroCave* tool, providing an interesting approach with multiple simultaneous views [KALF17].

Concerning 3D scatterplot projections of multidimensional data in particular, Etemadpour et al. evaluated a six-sided CAVE environment, observing better perception of distances between individual objects under this condition than when using a standard 2D screen [EML13]. In previous work, we have also focused specifically on this topic, exploring an HMD-based implementa-

tion [WRFN18]. We found that, for some datasets, both immersive and non-immersive 3D conditions allowed for a better perception of distances in the original higher-dimensional space. The immersive environment, however, required less navigation and effort to find information, and offered much higher perceptions of accuracy and engagement. Nonetheless, the use of flying navigation resulted in inefficient times and frequent user discomfort. Those results are used as baselines for some tasks in the present work (§5.3).

Navigation is indeed a key issue in immersive approaches. Most works adopt artificial metaphors such as flying [BH95, DDC*14, ZHF*16, LJKM*17], which frequently induce simulator sickness due to conflicts with the user's real perceived position. Others have also tried to employ physical movements, such as walking, as an alternative [SZK17], but this is generally very time and space consuming. Intermediate solutions, such as using physical movements like body leaning to control the artificial navigation, have also been proposed, but with limited success [ZHF*16]. We argue, however, that the best approach would actually be to render the data in smaller scale, at arm's reach [MBJS97], and just manipulate it with natural mid-air gestures to obtain different points of view.

The reproduction of the user's physical desk in the virtual environment, such as done in our work, was firstly seen in Zielasko et al.'s research [ZWB*17]. They later also experimented with the inclusion of the user's keyboard into the virtual scene [ZBM*17]. However, both these works still apply artificial flying navigation, with the desk flying coupled to the camera throughout the environment, making our concept and implementation fundamentally different. They address the issue of cybersickness by using user profiles, which would help to indicate when to limit certain system features. Bellgardt et al. also discussed the possibility of a seating immersive scenario [BPZ*17], but considered that it would sacrifice the level of immersion and realism, only being useful for short sessions. We disagree with this assessment, and argue that under an appropriate exploration paradigm, it could be highly immersive.

Cordeil et al. recently defined the concept of *spatio-data coordination* (SD), aiming to lower the user's cognitive workload when exploring information visualizations [CBL*17]. They argue for a one-to-one mapping of positions, directions and actions between the physical and virtual environments, and present a design space to categorize novel solutions. Our small-scale dataset rendering is consistent with their sketch of a virtual mid-air design for SD coordinated interaction.

Finally, we also borrow concepts from 3D user interfaces (3DUI) research. LaViola et al. presented a thorough discussion on 3D interaction techniques [LJKM*17]. In our scenario, the most relevant is the direct manipulation through simple virtual hands. Mine also discussed how interacting within arm's reach can take advantage of proprioception to provide a greater sense of position and orientation of manipulated objects. Body-relative interaction also provides higher precision and stronger stereopsis and head-motion parallax cues [MBJS97].

3. The VirtualDesk Metaphor and Prototype

The main foundations of our approach are the rendering of data at arm's reach for seated exploration (§3.1), the usage of the real


Figure 2: *In the VirtualDesk prototype, all system control and data manipulation are performed by tabletop tangible interaction (left) or controller-agnostic mid-air natural gestures, such as grabbing and tapping (right).*

user's desk for tangible interaction (§3.2), and the display of controls and coordinated 2D views (§3.3).

3.1. Interacting with Data

In the VirtualDesk prototype, all data manipulation is implemented by natural mid-air gestures, using direct interaction with virtual hands [LJKM*17] (see Figures 1 and 2). This is expected to minimize the user workload, given the intuitiveness of the actions and also the application of the sense of proprioception.

The main actions consist in *grabbing* the dataset and *tapping* data points. After grabbing the dataset with one hand, the user can move it and also rotate it around the hand position. Grabbing with two closed fists allows for the rotation and translation with relation to the central point between hands, and also the scaling of the dataset proportionally to the variation in distance between hands. Data points are selectable by quickly double tapping on their surfaces (see Figure 1 – center). This was chosen instead of single tapping to avoid the selection of undesired points in cluttered regions. Haptic feedback in the form of vibration when touching points contributes to the perception of a tangible interaction.

In this prototype, we opted to implement different actions for selection with each hand: while the right index finger activates a point (displaying its associated information or choosing it as answer in a task), the left one triggers the supportive action of highlighting a whole set of points for providing context to the user.

3.2. Tabletop Tangible Interaction

Following recent literature [ZWB*17], we decided to replicate the user's desk in the virtual environment. The use of tangible user interfaces (TUIs) is known to greatly benefit immersion [CBL*17]. The virtual desk is represented in an exact position (see §3.4) so as that, when the user touches the surface of the real table, his virtual hand touches the virtual one. We refer to this form of interaction as *tabletop*, to avoid confusion with the term *desktop*. Although the virtual desk is rendered larger than the real one, to provide a greater notion of space, a different marking keeps the user aware of the position of the actual desk (see Figure 1 – right).

Several controls are available on the virtual desk's surface in our prototype: buttons to reset the data points to the original position and scale, remove filters and change datasets. These buttons also provide haptic vibration to increase tangibility. Moreover, coordinated filters and visualization tools are also provided (§3.3). All

these components are shown in the frontal part of the desk, for easy access. An important note is that this segment must be free of obstacles (e.g. the user's keyboard) to avoid unintended collisions [ZBM*17].

By incorporating an element of the real world, VirtualDesk can also be described as a *mixed reality*, or *augmented virtuality* application [MK94]. We consider this to be a better fit than pure Augmented Reality (AR) to our purposes, since the remaining unnecessary and distracting surroundings can be eliminated, leading to a more immersive exploration experience. Despite not implemented in our current prototype, this also allows the analyst's environment to be entirely recreated to present extra information and to enable remote collaboration. Naturally, AR approaches would be potentially more convenient for combined exploration with other sources and collocated collaborations, and are also of interest. However, it should be noted that current AR HMDs are much less matured in comparison to their VR counterparts, rendering their evaluation still very difficult at this time. An interesting first assessment was presented by Bach et al. [BSB*17].

