

Hybrid Touch/Tangible Spatial 3D Data Selection

Lonni Besançon,^{1,3} Mickael Sereno,^{2,3} Lingyun Yu,⁴ Mehdi Ammi,⁵ and Tobias Isenberg²

¹Linköping Universitet, Norrköping, Sweden

²Inria, France

³Université Paris-Saclay, France

⁴University of Groningen, the Netherlands

⁵Université Paris 8, France




Figure 1: Tangible Brush technique: (a) drawing selection shape; (b) tangible shape extrusion; and (c) selection result.

Abstract

We discuss spatial selection techniques for three-dimensional datasets. Such 3D spatial selection is fundamental to exploratory data analysis. While 2D selection is efficient for datasets with explicit shapes and structures, it is less efficient for data without such properties. We first propose a new taxonomy of 3D selection techniques, focusing on the amount of control the user has to define the selection volume. We then describe the 3D spatial selection technique Tangible Brush, which gives manual control over the final selection volume. It combines 2D touch with 6-DOF 3D tangible input to allow users to perform 3D selections in volumetric data. We use touch input to draw a 2D lasso, extruding it to a 3D selection volume based on the motion of a tangible, spatially-aware tablet. We describe our approach and present its quantitative and qualitative comparison to state-of-the-art structure-dependent selection. Our results show that, in addition to being dataset-independent, Tangible Brush is more accurate than existing dataset-dependent techniques, thus providing a trade-off between precision and effort.

1. Introduction

Many visualization systems rely on exploratory data visualization and analysis [Tuk77] for the examination of previously unknown data. Such exploratory visualization approaches inherently rely on intuitive and effective interaction techniques that support a flexible specification of data analysis tasks and their parameters [Joh04, Kee10]. Essential for such exploratory analysis is the interactive selection of specific regions of interest [Wil96] to reveal their interesting patterns, properties, or internal structures. 2D regions are usually defined using picking, brushing, or lassos—easily achieved with a mouse/pen or on a touch[†] screen (both modalities provide the needed 2 degrees of freedom, DOF) or, more recently, by combining a brain-computer interface and gaze input [ÉAC*16].

3D datasets, however, often require the specification of 3D volumetric data subsets; an interactive task that is not easily supported by

mouse or touch input for general volumes. Existing techniques focused on creating constrained selections with these input paradigms (i. e., raycasting [AA09, OF03, FHZ96, GB06, LBGC07], automated 3D projection of 2D lassos [IMT99, YEII12, YEII16], and data-dependent picking [WVFH12]). These techniques either rely on contextual data analysis or pre-defined selection volumes and do not allow users to freely create their selection volume. While these interactive techniques are useful in many visualization contexts, they are limited when regions need to be selected which are difficult to define based on their structural or contextual features. These regions would then be missed by algorithms that rely on such features.

The recent development of mobile, spatially-aware devices, however, makes tangible manipulations a possible solution to the limit of 2 DOF touch input. We thus combine 2D touch input with 3D mobile manipulation and create a free-form, context-independent selection technique (Fig. 1). Our technique does not suffer from this drawback: it relies solely on user input and is context-independent. We use a Project Tango [Tan] tablet's 3D-self-tracking and touch

[†] Other approaches in the literature use the term *tactile input* (e. g., [BIAI17a, BIAI17b, HW78, PM03]) as a synonym for touch input.

screen and use a hybrid interaction paradigm [BIAI17b]: touch input specifies a 2D drawn shape and 3D tablet manipulations extend the 2D shape into 3D by “brushing” the 3D space with the 2D shape. Such techniques are essential because all fully or partially automated approaches use inherent data properties to facilitate a selection. A specific technique that is efficient with a given dataset may thus fail for another. Our approach thus complements existing automatic, interactive, and context-aware 3D selection techniques with explicit adjustments of the selection volume in a tangible spatial manner.

