

Hybrid Tactile/Tangible Interaction for 3D Data Exploration

Lonni Besançon, Paul Issartel, Mehdi Ammi, and Tobias Isenberg, *Senior Member, IEEE*




Fig. 1. Tangible and tactile interaction for 3D visualization: (a) tangible manipulation of a cutting plane in the visualization; (b) seed placement for particle tracing; and (c) tactile manipulation of a cutting plane in the visualization.

Abstract—We present the design and evaluation of an interface that combines tactile and tangible paradigms for 3D visualization. While studies have demonstrated that both tactile and tangible input can be efficient for a subset of 3D manipulation tasks, we reflect here on the possibility to combine the two complementary input types. Based on a field study and follow-up interviews, we present a conceptual framework of the use of these different interaction modalities for visualization both separately and combined—focusing on free exploration as well as precise control. We present a prototypical application of a subset of these combined mappings for fluid dynamics data visualization using a portable, position-aware device which offers both tactile input and tangible sensing. We evaluate our approach with domain experts and report on their qualitative feedback.

Index Terms—Interaction, tactile input, tangible input, 3D data visualization.

1 INTRODUCTION

Interactive data exploration has long been an essential aspect of the visualization of 3D datasets. Traditionally, researchers have been investigating both dedicated interactive visualization platforms such as immersive VR settings [12, 46, 77] and traditional workstations. While the former rely on dedicated 3D input devices such as wands, gloves, or 3D tracking, the latter make use of either desktop-based 3D input devices such as 3D mice or the traditional mouse+keyboard setup. Both of these interaction settings (VR and workstation) have a long tradition and continue to be important. Yet people have increasingly easy access to novel display and computation environments such as tablet computers and large displays. In addition to traditional ones, these offer new interaction paradigms such as tactile and tangible input.

Research has shown that these tactile and tangible input paradigms have many benefits for effective and efficient interaction, in particular for 3D data exploration (e.g., [19, 26, 35, 87]). Yet, they are quite different from each other: tactile input benefits from its directness and a resulting perception of control and precision of interaction [81, 87], while tangible input offers an integrated, multi-sensory, and intuitive 6DOF control due to its similarity to day-to-day interaction with real objects [17, 33, 48]. The development of portable position-aware devices offers opportunities to use a tablet for tangible input in addition to the usual tactile input. Indeed, a device capable of tracking its own position in 3D space and interacting with a digital environment fulfils the four requirements for tangible interfaces as defined by Ullmer and Ishii [74], while at the same time providing a display with tactile sensing.

One of the main benefits of both input paradigms is that the input sensing and the data display can be integrated into a single device—

Google’s Tango tablet even supports both input modalities. This sensor integration not only allows the devices to be used for data exploration by themselves, but also allows them to be integrated into traditional data exploration environments such as immersive settings (e.g., [12, 44, 77]). These devices allow us to take a large step toward an interaction continuum [29] in which different input and output modalities can be used for data exploration, depending on the setting and user needs.

Yet, it is still unclear how this transition between the different input modalities could and should be realized in practice, in particular due to the different characteristics of tactile and tangible inputs. While several mappings for the two paradigms have been explored in the past [54, 69], their respective benefits and challenges with respect to 3D data exploration remain uncertain. We thus investigate their integration in one device and the resulting possibilities and potential interaction mappings for common 3D data exploration tasks. Based on the analysis of a field-study and on follow-up interviews of five fluid dynamic researchers, we focus on a subset of the potential interaction mappings to provide interaction for common fluid-dynamic analysis tasks based on the Google’s Tango tablet as the interaction device.

Our contributions of this paper are thus threefold. First, we contribute an understanding of how the two interaction paradigms can be combined to benefit from their inherent characteristics in the context of 3D data exploration by discussing the design space for possible interaction mappings. Second, based on this understanding we propose a design of hybrid mappings to achieve common 3D visualization tasks. In particular, we focus on mappings that make tactile and tangible inputs complementary as well as investigate the resulting interaction accuracy and constrained control. Third, we evaluate a subset of these hybrid mappings and compare them to touch-only or tangible-only approaches in a qualitative evaluation with fluid dynamic experts. Our results inform the creation of hybrid and complementary mappings between tactile and tangible input and paves the way towards a more complete interaction continuum for scientific visualization.

2 CONTEXT AND RELATED WORK

Within the field of 3D interaction [6, 25, 38], our work relates to tactile, tangible, and mixed input techniques to explore 3D scenes. We review relevant work next with a focus on interactive 3D visualization.

- Lonni Besançon is with Inria Saclay and Univ. Paris Saclay, France; e-mail: lonni.besancon@gmail.com.
- Paul Issartel is with Univ. Paris Saclay, France; e-mail: paul.issartel@u-psud.fr.
- Mehdi Ammi is with Limsi/CNRS, France; e-mail: mehdi.ammi@limsi.fr.
- Tobias Isenberg is with Inria, France; e-mail: tobias.isenberg@inria.fr.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

2.1 Tactile Input and Its Use for 3D Data Exploration

Tactile input for interactive systems has been investigated for a long time [8] and has been popularized by the rise of mobile devices in the last decade. It has multiple advantages over other forms of input including an improved performance for certain tasks [43] while being compatible to mouse-based input for others [18, 62], its support of interaction collaboration awareness [27], its somesthetic feedback [59], its suitability for physically large displays [71], and its use as a communication channel when one is presenting visualizations to others [70].

The use of tactile input to control 3D scenes such as visualizations, however, requires a mapping from the 2D input surface to the 3D data space [28, 30]. Several basic interaction techniques have been proposed, including those for the exploration of 3D visualizations. For example, Coffey et al. [10] designed a set of tactile interaction mappings for virtual reality (VR) contexts, combining a large stereoscopic screen with a touch-enabled tabletop display using a world-in-miniature metaphor. Their widget-based interaction techniques let users navigate the 3D visualization (rotation, translation, scaling), position slicing planes, place annotations, select data subsets, and plan camera paths. Klein et al. [44] presented a similar design study, but this time for monoscopic-only projections of fluid dynamics data. They also provide navigation and cutting plane interaction techniques, using only a single surface. In addition, they investigate particle seeding in 3D vector fields and the support of collaboration. A third example of a design study for the exploration of scientific data is Lundström et al.'s [49] virtual surgery table, also intended for collaborative setups. It supports 6 DOF navigation as well as additional exploration techniques such as the exploration of slices from medical imaging datasets.

In addition to these systems and design studies, a number of additional tactile interaction techniques have been proposed. For example, Cohé et al.'s tBox [11], Reismann et al.'s 3D-RST [56], and Yu et al.'s FI3D [87] provide dataset navigation facilities. Fu et al.'s powers-of-10 ladder [19] provides scale navigation at different levels, and Yu et al. [85, 86] suggest context-aware spatial selection techniques.

Although these tactile interaction techniques offer the benefits described above, tactile input also has some issues compared to traditional input devices. In particular, finger fatigue [5], fat finger, and a higher sensitivity to noise in the input data [5, 73] are issues that one has to account for when selecting or designing a tactile interaction technique. It is also worth noticing that, although pinching is used and perceived as a natural gesture for zooming, the gesture cannot at the same time provide z -translations [22, 23]. In our work, we thus base the implementation on some of these methods, but also learn from the insights reported by their respective authors and propose slightly adjusted designs.

2.2 Tangible Input and Its Use for 3D Data Exploration

Tangible User Interfaces (TUI) and especially the class of TUIs called Graspable User Interfaces aim at taking advantage of people's natural skills for manipulating their physical environment [17, 32, 33]. While many TUIs use tangible props as both physical representation and means of interaction, several TUIs focus more on the input aspect [26, 34–37, 67] by considering the tangible props as *handles*. Tangible input inherently offers 6 integrated DOF per prop. In our case, we are thus restricted to a single set of 6 DOF as we work with a single device.

Tangible interaction has been shown to be more engaging [73] than other forms of input and to provide rich feedback [89]. There is also evidence that it requires very little mental effort to use [5]. It has been reported to be preferred by users when compared with other interaction means [5, 89], even though it is still unclear whether this reported preference was obtained because of a novelty effect [5]. Tangible interaction can be a source of fatigue and exit-errors [5] and users have reported feeling that it was less reliable than mouse or touch inputs [5, 89]. Furthermore, 3D data space navigation requires zooming which is something that a tangible handle does not naturally provide.