3.3. Coordinated 2D Views and Visualization Functionalities

Besides enabling tabletop tangible interaction, we also see the inclusion of the virtual desk as an opportunity to tackle another challenge in virtual environments: how to display and interact with texts and two dimensional information.

Two views associated to the main dataset were incorporated in the prototype as examples: a legend for categorical information and a map for spatial filtering (see Figure 1 – left and right). Both of them act as combinable coordinated filters, showing or hiding information in the main 3D view.

Additionally, an annotation panel was included as an example of possible extra analytical feature. This panel allows the user to change the color mapping of points to an uniform color, and then to mark individual points. These annotations could easily be persisted for future inspection either in VR or in a conventional display.

3.4. Technical Details and Choices

The VirtualDesk prototype was implemented using the Unity3D game engine and the Oculus Rift CV1 HMD (composed by two 1200x1080 stereoscopic displays). Adequate hardware was used to meet the recommendation of a frame rate around 90 FPS [YHD*14].

An important decision in the implementation was the selection of the hardware for the tracking of the user hands. Several related works that explored mid-air gestures in the past have employed the Leap Motion hand tracker [BFF*15, TLN17, ZWB*17]. Nonetheless, based on previous experience, we felt that this would not match the level of precision and comfort required for a satisfactory user experience, and opted instead to use the recently released Oculus Touch hand controllers. Although these controllers do not track the position of each finger, they are very precise in tracking the overall hand position based on the Constellation tracking technology. Moreover, they apply different touch and near-touch sensors

coupled with heuristics to determine the fingers positions. The official Unity Oculus Integration Package provided the hand models and the gesture mapping.

The Oculus Touch tracking was also used to implement the desk positioning. Upon the application start, the controllers are placed in a fixed location, and the virtual desk is then rendered in relation to their detected positions, resulting in a very accurate solution.

Another design choice was to not use any controller-specific tool, such as buttons, in any action – i.e., the actual controllers are completely abstracted by the users after they learn the gestures. The reason was twofold: we wanted to base interaction only on natural actions, and also to obtain a controller-agnostic framework, which could easily be adapted to any other tracking device. Although this also implicates in not benefiting from any controller-specific facilities, we are convinced that most required actions can be satisfactorily implemented by gestures, while more complex interactions will be clearer when assigned to the table controls.

4. User Study Evaluation

In order to assess how the VirtualDesk prototype would perform in comparison with conventionally used desktop-based approaches, a typical multidimensional projection use case was selected (§4.1), and a comparable condition implemented (§4.2). Finally, we collected a set of representative tasks (§4.3) and formulated the study hypotheses (§4.4) and design (§4.5).

4.1. Use Case

We decided to employ, in our evaluation, point cloud representations. These have multiple applications in information visualization and are also often associated with perceptual difficulties when shown in 3D. More specifically, we selected a multidimensional projection use case.

To this end, roll call voting data from the Brazilian Chamber of Deputies was collected. This data is particularly interesting for visualization because its projection to a lower dimension results in a political spectrum [SM06]. This is also a domain appealing to different kinds of public, with potential to engage participants in data exploration during the tests. Some of these datasets were also explored in our previous work on immersive exploration of dimensionally-reduced data [WRFN18], enabling those results to be used as baselines.

We extracted information about the votes of each deputy and the official vote instruction given by each party represented in the Chamber for every roll call in the last four four-year legislatures from the Brazilian Congress: 52nd (451 roll calls), 53rd (619 roll calls), 54th (428 roll calls) and 55th (493 roll calls). For each legislature, a voting matrix is constructed, where all deputies and parties are represented by M lines, and roll calls are represented by N columns. Each *(i, j)* cell is then attributed a value depending on the *ith* deputy or party vote on the *jth* roll call: -1 for "no", 1 for "yes" or 0 for abstention or absence. Following previous works [JB04, dBF15], Principal Component Analysis (PCA) is then applied, resulting in *min(N,M)* principal components [Hot33]. For visualization purposes, only the first three are considered.

In the VirtualDesk interface, party information is used for filtering in the legend view, and state information for filtering in the map view. In the point cloud, party information is encoded by colour, and the category of the point is encoded by shape: spheres for deputies and cubes for official party positions.

4.2. Comparable Condition

A key limitation in our previous study was the fact that the then-used desktop-based 3D condition was not really representative of a typical 3D scatterplot visualization tool. Its game-like design and interaction resulted in very high perception performance even in a monoscopic display due to the presence of multiple depth cues, but very inefficient task completion times.

Here, we decided to implement a new desktop-based comparable counterpart (*Desktop*) to the VirtualDesk condition based on the Rotate-Pan-Dolly paradigm, a very standard approach employed by almost all 3D modelling environments [JH15]. Depending on the mouse button being pressed (left or middle), mouse movement is mapped to either rotation around the dataset center or camera translation (*panning*). The scroll wheel can be used to *dolly*, or *zoom*, into the data. Additionally, we also allowed the rotation around any selected pivot point (by holding a keyboard key) in order to enable better local inspection, required in some tasks. The selection of data points is implemented by double-clicks with the left mouse button, while class highlight is associated to the right button. Perspective projection was used as an additional depth cue, increasing similarity to the immersive environment.

