



HAL
open science

Interactive Exploration of Three-Dimensional Scientific Visualizations on Large Display Surfaces

Tobias Isenberg

► **To cite this version:**

Tobias Isenberg. Interactive Exploration of Three-Dimensional Scientific Visualizations on Large Display Surfaces. Craig Anslow; Pedro Campos; Joaquim Jorge. Collaboration Meets Interactive Spaces, Springer, 2016, 978-3-319-45852-6. 10.1007/978-3-319-45853-3_6 . hal-01379938

HAL Id: hal-01379938

<https://inria.hal.science/hal-01379938v1>

Submitted on 12 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

This is an author-prepared version of the book chapter published by Springer as part of the 2016 book “Collaboration Meets Interactive Spaces,” edited by Craig Anslow, Pedro Campos, and Joaquim Jorge. The chapter can be found online at doi: 10.1007/978-3-319-45853-3_6. The text of this author version is virtually identical to the published one, albeit uses better typography and more hyperlinks (e. g., from the figures to online videos of the described interaction technique(s)). Notice also that many of the images contained in the publisher’s online version are of lower quality than those contained in this author-prepared version.

Interactive Exploration of Three-Dimensional Scientific Visualizations on Large Display Surfaces

Tobias Isenberg

Abstract The chapter surveys the different approaches investigated to interact with scientific visualizations on large surfaces such as tables and walls. The chapter particularly does not focus on VR-based interaction or tangible input but on those interaction techniques where the input is provided on the surface itself or where it is focused on the surface. In particular, tactile interaction techniques are covered and the challenges of gestural input as well as of combining touch input with stereoscopic rendering are discussed. Where possible, connections to collaborative interaction scenarios are pointed out, even though most publications to date focus on single-user interaction.

1 Introduction

Scientific visualization of data which has an implicit mapping to the 3D Euclidean space has traditionally been a domain for which interaction plays an important role. For example, the interactive exploration of 3D medical data, physical or astro-physical simulations, or models from structural biology has always been important as soon as the underlying graphics hardware had become powerful enough to support such interactive exploration. Initially, this interaction typically concentrated on navigation of 3D environments or the manipulation (translation, rotation, scaling) of parts of the visualization. Recently, however, researchers have started to focus on more flexible interaction techniques that facilitate advanced exploration of scientific datasets [51]. A large part of this work has explored surface-based interaction in which (one of) the data display(s) also serves as the main space where input is provided by the interacting person or people—a topic that has received an increasing amount of attention in the recent years for visualization in general [38, 39] and specifically for the exploration of 3D data [52].

Tobias Isenberg
Inria, France, e-mail: tobias.isenberg@inria.fr

Naturally, the use of interactive surfaces inherently supports the collaboration of people, in contrast to the single-person input of traditional interaction settings based on mouse, keyboard, or dedicated 3D input devices. While collaborative interaction is the core topic of this book, in 3D data visualization has the additional constraint that the data domain typically needs to use a single, linear mapping to a shared display surface so that a navigation of the 3D data representation, at least, can only be done by one person at a time. In this chapter we thus focus on single-user interaction techniques and point out the few approaches that have explored collaborative scenarios.

In addition, we focus on those interaction techniques that use the 2D input provided on interactive surfaces for interaction, but mention some special approaches that use both 2D input and immersive displays in our discussion. While there have been approaches for other forms of input, for example, through wands, gloves, or 3D tracking for use in immersive VR environments [26, 28, 57] as well as fully tangible interaction (e. g., [58]), our focus on directly capturing the input on an interactive surface has several advantages. First, the provided input is direct as the input location can directly correspond to a displayed data element, in contrast to wands and remote pointing devices as well as passive tangible props. Second, input on interactive surfaces does not require elaborate additional hardware (as for gloves) or 3D tracking setups (as in VR and tangible interaction), making the management and use of such data exploration systems easier and less expensive. Third, the properties and advantages of tactile interaction known from other forms of interaction [17] similarly apply, such as improved performance of tactile input [54], the support of awareness of collaborators [37], somesthetic information and feedback [67], and improved performance on physically large displays [77]. Moreover, recently it has been shown that certain interaction techniques such as based on tactile input on large displays can serve as a communication channel when presenting visualizations to an audience [76].

In the remainder of this chapter we first discuss the issue of the difference between the data space on the one side and input and output spaces on the other side. Next, we review basic interaction techniques for surface-based interaction with 3D data. Then we introduce a number of design studies for data exploration systems from various domains, before we conclude the chapter with a summary.¹

2 Data Space vs. Input and Output Spaces

The restriction to (typically planar) 2D input spaces for the control or exploration of 3D data spaces brings with it an important mapping problem: the need for mapping from the 2D input space to the 3D data space. In addition to this mapping, the normal visualization mappings, of course, also take place. In contrast to the mappings

¹ It is difficult to understand interaction techniques by simply reading about them or seeing traditional figures with snapshots. We thus hyperlink to videos of the discussed techniques where possible from the figures in the electronic version of the chapter to better illustrate the techniques.




Fig. 1 Surface-based 2D interaction with 3D data: a mental mapping from the input and the projected visualization to the 3D data space is required. Image © Tobias Isenberg, used by permission.

needed for the visualization of abstract, non-spatial data [18], for the visualization of spatial 3D data we already have an implicit mapping so that most mappings in the visualization pipeline concern data filtering, the assignment of abstract aspects to visual variables, and 3D projection and rendering (e. g., [79, Section 4.1]). Essential for interactive visualization in general and our specific case of 2D surface-based input, however, is that we also need to take the physical presentation of the generated visualization [47] into account.

2.1 2D Projected Viewing of 3D Data Visualizations

Unlike in immersive VR settings, much 3D visualization relies on the rendered visualization being projected to 2D and displayed on a “normal” screen. Typically, this screen is the display on which input is captured. In this case the input is collocated to the visual representation of the output, and the same *mental mapping* from visualization space to data space and vice versa can also be applied to the provided input (Fig. 1). In that sense the input is as direct as possible as the user never encounters a visual representation that has the same dimensionality as the data. If the input “display” is separate from the display that shows the projected visualization, then we have a situation similar to touch pad interaction or the use of digitizer tablets where, while the dimensionality of the displayed visualization and the provided input is the same, a *mental mapping* from input space to output space and from that to the data space is necessary, which makes this indirect interaction more difficult.

2.2 3D Stereoscopic Viewing of 3D Data Visualizations

In contrast to a projected display of the visualization one may also want to take advantage of the better depth perception and the resulting increased feeling of immersion of a stereoscopic presentation of the visualization. The use of surface-based input in such immersive virtual reality environments, however, presents additional challenges [82]. Here, the dimensionality of the presented visualization is the same as that of the data—the data is perceived at the same location as that of the data (with the exception of 3D manipulations of the visual representations that can also be understood as manipulations that are applied to the data itself).

The input, however, still is provided on a 2D surface to take advantage of the benefits mentioned above. This means that only in rare cases is the input actually performed where the user perceives the data to be manipulated. Moreover, it has been shown that such tactile (or pen-based) interaction suffers from the parallax between the two images that are shown for both of the eyes [13, 14, 24, 82, 83]. In addition, touch-through [19, 78, 81] and invisible wall problems make such an interaction setup problematic. Only in situations when the element to be accessed by the surface-based input is at a close distance to input screen do users perceive their input to directly control the manipulated elements [83].

Some solutions exist to address these problems, yet none are ideal. For example, Schmalstieg et al. [69] suggested to use transparent props, yet these are static and would not work well with 3D data space manipulations and time-dependent data. Hachet et al. [31] separated the touch surface from the stereoscopic display in their Toucheo system, but thus significantly restricted the space in which people can interact. Jackson et al. [45] use the touch surface of a table interface as the interaction reference frame on which widgets are placed, and input is provided not only through tactile sensing but also through over-the-surface means supported by 3D tracking (see Section 3.1 for a more detailed description of this technique). Butkiewicz and Ware [15, 16], finally, used a very specific setup that relies on a tilted setup, shallow-depth data, and a single “natural” interaction surface (see the more detailed description of this setup in Section 4).