Two general approaches have been proposed for tangible interaction in a visualization context: either a tangible object serves as an input device while the result of interaction is displayed on an external screen, or the result is displayed on the tangible object itself. In the first category, Hinckley et al. [26] use a rectangular acrylic surface as a

slicing tool for neuro-surgical datasets and display the sliced dataset on an classical external display. They combine this approach with a head prop to manipulate the dataset as well as clutching through a foot pedal (for the head-prop) and a button directly located on the plane (for the cutting plane). Similarly, De Guzman et al. [13] use a fork as a metaphor for the slicing plane to help children navigate through a 3D virtual model of the human body. The fork can be attached to a mechanical arm (thus providing a form of clutching) or can be freely manipulated by hand. They represent the dataset by either a 2D or 3D physical body model but its orientation in physical space is not linked to that of the virtual dataset. In contrast, Qi et al. [55] provide a generic shape (a cube) to represent their dataset and propose cutting plane with either a tracked pen (that defines the plane's normal) or a tracked square frame similar to De Guzman et al.'s solution [13]. Like Guzman et al., Mulder et al.'s [52] approach relies on a fixed tracking system and a fixed display setup and maps the orientation of a tracked pen to that of a virtual slicing plane, while Schkolne et al. [60] use tracked props to position elements in 3D space for molecular visualization. A more lightweight tracking approach is used by Jackson et al. [37] whose tangible interface consists of a printed 2D barcode that can be rolled into a pen-like object whose location and orientation can be easily captured using a camera—to be used, for example, for fibertract exploration. While Issartel et al.'s [35] data exploration system also relies on simple fiducial tracking to control the cutting plane with a tangible pen, they integrate this approach in a portable and affordable augmented reality setup that uses a tablet computer to view the dataset and track the interaction props. As an alternative to pen-based control, they also investigate the use of a cutting plane that is slightly offset from the tablet toward the data. The last two approaches have the benefit that a dedicated 3D tracking system is no longer necessary and thus avoid the calibration and maintenance issues that otherwise affect such setups. We use a similar approach in form of a spatially-aware tablet computer which can track its location and orientation in space.

The second approach for tangible interaction with 3D visualizations is to display (at least a part of) the data on the tangible device itself. For example, Spindler et al. [68] track the location of a small tangible surface (the PaperLens) over a tabletop display and project visual information on this surface as it is used to slice through the data. Song et al. [67] use a tablet computer instead as a tangible cutting plane and show the resulting data slice both on the tablet and on a large vertical display. On the other hand, Bertelsen et al. [4] use a full monitor mounted on a mechanical arm both for tracking and support, and show the respective data slice on the display depending on its location and orientation. In a way, all of these are variations of an earlier approach by Konieczny et al. [45] that also allowed users to bend the tangible cutting plane and show the appropriate intersection with the data.

While we believe that all these techniques are indeed useful, they fail to provide other means of input in addition to the tangible 6 DOF control without the use of additional devices. We thus base our design on displaying a view on the tangible, but go beyond tangible interaction by using a position-aware table to also provide tactile input possibilities.

2.3 Combinations of Tactile and Tangible Interaction

From the respective challenges and benefits of both input modalities and results from a previous study [5], it appears that the touch and tangible modalities are complementary paradigms and can be combined. Past work that combines tactile and tangible input in a single system is largely found in tangible additions to tabletop displays where props are tracked using fiducial markers. Good examples of this approach are Jorda et al.'s [40] prototype for live music performance as well as Al-Megren and Ruddle's [1] setup for abstract data analysis. An example of such a combination for spatial 3D visualization is Sultanum et al.'s [69] table-based system for exploring geologic reservoir data. They use tangible props for detailed data read-out and parameterizing a focus+context view, while tactile input is used for regular data navigation as well as for dedicated exploration techniques such as dataset splitting and layer peeling. However, most existing applications (mainly based on TUIO [41]), do not take full advantage of the physicality of tangible interaction as the tangible props remain on the tabletop—only their

2D position and orientation are used. We investigate, in contrast, the combination of tactile with full 6 DOF tangible input.

A previous step in this direction was Olwal et al.'s [54] use of a spatially-aware small display device on a large tabletop surface. The small device was tracked in 2D and could thus show a section of the data displayed on the tabletop, but at much higher resolution. Tactile input was possible both on the tabletop display and the small device to explore the data. Taking this concept into 3D space, López et al. [48] investigated the use of tactile input on a mobile device for 3D visualization, with both a stereoscopic view of the data and the mobile device's monoscopic view. In their study, they included a tangible interaction mode in which the tablet's orientation controlled the data view, yet restricted to 3D orientations. We extend such interaction to use the tangible device's full physicality, including full 6 DOF interaction.

Our approach also builds on earlier work by Watsen et al. [80] questioning whether 2D or 3D interaction is best to integrate tactile PDA interaction in CAVE environments. Similarly, as early as 1993, Fitzmaurice [16] and later Rekimoto et al. [57] proposed to use position-aware PDAs as a magic lens to interact with the physical world. Schmalstieg et al. [61] used transparent props and a pen that were tracked to augment the interaction space of a virtual table. Later, Miguel et al. [50] used a tracked PDA to facilitate interaction in CAVEs: users moved the PDA in 3D to get a suitable "captured" view and then selected a 3D object with a tactile input. Similarly, Yee's peephole displays [83] combined position-aware displays with pen input and applied them to three different applications scenarios. We mainly based our approach on the last three but, instead of using pen interaction we explore tactile interaction in the context of 3D manipulations—as previously done by Tsang et al. [72]. Recently, Bergé et al. [3] compared tactile and tangible interaction with a smartphone to explore 3D public displays.

3 TACTILE AND TANGIBLE DATA EXPLORATION

To better understand a potential integration of touch and tangible inputs we start by discussing the interaction tasks needed for spatial 3D data visualization, then analyze the resulting design space for the two input modalities, and finally motivate our prototypical implementation.

3.1 Interaction Tasks in Spatial 3D Data Visualization

Interaction tasks for the exploration of data visualizations have been analyzed in detail in the past (e. g., [7, 58, 64, 84]). Shneiderman [64], for instance, describes abstract tasks such as getting an overview, zooming, filtering, finding and selecting details or data subsets, discovering relationships, and interacting with the data exploration history. While others [7, 58] discuss the multiple levels of granularity in tasks concepts for visualization, Yi et al. [84] provide a synthesis of abstract tasks based on a literature analysis. They include exploration, selection, reconfiguration, encoding, abstraction/elaboration, filtering, and discovery of relationships as well as an "other" category.

Applied to visual representations of 3D spatial data (rather than abstract data), the first four of these abstract tasks can be loosely mapped to 3D navigation (translation, rotation, zoom), data selection, and parameterization of the visualization mapping. Aspects of reconfiguration and encoding that change the spatial mapping of visual representation are rare due to the spatial data's inherent mapping to 3D space. Abstraction/elaboration, filtering, and relationship discovery also exist as well as "other" interaction tasks. For example, the use of a cutting plane, a drilling probe, or isosurface rendering could be seen as a form of data abstraction, while seed point placement adds a visual encoding of vector field's dynamic aspects. Overall, we thus need the following fundamental interaction techniques in most visualization systems of 3D spatial data [42]:

- 3D data space/view navigation: 3 DOF translation, 3 DOF rotation, 1 DOF uniform zooming; potentially with the possibility of constraining the interaction to specific DOF and/or align them to specific data dimensions,
- visualization styles/types adjustment and parameterization: selection of volumetric, iso-surface, or vector-based representations and their parameters,

- positioning/manipulating data exploration objects such as cutting planes (3 DOF) or probes such as drilling cores (2 DOF),
- 3D picking or selection of data subsets for further analysis,
- specifying/manipulating 3D points and other primitives for particle seeding, picking, or path planning,
- generating data read-outs or measurements, and
- temporal navigation.

Beyond these fundamental tasks, a smaller or larger set of other techniques may be needed depending on the data and application domain. However, the described set can be seen as a common set that is needed in most applications including our fluid dynamics application domain.

Several mappings for such data exploration tasks have been proposed for both input modalities (see Sect. 2). We can now analyze how the different tasks can be mapped to the two modalities, thus creating a design space for tactile and tangible control of 3D spatial visualizations.

3.2 The Design Space for Tactile and Tangible Control

Table 1 shows the design space and gives examples for specific types of control based on the related work. For navigation tasks we assume that the tangible interaction device is able to display a visualization of the data on its own touch-enabled screen, such a tablet computer, in order to investigate combinations of both input paradigms. However, we also include cases where dedicated tangible props are used for additional interactions as these could be combined with the tablet device. For tactile input we specifically point out the minimum number of hands necessary for a particular technique. For example, even though Coffey et al.'s Slice WIM widgets [10] were demonstrated in their video in a bimanual fashion, the widget could also be used with several fingers of the same hand and we thus classify it as unimanual interaction. In contrast, while Yu et al.'s FI3D [87] can be used in a unimanual fashion, their constrained interaction modes necessarily require bimanual input due to the widget's design.