This condition is explored in a Full HD 22" monoscopic display. The screen was divided into two areas: a 3D dataset view and a 2D menu panel, with all the components from the VirtualDesk's surface (see Figure 3). A 65%/35% screen space distribution was defined through empirical testing to maximize the dataset view without compromising the legibility and usability of the menus.
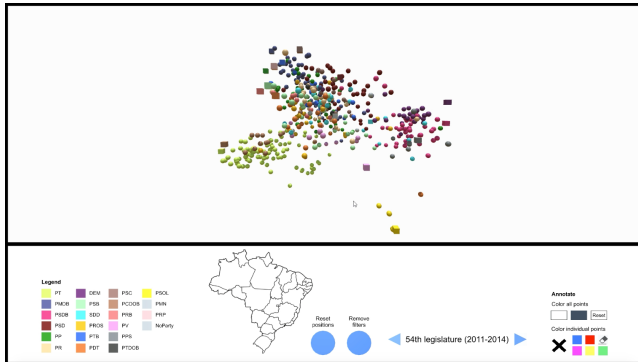

Figure 3: *The desktop-based implementation provides all the same functionalities as the immersive environment, but employing a two-panels interface and Rotate-Pan-Dolly interaction for the 3D point cloud exploration.*

4.3. Tasks

In order to assess our two visualization conditions under a variety of usage patterns, we selected a set of 9 different tasks, divided into four categories. These tasks were based on both relevant task taxonomies [SG17, ELCF15] and previous related evaluation studies [BSB*17, EML13].

Point-based distance perception tasks. Much information in point cloud representations is encoded through pairwise distances. In multidimensional projections, they quantify the similarities or differences between points in the original data space. Tasks were defined relatively to different competencies: perception of near, medium and far distances, and of different shape encodings.

T1 *Selection of a deputy's closest deputy*. In this near-distance perception task, the user is requested to select the closest deputy (sphere) to a given one.

T2 *Selection of a deputy's closest party*. In a more difficult variation of the previous task (since deputies are usually positioned between multiple parties), the user is requested to select the closest party (cube) to a given deputy.

T3 *Selection of a party's furthest deputy*. In this simplified outlier identification task, the user must select the member of a given party who is furthest located from the official party position.

T4 *Selection of a party's closest party*. Also a variation of T1, but exploring different competencies since parties are more distributed on the spectrum.

In terms of abstract scatterplot tasks, these refer to object comparison, neighbourhood exploration and distances understanding [SG17]. To enable the comparison with our previous study, where these four tasks were also used, the same selection of question points was repeated. These had been selected randomly, forming different sets of points from the *54th legislature* dataset that are repeated only once by a unique participant in each condition (all questions answered by a user in one condition will also be answered by a different user in the other). This maximizes the exploration of different possible situations in the data, and cross validates the results [GGR*16].

Class-based density perception tasks. Class or cluster density is another important factor in point cloud representations, indicating group cohesion. This is one of the behaviour comparison tasks recommended by Etemadpour et al. [ELCF15], and also a key scatterplot analysis task (numerosity/density comparison) in Sarikaya et al.'s taxonomy [SG17].

T5 *Density comparison between two parties*. In this task, the user must choose which of two simultaneously visualized given parties is the densest one.

T6 *Density comparison over time*. In this variation of the previous task, the user must choose one between two time periods, which cannot be simultaneously viewed, when the given party was denser.

For each task, two different sets of 3 questions were selected, and were alternated between conditions. In T5, pairs of parties in the *54th legislature* are compared, and in T6, the same party is compared between the *53rd* and *54th legislatures*. Only large parties with at least 20 deputies were considered.

Clustering task. Clustering is a typical pattern identification [ELCF15] or known motif search [SG17] task in point clouds.

T7 *Estimation of the number of clusters in a given point cloud.*

This task requires the inspection of different orthogonal points of view [BSB*17].

One of the four collected datasets was presented for each user in each condition, in varying orders.

Interaction tasks. Interaction efficiency is a main concern in 3D representations, also having been evaluated, for example, in Bach et al.'s recent AR proposal [BSB*17]. Tasks were designed to assess the interaction with the tabletop 2D views and the 3D data points.

T8a *Filtering of a party-state combination*. The user is requested to select, as quickly as possible, the correspondent filters in the tabletop views.
T8b *Selection of all remaining deputies*. Continuing the previous task, the user is requested to select, as fast as possible, all the remaining points in the 3D scatterplot.

Six different party-state combinations were selected in the *54th legislature* dataset, with 3 being presented in each condition in varying orders. To maximize the representation of real use scenarios, states of different sizes on the map and parties in different positions in the legend were selected. Pairs were also carefully selected to result in the selection of different numbers of points (3, 6 or 9).

4.4. Hypotheses

The following specific hypotheses were defined for evaluation.

H1 Easier data manipulation, proprioception and stereopsis combined will result in enhanced perception of distances and densities in the VR condition.
H2 Consolidated mouse-based interaction will still be quicker and more accurate for the selection tasks.
H3 Natural embodied interaction will decrease user mental workload and increase subjective perceptions of accuracy and engagement.
H4 The VirtualDesk metaphor will be more comfortable and efficient, both in time and task correctness, than our previous flying immersive exploration approach [WRFN18].

4.5. Experiment Design

A population of 24 undergraduate Computer Science students (20 male/4 female, mean age 23.7, SD 2.7) was invited to perform all tasks in the two compared conditions, in a *within-subjects* protocol. Half the users presented some visual condition and wore glasses in combination with the HMD. Twenty-two of the users reported no or low previous experience with HMDs, and 20 of them had no previous contact at all with the Oculus Touch controllers. Nonetheless, 22 reported at least average experience with 3D computer games, 21 with gamepads and 16 with motion controllers in general.

In the beginning of each condition, users were always presented a tutorial, which guided them through all system functions and exercised the different forms of interaction. Then, they proceeded to execute the tasks, which were always introduced by text accompanied by an illustrative icon (on-screen or close to the surface of the VirtualDesk). Participants were allowed to raise questions at any moment. The condition order was always alternated to compensate for the fact that, in the second condition, tasks would already be familiar to the users, but the task order was always kept the same to avoid confusion. Tasks were also distributed according to their increasing needs for interaction, so that previous tasks contribute to the familiarization with the system. Tasks always started in a data overview position. In Desktop, the monitor was positioned approximately 50cm in front of the users. In VirtualDesk, the center of the point cloud was initially positioned approximately 60cm in front of the users, and points rendered with a 1.5cm diameter.