In addition to these hardware solutions, also some software-based interaction designs were proposed to alleviate the parallax problem. For example, Valkov et al. [81] suggested to move objects toward the touch surface as a user reaches out to them. Giesler et al. [30] used “void shadows” that connect the objects behind a touch surface with it and which serve as interaction proxies. These interaction designs, however, may cause problems in a visualization environment as the data display itself should not be obscured and often no dedicated objects exist. As an alternative, people thus also examined hybrid settings that separate the stereoscopic display from the input surface as discussed next.




Fig. 2 Static mapping between input surface and stereoscopic data display in Coffey et al.'s [21, 22] Slice WIM setup. Image courtesy of and © Daniel F. Keefe, used by permission.

2.3 Hybrid View Settings for 3D Data Visualization

In normal PC-based settings we are quite used to having the surface that displays the data or object with which we interact (i. e., the computer screen) to be different from the surface on which we provide input (i. e., the table on which the keyboard and mouse are located). Humans are able to deal with such separation and are ready to make a mental mapping from one surface to the other if the mapping only contains translations [8, 10, 85] and simple rotations [1, 2, 25], also in immersive environments [84]. The same concept has also been used for the surface-based interaction with stereoscopic displays. For example, Coffey et al. [20, 21, 22] use a vertical display to show the stereoscopic content of a visualization, while capturing tactile input on a horizontal surface (see Fig. 2). Both surfaces are physically connected perpendicularly, and are visually connected to each through a stereoscopic world-in-miniature (WIM) display of the data as well as shadows that this WIM casts on the tactile input surface. Similar static hybrid setups at interactive 3D visualization have been used by Bogorin et al. [11] and Novotný et al. [64].

In addition to such static setups, modern smart phones and tablet computers also facilitate interaction styles where the input is provided on a mobile surface [72], while the data is still visualized stereoscopically. This scenario, however, poses additional challenges as the mapping from the input provided and data displayed on the mobile surface to the stationary (and typically large) stereoscopic surface constantly changes as the interacting person is moving around. López et al. [60] recently analyzed



Fig. 3 Stereoscopic data exploration with mobile control. Image from [60], © IEEE, used by permission.

the interaction modes in this situation (e. g., Fig. 3). In particular, they described the mapping issues between the different frames of reference that arise when the interacting person moves beyond a small distance around their initial position—due to the issues humans have with such mappings if more than simple translations and a single rotation is involved as described above. Based on this analysis they suggested an data exploration workflow that allows users to move in the physical space and transition between the different interaction modes, synchronizing the two views when the reference frames can no longer be easily mentally integrated.

3 Basic Interaction Techniques

Based on these viewing and interaction settings we can now turn to the specific interaction techniques used for surface-based 3D data exploration and analysis. In doing this, we concentrate on those techniques that use planar, monoscopic 2D surfaces as input spaces because most of the special cases mentioned in Section 2.2 do not play a role in 3D visualization applications so far. In this chapter we first describe basing interaction techniques that are common to many of not most 3D visualization applications, in particular those for 3D navigation and for data selection. Next, in Section 4, we then discuss a number of systems and design studies that combine several interaction techniques for a more comprehensive interaction with and exploration of data.



Fig. 4 Interaction sequence from the control of the orientation and location of a map orientation in 3D space using 3D-RST [66]. Images courtesy of and © Jason L. Reisman, Philip L. Davidson, and Jefferson Y. Han, used by permission.

3.1 Data Navigation

Navigation techniques for 3D environments have been investigated for a long time [12, 36]. Also surface-based 3D interaction techniques are not only a domain of visualization [46]. Here we mention a number of techniques that are, in principle applicable to interactive 3D visualization, even if the techniques were not initially designed with this application in mind. However, in general 3D interaction the focus often is placed on the *manipulation of individual 3D objects* within a larger space, such as moving furniture items around in a virtual environment that shows a new interior design. In 3D visualization, however, we rarely manipulate individual objects but rather navigate in the 3D data space to look at specific aspects of the data more closely. Nevertheless, many generic surface-based 3D navigation techniques can be used in visualization by using them to *affect the “data space”* of the visualization.²

One of these interaction techniques for general 3D shapes is the 3D-RST approach by Reisman et al. [66] that is inspired by the common two-finger pinching interaction. RST stands for rotation, scaling, and translation and, in the 2D case, allows users to perform these transformations for 2D shapes in their native 2D space [35]. This interaction relies on the principle that the interaction points are “sticky” [33, 66], i. e., that they stay connected to the same location on the manipulated object for the

² In interactive 3D visualization there may indeed be some dedicated objects to be moved such as cutting planes and particle sources. Nevertheless, for such objects often dedicated interaction techniques are used as explained in the remainder of the chapter.




Fig. 5 Illustration of the interaction with the tBox [23] which is shown overlaid on a 3D object. Image courtesy of and © Martin Hachet, used by permission.

entire interaction. Reisman et al.’s [66] 3D-RST³ extends this general principle to the reorientation and translation of 3D shapes, using the 2D screen-space locations of multiple interaction points to constrain the mapping. Of course, this mapping is easily either under-constrained or over-constrained:

- one to two fingers provide only ≤ 4 DOF, while 6 DOF are needed to specify the location and orientation of a 3D shape, while
- four or more fingers provide ≥ 8 DOF for the same necessary 6 DOF.

Reisman et al. solve this problem in screen-space using an energy minimization approach to find the best possible mapping despite the possibly under- or over-constrained input (Fig. 4). They demonstrate how their technique can be applied to many types of surfaces including terrain renderings, and it is not difficult to envision to apply the same technique to other planar elements in visualization applications such as cutting planes.

An alternative interaction mapping was designed by Liu et al. [59] who integrate both the 4DOF x -/ y -/ z -translation plus z -rotation with the 2DOF x -/ y -rotation. Users can seamlessly switch between these two modalities by either moving both fingers at the same time or by leaving one finger static on the tactile surface. Similar to Reisman et al.’s [66] approach, this technique allows Liu et al. [59] to facilitate flexible and fluid interactions with 3D shapes.

While Reisman et al.’s [66] 3D-RST and Liu et al.’s [59] two-finger technique facilitate a flexible and fluid type of interaction, it always affects all 6 DOF (for

³ 3D-RST is a somewhat inappropriate name as Reisman et al.’s [66] technique is constrained to translations and rotations. The scale always remains constant with this technique. In fact, a technique that is entirely based on “sticky” contact control cannot affect both z -distance and object scale at the same time, the two properties are visually ambiguous (see also Hancock et al.’s [33] “Sticky Tools” interaction mapping and its application to Sandtray therapy [34]).




Fig. 6 Widget-based 3D navigation using FI3D, using different mappings on the man data view and on the FI3D frame. Image from [90], © IEEE, used by permission.

Reisman et al.'s [66] technique) or 2 DOF resp. 4 DOF (for Liu et al.'s [59] approach) of the output simultaneously. In visualization applications, however, it is often necessary to single out specific DOF for the interaction to be able to constrain which aspects of a visualization are affected.

To address this problem, Au et al. [6] describe a set of gestures to single out specific DOFs to control individually. One problem with such a gestural approach is that the gesture set has to be learned and is not easily discoverable. Cohe et al. [23] describe a similar constrained interaction with their tBox technique (Fig. 5) for up to 9 DOF control. This widget-based approach shows a box-shaped interaction widget overlaid on the rendering, whose orientation is tied to that of the shown 3D scene or object. Manipulations can now be applied based on where on the widget input is provided. For example, single inputs on the cube's sides provide single-axis rotations, while single inputs on a cube edge start translations along the axis parallel to the edge. Uniform scaling is possible using pinching on a cube side, non-uniform scaling by pinching on two opposite cube edges. These interactions allow users to constrain their manipulations to only single-DOF control, and studies [60] indicate that the tBox provides people with an increased feeling of precision for the 3D interaction.