As Table 1 shows, tangible interaction based on physically moving a mobile device in 3D space facilitates a direct mapping of up to 6 input DOF to the respective output DOF. Scaling would need to be supported separately, but mode switches facilitate the control of data space navigation, cutting plane manipulation, seed point placement, spatial selection, and data readout. Tactile interaction, in contrast, uses a variety of mappings and widgets that are either controlled uni-manually or bimanually. This essential distinction affects our hybrid interaction design because the need to hold the device during interaction severely restricts potential input from the non-dominant (carrying) hand [79].

In all manipulation interactions, precision can be controlled by setting specific control-display gains either explicitly (e. g., through physical or virtual widgets) or implicitly (e. g., velocity of the tangible device/finger or distance to the starting point). In addition, precise interaction is supported when specific interaction constraints are observed. Such constraints can also be specified though the use of specific widgets (physical [39] or virtual [11, 19, 87]), or, in the tangible case, by using the device's own orientation (e. g., constraining the manipulation along the nearest data axis) or its position (e. g., translating along the axis with the longest distance from the starting point of the interaction).

After having discussed the individual mappings, we can now discuss how a combination of tactile and tangible input can be achieved. We concentrate on the spatial direct manipulation tasks because tasks such as style setting and temporal navigation, on a tablet-based interface, are most flexibly realized using a tactile widget. Table 2 summarizes some of the possible combinations. In contrast to the interactions in Table 1, we place here particular emphasis on ensuring that combinations do not conflict with each other. Table 2 shows that hybrid techniques exclusive to 3D data space navigation or cutting plane manipulation are possible, but do not make much sense. However, with situations requiring the two tasks it becomes meaningful to assign the two input modalities to one of the two mapping alternatives.

Additional challenges arise when an additional 3D point needs to be specified—either for picking/seed point placement or for data readout—because both input modalities are already mapped to either view or cutting plane manipulation. If this point is specified using the tactile

Table 1. Design space for tangible and tactile control of 3D visualizations. For the uni- and bimanual usage we specify the lowest mode that a technique can be used in, not how it was demonstrated in publication videos.

task	DOF	tangible	tactile
3D data space navigation	7	6 input DOF mapped directly to 6 output DOF; absolute or relative motion, rate control [36] (uniform scale cannot be provided by tangible interaction)	widget-based mapping: unimanual [10, 11, 19, 87] or bimanual [19, 87]; purely posture-based mapping [56] (uni- or bimanual)
cutting plane manipulation	3	same as 3D data space navigation	same as 3D data space navigation + posture/widget combinations [44]
integrated data space + cutting plane manipulation	7+3	same as 3D data space navigation; through either explicit mode switches on the tangible device or multiplexed TUI interaction with more than one device	same as 3D data space or cutting plane navigation, using widget- or posture-based mode selection (e. g., [10, 44])
style setting / parameterization	n/a	dedicated tangibles (e. g., [51, 53, 76])	typically via separate widgets (e. g., [44, 69]; typically uni-manually)
picking / seed point placement	3	same as 3D data space navigation (direct 3DOF pointing) + physical or virtual button to activate	tactile 3D positioning such as balloon positioning [2] (e. g., [10]) or 2DOF pointing by casting a ray on a cutting plane [44] (both uni-manually)
spatial selection	n/a	“tangible brush” with button-based moding and additional size controls	context-aware selection techniques based on 2D projected view (e. g., [85, 86], bi-manually)
data read-out	3	same as 3D data space navigation (direct 3DOF pointing) + physical or virtual button to activate	2 DOF pointing based on 2D projection, ray-casting on a cutting plane, or view-aware picking [82] (uni-manually)
temporal navigation	1	tangible sliders [39, 75]	slider-based widget (uni-manually)

input, positioning a cutting plane is necessary in order to perform a ray-casting so that the finger’s position can be interpolated to a 3D position. In this case, the positioning of the cutting plane and dataset are assigned to touch or tangible input. In the alternative case of tangible specification, the cutting plane is not necessary to provide a 3D location, so the tactile input can control the data view.

Similar difficulties arise when considering spatial selection, as this interaction also has to be mapped either to the tangible or to the tactile modality. Here, no additional cutting plane is necessary. In the first of these cases, the tactile modality can thus be used to specify the data view while the tangible is used as a “tangible brush.” Physical or virtual buttons can be used to activate the brush mode and control its size. In the latter case, the view on the data can be adjusted with the tangible input, while the tactile input is used to specify context-aware selections.

3.3 Field Study and Prototypical Implementation

This design space can now inform the creation of actual hybrid interaction mappings for a specific domain. In this paper we focus on supporting data exploration for fluid dynamics researchers. In order to better understand their needs, and following Shneiderman’s [65] recommendations, we carried out a field study with five experts (3 males; 2 females; ages 22–44; mean of 13.6 years of professional experience). For this purpose we visited them in their lab and individually observed their normal working procedures as they analyzed new datasets. Each

Table 2. Design space for hybrid tactile and tangible interaction.

task	hybrid tangible + tactile control
3D data space navigation	tangible and tactile input mapped separately to location and orientation; mapping of subsets (e. g., only x -/ y -rotation to tactile) also possible
cutting planes	either tangible or tactile control, not both simultaneously
integr. data space + cutting plane	tangible and tactile input mapped separately to data space navigation and cutting plane manipulation
picking / seed point placement	as with integrated data space + cutting plane interaction, specifying the 3D point can be done with either tangible or tactile input: use of either an explicit user-controlled mode switch for tangible input or of the intersection of the finger’s position with the cutting plane’s position (ray-casting) for tactile input
spatial selection	e. g., tangible input to set data view and tactile input for context-aware selection; or tactile input to set the view and tangible input for selection using a “tangible brush” metaphor
data read-out	same as picking/seed point placement

observation session was video-recorded for further analysis and was followed up with a semi-structured interview in which we asked about the steps involved in dataset analysis and the interaction features that participants thought were lacking in current software. One result from this field study was the realization that an essential part of understanding new datasets relies on being able to manipulate cutting views of the data. Fluid dynamics researchers do not perform much translations on the data itself, but frequently rotate it in order to get better views and then use several cutting planes to get an understanding of its internal structure. When analyzing new datasets, experts first want to obtain a general understanding of the dataset—particularly through cutting planes—and then focus on understanding how the flows evolve spatially and temporally. The latter can be evaluated thanks to particle seeding.¹

Based on this analysis, we decided to place our focus on several tasks in 3D space as the most fundamental building blocks for 3D data exploration. In line with the described design space, we decided to support the following tasks: 3D data space navigation, cutting plane manipulation, and seed point placement. Our work particularly focused on three points: (a) supporting these tasks in a way that the interaction mappings do not conflict with each other, (b) providing not only hybrid but also both purely tactile and purely tangible mappings so that they can be compared to each other, and (c) adding support for control-display gain control and DOF constraints in our mappings to provide a fine-grained control of manipulation.

Based on Table 1, we used a simple 1:1 mapping for the tangible-only case in which translations and rotations are captured by the Tango tablet and directly mapped to respective manipulations of the dataset. Zooming was not possible in the tangible-only setting. For the tactile-only case we used a simple posture-based mapping that assigns motions of one finger to arcball rotations (e. g., [19, 44]), motions of two fingers to 2D rotations, scaling, and translations (2D-RST, e. g., [87]). These mappings were chosen in part because unimanual input is easily supported and because they relate to common 2D tactile interaction mappings. Since many people are also used to translating items only using two fingers placed and kept next to each other, our 2D-RST mapping only becomes active if the distance between two simultaneous touches grows beyond a threshold (i. e., 300 pixels, equivalent to 29 mm). Due to the scaling gesture, translation along the z -axis was not possible with this mapping but could be achieved through combinations of x -/ y -translations and arcball rotations.

For hybrid mappings we implemented four combinations. In the first two, we explore how a combination of tangible and tactile input can be used for either view navigation or cutting plane manipulation alone

¹We contrast the traditional interaction software currently used by experts in the discussion section of this paper.




Fig. 2. Our interface controlling the tactile input and the mode switching for the tangible control.


(Fig. 1(a), (b)). For this study, instead of mapping dataset orientation to one modality and location to another, as suggested in the first row of Table 2, we mapped both *both* location and orientation to a single modality. In the first mapping both modalities controlled the data volume, and in the second mapping both controlled the cutting plane. This way a temporal multiplexing made it possible to switch between the two input modalities for a given interaction, allowing us to investigate which mappings would be preferred by participants. In the other two combinations we explored the mappings mentioned in row 3 of Table 2: either tactile input mapped to view manipulation and tangible input mapped to the cutting plane manipulation, or the other way around.