For tasks T1-T4, one point in the cloud is shown blinking, and the user must select another point as answer. In T5 and T6, relevant parties are already shown highlighted (i.e. with the remaining points semitransparent), and the user must select the party cube correspondent to his answer. In T7, all points are shown in black to facilitate the perception of clusters and avoid confusion with classes, and the answer is given by an incremental counter positioned in the lower panel of the screen or near the surface of the desk. In T8, the acronym of the party to be filtered is shown on the task display, and the state is marked in red on the map (so as to avoid interference of varying previous geographical knowledge). This is the only task where reading is performed during the execution, and so font size was made large to minimize the effect of the participants' varying levels of visual acuity. For T1-T7, users were asked to be accurate and, for T8, to be fast. Following previous experiences, we blocked semantically impossible answers (e.g. a party outlier that is not from the given party), so as to reduce noise resulting from accidental clicks or misunderstandings. When this is the case, the user hears a negative audio feedback. Upon an acceptable answer, a positive sound is played, the image briefly fades and the data returns to its original overview position.

After each condition, users were asked to fill standardized questionnaires and answer general questions. In both parts, the SUS questionnaire was applied to assess system usability [Bro96], while the NASA Raw TLX was applied to assess user workload [Har06]. SSQ was applied to evaluate simulator sickness, comparing reported levels of sixteen different symptoms pre and post VR exposure [KDC*03]. IPQ was also applied post VR exposure to assess the level of presence experienced by users in the virtual environment [SFR01].

The complete experiment took approximately 40 minutes. The accompanying videos to this paper illustrate all system functionalities and how tasks were executed under both conditions.

5. Results

Results from the user study evaluation are reported here in terms of task performance (§5.1), user feedback (§5.2) and a comparison with our previous experiment (§5.3). Significance under the adequate statistical tests is indicated in the text and figures as follows: (*) for $p < 0.05$, (**) for $p < 0.01$ and (***) for $p < 0.001$. Z-values and effect sizes (r) are also reported [Pal13].

5.1. Task Performance

Task performance was assessed in terms of both task completion times and error rates. Since, here, we are concerned only with the

Figure 4: *Average task completion times for all tasks and conditions, with standard deviations indicated by error bars. For T8b, reported times are normalized per selected point. The immersive environment was only significantly slower in tasks which required higher amounts of interaction with the tabletop controls.*


Figure 5: *Average error rates for all tasks and conditions, with standard deviations. For the interaction tasks, errors are given by the number of unintended selections. All perception tasks were performed equally well or better in the immersive environment. Point selection, however, was more accurate in Desktop.*

correct perception of the actual representation, and not with the dimensionality reduction accuracy, all tasks are evaluated considering the lower-dimensional space only. For a discussion on dataset-dependent information gain when projecting data to 3D instead of 2D, we refer to our previous work [WRFN18].

For distance perception tasks T1-T4, pairwise Euclidean distances were computed to determine the correct answers. For density perception tasks T5-T6, we followed Etemadpour et al.'s approach based on the inverse of the average edge length in the Euclidean Minimum Spanning Tree of a class [EML13]. For the clustering task T7, ground truth was computed by the X-Means algorithm, and varied between 2 and 3 clusters [PM00]. Interaction tasks T8a and T8b, on the other hand, are assessed in terms of unintended selections. Times for T8b exclusively are averaged per selected point. Figures 4 and 5 present these results. Since parametric requirements were not met by multiple samples, paired Wilcoxon signed-rank tests were used to determine statistical significances.

Significant differences in time were only found in tasks T5-T8, in which cases the immersive condition was slower. With the exception of T5 ($***$, $z = -3.9$, $r = .56$), all of these were tasks with higher requirement of interaction with the tabletop controls. In T6 ($***$, $z = -4.1$, $r = .59$), the dataset needed to be changed; in T7 ($**$, $z = -3.1$, $r = .44$), the user answer was inputed through an incremental counter on the desk; in T8a ($**$, $z = -2.8$, $r = .41$), filters should

be applied. We believe this is partially related to the fact that some users experienced difficulties with tabletop interaction due to hand sizes (see §6.2). Moreover, the mouse interaction was already expected (H2) to be faster due to its consolidated usage. As opposed to the desktop-based condition, controls and data did not share the user's field of view in the immersive condition, what also required additional time. It is important to note that users were asked to be precise and not fast in tasks T1-T7. Considering task T8b ($***$, $z = -4.2$, $r = .61$), the point selection time was found to be 43% higher in the VR setup (3.3 vs 2.3s per point), confirming H2.

Hypothesis H1 could be partially accepted, given that tasks T1 – distances between spheres – ($*$, $z = -2$, $r = .28$) and T6 – density comparison over time – ($*$, $z = -2$, $r = .29$) obtained significantly smaller error rates in the VirtualDesk condition. The VR setup was also never significantly worse than the desktop-based condition in terms of perception. It was, however, more inaccurate in terms of point selection in task T8b ($**$, $z = -2.9$, $r = .42$): despite having an extra degree of freedom (DoF), users selected almost three times more unintended points with the virtual finger than with the mouse. We believe this was particularly problematic in cluttered areas of the representation, where it was difficult not to hit adjacent points during selection, especially considering that users had still not mastered the double tap action.

Finally, an interesting difference was observed in terms of

Figure 6: *Average accumulated dataset rotations per task question in degrees. This form of exploration was performed 5.8 times more in VirtualDesk, probably due to the intuitiveness of the grabbing action. This increased task accuracy with minimal time overhead.*

Figure 7: *User agreements with different assertions, ranging from completely disagree (dark red) to completely agree (dark green), for Desktop (D) and VirtualDesk (VD). Intuitive embodied data manipulation gestures were well received and allowed easy and rapid inspection and information finding in any region of the dataset.*

dataset rotations. These were performed, on average, 5.8 times more in the immersive condition, probably due to the intuitiveness of the grabbing action. Average accumulated data rotations per task question were 190.4 degrees (SD 179.6) in Desktop and 1,114.8 degrees (SD 935.4) in VirtualDesk. Considering that the observation from different points of view is fundamental in the comprehension of a 3D point cloud, this also partially explains VirtualDesk's advantage in perception tasks such as T1. Figure 6 presents results per task. Also note that, despite large differences in T1-T4, these tasks did not present any differences in time. In T7, however, the large difference in exploration between conditions (7.3x) may contribute to explain the difference in completion times. This probably was not reflected on answer accuracy, though, because questions turned out to be very easy, with only 2 or 3 clusters per dataset.