While such precise control is essential, a flexible and fluid 3D navigation may also be important. Unfortunately, Reisman et al.'s [66] 3D-RST which provides such flexibility only facilitates 6 DOF control. So Yu et al. [90] conceived FI3D (Fig. 6), a widget-based 3D navigation approach that allows researchers to control up to 7 DOF (3D rotations, 3D translations, and uniform scaling). The approach is based on a widget placed around a central data display, with which controls the interaction mode based on the location of where an input starts as well as its direction. In the center, Yu et al. map x/y -translation as well as 2D RST manipulation. Interactions

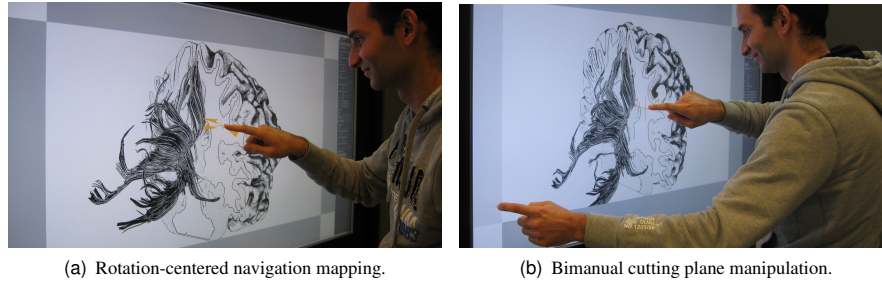


Fig. 7 Aspects of the FI3D-based interaction mapping can and should depend on the data. For data that does not require much zooming, the main interaction is trackball rotation and is thus used for the one-finger mapping, while x/y -translation is accessed from the FI3D frame (a). In addition, dedicated interaction techniques such as cutting plane manipulation are supported (b). Images from [90], © IEEE, used by permission.

started in the frame, initially moving along it start rotations around the z -axis, while interactions from the frame initially into the central data view start trackball rotations. Separate regions are used for uniform scaling and z -translations, and bi-manual techniques allow users to constrain their rotation input also to the x - and y -axis as well as allow to provide different 2D rotation centers.

Yu et al. [90] also mention that the specific mappings should depend on the data being shown. Their initial mapping works well for data such as astro-physical particle simulations that does not have an inherent center point and that require a lot of scale changes to explore different aspects—in such cases translations are more important than rotations and, hence, are mapped to one-point input in the central view. Other types of data such as brain scans, however, may have an inherent center and may not require much scaling—in such cases rotations are more important than translations. Yu et al. thus also demonstrate that the two mappings can be flipped (Fig. 7(a)), and that additional functions such as cutting plane manipulations and fibertract selection can be realized using bimanual interaction (Fig. 7(b)).

In addition to these generic navigation techniques, some techniques with domain-specific constraints have also been created. In the context of visualization, two should be mentioned at this point. The first one specifically supports the navigation of 3D astronomical datasets such as models of the solar system, its local neighborhood, the Milky Way, and the spatial arrangement of multiple galaxies. Such a setup has two major constraints: it (a) primarily requires rotations (as opposed to translations) and it (b) requires scaling across multiple orders of magnitude. For this purpose, Fu et al. [29] provide a set of interaction techniques that combine spherical navigation based on a trackball metaphor as well as their unique “powers-of-ten-ladder” for multi-scale navigation (Fig. 8). The latter is invoked with two fingers touching the interaction surface in a vertical arrangement, and then a second hand can initiate scale navigation, either initiating small-scale changes (Fig. 8(b)) or large-scale changes (Fig. 8(c))—depending on the distance of the input to the basis of the widget.

Another domain-specific tactile surface-based navigation technique was introduced by Sultanum et al. [75] for the exploration of geological outcrops. Such 3D




Fig. 8 Powers-of-10-Ladder [29]. Depending on where input is provided with respect to the basis of the interaction widget, either (b) small-scale or (c) large-scale changes to the zoom level are being introduced. Images courtesy of and © Chi-Wing Fu, used by permission.

datasets are captured, for example, by LiDAR scans and reveal information on geological layers as part of geological surveys. To provide 3D navigation with such outcrops, Sultanum et al. facilitate a first-person fly-by navigation strategy by defining a separate navigation surface to which the camera is constrained. The scientist can then explore the dataset through tactile interaction gestures, controlling the remaining camera parameters such as panning its x - y -location, zooming, and tilting.

While all previous techniques restrict themselves to controlling the 3D scene or objects solely based on 2D input captured on the tactile surface, Jackson et al. [45] augment this 2D input with additional information based on the posture of the interacting hands (Fig. 9). Their “nailing down multi-touch” set of interaction techniques thus allows users to tilt, bend, or twist objects or datasets within the 3D space, supported by a stereoscopic data display. As they specifically treat the interaction surface as the location where the interaction control widgets are placed (as can be seen in Fig. 9), this interaction style does not cause many problems despite the previously discussed issues of tactile interaction with stereoscopically displayed scenes (Section 2.2).

In addition to the navigation techniques we discussed so far, several other 3D interaction techniques for surface-based 3D navigation have been designed (see, e. g., [12, 36, 46]). Most of them, however, rely on the manipulation of individual 3D objects within the 3D space in a way that is not very well suitable for 3D interaction with data visualizations [40].

3.2 Data Picking and Data Selection

In addition to 3D navigation, a second interaction technique that is essential for the exploration of 3D data visualizations is the selection of sub-elements of the depicted datasets. While selection techniques for 2D [87] and 3D [4, 5, 7] environments have




Fig. 9 Nailing down multi-touch interaction [45] for providing additional tilting, bending, or twisting of the 3D data. Image courtesy of and © Daniel F. Keefe, used by permission.

long been studied, the picking or selection within 3D datasets is not as straightforward as one may think. A first challenge lies in the problem that either no explicit objects make up the dataset (for example, in sampled data of a continuous domain such as volumetric datasets) or that the explicit data objects are too small or narrow to be easily captured by traditional picking or selection techniques (such as in point-based line-based datasets). A second issue arises from the fact that interaction can only be recorded on the two-dimensional input surface, so this input is not able to fully constrain the intended three-dimensional selection.

The ultimate challenge in 3D data selection is thus to effectively and intuitively specify that sub-space of the dataset that contains the elements to be further processed. While data filtering is one approach to arrive at such selections sub-spaces, the characteristics of the intended selections may not be known ahead of time or there may not even be data aspects that would allow such an effective filtering during exploratory data analysis. Below we review a number of spatial input techniques that specify intended selections based on spatial input. While none of the techniques we review in this section were specifically created for the application to surface-based interaction, they all work particularly in surface-based interaction contexts due to their spatial input character and direct manipulation paradigms they support: spatial selection input can thus be directly specified with respect to the displayed data.

A fundamental interaction technique in this context is picking. While the picking of individual objects is simple using ray-pointing and similar techniques, picking in continuous data such as medical volume scans or physical simulations is far from straightforward. For this purpose Wiebel et al. [86] created their What You See Is What You Pick interaction technique (Fig. 10), a structure- and view-aware picking approach that takes the data along the picking ray as well as the transfer function into




Fig. 10 What you see is what you pick interaction [86]. Once a location is picked in the projected view of a volumetric visualization (left), its 3D location is deduced and a cutting plane can be placed accordingly, facilitating the exploration of this cutting plane in the context of the volume rendering (center) or by itself (right). Image courtesy of and © Alexander Wiebel, used by permission.

account. Specifically, they extract the section along the ray that constitutes the largest jump in accumulated opacity, corresponding to the feature that is visually dominant at the picked 2D position. For this section they then select either its front or center, depending on the user's preference.