We also wanted to investigate at least one mapping that requires the specification of a 3D point. We thus decided to explore 3D seed point placement (Fig. 1(c), row 4 in Table 2). From the different alternatives we implemented tactile input for specifying a seeding point because the act of touching is a good metaphor for the placement of objects. Consequently, we realized the two alternative mappings for integrated view specification and cutting plane arrangement with tactile and tangible input—using the considerations discussed in Sect. 3.2—and based the seeding on the specified cutting plane.

To facilitate the necessary mode switching, we implemented the interface shown in Fig. 2. Menus on the top allow users to load datasets and change general settings. The buttons on the lower right constrain the interaction to a particular coordinate axis of the dataset. The slider on the right manipulates the control display gain factor associated with both tactile and tangible input to provide experts with an explicit control over the accuracy of their interactions. The buttons on the lower left control the mapping, allowing people to enable or disable the two input modalities and to map them to data or cutting plane manipulation. These are system-controlled states as user-controlled moding is already used for seeding and to activate tangible input through clutching. Tangible clutching is achieved by pressing and releasing the red button on the upper left (i. e., located beneath the left thumb of the user when holding a tablet ‘normally’). Similarly, seeding point placement is achieved by placing the finger on the screen when the seed button is pressed.

4 OBSERVATIONAL STUDY WITH EXPERT USERS

To better understand the combined touch-tangible interaction and the use of the different possible mappings in practice, we conducted a second observational study with domain experts, this time with our actual prototype implementation. We were interested in their general opinion about such a hybrid interaction style, their understanding of the mappings, the way they transition between different modalities, how well the chosen mappings support their data exploration goals, and how they made use of those interaction capabilities in practice. We used an observational strategy like several visualization researchers before us (e. g., [19, 44, 48, 49, 69]). We specifically decided against a classical usability study for several reasons. First, interaction mappings such as the ones we study are highly complex and are not easily studied by means of completion times and error metrics. Second, our pool of experts confirmed they do not consider data exploration and understanding as a task to be completed as fast as possible—in fact, a slower technique



(a) Isosurface view.

(b) Volumetric view.

Fig. 3. Two views of the FTLE dataset used in the study.

may be equally good or better when trying to understand unknown data. Moreover, as emphasized by Carpendale [9] and Greenberg and Buxton [21], classical quantitative studies can prevent the desired insights from users on the suitability of the different interaction techniques, e. g., by muting creative ideas or meaningful critique. Based on Lam et al.’s [31, 47] categorization of evaluation strategies, we thus conducted a combination of a User Experience and VDAR evaluation: We asked our experts about how our techniques support their data exploration needs to understand how they can be improved and/or integrated in their work practice. Both of these evaluations have previously been conducted by means of observations and questionnaires/interviews (e. g., [15, 20, 66]).

4.1 Participants

We recruited 7 researchers (all male; ages 23–61 years, mean: 35.7, median: 32, and SD: 13.1) from a fluid dynamics lab whose work focuses on volumetric flow data. Our unpaid volunteers had 1–38 years (mean: 12.9, median: 9, and SD: 12.9) of post-Master’s professional experience. All participants were used to interacting with 3D datasets in the course of their work using typical mouse+keyboard interaction. All were familiar with tactile interaction on their smartphones, and two had previously participated in tactile interaction experiments for 3D visualization. Only one reported to be familiar with the term tangible interaction and had been using such techniques before in experiments for classical 3D manipulations and 3D scientific visualization.

4.2 Apparatus

Our setup included the 7 inch Google Tango tablet² (370 g, 1920 × 1200 pixel resolution) and a 55 inch (139.7 cm diagonal) vertical screen with a 3840 × 2160 pixel resolution. Users were asked to stand in front of the large display throughout the experiment. The external screen showed larger views of the data as well as additional visualization elements in order to address the occlusion issue of tactile interaction [24, 63]. As dataset, we used a domain-specific Finite Time Lyapunov Exponent (FTLE) scalar field (Fig. 3) with its associated vector field. Our implementation uses the VTK library³ to load and process the datasets. The dataset was rendered using OpenGL ES 2.0 on the tablet and OpenGL 3.0 on the vertical display. Communication between the devices used the UDP protocol and we sent absolute transformation matrices to ensure that packet loss would not be critical for the display/tablet synchronization. Elaborate computations and visualizations were restricted to the vertical display. For instance, the input information for particle tracing/seeding was captured by the tablet but processed and rendered by the vertical display’s computer, and the external display also showed a 2D view of the slicing plane to ease the understanding of the sliced data. Our prototype is modular: the vertical display is handled by a PC running Ubuntu, while the tablet code can be adjusted to be fully functional with any position-aware device. Others can thus build on our work to create other hybrid interaction mappings with other devices.

²<https://www.google.com/atap/project-tango/>

³<http://www.vtk.org/>

4.3 Study Design and Tasks

We started by telling participants the purpose of the study, the setup, and the handling of the tablet. The study was divided into a training stage, two main tasks, and a final session consisting of a questionnaire and an interview. Overall, each study session took between 18 and 50 minutes, not counting the questionnaire and interview parts.

Training. Using a simple 3D shape as a training dataset, we introduced the participants to all interaction techniques and gave details on the user interface in the form of a tutorial. Participants were also encouraged at this point to ask questions to the experimenter. This training session is not taken into account in our analysis of the results. Because we wanted to understand the entire interaction spectrum between tactile-only and tangible-only input, we encouraged the participants to explore all possible mapping combinations and ensured that this happened at least during the training stage.

First Task: Exploration. Once the participants declared that they were ready for the actual study, we switched to the FTLE dataset. We asked them to navigate the data and to try to understand it—as they would normally do in their work—by using any interaction technique provided by our interface. While interacting with the data, the participants were encouraged to ask questions and especially to think-aloud: explain what they liked, what they disliked, what they did not understand, or when they felt that unexpected things happened. We took notes about the oral insights given by participants. These notes were completed after each session based on video recordings and log data.

Second Task: Particle Tracing. After a thorough exploration of the data using the different mappings, the experimenter enabled the interaction mode for particle tracing⁴ and explained how it could be performed. Participants were then asked to use it to gain additional insights on the datasets. Similar to the previous section, they were encouraged to explore all the different interaction mappings and to use the think-aloud protocol to allow us to capture their reasoning and preferences. In both the 3D navigation and the particle-based exploration parts, participants were free to explore the dataset until they had gathered enough knowledge.

Questionnaires. In the final part of each session we asked participants to fill out a questionnaire that asked about the effectiveness, usability, and intuitiveness of each interaction mapping. Participants were asked to provide their opinions using 5-point Likert scales. In addition, we asked about their overall preference for any of the different interaction mappings and the rationale for their choice. Finally, we conducted a semi-structured interview to discuss their experience in order to understand the pros and cons of each technique, why and how they would like to use them, and possible improvements. We based the interview on Lam et al.'s [47] main questions for User Experience (UE) evaluation. An experimenter took notes which were augmented from the captured video of the session.

5 QUANTITATIVE RESULTS

Next we report the quantitative data we captured during the study. Due to variations in the overall study duration we report time ratios as suggested by Dragicevic [14]. Even though our participant pool was small, it was composed of domain experts and observational studies of interactive systems with experts can yield important insights (e. g., [49]). Nevertheless, we report our results using estimation techniques with simple effect sizes and confidence intervals (instead of p -value statistics), consistent with recent APA recommendations [78]. These quantitative results are no statistical proof and serve to support the qualitative feedback we obtained from our observations and interviews.

5.1 Relative Interaction Times

Fig. 4 shows the fraction of time the participants spent using tangible, tactile, and hybrid input mappings. There is strong evidence that, after the training, participants predominantly used the hybrid mapping (86% of the time on average). In addition to this overall usage, we also

⁴Particle tracing/seeding is a technique to explore vector fields datasets by generating a number of particles at a given 3D location. Each of them then follows the vector fields for this starting point.




Fig. 4. Ratio of time spent interacting in the different conditions. Error bars are 95% confidence intervals.




Fig. 5. Ratio of time spent using the tactile and tangible conditions while using the hybrid interaction. Error bars: 95% CIs.




Fig. 6. Distribution of time spent interacting with the different plane/data associations while using the hybrid interaction. Error bars: 95% CIs.

examined the distribution of tactile and tangible inputs in the hybrid interaction mappings (Fig. 5). The results show that our participants, on average, spent approx. 74% of their time interaction with the tangible and only 26% with the tactile input. We hypothesize that the flexible manipulations of tangible interaction is the reason why it is mostly used, which was confirmed by our interviews. Finally, we were interested in the use of the four different hybrid mappings described in Sect. 3.3 for view and/or cutting plane manipulation. Fig. 6 reports the relative times spent in each of the four mappings. The ratios of time spent with each combination of data and plane manipulation (Fig. 6) do not provide evidence for a dominant mapping. It appears, however, that mapping tactile input to the cutting plane and tangible input to the dataset was not used much by participants. We conjecture that this result was caused by the tablet being—due to its shape—an easier metaphor for the cutting plane than the data volume, leading participants to more directly associate it with the cutting plane rather than with the dataset when selecting the mappings. This was confirmed by participants.