5.2. User Feedback

The general subjective feedback received from users in post-test interviews was very positive, especially regarding the use of 3D interaction for data manipulation. In terms of usability, both conditions were well rated in the SUS questionnaire. VirtualDesk obtained a 77.2 mean score (standard deviation 16.4) and, Desktop, 72.8 (SD 20.2), but differences were not significant.

For task workload, nonetheless, VirtualDesk's NASA Raw TLX score was significantly higher (*): 30.9 (SD 14.7) compared to 23.2 (SD 15.4). This was especially influenced by two workload components: Physical Workload (37.4 vs 9.7) (***) and Effort (36.1 vs 23.6) (*). This is understandable considering that users were observed to move their left and right hands on average 2.4m (SD 1.2m) and 4.3m (SD 1m) per task question, respectively. Mental Workload was scored at 26.3 (SD 16.2) against 22.2 (SD 20) of Desktop, without statistical difference, partially contradicting H3.

Concerning the immersive environment, SSQ scores were very satisfactory, averaging only 2.18 (SD 9.0), symptoms which can be considered negligible [KDC*03]. No user reported discomfort during or after the tasks. In terms of presence, VirtualDesk was rated in the IPQ (6 points scale) 4.7 (SD 0.88) for Spatial Presence, 4.07 (SD 1.06) for Involvement, 3.11 (SD 0.79) for Experienced Realism and 5.41 (SD 1.17) for the General Item (feeling of "being

there"). We provide these results in the expectation of serving as a baseline for future setups. It is important to note that participants were allowed to communicate with the experimenter at any time, keeping them aware of the external environment.

Analysing the users' agreements to different assertions (see Figure 7), it becomes clear what were the strengths and weaknesses of our prototype. 46% more participants agreed that it was easy to find information in VirtualDesk. This is probably closely related to the embodied data manipulation, which was not considered difficult by any user. By executing instinctive grabbing and scaling actions, users could easily and rapidly inspect any region of the dataset, as opposed to combining several Rotation-Pan-Dolly actions in the desktop-based version. This was probably what most impressed participants in the experience. Pointing data interaction was rated similarly in both versions, what is very positive considering that the quick double-tap metaphor had just been learned for the experiment, while double-mouse clicking is an universal action. On the other hand, difficulties with the tabletop interaction were the main system weakness: six participants experienced difficulties due to their real hands being larger than the fixed-size virtual model employed (see §6.2).

In a ranking question after the completion of both parts of the test, VirtualDesk was selected by all participants as the most engaging condition, and by 21 (87.5%) as the most intuitive. This was already expected, and is partially related to the novelty of VR, but also to the experienced immersion and the use of natural gestures for interaction. More importantly, VirtualDesk was perceived by 15 participants (62.5%) as the fastest technique. Both conditions tied in terms of accuracy, with 12 users choosing each. When asked, many reported that Desktop was most accurate for selection, but VirtualDesk for manipulation.

5.3. Comparison with Flying Navigation

One of the main motivations for this work came from the observation of high levels of user discomfort and impractical completion times in a previous study employing Flying navigation across a large-scaled point cloud [WRFN18].

Figure 8 contrasts results between the new and old paradigms for T1-T4 (tasks present in both studies), in terms of completion times

(a) Time

(b) Error rates

Figure 8: *Comparison between VirtualDesk and a previous implementation employing Flying navigation. Embodied data manipulation resulted in up to 51% shorter average completion times and 30% smaller error rates.*

(left) and error rates (right). Given that questions for these tasks are not repeated more than once in each condition, only the first 24 users from the previous study and the same dataset are considered, since it would be unfair to consider ones who performed potentially easier or more difficult questions. Mann–Whitney U tests for independent samples were used to compare results.

As expected, VirtualDesk was more time-efficient, and all tasks were executed significantly faster, reaching a 51% improvement in T2. This is explained by embodied manipulation allowing much more agile access to different parts of the point cloud than slowly flying through it. In terms of task performance, collected data has no statistical significance. Although all tasks consistently achieved lower error rates under the new approach (25%, 30%, 10% and 18% reductions, respectively), we are aware that this can be due to a random factor. Just as a speculation, we attribute this to the added notion of proprioception, and the stronger stereopsis and head-motion parallax cues at short distances, as discussed by Mine [MBJS97]. However, new tests should be conducted to verify this.

Another key result, in our opinion, is shown in Figure 9. Despite very similar VR exposure times in both studies, the average SSQ score in VirtualDesk was 7 times lower than in the artificially navigated version. Moreover, while in that study 40% of the users had experienced very significant discomfort levels (scores >= 20), now 83% perceived only negligible or minimal symptoms, and the maximum individual score was 18.3. This completed the confirmation of hypothesis H4.


Figure 9: *Due to its more natural and comfortable navigation paradigm, VirtualDesk achieved a 7x smaller SSQ score than the previous Flying approach, despite very similar VR exposure times.*

6. Discussion

6.1. Findings

Results from the user study confirmed our main intuitions in the proposal of this approach. VirtualDesk performed equally well or better across all analytical tasks, both in comparison to a standard Desktop interface and to a previous immersive implementation with Flying navigation (confirming hypotheses H1 and H4). The added time overhead with relation to Desktop was only significant in tasks with higher requirements for tabletop interaction (which demanded viewpoint change and also imposed certain difficulties for some users), and was generally only of a few seconds. This was despite the fact that data exploration in terms of dataset rotation was found to be 5.8 times higher in the immersive condition.