While this structure- and view-aware picking technique can only yield a single 3D position within the 3D data, it is sometimes also necessary to select a whole 3D subspace, such as to be able to do carry out a specific data analysis of visually interesting features. For this purpose several structure- and view-aware spatial selection techniques have been created. Based on the initial work by Owada et al. [65], Yu et al. [88] introduced their CloudLasso that bases the selection on an interactively drawn 2D selection lasso and the analysis of a scalar value such as density based on which a selection volume is extrapolated (Fig. 11 shows the use of the CloudLasso selection in a surface-based interaction context).

Structure-aware and view-dependent selection techniques such as CloudLasso, however, have the disadvantage that they (a) select everything along the line where the chosen scalar threshold is surpassed, which can lead to multiple selected but unconnected components. Moreover, they (b) do not take the shape of drawn lasso itself into account other than to use it as a 2D cut-off constraint for the 3D selection. To address both issues, Yu et al. [89] extended their initial approach and described the CAST family of context-aware selection techniques (Fig. 12). In particular, they describe SpaceCAST (Fig. 12(a)), that works similar to CloudLasso but selects that connected component whose outline is most similar to the drawn lasso. Next, they created TraceCast (Fig. 12(b)) that relaxes the constraint that the lasso cuts off the selection with respect to the 2D view, facilitating the easy selection of complex structures. Finally, they describe PointCAST (Fig. 12(c)) which only requires a point as an input to facilitate the selection of small clusters. An alternative is to pick that cluster from the selection that has the largest 2D projection [70] but this approach does not take the actual shape of the selection lasso into account.

The approaches discussed so far work well for point-based or scalar 3D data (i. e., particle clouds and volumetric data), but other 3D data types require different approaches. In particular, line based data such as streamlines and similar or fiber





Fig. 11 Bimanual CloudLasso selection within an astronomical particle dataset. Image from [88], © IEEE, used by permission.



(a) SpaceCAST selection.

(b) TraceCast selection.

(c) PointCAST selection.

Fig. 12 CAST selection for 3D particle data which can be used in a similar surface-based environment as shown in Fig. 11. Image from [89], © IEEE, used by permission.

tracts are too long for it to be possible to effectively select subsets with a volume-based technique. While it is possible to use dedicated input hardware for a spatial selection of line-based data (e. g., [43, 44, 53]), solely surface-based spatial selection techniques are less common and we are only aware of two approaches. Akers' [3] fibertract selection uses the shape of a drawn selection mark to guide the selection of 3D neurologic pathways in a structure-aware fashion, while Coffey et al.'s [22] Slice WIM widget facilitates the selection of flowline bundles by drawing a selection lasso on a plane that was previously placed roughly perpendicular to the flow. Tong et al.'s [80] interaction techniques, in contrast, are not used for streamline selection but allow users to specify spatial lenses though tactile input that then reveal hidden parts of a streamline dataset.

3.3 Summary

The discussed basic navigation and selection techniques demonstrate that 3D navigation and selection can effectively be carried out also in a surface-based interaction context. While other interaction mappings for these data exploration tasks will certainly be explored in the future, the existing ones already provide a good selection for the practical implementation of surface-based visualization tools. In the survey we intentionally did not cover, however, techniques for the manipulation of 3D data elements because data visualization of scanned or simulated data typically does not require the *manipulation of the data*, but focuses on the *exploration of the data*. Yet, many additional interaction tasks for data exploration also have to be supported [52] such as particle seeding at 3D locations, cutting plane manipulation, path planning, data value probing, and visualization parameter adjustment. Some of these have already been explored for surface-based interaction settings, such as the visual exploration of different settings for volume rendering transfer functions [50]. Most of them, however, have been explored within the context of a specific application domain or design study. We thus review several of such existing surface-based 3D data visualization systems⁴ next.

4 Systems and Design Studies

Surface-based data exploration systems have been created for a variety of target audiences including museum visitors [32], scientific researchers in domains such as fluid mechanics and oceanography, medical doctors and researchers, and exploration geologists. As we focus in this chapter on the interaction techniques, we loosely group them by similar interaction characteristics, rather than chronologically, by the mentioned application domains, or intended target audiences.

Marton et al. [62] describe IsoCam (Fig. 13), a touch-based system for the exploration of 3D scans of archeological artifacts in a museum context. As the target audience for this interactive system is almost exclusively untrained in 3D navigation, they use a constrained navigation that provides a robust exploration of large 3D virtual reconstructions. They use an indirect navigation approach [56, 63]: gestural input on the touch surface is mapped to changes in the visualization that is shown on a remote screen. Specifically, they provide 2D navigation along an iso-surface of the distance field of the depicted objects (related to Sultanum et al.'s [75] work), zooming to change the distance to the object (i. e., the iso-value), and twisting to change the camera orientation. In addition, additional information can be accessed about the objects on demand. The interesting aspects about this interaction system is that it was deployed into the real world, with a non-expert target audience. The way

⁴ The classification into interaction technique and interactive system is not always crystal clear—we used our best judgment to differentiate between the two groups.



Fig. 13 IsoCam interaction with massive cultural heritage models [62] at the Digital Mont'e Prama installation by CRS4 Visual Computing at the National Archaeological Museum of Cagliari, Italy. Image courtesy of and © Alberto Jaspe Villanueva, used by permission.




Fig. 14 Song et al.'s [71] system that combines tactile and tangible interaction on for surface-based medical visualization. Images courtesy of and © Peng Song, Wooi Boon Goh, and Chi-Wing Fu, used by permission.

Marton et al. thus constrained the degrees of freedom for navigation can thus be an inspiration for future interactive systems “for the masses.”

Song et al. [71] combine a large monoscopic touch wall with a mobile secondary touch and orientation input device for the visualization of volumetric data (Fig. 14). They explore a combination of direct interaction on the large surface with manipulations of a cutting plane by means of the mobile device. They explore a number of combinations of tangible and tactile interaction techniques on the remote device to enable users to translate, reorient, zoom, and annotate the remotely shown visualization. The interesting aspect of this design study is the combination of tangible interaction with tactile input to create a larger interaction vocabulary including constrained

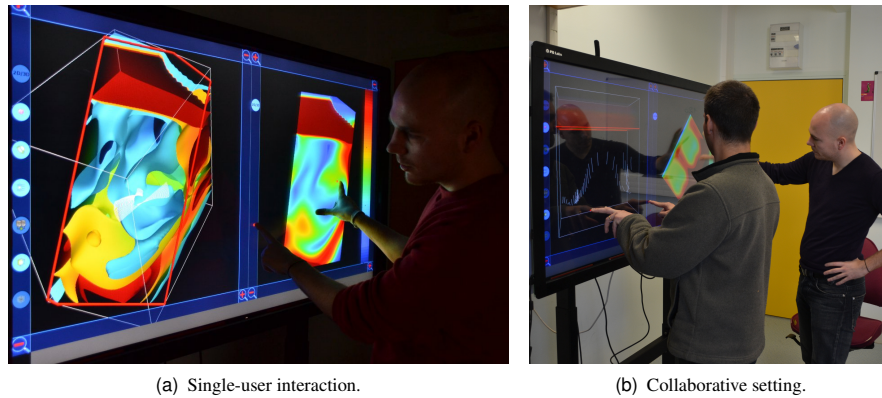


Fig. 15 Design study of a surface-based system for the exploration of fluid flow simulations [55]. Image (a) from [52], © IEEE, used by permission. Image (b) from [55], © Eurographics Association, used by permission.

interactions and alternative mappings for the same interaction. The interaction with the large display then facilitates the interaction with the data space itself, similar to some of those described in Section 3.1. A small user study conducted by Song et al. suggests that their hybrid interaction can outperform a PC-based interface with equivalent controls with respect to interaction time.