5.2 User Preferences and Assessments

We asked experts to rate the tactile-only, tangible-only, and hybrid interaction modes using 5-point Likert scales according to several usability criteria. In addition, we asked them to rank the three different techniques by preference. The results, collected using a questionnaire, are shown in Table 3. Overall, our participants agreed that they were able to accomplish what they wanted, quickly achieve their goals, and that the three techniques were intuitive enough. It appears, however, that they needed some mental effort to use the techniques, an effect that could be attributed to the mappings being new to them.

6 QUALITATIVE OBSERVATIONS AND DISCUSSION

Our main goal is to understand the role of hybrid interaction in the context of 3D visualization. We thus not only need to evaluate the tested interaction techniques by themselves but also have to discuss how they compare to and can be integrated with the usual PC/workstation-based environment that experts typically use. Most rely on tools such as Paraview⁵ or MATLAB.⁶ Based on our initial field study as well as the present qualitative observations and, in particular, the comments of the participants from the think-aloud protocol we can make this comparison. We discuss the two general approaches in this section, together with the general comments from our participants on the interaction mappings and their reasoning for using particular interaction styles.

⁵<http://www.paraview.org/>

⁶<http://www.mathworks.com/products/matlab/>

Table 3. Results from Likert-based ratings for different statements, with values ranging from 1 (completely disagree) to 5 (completely agree). Ranking is expressed from 1 to 3, 1 standing for the preferred technique

statement	factor	mean	median	SD
I could do what I wanted.	tangible	3.57	4.00	0.79
	tactile	3.57	4.00	0.98
	hybrid	4.14	4.00	0.90
I could achieve my goals quickly.	tangible	3.57	3.00	0.79
	tactile	3.86	4.00	1.07
	hybrid	4.14	4.00	0.90
It could be used without much information.	tangible	3.57	4.00	0.79
	tactile	4.00	4.00	0.81
	hybrid	4.00	4.00	0.58
It required a lot of mental effort to use.	tangible	3.14	3.00	0.69
	tactile	2.71	2.00	1.11
	hybrid	3.57	3.00	1.13
overall ranking (1 to 3, 1=best)	tangible	2.57	3.00	0.53
	tactile	2.28	2.00	0.76
	hybrid	1.14	1.00	0.38

6.1 Preferences

Overall, the hybrid interaction was preferred by most participants as seen in Table 3—only one participant did not name it as his first choice but as his second. Participants reported that our approach “was way faster and way more natural. It was more engaging, making people want to try more things, in particular when using the seeding.” They reported that the tangible interaction “adds more options once you get used to it and also allows for axis-constrained interaction. It is more natural and provides more DOFs than a mouse,” and finally reported that our approach “is easier to use. Less powerful for now in the different options that it gives than the traditional Paraview interface, but it is only an implementation problem and could easily be solved. If it were, it would be used for teaching or sharing knowledge. It is more enjoyable to use, even though the mouse seems to provide more precision.” One participant said that the system, “is a reproduction of what [he] can do with a PC but with better interactions.” This statement was mirrored by P6 who stated that he “could achieve complex rotations or translation difficult to perform with a mouse.” Another participant (P2) stated that the prototype could be used as an “extra visualization tool” in the sense that he could carry most of the primary analysis with it, before switching to a PC/workstation for a more in-depth analysis that requires scripts, scatterplots, and mathematical functions that are better created with and manipulated on a PC/workstation.

6.2 Mappings

Overall, the two ‘redundant’ mappings (that associate both input modalities to the same interaction target) and the the tangible to plane/tactile to data mapping were used similarly often. Which of the three mappings was most frequently used still varied quite a lot between participants, and it appears that the choice between them is guided by personal preference—we could not identify particular reasons for a given usage pattern. When asked to explain why they used one over the other, participants said that it would correspond to their way of thinking. This leads, however, to the surprising observation that, contrary to what we expected, participants actually used ‘redundant’ mappings a lot. They all reported that the tactile modality was a way for them to adjust and fine-tune the final view/cutting-plane following large-grained changes with the tangible modality.

We believe, however, that a single study with one class of domain experts is not sufficient to fully understand the potential of these mappings. It would thus be useful to conduct a study with users from other backgrounds to better evaluate how tactile and tangible inputs can be combined when manipulating a single space or object. From our obser-

ations, it is also difficult to conclude whether a participant preferred to have a static cutting plane and a moving dataset, or the opposite. We conjecture that this is a matter of each person’s mental model, which may be related to the type of datasets they usually manipulate.

6.3 Accuracy and constrained interaction

Both tangible and tactile input are inherently susceptible to noise and imprecise input, which is partially exacerbated by the integrated control of several DOF at the same time. We thus specifically considered how to support constrained input and control-display gain in order to improve manipulation accuracy.

Axis-Constrained Interaction. The possibility to constrain interaction to a given axis was much appreciated by participants, in particular for mappings that relied on tangible input. Participants reported that the tangible condition was a good way to freely explore the dataset, but that being able to constrain the input allowed them to achieve interactions they usually perform in PC/workstation-based environments. Indeed, in our field study we observed that, when using Paraview, experts often place a cutting plane perpendicular to one axis and translate it along this axis through the entire dataset. The use of axis-based input constraints can thus be considered as an important element of both tangible and tactile input mappings, confirming suggestions from previous work [44, 49]. Moreover, P5 and P6 suggested to add buttons for placing data or planes at precise configurations without any manipulation to provide a better way to compare views with Paraview.

Accuracy. Even though participants could change the gain factor associated with the tangible and touch interaction, we noticed that this was rarely used. Indeed, P1, P2, P4, and P6 adjusted the gain factor twice, while Participants 3, 5 and 7 only used this feature once. P2, P3, and P4 reported, however, that they used the tactile modality to adjust orientations or positions obtained with the tangible modality. P2, P3, P4 and P6 especially mentioned the lack of accuracy they felt in the tangible condition, even though they did not use the widget we provided to control the interaction’s gain factor. This contrasts recent work which found a similar level of accuracy between tangible and tactile control for 3D manipulations [5]. From our interviews it also appears that having an explicit widget for setting the control-display gain was not appropriate, confirming the hypothesis stated by Issartel et al. [36]. One participant stated that he knew the widget could be used but did not think about using it and rather used the tactile modality to adjust the positioning. Perhaps we should thus focus on providing more implicit ways of controlling the gain factor while interacting. One participant suggested a *rate-control* approach—which is inappropriate for isotonic devices not providing a self-centering mechanism [88].

6.4 Particle Seeding

The ability to seed particles in the data volume to explore its vector component was also appreciated by the experts. Six participants mentioned this feature first among the aspects they particularly appreciated. Participants compared particle seeding to the PC/workstation-based exploration tools, where they place particle sources by editing a script that is then executed. To be able to adjust the particle source they need to edit the script and re-run it. In contrast to this rather crude form of data exploration, our approach allowed them to interactively adjust the placement based on the location and orientation of the tablet (when using tangible control) or the cutting plane (for tactile input). All experts reported that this technique was “engaging,” “easy to use,” and that it could be “easily and greatly” improved with a few additional visual features such as depth cues or colors to represent temperature, speed, etc.

6.5 Use of a Separate Vertical Display

Five experts reported that the large display was not necessary in a non-collaborative setup. They reported that the screen of the tablet was large enough. Indeed, we observed that these participants actually spent most of their time looking at the tablet. This is surprising since previous observations of a tablet-based interaction with a large vertical screen [48] (albeit using a stereoscopic view) showed a preference of participants to focus on the vertical display. One of them explained that he looked at the tablet because that was “were the interaction was

happening.” Even in the supposedly eye-free tangible condition these three participants kept looking at the tablet. Participants 4–6, however, used the external display a lot, saying that they would always have access to an external display at their workplace anyway. Still, they did not think that the display had to be this large and that a typical 24” display would be enough. Yet, all of them mentioned that a large display was necessary for demonstration/presentation or collaborative analysis and that, with a larger screen, the tablet could be used alone as an on-the-go device to integrate in their working procedures.