Tasks T1 (identification of the closest point) and T6 (density comparison over time) in particular presented significant error rate decrease under immersive data exploration. We believe this was related, respectively, to the easier inspection of local areas using 3D interaction, and to a possibly longer persistent obtained mental model of the data in the immersive condition.

The desktop-based 3D condition also performed well across tasks, as had already been observed in our previous study. As discussed by Ware, structure-from-motion cues enable the perception of point positions even without stereopsis. We are convinced that the comparison between the two implementations was fair, and most participants reported that each condition had its pros and cons. In particular, interaction tasks in this condition were still found to be quicker and less error-prone (as expected in H2).

Subjective feedback indicated that VirtualDesk was perceived as quicker, more intuitive and engaging, partially confirming H3. Nonetheless, the mental workload, as measured by the NASA Raw TLX questionnaire, did not present significant variation, and the overall workload increased due to higher inherent physical workload and perceived required effort to achieve the task goals.

In terms of interaction gestures, one of our main mistakes, in retrospect, was to assign different selection behaviours to the left and right hands. Despite being familiarized with them in the tutorial phase, even right-handed users intuitively constantly tried to se-

lect points with their left hands when they were closer. Meanwhile, the double tap gesture for point selection (as opposed to some controller-dependent action such as button clicking [BSB*17]), though difficult to master at the beginning for many users, did not affect the correspondent ratings (Figure 7), and we believe that, in the long term, would be more intuitive and efficient and reduce workload. Alternative object selection techniques, coupled with disambiguation mechanisms, could also be considered for evaluation in this context to minimize the unintended selections – a thorough review was presented by Argelaguet and Andujar [AA13].

6.2. Limitations

The main limitation of the present users evaluation study was the fact that it has only been performed in one specific use case (multidimensional roll call data projected to 3D) and one information representation (point clouds). Nonetheless, we believe this was adequate for our current purposes, which were to investigate and demonstrate the potential of a so-far atypical immersive data exploration paradigm, rather than to propose its mediate adoption in data analysis. It is also important to emphasize that our evaluation is admittedly only concerned with the benefits of different factors, such as stereopsis, tangible interaction, proprioception and embodied data manipulation when combined, and not individually, what could also be assessed in future specific studies.

Considering the prototype implementation, the main identified limitation was that the virtual hand models, obtained from the Unity Oculus Integration Package, were not adjusted accordingly to the participants' real hand sizes. This resulted in difficulties for at least six users who had larger hands and faced difficulties to reach the virtual desk surface despite being touching the real desk. This was always solved by slightly changing the controller position in the user's hand, but negatively affected their overall perception of interaction ease (see Figure 7) and partially compromised the evaluation of this aspect of the prototype. We intend to circumvent this limitation in a future version. Nevertheless, we believe the choice for the Oculus Touch controllers instead of other hand tracking solutions was appropriate, and resulted in highly realistic and precise modelling of the hands and hand gestures, what contributed to immersion and user engagement.

Another key limitation in our comparative study is the limited training provided to users in the new technique, as opposed to the ubiquitous familiarity of 2D user interaction, a common issue in the evaluation of novel approaches. Note that only 8% of our users had average previous experience with VR HMDs. After working on this prototype for several months, we are convinced that, upon longer training, interaction times and error rates in tasks T8a and T8b become much lower. In order to demonstrate this, however, long term evaluations will be needed. One such attempt was recently shown by Bach et al. [BSB*17], who reassessed 6 participants across five daily sessions. They observed speed improvements in 22 out of the 24 task-participant combinations, being 9 with significance. However, no significant improvements in precision were observed, and they noted that 5 sessions might have been too few. Finally, our user comfort results refer to an average 14.3 minutes VR session, and it is still unknown how this would change for longer exposures, also requiring further studies.

6.3. Perspectives

Despite being used as two alternative conditions in our study, we believe that one of the main perspectives for the VirtualDesk prototype is its combination with the non-immersive counterpart implementation. This paradigm allows for a direct mapping between all immersive environment contents and a two-panels 2D interface, which reproduce a 2D projection of the 3D data, or even a simpler 2D representation, in one part, and the surface of the desk in the other. Combining immersive and conventional data exploration environments becomes thus straightforward. This way, for example, annotations introduced in VirtualDesk could easily be persisted for further inspection in the monocular display if necessary [KALF17]. Moreover, through a task simulation approach [WRFN18], the analyst could quantify the contribution of the third dimension to the particular task or dataset in focus and, weighting the time trade-off, choose the most convenient exploration mode.

Taking into account our observed results for user comfort and low required completion times, we believe this is one of the first implementations for immersive exploration of abstract information actually convenient for real world usage, requiring only minor improvements. It also has the key advantage of being easily integrable into an analyst's work environment [ZWB*17].

From a research point of view, perspectives for further investigation include the inspection of larger datasets and different representations, and also multiple datasets or representations combined at the same time. Moreover, other interesting directions would be to expand the virtual environment to include extra information beyond the user's desk, and to support user collaboration.

7. Conclusion

In this work, we expanded the discussion into immersive information visualizations, implementing and evaluating a different exploration paradigm which we believed would be more fit for real usage. The VirtualDesk metaphor is based mainly on embodied natural manipulation and interaction with data rendered at arm's reach and tabletop tangible interaction with controls and 2D coordinated views positioned on the surface of a virtual desk, whose position is synchronized with the analyst's real desk. In a comparative user study employing a multidimensional projection use case, error rates for a series of perception tasks were always equal or lower than in a conventional desktop interface and in a previous immersive implementation with flying navigation. The immersive environment also contributed to higher subjective perceptions of efficiency and engagement and much higher data exploration, while incurring minimal time overhead and generating almost no simulator sickness symptoms, typical issues in previous studies.

As future work, we intend to improve the VirtualDesk prototype based on the user study participants' feedback, and also test it under different conditions, including different datasets and representations. We believe that strong candidates to benefit from this exploration metaphor would be node-link diagrams and space-time cube representations. Long term tests and longer VR exposure times will also be important to assess its real applicability. Lastly, we encourage further studies of alternative proposals, and provided here results which can be used as baselines.