Klein et al. [55], instead, concentrate on capturing input only on a single large display, but design their system for the exploration of fluid simulations to provide similar interaction techniques (Fig. 15): the navigation of the overall dataset, cutting plane manipulation, data value probing, particle seeding, visualization parameterization, and temporal exploration. Their interaction design is largely based on the FI3D widget [90] but adds additional interaction mappings for manipulating the other mentioned elements, typically in a widget-based bimanual fashion. The interesting aspect of this system is its dual use of the cutting plane: it is not only used in the normal sense to cut parts of the volumetric visualization but also is shown in an unprojected and undistorted way to provide a 2D input space to specify locations or regions for seed point placement. Moreover, the system was designed to allow up to two people to interact at the same time due to the split-screen setup (Fig. 15(b)). Klein et al. also carried out a qualitative evaluation with domain experts from fluid mechanics which showed that the collaboration aspect of their design was liked by most participants, and that despite the use of a vertical display this collaboration worked well—to some degree contradicting previous work on the subject [68]. In addition, the evaluation also revealed the need of precise, constrained, and/or tightly controlled interactions, in addition to the fluid and flexible navigation techniques implemented in the system.

The systems discussed so far relied on vertical displays of the data—as it is common in many visualizations of spatial 3D datasets. Some types of data, however, inherently favor a horizontal mapping, such as surgery-based medical visualizations. Lundström et al. [61] thus use such a horizontal setup to create their virtual surgery



Fig. 16 The virtual surgery table [61], a horizontal interactive surface for the analysis of medical datasets. Image courtesy of and © Sectra, used by permission.

table (Fig. 16). Similar to the setup described before, the virtual surgery table also was designed with collaboration in mind, but in this case with collaborators located around the table’s horizontal surface (Fig. 16(b)). To specifically support this collaboration, they introduce “movable alternator pucks” which allow doctors to switch between the different interaction modalities in the system in a user-controlled fashion. For the main navigation interactions they use a typical 6DOF one- and two-finger mapping that supports x/y -panning, $x/y/z$ -rotation, and uniform zoom. Lundström et al. also conducted an observational study with domain experts (five medical doctors). This study provided numerous insights on the usability of the design, its clinical usefulness, and needed additional features for the system to be used in practice. In particular, the study demonstrated that a system such as the virtual surgery table is particularly useful for the analysis of complex cases, an insight that may also be possible to extend to other application domains such as data analysis by scientists in other domains. The participants also reported that a pure 3D interaction is not sufficient—the possibility to view and interact with additional 2D views such as traditional slices is needed. In the meantime, the research on the virtual surgery table has led to the founding of a company (Sectra) which has continued to develop the system into a product that is now actively being marketed (as it is evident in the pictures shown in Fig. 16), and other companies (e. g., Anatomage) are offering similar setups.

The horizontal form factor has also been used by Sultanum et al. [73, 74] for their system to support the analysis of geologic reservoir data (Fig. 17). Their system is based on volumetric datasets that capture geological features such as seismic data, different surface layers, permeability levels, oil saturation levels, etc. Sultanum et al.’s system then allows geologists to explore these different aspects of the model, both by looking at the different data attributes (e. g., using physical property cards [73]—similar to Lundström et al.’s [61] virtual alternator pucks) as well as by manipulating the volumetric model itself. The latter is done by interactions such as




Fig. 17 Data exploration techniques for reservoir visualization [73, 74]. Images courtesy of and © Nicole Sultanum, Sowmya Somanath, Ehud Sharlin, and Mario Costa Sousa, used by permission.

splitting (Fig. 17(b)), zipping it (Fig. 17(c), (d)), or layer peeling (Fig. 17(e)). The interaction by means of tangible objects also facilitates data readout (Fig. 17(a)) and focus+context visualizations (Fig. 17(f)). One particularly interesting aspect of this interaction design is the large set of interaction techniques that are mapped in a non-conflicting way, enabling both navigation/view correction as well as several dataset manipulations using a coherent interaction design. Sultanum et al. not only base their work on observational sessions with the target users [73] but also conducted an evaluation of their final design [74] with domain experts. This last study revealed that, while the participants liked the overall system design with its flexible and fluid data exploration, they too asked for specific precise views such as 2D projections, similar to what was reported by Klein et al. [55] and Lundström et al. [61].

The systems discussed so far rely on the interaction with 2D projections of the 3D data visualizations—largely due to the interaction problems that arise from the combination of touch input and stereoscopic projection (as discussed in Section 2.2). We are only aware of two systems that use interaction setups that include stereoscopic projections, the hybrid Slice WIM setup by Coffey et al. [21, 22] for medical data analysis and the purely stereoscopic setup by Butkiewicz and Ware [15, 16] for oceanographic data. Coffey et al.’s [21, 22] Slice WIM (Fig. 2) combines a large vertical stereoscopic projection of the explored 3D data with a horizontal tactile input surface. The core aspect of their system is the use of a stereoscopically displayed world-in-miniature visualization of the entire dataset using the large vertical display that also casts a shadow/projection onto the horizontal interaction surface. This connects both views and allows users to mentally map their 2D input into the stereoscopic 3D scene. The input itself is based on an interaction widget that allows the user to navigate the 3D view, manipulate exploration elements such as cutting planes, create curves in 3D space for path planning, and select subsets of the data such as bundles of flowlines as mentioned in Section 3.2.

In contrast to the previous hybrid setup, Butkiewicz and Ware [15, 16] use a purely stereoscopic data display that is unique in several ways. Their setup includes a slightly




Fig. 18 Purely stereoscopic setup for the exploration of oceanographic data [15, 16]. Image courtesy of and © Thomas Butkiewicz, used by permission.

tilted stereoscopic screen on which also the tactile input is provided (captured using a depth camera) and a data view that is shown at a similarly tilted angle (Fig. 18). In addition, they designed their system for the exploration of oceanographic data. All these aspects together make it possible that the stereoscopic display does not conflict with tactile interaction because (a) the oceanographic data is usually rather shallow if viewed at a scale of large water bodies or oceans, (b) this data has an inherent surface with which to interact—the water surface, and (c) the tilted view of the data projection is quite similar to that of the display setup. This means that the touch surface can be easily placed roughly at the inherent interaction surface (i. e., with zero parallax), without flexible 3 DOF rotations being necessary that would disrupt this ideal alignment. Using this setup they allow 2 DOF translations of the visualization, zooming, the placement of dye poles that allow the exploration of dynamic aspects of the displayed flow simulation. Butkiewicz and Ware also use specific precise interaction techniques such as two-handed mappings for the exact placement of dye poles.

Other interaction techniques with stereoscopic 3D data displays and visualizations are being investigated such as the use of a grasping-based metaphor [27] or the use of the monoscopic display in a hybrid setup as a mobile input device [60]. Yet, none of these approaches has yet led to a complete design study so we do not describe them in detail in this list of systems.

5 Conclusion and Open Research Challenges

With this survey we have demonstrated that the surface-based exploration of 3D visualizations is not only an active field of research but also has led to approaches that support the basic interaction techniques including 3D navigation, selection, data manipulation, seed particle placement, and more. The surface-based interaction benefits from the ability to provide spatial input directly at the location where the data is shown, thus inherently supports direct manipulation which is essential for the exploration of scientific data.

Yet, several research questions remain open for the field of surface-based exploration of three-dimensional, spatially explicit scientific data. For example, it seems clear that tactile interaction will only become another interaction modality to explore data, it will by no means replace existing approaches such as VR-based environments or traditional workstation settings. This means that the integration of surface-based data exploration into a **practical workflow for domain experts** is a pressing issue with the goal of providing an **interaction continuum** [41] in which the data analysts can choose the best interaction paradigm for the situation as well as easily transition between different paradigms as necessary.