6.6 Integration into the Workflow

We wanted to ensure that our interaction technique would be useful in improving the workflow and we asked participants whether they would use it in their everyday work. P6 reported that with a longer learning phase “it would be a nice tool and that [he] was interested and willing to use it whenever [he] needs to roughly analyze 3D data.” P5 stated that he “would use it if it were improved with some extensions” (which we mention below). Overall, only P1 did not consider the tool as being fit for his practical work yet. Other participants welcomed the opportunity and imagined using it for presentations, collaboration, or teaching purposes. When asked whether they could integrate it with their classical interactions, five participants stated they would gladly use it if it synchronizes with their desktop station. Four participants also mentioned they would be interested in combinations of our hybrid technique with a mouse-based approach which worked well with classical software (to be used for further analysis). This way they could avoid a “going-back-and-forth” behavior and actually directly use our prototype for highlighting and selecting data—as it is “easier and more complete”—to further study it with mouse and scripts. We believe that these remarks highlight even further the need for an interaction continuum between new and classical interaction paradigms.

6.7 Suggested Extensions

Like many other qualitative studies, our experiment led to several suggestions for improvement. We discuss the most important ones:

- *provide separate “frozen” views on the large display*: participants wanted to be able to save frozen views on the large display and be able to interact with them at a later point (possibly using tactile interaction on the large display);
- *include the possibility to navigate in time*: datasets in scientific fields are often time-dependent and participants reported that in the seeding interaction they would like to see the influence of time on the particle propagation—a classical slider implementation can easily be added, but other possibilities such as tangible sliders [39] may also be interesting in a hybrid interaction context;
- *add widgets to obtain specific, well-defined views and cutting plane positions*: experts reported that they would like to have widgets to set data and cutting plane to specific positions as it is offered by their traditional PC-based software—a feature that they frequently use; and
- *consider interaction mappings for heterogeneous data*: while our dataset was relatively box-shaped with $100 \times 80 \times 54$ samples, extremely thin and long dataset may require adjusted interactions that better fit their overall shape; based on our study we could envision to address this issue by mapping large-scale motions to the tangible interaction while small-scale motions would be mapped to the tactile interaction.

6.8 Limitations

While our design space exploration and observational study have demonstrated that hybrid tactile/tangible interaction can support 3D data exploration tasks, our work still has some limitations that we want to summarize here. As previously mentioned, our field-study and observational study only focused on one class of domain experts—fluid dynamic namely. Even though fluid dynamic visualization is a scientific field in which experts are often confronted with the manipulation of multiple DOFs, there are other domains to consider that could give complementary/different insights on our hybrid interaction design space

exploration. Similarly, during both studies, experts were faced with a rather homogeneous dataset thus avoiding issues that could be encountered with extremely heterogeneous datasets. Our observational study also does not explain why certain mappings worked or were preferred, in particular due to the small number of participants. Similarly, our study does not explain how tactile and tangible input—when assigned to a single target—can be efficiently combined. Further studies with more participants are needed to sketch different possibilities.

7 CONCLUSION

We explored the design space for the combination of tactile and tangible inputs to control 3D spatial visualizations. We proposed several mappings and studied a subset of these mappings to understand how the combination of touch and tangible inputs could benefit fluid dynamics experts in their 3D visualization tasks.

Our observational study showed that our participants appreciated our prototype and that they found it better suited for primary 3D visualization tasks than a traditional mouse-and-keyboard setup. We saw surprising effects such as the participants using interaction mappings we initially did not think would be very useful. Our participants especially appreciated that a complex seed point placement task could easily be achieved by combining a tangible manipulation of the cutting plane and a ray-casting with the tactile input, thus demonstrating the potential of hybrid tactile-tangible interactions. The flexible seed point placement enabled them to use an exploratory data analysis style of vector fields. Participants also appreciated having the ability to constrain input to specific DOF. Our work is thus a first step towards combining these two complementary input techniques and providing an alternative data exploration platform for scientific visualization. Most importantly, however, we demonstrated that with current hardware it is possible to realize hybrid tactile/tangible interaction techniques that support fundamental 3D data exploration tasks and that no longer rely on external 3D tracking. Without the need for constant maintenance, calibration, and support that such hybrid interaction would normally require, it is thus now possible to make the proposed interactive data exploration techniques available to researchers in various domains.

Even though we focused on fluid dynamic experts, some of our findings are generalizable. First, hybrid interaction is easily understood, is largely preferred by participants, and is the most frequently used interaction mode. Second, it appears that the tablet, due to its shape, is a good metaphor for cutting plane manipulations, which affects the interaction mapping preference. Third, the combination with a large display does not seem to be necessary—regular displays (or none) can be enough. Finally, explicit accuracy adjustments are not useful for tactile/tangible interaction, other accuracy controls are needed. Moreover, while in our work we chose an observational approach that led to these insights, specific application domains may warrant quantitative studies, focusing on a task or domain that requires certain levels of accuracy or has time constraints. For example, doctors in an ER need to analyze medical scans quickly, surgeons need to understand medical images very precisely to plan or adjust a surgery. For such application domains we may thus need to reconsider the interaction design.

Naturally, several questions remain open. For example, even though experts reported a lack of accuracy—particularly in the tangible condition—surprisingly few of them actually used our display-control gain factor widget to adjust the sensibility of the interaction. A follow-up study should thus investigate how to better control the accuracy for both interaction modalities. Future work also includes the exploration of additional 3D visualization tasks and how they can better be supported with hybrid tangible-tactile interaction—focusing, in particular, on data read-out and temporal navigation. We also wish to investigate the use of our prototype in a VR environment to check that it can also improve interaction in this context. Our discussion of the design space exploration also revealed that some tasks such as spatial selection are currently only supported by tactile or mouse-based interaction. These tasks may also benefit from having tangible and hybrid counterparts. We will also further improve our interface to better suit the needs of domain experts and to thus be able to further explore its usage in a practical context. One particularly interesting aspect to investigate is

the support of collaborative work in different configurations, such as a joint data exploration by two or three colleagues or by a small team. In this particular case, the combination of a small device with a large vertical display (potentially showing a stereoscopic view) would be needed, and the latter could also be used to provide tactile input.

ACKNOWLEDGMENTS

We wish to thank Andre Spritzer for his insightful comments and our participants for their insightful feedback and their time.