Acknowledgements

References

[AA13]   ARGELAGUET F., ANDUJAR C.: A survey of 3d object selection techniques for virtual environments. *Computers & Graphics 37*, 3 (2013), 121–136. 10

[BFF*15]   BURGESS R., FALCÃO A. J., FERNANDES T., RIBEIRO R. A., GOMES M., KRONE-MARTINS A., DE ALMEIDA A. M.: Selection of large-scale 3d point cloud data using gesture recognition. In *Doctoral Conference on Computing, Electrical and Industrial Systems* (2015), Springer, pp. 188–195. 4

[BH95]   BOWMAN D. A., HODGES L. F.: *User interface constraints for immersive virtual environment applications.* Tech. rep., Georgia Institute of Technology, 1995. 3

[BPZ*17]   BELLGARDT M., PICK S., ZIELASKO D., VIERJAHN T., WEYERS B., KUHLEN T. W.: Utilizing immersive virtual reality in everyday work. In *3rd Workshop on Everyday Virtual Reality* (2017). 3

[Bro96]   BROOKE J.: Sus-a quick and dirty usability scale. *Usability Evaluation in Industry 189*, 194 (November 1996), 4–7. 6

[BSB*17]   BACH B., SICAT R., BEYER J., CORDEIL M., PFISTER H.: 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* (2017). 2, 4, 5, 6, 10

[CBL*17]   CORDEIL M., BACH B., LI Y., WILSON E., DWYER T.: A design space for spatio-data coordination: Tangible interaction devices for immersive information visualisation. In *Proceedings of IEEE Pacific Visualization Symposium (Pacific Vis)* (2017). 3

[CCC*15]   CHANDLER T., CORDEIL M., CZAUDERNA T., DWYER T., GLOWACKI J., GONCU C., KLAPPERSTUECK M., KLEIN K., MARRIOTT K., SCHREIBER F., ET AL.: Immersive analytics. In *2015 Big Data Visual Analytics (BDVA)* (September 2015), IEEE, pp. 1–8. doi:10.1109/BDVA.2015.7314296. 2

[CCD*17]   CORDEIL M., CUNNINGHAM A., DWYER T., THOMAS B. H., MARRIOTT K.: Imaxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (2017), ACM, pp. 71–83. 2

[CDK*17]   CORDEIL M., DWYER T., KLEIN K., LAHA B., MARRIOTT K., THOMAS B. H.: Immersive collaborative analysis of network connectivity: Cave-style or head-mounted display? *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (January 2017), 441–450. doi:10.1109/TVCG.2016.2599107. 2

[CML*11]   COFFEY D., MALBRAATEN N., LE T., BORAZJANI I., SOTIROPOULOS F., KEEFE D. F.: Slice wim: A multi-surface, multi-touch interface for overview+detail exploration of volume datasets in virtual reality. In *Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2011), ACM, pp. 191–198. 2

[dBF15]   DE BORJA F. G., FREITAS C. M.: Civisanalysis: Interactive visualization for exploring roll call data and representatives' voting behaviour. In *Proceedings of SIBGRAPI* (August 2015), IEEE, pp. 257–264. doi:10.1109/SIBGRAPI.2015.34. 4

[DDC*14]   DONALEK C., DJORGOVSKI S. G., CIOC A., WANG A., ZHANG J., LAWLER E., YEH S., MAHABAL A., GRAHAM M., DRAKE A., ET AL.: Immersive and collaborative data visualization using virtual reality platforms. In *2014 IEEE International Conference on Big Data (Big Data)* (October 2014), IEEE, pp. 609–614. doi:10.1109/BigData.2014.7004282. 2, 3

[DSH*15]   DROUHARD M., STEED C. A., HAHN S., PROFFEN T., DANIEL J., MATHESON M.: Immersive visualization for materials science data analysis using the oculus rift. In *2015 IEEE International Conference on Big Data (Big Data)* (October 2015), IEEE, pp. 2453–2461. doi:10.1109/BigData.2015.7364040. 2

[ELCF15]   ETEMADPOUR R., LINSEN L., CRICK C., FORBES A.: A user-centric taxonomy for multidimensional data projection tasks. In *IVAPP* (2015), pp. 51–62. 5

[EML13]   ETEMADPOUR R., MONSON E., LINSEN L.: The effect of stereoscopic immersive environments on projection-based multidimensional data visualization. In *IV* (2013), pp. 389–397. 2, 5, 7

[GGR*16]   GRACIA A., GONZÁLEZ S., ROBLES V., MENASALVAS E., VON LANDESBERGER T.: New insights into the suitability of the third dimension for visualizing multivariate/multidimensional data: A study based on loss of quality quantification. *Information Visualization 15*, 1 (2016), 3–30. doi:10.1177/1473871614556393. 2, 5

[GHAWK16]   GARCÍA-HERNÁNDEZ R. J., ANTHES C., WIEDEMANN M., KRANZLMÜLLER D.: Perspectives for using virtual reality to extend visual data mining in information visualization. In *2016 IEEE Aerospace Conference* (March 2016), IEEE, pp. 1–11. doi:10.1109/AERO.2016.7500608. 2

[Har06]   HART S. G.: Nasa-task load index (nasa-tlx); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting* (2006), vol. 50, Sage Publications Sage CA: Los Angeles, CA, pp. 904–908. 6

[Hot33]   HOTELLING H.: Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology 24*, 6 (1933), 417. 4

[HZBK08]   HALPIN H., ZIELINSKI D. J., BRADY R., KELLY G.: Exploring semantic social networks using virtual reality. In *International Semantic Web Conference* (Berlin, Heidelberg, 2008), Springer, pp. 599–614. doi:10.1007/978-3-540-88564-1_38. 2

[JB04]   JAKULIN A., BUNTINE W.: Analyzing the us senate in 2003: Similarities, networks, clusters and blocs. *Preprint. Available at http://kt.ijs.si/aleks/Politics/us_senate. pdf* (December 2004). 4