To arrive at such a continuum, we have to continue the work on better **understanding the suitability of different input paradigms** for different data exploration tasks to be able to use tactile, tangible, haptic, traditional, or other sensing as it works best (e. g., [9, 48, 49]). The future continuum can then also include data exploration environments that effectively integrate tactile, surface-based interaction with other visual or interaction paradigms such as stereoscopic views (e. g., [15, 16, 21, 22, 60]) or tangible input (e. g., [71, 74]).

Connected to this issue of creating an interaction continuum is the challenge of providing **coherent interaction designs** even for a single interaction paradigm. The techniques reviewed in Section 3, in particular, often focus on a single type of manipulation only. Their interaction mappings are thus relatively flexible and without many external constraints. Yet, in practice analysts require a whole toolkit of data exploration techniques. Section 4 provided some examples for how several interaction techniques can be combined using coherent mappings. In practice, however, it is likely that many more techniques are needed so that more work is necessary to understand how to best integrate a large set of interaction techniques within the context of surface-based data exploration.

Another open research question is the issue of the applicability of widget-based or of gestural specifications of the interaction intents. In our discussion we, on purpose and with only a few exceptions [6, 59], did not mention the use of gestural interaction techniques—it turns out that almost all published interaction techniques rely on interaction mappings that are clearly specified based on the location of input points with respect to the data or interaction widgets. Such posture-based interaction [42] has the benefit that any input point motion can directly be interpreted as data manipulations. This interaction specification paradigm thus avoids complex mode specification and allows users of respective systems to concentrate on the data analysis. Yet, in some situations such as the initiation of specific data exploration

actions the use of gestures may be useful, so this question of **widget-based vs. gestural interaction control** still needs to be further explored in the future.

From our survey it also became apparent that the data exploration scenarios that exist today focus primarily on single-user interaction. Only two of the discussed systems [55, 61] were specifically designed with **collaboration and parallel input** in mind. However, collaboration is still possible with the other systems when the different collaborators are taking turns. It would thus be nice to see more work in the future that specifically explores collaborative settings, both those that require turn-taking and those that allow parallel work.

Acknowledgments

I would like to thank all the authors of the original papers who kindly provided me with images for their techniques in order for me to be able to illustrate the described interaction techniques and systems. Also, I would like to thank all people who have collaborated with me in the past on surface-based interaction techniques for 3D visualization, in particular my past and present students who worked with me on this topic including Lonni Besançon, Tijmen Klein, David López, Pjotr Svetachov, Thomas ten Cate, and Lingyun Yu.

References

1. Abeele, S., Bock, O.: Mechanisms for sensorimotor adaptation to rotated visual input. *Experimental Brain Research* **139**(2), 248–253 (2001). DOI 10.1007/s002210100768
2. Abeele, S., Bock, O.: Sensorimotor adaptation to rotated visual input: Different mechanisms for small versus large rotations. *Experimental Brain Research* **140**(4), 407–410 (2001). DOI 10.1007/s002210100846
3. Akers, D.: CINCH: A cooperatively designed marking interface for 3D pathway selection. In: *Proc. UIST*, pp. 33–42. ACM, New York (2006). DOI 10.1145/1166253.1166260
4. Argelaguet, F., Andujar, C.: Efficient 3D pointing selection in cluttered virtual environments. *IEEE Computer Graphics and Applications* **29**(6), 34–43 (2009). DOI 10.1109/MCG.2009.117
5. Argelaguet, F., Andujar, C.: A survey of 3D object selection techniques for virtual environments. *Computers & Graphics* **37**(3), 121–136 (2013). DOI 10.1016/j.cag.2012.12.003
6. Au, O.K.C., Tai, C.L., Fu, H.: Multitouch gestures for constrained transformation of 3D objects. *Computer Graphics Forum* **31**(2pt3), 651–660 (2012). DOI 10.1111/j.1467-8659.2012.03044.x
7. Banic, A.: Selection classification for interaction with immersive volumetric visualizations. In: *Human Interface and the Management of Information: Information and Knowledge Design and Evaluation, LNCS*, vol. 8521, pp. 10–21. Springer, Cham/Heidelberg (2014). DOI 10.1007/978-3-319-07731-4_2
8. Bedford, F.: Perceptual learning. In: D.L. Medin (ed.) *The Psychology of Learning and Motivation*, vol. 30, pp. 1–60. Academic Press, Inc., New York (1993). DOI 10.1016/S0079-7421(08)60293-5
9. Besançon, L., Issartel, P., Ammi, M., Isenberg, T.: Usability comparison of mouse, touch and tangible inputs for 3D data manipulation. *Tech. Rep. 1603.08735*, arXiv.org (2016)

10. Bock, O.: Basic principles of sensorimotor adaptation to different distortions with different effectors and movement types: A review and synthesis of behavioral findings. *Frontiers in Human Neuroscience* 7(81) (2013). DOI 10.3389/fnhum.2013.00081
11. Bogorin, M.A., Luderschmidt, J., Dörner, R., Geiger, C.: ViNet – Interaction with information visualizations in VR applications via multi-touch and tangible user interfaces. In: Proc. 6th GI Workshop “Virtuelle und Erweiterte Realität”, pp. 211–222. Shaker-Verlag, Aachen (2009)
12. Bowman, D.A., Kruijff, E., LaViola Jr., J.J., Poupyrev, I.: 3D User Interfaces: Theory and Practice. Addison-Wesley, Boston (2005)
13. Bruder, G., Steinicke, F.: 2.5D touch interaction on stereoscopic tabletop surfaces. In: Proc. ISIS3D, pp. 1–4 (2013)
14. Bruder, G., Steinicke, F., Stürzlinger, W.: To touch or not to touch? Comparing 2D touch and 3D mid-air interaction on stereoscopic tabletop surfaces. In: Proc. SUI, pp. 9–16. ACM, New York (2013). DOI 10.1145/2491367.2491369
15. Butkiewicz, T., Ware, C.: Exploratory analysis of ocean flow models with stereoscopic multi-touch. In: IEEE Visualization Posters (2011)
16. Butkiewicz, T., Ware, C.: Multi-touch 3D exploratory analysis of ocean flow models. In: Proc. OCEANS, pp. 746–755. IEEE, Los Alamitos (2011)
17. Buxton, B.: Multi-touch systems that I have known and loved. Website: <http://www.billbuxton.com/multitouchOverview.html> (2007). Updated last on June 12, 2014, visited in March 2016
18. Card, S.K., Mackinlay, J.D., Shneiderman, B.: Readings in Information Visualization: Using Vision to Think. Morgan Kaufmann, San Francisco (1999)
19. Chan, L.W., Kao, H.S., Chen, M.Y., Lee, M.S., Hsu, J., Hung, Y.P.: Touching the void: Direct-touch interaction for intangible displays. In: Proc. CHI, pp. 2625–2634. ACM, New York (2010). DOI 10.1145/1753326.1753725
20. Coffey, D., Korsakov, F., Keefe, D.F.: Low cost VR meets low cost multi-touch. In: Proc. ISVC, vol. 2, pp. 351–360. Springer, Berlin, Heidelberg (2010). DOI 10.1007/978-3-642-17274-8_35
21. Coffey, D., Malbraaten, N., Le, T., Borazjani, I., Sotiropoulos, F., Erdman, A.G., Keefe, D.F.: Interactive Slice WIM: Navigating and interrogating volume datasets using a multi-surface, multi-touch VR interface. *IEEE TVCG* 18(10), 1614–1626 (2012). DOI 10.1109/TVCG.2011.283
22. 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: Proc. I3D, pp. 191–198. ACM, New York (2011). DOI 10.1145/1944745.1944777
23. Cohé, A., Dècle, F., Hachet, M.: tBox: A 3D transformation widget designed for touch-screens. In: Proc. CHI, pp. 3005–3008. ACM, New York (2011). DOI 10.1145/1978942.1979387
24. Colley, A., Häkkinä, J., Schöning, J., Daiber, F., Steinicke, F., Krüger, A.: Touch the 3rd dimension! Understanding stereoscopic 3D touchscreen interaction. In: Computer-Human Interaction: Cognitive Effects of Spatial Interaction, Learning, and Ability (extended papers of OzCHI 2013), *LNCS*, vol. 8433, pp. 47–67. Springer, Berlin/Heidelberg (2015). DOI 10.1007/978-3-319-16940-8_3
25. Cunningham, H.A.: Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. *Journal of Experimental Psychology: Human Perception and Performance* 15(3), 493–506 (1989). DOI 10.1037/0096-1523.15.3.493
26. Cutler, L.D., Fröhlich, B., Hanrahan, P.: Two-handed direct manipulation on the responsive workbench. In: Proc. SI3D, pp. 107–114. ACM, New York (1997). DOI 10.1145/253284.253315
27. Daiber, F.: Interaction with stereoscopic data on and above multi-touch surfaces. In: ITS Doctoral Colloquium, pp. 2:1–2:4. ACM, New York (2011). DOI 10.1145/2076354.2076428
28. van Dam, A., Forsberg, A.S., Laidlaw, D.H., LaViola, J.J., Simpson, R.M.: Immersive VR for scientific visualization: A progress report. *IEEE Computer Graphics and Applications* 20(6), 26–52 (2000). DOI 10.1109/38.888006
29. Fu, C.W., Goh, W.B., Ng, J.A.: 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