REFERENCES

- [1] S. Al-Megren and R. A. Ruddle. Comparing tangible and multi-touch interaction for interactive data visualization tasks. In *Proc. TEI*, pp. 279–286. ACM, New York, 2016. doi: 10.1145/2839462.2839464
- [2] H. Benko and S. Feiner. Balloon selection: A multi-finger technique for accurate low-fatigue 3D selection. In *Proc. 3DUI*, pp. 79–86. IEEE Computer Society, Los Alamitos, 2007. doi: 10.1109/3DUI.2007.340778
- [3] L.-P. Bergé, M. Serrano, G. Perelman, and E. Dubois. Exploring smartphone-based interaction with overview+detail interfaces on 3D public displays. In *Proc. MobileCHI*, pp. 125–134. ACM, New York, 2014. doi: 10.1145/2628363.2628374
- [4] A. Bertelsen, P. Irrazaval, and R. F. Cadiz. Volume visualization using a spatially aware mobile display device. *Computerized Medical Imaging and Graphics*, 36(1):66–71, June 2012. doi: 10.1016/j.compmedimag.2011.06.003
- [5] L. Besançon, P. Issartel, M. Ammi, and T. Isenberg. Usability comparison of mouse, touch and tangible inputs for 3D data manipulation. arXiv.org preprint 1603.08735, 2016.
- [6] D. A. Bowman, E. Kruijff, J. J. LaViola, Jr., and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley, Boston, 2005.
- [7] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, Dec. 2013. doi: 10.1109/TVCG.2013.124
- [8] B. Buxton. Multi-touch systems that I have known and loved. Website: <http://www.billbuxton.com/multitouchOverview.html>, Jan. 2007. Updated last on June 12, 2014, visited in March 2016.
- [9] S. Carpendale. Evaluating information visualizations. In *Information Visualization: Human-Centered Issues and Perspectives*, vol. 4950 of LNCS, pp. 19–45. Springer, Berlin, 2008. doi: 10.1007/978-3-540-70956-5_2
- [10] D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiropoulos, A. G. Erdman, and D. F. Keefe. Interactive Slice WIM: Navigating and interrogating volume datasets using a multi-surface, multi-touch VR interface. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1614–1626, 2012. doi: 10.1109/TVCG.2011.283
- [11] A. Cohé, F. Dècle, and M. Hachet. tBox: A 3D transformation widget designed for touch-screens. In *Proc. CHI*, pp. 3005–3008. ACM, New York, 2011. doi: 10.1145/1978942.1979387
- [12] L. D. Cutler, B. Fröhlich, and P. Hanrahan. Two-handed direct manipulation on the responsive workbench. In *Proc. SIGD*, pp. 107–114. ACM, New York, 1997. doi: 10.1145/253284.253315
- [13] E. De Guzman, F. Ho-Ching, T. Matthews, T. Rattenbury, M. Back, and S. Harrison. EEWWW!!: Tangible interfaces for navigating into the human body. In *CHI Extended Abstracts*, pp. 806–807. ACM, New York, 2003. doi: 10.1145/765891.766003
- [14] P. Dragicevic. My technique is 20% faster: Problems with reports of speed improvements in HCI. Research Report hal-00739237, Inria, Oct. 2012.
- [15] T. Dwyer and D. R. Gallagher. Visualising changes in fund manager holdings in two and a half-dimensions. *Information Visualization*, 3(4):227–244, 2004. doi: 10.1057/palgrave.ivs.9500081 10.1057/palgrave.ivs.9500081
- [16] G. W. Fitzmaurice. Situated information spaces and spatially aware palm-top computers. *Comm. ACM*, 36(7):39–49, 1993. doi: 10.1145/159544.159566
- [17] G. W. Fitzmaurice. *Graspable User Interfaces*. PhD thesis, University of Toronto, Canada, 1996.
- [18] C. Forlines, D. Wigdor, C. Shen, and R. Balakrishnan. Direct-touch vs. mouse input for tabletop displays. In *Proc. CHI*, pp. 647–656. ACM, New York, 2007. doi: 10.1145/1240624.1240726
- [19] C.-W. Fu, W.-B. Goh, and J. A. Ng. Multi-touch techniques for exploring large-scale 3D astrophysical simulations. In *Proc. CHI*, pp. 2213–2222. ACM, New York, 2010. doi: 10.1145/1753326.1753661
- [20] L. Grammel, M. Tory, and M.-A. Storey. How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):943–952, 2010. doi: 10.1109/TVCG.2010.164
- [21] S. Greenberg and B. Buxton. Usability evaluation considered harmful (some of the time). In *Proc. CHI*, pp. 111–120. ACM, New York, 2008. doi: 10.1145/1357054.1357074
- [22] M. Hancock, T. ten Cate, and S. Carpendale. Sticky Tools: Full 6DOF force-based interaction for multi-touch tables. In *Proc. ITS*, pp. 145–152. ACM, New York, 2009. doi: 10.1145/1731903.1731930
- [23] M. Hancock, T. ten Cate, S. Carpendale, and T. Isenberg. Supporting sandtray therapy on an interactive tabletop. In *Proc. CHI*, pp. 2133–2142. ACM, New York, 2010. doi: 10.1145/1753326.1753651
- [24] M. S. Hancock and K. S. Booth. Improving menu placement strategies for pen input. In *Proc. GI*, pp. 221–230. CHCCS, Toronto, 2004.
- [25] C. Hand. A survey of 3D interaction techniques. *Comp. Graph. Forum*, 16(5):269–281, Dec. 1997. doi: 10.1111/1467-8659.00194
- [26] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proc. CHI*, pp. 452–458. ACM, New York, 1994. doi: 10.1145/191666.191821
- [27] E. Hornecker, P. Marshall, N. S. Dalton, and Y. Rogers. Collaboration and interference: Awareness with mice or touch input. In *Proc. CSCW*, pp. 167–176. ACM, New York, 2008. doi: 10.1145/1460563.1460589
- [28] T. Isenberg. Position paper: Touch interaction in scientific visualization. In *Proc. DEXIS*, pp. 24–27. Inria, Rocquencourt, France, 2011.
- [29] T. Isenberg. An interaction continuum for visualization. In *VIS Workshop on “Death of the Desktop: Envisioning Visualization without Desktop Computing”*, 2014.
- [30] T. Isenberg. Interactive exploration of 3D scientific visualizations on large display surfaces. In *Collaboration Meets Interactive Surfaces and Spaces: Theory and Practice*. Springer, Berlin/Heidelberg, 2016. To appear.
- [31] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2818–2827, 2013. doi: 10.1109/TVCG.2013.126
- [32] H. Ishii. Tangible bits: Beyond pixels. In *Proc. TEI*, pp. xv–xxv. ACM, New York, 2008. doi: 10.1145/1347390.1347392
- [33] H. Ishii and B. Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proc. CHI*, pp. 234–241. ACM, New York, 1997. doi: 10.1145/258549.258715
- [34] P. Issartel, F. Guéniat, and M. Ammi. A portable interface for tangible exploration of volumetric data. In *Proc. VRST*, pp. 209–210. ACM, New York, 2014. doi: 10.1145/2671015.2671130
- [35] P. Issartel, F. Guéniat, and M. Ammi. Slicing techniques for handheld augmented reality. In *Proc. 3DUI*, pp. 39–42. IEEE Computer Society, Los Alamitos, 2014. doi: 10.1109/3DUI.2014.6798839
- [36] P. Issartel, F. Guéniat, T. Isenberg, and M. Ammi. Analysis of locally coupled 3D manipulation mappings based on mobile device motion. arXiv.org preprint 1603.07462, Mar. 2016.
- [37] B. Jackson, T. Y. Lau, D. Schroeder, K. C. Toussaint, and D. F. Keefe. A lightweight tangible 3D interface for interactive visualization of thin fiber structures. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2802–2809, Dec. 2013. doi: 10.1109/TVCG.2013.121
- [38] J. Jankowski and M. Hachet. A survey of interaction techniques for interactive 3D environments. In *Eurographics State of the Art Reports*, pp. 65–93. Eurographics Association, Goslar, Germany, 2013. doi: 10.2312/conf/EG2013/stars/065-093
- [39] Y. Jansen, P. Dragicevic, and J.-D. Fekete. Tangible remote controllers for wall-size displays. In *Proc. CHI*, pp. 2865–2874. ACM, New York, 2012. doi: 10.1145/2207676.2208691
- [40] S. Jordà, G. Geiger, M. Alonso, and M. Kaltenbrunner. The reacTable: Exploring the synergy between live music performance and tabletop tangible interfaces. In *Proc. TEI*, pp. 139–146. ACM, New York, 2007. doi: 10.1145/1226969.1226998
- [41] M. Kaltenbrunner, T. Bovermann, R. Bencina, and E. Costanza. TUIO: A protocol for table-top tangible user interfaces. In *Proc. Workshop on Gesture in Human-Computer Interaction and Simulation*, 2005.
- [42] D. F. Keefe and T. Isenberg. Reimagining the scientific visualization interaction paradigm. *IEEE Computer*, 46(5):51–57, May 2013. doi: 10.1109/MC.2013.178
- [43] K. Kin, M. Agrawala, and T. DeRose. Determining the benefits of direct-touch, bimanual, and multifinger input on a multitouch workstation. In *Proc. GI*, pp. 119–124. CIPS, Mississauga, ON, Canada, 2009.
- [44] T. Klein, F. Guéniat, L. Pastur, F. Vernier, and T. Isenberg. A design study of direct-touch interaction for exploratory 3D scientific visualization. *Comp. Graph. Forum*, 31(3pt3):1225–1234, June 2012. doi: 10.1111/1467-8659.2012.03115x