[JH15]   JANKOWSKI J., HACHET M.: Advances in interaction with 3d environments. In *Computer Graphics Forum* (2015), vol. 34, Wiley Online Library, pp. 152–190. 5

[KALF17]   KEIRIZ J. J., AJILORE O., LEOW A. D., FORBES A. G.: Immersive analytics for clinical neuroscience. In *2nd Workshop on Immersive Analytics* (2017), IEEE. 2, 10

[KDC*03]   KENNEDY R. S., DREXLER J. M., COMPTON D. E., STANNEY K. M., LANHAM D. S., HARM D. L.: Configural scoring of simulator sickness, cybersickness and space adaptation syndrome: Similarities and differences. *Virtual and adaptive environments: Applications, implications, and human performance issues* (2003), 247. 6, 8

[KH14]   KUHLEN T. W., HENTSCHEL B.: Quo vadis cave: does immersive visualization still matter? *IEEE Computer Graphics and Applications 34*, 5 (September 2014), 14–21. doi:10.1109/MCG.2014.97. 2

[KMLM16]   KWON O.-H., MUELDER C., LEE K., MA K.-L.: A study of layout, rendering, and interaction methods for immersive graph visualization. *IEEE Transactions on Visualization and Computer Graphics 22*, 7 (July 2016), 1802–1815. doi:10.1109/TVCG.2016.2520921. 2

[LJKM*17]   LAVIOLA JR J. J., KRUIJFF E., MCMAHAN R. P., BOWMAN D., POUPYREV I. P.: *3d user interfaces: theory and practice.* Addison-Wesley Professional, 2017. 3

[MBJS97]   MINE M. R., BROOKS JR F. P., SEQUIN C. H.: Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 19–26. 2, 3, 9

[MK94] MILGRAM P., KISHINO F.: A taxonomy of mixed reality visual displays. *IEICE Transactions on Information and Systems 77*, 12 (1994), 1321–1329. 4

[Mun14] MUNZNER T.: *Visualization analysis and design*. CRC press, 2014. 1

[Pal13] PALLANT J.: *SPSS survival manual*. McGraw-Hill Education (UK), 2013. 6

[PEP∗11] POCO J., ETEMADPOUR R., PAULOVICH F. V., LONG T., ROSENTHAL P., OLIVEIRA M. C. F. D., LINSEN L., MINGHIM R.: A framework for exploring multidimensional data with 3d projections. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 1111–1120. 2

[PM00] PELLEG D., MOORE A. W.: X-means: Extending k-means with efficient estimation of the number of clusters. In *17th International Conference on Machine Learning* (2000), Morgan Kaufmann, pp. 727–734. 7

[SBSSB15] SHOVMAN M., BOWN J., SZYMKOWIAK A., SCOTT-BROWN K. C.: Twist and learn: Interface learning in 3dof exploration of 3d scatterplots. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), ACM, pp. 313–316. 2

[SFR01] SCHUBERT T., FRIEDMANN F., REGENBRECHT H.: The experience of presence: Factor analytic insights. *Presence: Teleoperators and virtual environments 10*, 3 (2001), 266–281. 6

[SG17] SARIKAYA A., GLEICHER M.: Scatterplots: Tasks, data, and designs. *IEEE Transactions on Visualization and Computer Graphics* (2017). 5

[SM06] SPIRLING A., MCLEAN I.: The rights and wrongs of roll calls. *Government and Opposition 41*, 4 (2006), 581–588. `doi:10.1111/j.1477-7053.2006.00204.x`. 4

[SZK17] SIMPSON M., ZHAO J., KLIPPEL A.: Take a walk: Evaluating movement types for data visualization in immersive virtual reality. In *2nd Workshop on Immersive Analytics* (2017), IEEE. 3

[TLN17] THEART R. P., LOOS B., NIESLER T. R.: Virtual reality assisted microscopy data visualization and colocalization analysis. *BMC Bioinformatics 18*, 2 (2017), 64. 4

[UKF∗17] USHER W., KLACANSKY P., FEDERER F., BREMER P.-T., KNOLL A., ANGELUCCI A., PASCUCCI V.: A virtual reality visualization tool for neuron tracing. *IEEE Transactions on Visualization and Computer Graphics* (2017). 2

[WM08] WARE C., MITCHELL P.: Visualizing graphs in three dimensions. *ACM Transactions on Applied Perception (TAP) 5*, 1 (2008), 2. 1, 2

[WRFN18] WAGNER FILHO J. A., REY M. F., FREITAS C. M. D. S., NEDEL L.: Immersive visualization of abstract information: An evaluation on dimensionally-reduced data scatterplots. In *Proceedings of the 25th IEEE Conference on Virtual Reality and 3D User Interfaces* (March 2018). 2, 3, 4, 6, 7, 8, 10

[YHD∗14] YAO R., HEATH T., DAVIES A., FORSYTH T., MITCHELL N., HOBERMAN P.: Oculus vr best practices guide. *Oculus VR* (2014). 4

[ZBM∗17] ZIELASKO D., BELLGARDT M., MEISSNER A., HAGHGOO M., HENTSCHEL B., WEYERS B., KUHLEN T. W.: buenosdias: Supporting desktop immersive analytics while actively preventing cybersickness. In *2nd Workshop on Immersive Analytics* (2017), IEEE. 3, 4

[ZHF∗16] ZIELASKO D., HORN S., FREITAG S., WEYERS B., KUHLEN T. W.: Evaluation of hands-free hmd-based navigation techniques for immersive data analysis. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)* (March 2016), IEEE, pp. 113–119. `doi:10.1109/3DUI.2016.7460040`. 3

[ZWB∗17] ZIELASKO D., WEYERS B., BELLGARDT M., PICK S., MEIBNER A., VIERJAHN T., KUHLEN T. W.: Remain seated: towards fully-immersive desktop vr. In *2017 IEEE 3rd Workshop on Everyday Virtual Reality (WEVR)* (March 2017), IEEE, pp. 1–6. `doi:10.1109/WEVR.2017.7957707`. 2, 3, 4, 10