30. Giesler, A., Valkov, D., Hinrichs, K.: Void shadows: Multi-touch interaction with stereoscopic objects on the tabletop. In: Proc. SUI, pp. 104–112. ACM, New York, NY, USA (2014). DOI 10.1145/2659766.2659779
31. Hachet, M., Bossavit, B., Cohé, A., de la Rivière, J.B.: Toucheo: Multitouch and stereo combined in a seamless workspace. In: Proc. UIST, pp. 587–592. ACM, New York (2011). DOI 10.1145/2047196.2047273
32. Hachet, M., de la Rivière, J.B., Laviolle, J., Cohé, A., Cursan, S.: Touch-based interfaces for interacting with 3D content in public exhibitions. *IEEE Computer Graphics and Applications* **33**(2), 80–85 (2013). DOI 10.1109/MCG.2013.34
33. Hancock, M., ten Cate, T., Carpendale, S.: 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
34. Hancock, M., ten Cate, T., Carpendale, S., Isenberg, T.: Supporting sandtray therapy on an interactive tabletop. In: Proc. CHI, pp. 2133–2142. ACM, New York (2010). DOI 10.1145/1753326.1753651
35. Hancock, M.S., Carpendale, S., Vernier, F.D., Wigdor, D., Shen, C.: Rotation and translation mechanisms for tabletop interaction. In: Proc. Tabletop, pp. 79–88. IEEE Computer Society, Los Alamitos (2006). DOI 10.1109/TABLETOP.2006.26
36. Hand, C.: A survey of 3D interaction techniques. *Computer Graphics Forum* **16**(5), 269–281 (1997). DOI 10.1111/1467-8659.00194
37. Hornecker, E., Marshall, P., Dalton, N.S., Rogers, Y.: Collaboration and interference: Awareness with mice or touch input. In: Proc. CSCW, pp. 167–176. ACM, New York (2008). DOI 10.1145/1460563.1460589
38. Isenberg, P., Isenberg, T.: Visualization on interactive surfaces: A research overview. *i-com* **12**(3), 10–17 (2013). DOI 10.1524/icom.2013.0020
39. Isenberg, P., Isenberg, T., Hesselmann, T., Lee, B., von Zadow, U., Tang, A.: Data visualization on interactive surfaces: A research agenda. *IEEE Computer Graphics and Applications* **33**(2), 16–24 (2013). DOI 10.1109/MCG.2013.24
40. Isenberg, T.: Position paper: Touch interaction in scientific visualization. In: Proc. DEXIS, pp. 24–27 (2011)
41. Isenberg, T.: An interaction continuum for visualization. In: Proc. IEEE VIS Workshop on “Death of the Desktop: Envisioning Visualization without Desktop Computing” (2014)
42. Isenberg, T., Hancock, M.: *Gestures* vs. postures: ‘Gestural’ touch interaction in 3D environments. In: Proc. 3DCHI, pp. 53–61 (2012)
43. Jackson, B., Coffey, D., Keefe, D.F.: Force Brushes: Progressive data-driven haptic selection and filtering for multi-variate flow visualizations. In: Short Paper Proc. EuroVis, pp. 7–11. Eurographics Association, Goslar, Germany (2012). DOI 10.2312/PE/EuroVisShort/EuroVisShort2012/007-011
44. Jackson, B., Lau, T.Y., Schroeder, D., Toussaint Jr., K.C., Keefe, D.F.: A lightweight tangible 3D interface for interactive visualization of thin fiber structures. *IEEE Transactions on Visualization and Computer Graphics* **19**(12), 2802–2809 (2013). DOI 10.1109/TVCG.2013.121
45. Jackson, B., Schroeder, D., Keefe, D.F.: Nailing down multi-touch: Anchored above the surface interaction for 3D modeling and navigation. In: Proc. Graphics Interface, pp. 181–184. Canadian Information Processing Society, Toronto (2012). DOI 10.20380/GI2012.23
46. Jankowski, J., Hachet, M.: 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
47. Jansen, Y., Dragicevic, P.: An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics* **19**(12), 2396–2405 (2013). DOI 10.1109/TVCG.2013.134
48. Jansen, Y., Dragicevic, P., Fekete, J.D.: Tangible remote controllers for wall-size displays. In: Proc. CHI, pp. 2865–2874. ACM, New York (2012). DOI 10.1145/2207676.2208691
49. Jansen, Y., Dragicevic, P., Fekete, J.D.: Evaluating the efficiency of physical visualizations. In: Proc. CHI, pp. 2593–2602. ACM, New York (2013). DOI 10.1145/2470654.2481359