- [45] J. Konieczny, C. Shimizu, G. Meyer, and D. Colucci. A handheld flexible display system. In *Proc. Visualization*, pp. 591–597. IEEE Computer Society, Los Alamitos, 2005. doi: 10.1109/VISUAL2005.1532846
- [46] W. Krüger and B. Fröhlich. The responsive workbench. *IEEE Computer Graphics and Applications*, 14(3):12–15, May 1994. doi: 10.1109/38.279036
- [47] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics*, 18(9):1520–1536, Sept. 2012. doi: 10.1109/TVCG.2011.279
- [48] D. López, L. Oehlberg, C. Doger, and T. Isenberg. Towards an understanding of mobile touch navigation in a stereoscopic viewing environment for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 22(5):1616–1629, May 2016. doi: 10.1109/TVCG.2015.2440233
- [49] C. Lundström, T. Rydell, C. Forsell, A. Persson, and A. Ynnerman. Multi-touch table system for medical visualization: Application to orthopedic surgery planning. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), Dec. 2011. doi: 10.1109/TVCG.2011.224
- [50] M. M. Miguel, T. Ogawa, K. Kiyokawa, and H. Takemura. A PDA-based see-through interface within an immersive environment. In *Proc. ICAT*, pp. 113–118. IEEE, Los Alamitos, 2007. doi: 10.1109/ICAT.2007.41
- [51] D. J. Moore, R. Want, B. L. Harrison, A. Gujar, and K. Fishkin. Implementing phicons: Combining computer vision with infrared technology for interactive physical icons. In *Proc. UIST*, pp. 67–68. ACM, New York, 1999. doi: 10.1145/320719.322585
- [52] J. D. Mulder and R. Van Liere. The personal space station: Bringing interaction within reach. In *Proc. VRIC*, vol. 2002, pp. 73–81, 2002.
- [53] S. Oh and W. Woo. Manipulating multimedia contents with tangible media control system. In *Proc. ICEC*, pp. 57–67. Springer, Berlin/Heidelberg, 2004. doi: 10.1007/978-3-540-28643-1_9
- [54] A. Olwal and S. Feiner. Spatially aware handhelds for high-precision tangible interaction with large displays. In *Proc. TEI*, pp. 181–188. ACM, New York, 2009. doi: 10.1145/1517664.1517705
- [55] W. Qi and J.-B. Martens. Tangible user interfaces for 3D clipping plane interaction with volumetric data: A case study. In *Proc. ICMI*, pp. 252–258. ACM, New York, 2005. doi: 10.1145/1088463.1088507
- [56] J. L. Reisman, P. L. Davidson, and J. Y. Han. A screen-space formulation for 2D and 3D direct manipulation. In *Proc. UIST*, pp. 69–78. ACM, New York, 2009. doi: 10.1145/1622176.1622190
- [57] J. Rekimoto and K. Nagao. The world through the computer: Computer augmented interaction with real world environments. In *Proc. I3D*, pp. 29–36. ACM, New York, 1995. doi: 10.1145/215585.215639
- [58] L. Ren, J. Cui, Y. Du, and G. Dai. Multilevel interaction model for hierarchical tasks in information visualization. In *Proc. VINCI*, pp. 11–16. ACM, New York, 2013. doi: 10.1145/2493102.2493104
- [59] G. Robles-De-La-Torre. The importance of the sense of touch in virtual and real environments. *IEEE MultiMedia*, 13(3):24–30, July–Sept. 2006. doi: 10.1109/MMUL.2006.69
- [60] S. Schkolne, H. Ishii, and P. Schroder. Immersive design of DNA molecules with a tangible interface. In *Proc. Visualization*, pp. 227–234. IEEE Computer Society, Los Alamitos, 2004. doi: 10.1109/VISUAL2004.47
- [61] D. Schmalstieg, L. M. Encarnação, and Z. Szalavári. Using transparent props for interaction with the virtual table. In *Proc. I3D*, pp. 147–153. ACM, New York, 1999. doi: 10.1145/300523.300542
- [62] A. Sears and B. Shneiderman. High precision touchscreens: Design strategies and comparisons with a mouse. *International Journal of Man-Machine Studies*, 34(4):593–613, Apr. 1991. doi: 10.1016/0020-7373(91)90037-8
- [63] B. Shneiderman. Touchscreens now offer compelling uses. *IEEE Software*, 8(2):93–94, 107, Mar. 1991. doi: 10.1109/52.73754
- [64] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. VL*, pp. 336–343. IEEE Computer Society, Los Alamitos, 1996. doi: 10.1109/ML.1996.545307
- [65] B. Shneiderman. *Designing the user interface: strategies for effective human-computer interaction*. Pearson Education Inc., 2010.
- [66] H. Song, E. P. Curran, and R. Sterritt. Multiple foci visualisation of large hierarchies with flextree. *Information Visualization*, 3(1):19–35, 2004. doi: 10.1057/palgrave.ivs.9500065
- [67] P. Song, W. B. Goh, C.-W. Fu, Q. Meng, and P.-A. Heng. WYSIWYF: Exploring and annotating volume data with a tangible handheld device. In *Proc. CHI*, pp. 1333–1342. ACM, New York, 2011. doi: 10.1145/1978942.1979140
- [68] M. Spindler and R. Dachselt. PaperLens: Advanced magic lens interaction above the tabletop. In *Proc. ITS*, pp. 69–76. ACM, New York, 2009. doi: 10.1145/1731903.1731920
- [69] N. Sultanum, S. Somanath, E. Sharlin, and M. C. Sousa. “Point it, split it, peel it, view it”: Techniques for interactive reservoir visualization on tabletops. In *Proc. ITS*, pp. 192–201. ACM, New York, 2011. doi: 10.1145/2076354.2076390
- [70] E. Sundén, A. Bock, D. Jönsson, A. Ynnerman, and T. Ropinski. Interaction techniques as a communication channel when presenting 3D visualizations. In *Proc. 3DVis*, pp. 61–65. IEEE Computer Society, Los Alamitos, 2014. doi: 10.1109/3DVis.2014.7160102
- [71] D. S. Tan, D. Gergle, P. Scupelli, and R. Pausch. Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction*, 13(1):71–99, Mar. 2006. doi: 10.1145/1143518.1143521
- [72] M. Tsang, G. W. Fitzmaurice, G. Kurtenbach, A. Khan, and B. Buxton. Boom chameleon: Simultaneous capture of 3D viewpoint, voice and gesture annotations on a spatially-aware display. In *Proc. UIST*, pp. 111–120. ACM, New York, 2002. doi: 10.1145/571985.572001
- [73] P. Tuddenham, D. Kirk, and S. Izadi. Graspables revisited: Multi-touch vs. tangible input for tabletop displays in acquisition and manipulation tasks. In *Proc. CHI*, pp. 2223–2232. ACM, New York, 2010. doi: 10.1145/1753326.1753662
- [74] B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM Systems Journal*, 39(3.4):915–931, 2000. doi: 10.1147/sj.393.0915
- [75] B. Ullmer, H. Ishii, and R. J. Jacob. Tangible query interfaces: Physically constrained tokens for manipulating database queries. In *Proc. of INTERACT*, vol. 3, pp. 279–286, 2003.
- [76] B. Ullmer, R. Sankaran, S. Jandhyala, B. Tregre, C. Toole, K. Kallakuri, C. Laan, M. Hess, F. Harhad, U. Wiggins, and S. Sun. Tangible menus and interaction trays: Core tangibles for common physical/digital activities. In *Proc. TEI*, pp. 209–212. ACM, New York, 2008. doi: 10.1145/1347390.1347436
- [77] A. van Dam, A. S. Forsberg, D. H. Laidlaw, J. J. LaViola, and R. M. Simpson. Immersive VR for scientific visualization: A progress report. *IEEE Computer Graphics and Applications*, 20(6):26–52, Nov./Dec. 2000. doi: 10.1109/38.888006
- [78] G. R. VandenBos, ed. *Publication Manual of the American Psychological Association*. APA, Washington, DC, 6th ed., 2009.
- [79] J. Wagner, S. Huot, and W. Mackay. BiTouch and BiPad: Designing bimanual interaction for hand-held tablets. In *Proc. CHI*, pp. 2317–2326. ACM, New York, 2012. doi: 10.1145/2207676.2208391
- [80] K. Watson, R. Darken, and M. Capps. A handheld computer as an interaction device to a virtual environment. In *Third International Immersive Projection Technology Workshop (IPT)*, Stuttgart, Germany, 1999.
- [81] D. Watson, M. Hancock, R. L. Mandryk, and M. Birk. Deconstructing the touch experience. In *Proc. ITS*, pp. 199–208. ACM, New York, 2013. doi: 10.1145/2512349.2512819
- [82] A. Wiebel, F. M. Vos, D. Foerster, and H.-C. Hege. WYSIWYP: What you see is what you pick. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2236–2244, Dec. 2012. doi: 10.1109/TVCG.2012.292
- [83] K.-P. Yee. Peephole displays: Pen interaction on spatially aware handheld computers. In *Proc. CHI*, pp. 1–8. ACM, New York, 2003. doi: 10.1145/642611.642613
- [84] J. S. Yi, Y. a. Kang, J. Stasko, and J. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, Nov. 2007. doi: 10.1109/TVCG.2007.70515
- [85] L. Yu, K. Efstathiou, P. Isenberg, and T. Isenberg. Efficient structure-aware selection techniques for 3D point cloud visualizations with 2DOF input. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2245–2254, Dec. 2012. doi: 10.1109/TVCG.2012.217
- [86] L. Yu, K. Efstathiou, P. Isenberg, and T. Isenberg. CAST: Effective and efficient user interaction for context-aware selection in 3D particle clouds. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):886–895, Jan. 2016. doi: 10.1109/TVCG.2015.2467202
- [87] L. Yu, P. Svetachov, P. Isenberg, M. H. Everts, and T. Isenberg. FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1613–1622, Nov./Dec. 2010. doi: 10.1109/TVCG.2010.157
- [88] S. Zhai. User performance in relation to 3D input device design. *ACM SIGGRAPH Computer Graphics*, 32(4):50–54, Nov. 1998. doi: 10.1145/307710.307728
- [89] O. Zuckerman and A. Gal-Oz. To TUI or not to TUI: Evaluating performance and preference in tangible vs. graphical user interfaces. *International Journal of Human-Computer Studies*, 71(7–8):803–820, July 2013. doi: 10.1016/j.ijhcs.2013.04.003