50. Jönsson, D., Falk, M., Ynnerman, A.: Intuitive exploration of volumetric data using dynamic galleries. *IEEE Transactions on Visualization and Computer Graphics* **22**(1), 896–905 (2016). DOI 10.1109/TVCG.2015.2467294
51. Keefe, D.F.: Integrating visualization and interaction research to improve scientific workflows. *IEEE Computer Graphics and Applications* **30**(2), 8–13 (2010). DOI 10.1109/MCG.2010.30
52. Keefe, D.F., Isenberg, T.: Reimagining the scientific visualization interaction paradigm. *IEEE Computer* **46**(5), 51–57 (2013). DOI 10.1109/MC.2013.178
53. Keefe, D.F., Zeleznik, R.C., Laidlaw, D.H.: Tech-note: Dynamic dragging for input of 3D trajectories. In: *Proc. 3DUI*, pp. 51–54. IEEE Computer Society, Los Alamitos (2008). DOI 10.1109/3DUI.2008.4476591
54. Kin, K., Agrawala, M., DeRose, T.: Determining the benefits of direct-touch, bimanual, and multifinger input on a multitouch workstation. In: *Proc. Graphics Interface*, pp. 119–124. CIPS, Toronto (2009). DOI 10.20380/GI2009.16
55. Klein, T., Guéniat, F., Pastur, L., Vernier, F., Isenberg, T.: A design study of direct-touch interaction for exploratory 3D scientific visualization. *Computer Graphics Forum* **31**(3), 1225–1234 (2012). DOI 10.1111/j.1467-8659.2012.03115.x
56. Knoedel, S., Hachet, M.: Multi-touch RST in 2D and 3D spaces: Studying the impact of directness on user performance. In: *Proc. 3DUI*, pp. 75–78. IEEE Computer Society, Los Alamitos (2011). DOI 10.1109/3DUI.2011.5759220
57. Krüger, W., Fröhlich, B.: The responsive workbench. *IEEE Computer Graphics and Applications* **14**(3), 12–15 (1994). DOI 10.1109/38.279036
58. Kruszyński, K.J., van Liere, R.: Tangible props for scientific visualization: Concept, requirements, application. *Virtual Reality* **13**(4), 235–244 (2009). DOI 10.1007/s10055-009-0126-1
59. Liu, J., Au, O.K.C., Fu, H., Tai, C.L.: Two-finger gestures for 6DOF manipulation of 3D objects. *Computer Graphics Forum* **31**(7), 2047–2055 (2012). DOI 10.1111/j.1467-8659.2012.03197.x
60. López, D., Oehlberg, L., Doger, C., Isenberg, T.: 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 (2016). DOI 10.1109/TVCG.2015.2440233
61. Lundström, C., Rydell, T., Forsell, C., Persson, A., Ynnerman, A.: Multi-touch table system for medical visualization: Application to orthopedic surgery planning. *IEEE Transactions on Visualization and Computer Graphics* **17**(12) (2011). DOI 10.1109/TVCG.2011.224
62. Marton, F., Rodriguez, M.B., Bettio, F., Agus, M., Villanueva, A.J., Gobbetti, E.: IsoCam: Interactive visual exploration of massive cultural heritage models on large projection setups. *Journal on Computing and Cultural Heritage* **7**(2), 12:1–12:24 (2014). DOI 10.1145/2611519
63. Moscovich, T., Hughes, J.F.: Indirect mappings of multi-touch input using one and two hands. In: *Proc. CHI*, pp. 1275–1284. ACM, New York (2008). DOI 10.1145/1357054.1357254
64. Novotný, M., Lacko, J., Samuelčík, M.: Applications of multi-touch augmented reality system in education and presentation of virtual heritage. *Procedia Computer Science* **25**, 231–235 (2013). DOI 10.1016/j.procs.2013.11.028
65. Owada, S., Nielsen, F., Igarashi, T.: Volume catcher. In: *Proc. I3D*, pp. 111–116. ACM, New York (2005). DOI 10.1145/1053427.1053445
66. Reisman, J.L., Davidson, P.L., Han, J.Y.: 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
67. Robles-De-La-Torre, G.: The importance of the sense of touch in virtual and real environments. *IEEE MultiMedia* **13**(3), 24–30 (2006). DOI 10.1109/MMUL.2006.69
68. Rogers, Y., Lindley, S.: Collaborating around vertical and horizontal large interactive displays: Which way is best? *Interacting with Computers* **16**(6), 1133–1152 (2004). DOI 10.1016/j.intcom.2004.07.008
69. Schmalstieg, D., Encarnação, L.M., Szalavári, Z.: Using transparent props for interaction with the virtual table. In: *Proc. I3D*, pp. 147–153. ACM, New York (1999). DOI 10.1145/300523.300542

70. Shan, G., Xie, M., Li, F., Gao, Y., Chi, X.: Interactive visual exploration of halos in large-scale cosmology simulation. *Journal of Visualization* **17**(3), 145–156 (2014). DOI 10.1007/s12650-014-0206-5
71. Song, P., Goh, W.B., Fu, C.W., Meng, Q., Heng, P.A.: 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
72. Steinicke, F., Hinrichs, K.H., Schöning, J., Krüger, A.: Multi-touching 3D data: Towards direct interaction in stereoscopic display environments coupled with mobile devices. In: Proc. AVI Workshop on Designing Multi-Touch Interaction Techniques for Coupled Public and Private Displays, pp. 46–49 (2008)
73. Sultanum, N., Sharlin, E., Sousa, M.C., Miranda-Filho, D.N., Eastick, R.: Touching the depths: Introducing tabletop interaction to reservoir engineering. In: Proc. ITS, pp. 105–108. ACM, New York (2010). DOI 10.1145/1936652.1936671
74. Sultanum, N., Somanath, S., Sharlin, E., Sousa, M.C.: “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
75. Sultanum, N., Vital Brazil, E., Costa Sousa, M.: Navigating and annotating 3D geological outcrops through multi-touch interaction. In: Proc. ITS, pp. 345–348. ACM, New York (2013). DOI 10.1145/2512349.2512396
76. Sundén, E., Bock, A., Jönsson, D., Ynnerman, A., Ropinski, T.: Interaction techniques as a communication channel when presenting 3D visualizations. In: Proc. 3DVis, pp. 61–65. IEEE, Los Alamitos (2014). DOI 10.1109/3DVis.2014.7160102
77. Tan, D.S., Gergle, D., Scupelli, P., Pausch, R.: Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction* **13**(1), 71–99 (2006). DOI 10.1145/1143518.1143521
78. Teather, R.J., Stuerzlinger, W.: Pointing at 3D targets in a stereo head-tracked virtual environment. In: Proc. 3DUI, pp. 87–94 (2011). DOI 10.1109/3DUI.2011.5759222
79. Telea, A.C.: *Data Visualization: Principles and Practice*, 2nd edn. A K Peters/CRC Press (2015)
80. Tong, X., Chen, C.M., Shen, H.W., Wong, P.C.: Interactive streamline exploration and manipulation using deformation. In: Proc. PacificVis, pp. 1–8. IEEE, Los Alamitos (2015). DOI 10.1109/PACIFICVIS.2015.7156349
81. Valkov, D., Giesler, A., Hinrichs, K.H.: Imperceptible depth shifts for touch interaction with stereoscopic objects. In: Proc. CHI, pp. 227–236. ACM, New York (2014). DOI 10.1145/2556288.2557134
82. Valkov, D., Steinicke, F., Bruder, G., Hinrichs, K.: 2D touching of 3D stereoscopic objects. In: Proc. CHI, pp. 1353–1362. ACM, New York (2011). DOI 10.1145/1978942.1979142
83. Valkov, D., Steinicke, F., Bruder, G., Hinrichs, K.H., Schöning, J., Daiber, F., Krüger, A.: Touching floating objects in projection-based virtual reality environments. In: Proc. EGVE/EuroVR/VEC, pp. 17–24. Eurographics Association, Goslar, Germany (2010). DOI 10.2312/EGVE/JVRC10/017-024
84. Ware, C., Arseneault, R.: Frames of reference in virtual object rotation. In: Proc. APGV, pp. 135–141. ACM, New York (2004). DOI 10.1145/1012551.1012576
85. Welch, R.B.: *Perceptual Modification: Adapting to Altered Sensory Environments*. Academic Press, New York (1978). DOI 10.1016/B978-0-12-741850-6.50002-1
86. Wiebel, A., Vos, F.M., Foerster, D., Hege, H.C.: WYSIWYP: What you see is what you pick. *IEEE Transactions on Visualization and Computer Graphics* **18**(12), 2236–2244 (2012). DOI 10.1109/TVCG.2012.292
87. Wills, G.J.: Selection: 524,288 ways to say “This is interesting”. In: Proc. InfoVis, pp. 54–60. IEEE Computer Society, Los Alamitos (1996). DOI 10.1109/INFVIS.1996.559216
88. Yu, L., Efstathiou, K., Isenberg, P., Isenberg, T.: Efficient structure-aware selection techniques for 3D point cloud visualizations with 2DOF input. *IEEE Transactions on Visualization and Computer Graphics* **18**(12), 2245–2254 (2012). DOI 10.1109/TVCG.2012.217
89. Yu, L., Efstathiou, K., Isenberg, P., Isenberg, T.: 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 (2016). DOI 10.1109/TVCG.2015.2467202

90. Yu, L., Svetachov, P., Isenberg, P., Everts, M.H., Isenberg, T.: FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces. *IEEE Transactions on Visualization and Computer Graphics* **16**(6), 1613–1622 (2010). DOI 10.1109/TVCG.2010.157