Our contributions are threefold. First, we classify 3D selection techniques based on input and shape control, and derive a design space for context-free 3D selection techniques. Second, we propose a new hybrid touch/tangible technique for 3D spatial selection using a spatially-aware device to give users manual control of the 3D selection volume they create. Third, we compare this new technique to a state-of-the-art and structure-dependent technique. We report and discuss quantitative and qualitative findings of this comparison.

2. Context and Related Work

The ability to interact with the data representation is an essential feature of all visualization systems [Tuk77, Mun14] and has been considered to be a crucial and a major challenge of the visualization community for 3D representation from the 1960s [Sut66] to recent days [Hib99, Rhe02, Joh04, TM04, Kee10, KI13, BCD*18, WBG*19]. We thus discuss related work within the field of 3D interaction [Han97, JH13, Ise16, LKM*17], where our work relates to 3D selection techniques and to hybrid input paradigms used in 3D exploratory data analysis.

2.1. 3D Object Selection in VR Environments

Many 3D object selection techniques in virtual environments rely on 3D ray-casting [AA13]: a ray, cast from the user’s hand, selects the first object it hits. It suffers from three major limitations. First, the selected object may not be visible from the user’s view, while no obstructions exist along the selection ray, leading to unintentional selections. This *eye-hand visibility mismatch* is tackled by *raycasting from the eye* [AAT08], while other techniques such as the selection based only on eye movements [CSD03, TJ00] circumvent it entirely. Second, a desired object can be (partially) obscured by third-party objects, making first-hit-raycasting selection difficult. Some variations address this issue too [OF03, GB06, VGC07]. Finally, raycasting is often too imprecise to select distant/small objects. To solve this, *IntenSelect* [DHKP05] uses a scoring system, while *adaptive pointing* [KGDR09] improves the precision of pointing devices. Cone-selections [FHZ96, LG94, SP04, Ste06] can avoid the precision issue: objects get selected if they are (partially) contained within the cone. Several items can, however, be located within the cone, thus making single-item selection difficult. To circumvent this, Schmidt et al. [SBB*06] investigated a probabilistic approach, while Olwal et al. [OBF03] relied on statistical geometry to better capture the item intended by the user. Other primitives [VGC07, WHB06, ZBM94] are also used. As an alternative, Mine et al. [MBS97] used hand-held widgets (3D objects appearing in the user’s virtual hand), while Poupyrev et al.’s *Go-Go* technique [PBWI96] interactively grow the user’s virtual hand to reach distant objects.

All these approaches have been extensively used in VR environments to select a pre-defined shape or object in the environment. They can be adapted for AR environments with varying performances [KOM*03, LBGC07, OBF03]. However, some selection tasks require specific, yet arbitrary, 3D selection as described next.

2.2. 3D Selection for Exploratory Data Analysis

Exploratory data analysis [Tuk77] is one such case where users may want to focus on specific *regions of interest* (ROI) that are not necessarily predefined or distinct. The data is often unsegmented and requires the user to perform a segmentation of the volume which is complicated and time-consuming. Researchers thus proposed to manipulate specific volumes in the 3D space [ASM*04, SAM*05] (such as a box) to specify the ROI but that can also include undesired objects since target region is usually not cuboidal [LBCW05a, YEII12].

For such datasets, some approaches rely on specific hardware. For instance, Harders et al. [HWS02] use a 3D mouse with force feedback to facilitate the segmentation of linear structures. Similarly, Malmberg et al. [MVN06] use a haptic device and stereoscopic rendering to allow users to draw 3D curves based on the 2D live-wire method. This idea was improved with *Spotlight* [THA10] which adds visual guidance to improve the quality of the segmentation. A similar setup is used by Nystrom et al. [NMVB09]. Finally, Jackson et al. [JLS*13] use a rolled paper as a tangible prop to facilitate selection of thin fiber structures. We also allow users to provide 3D tangible input but we combine it with an initial touch input.

Other approaches such as Wiebel et al.’s [WVFH12] are based on the data itself, the transfer function, and the way the volumetric rendering is perceived by the user. However, these techniques heavily rely on the ability to make sense of the data computationally which limits their applications to specific goals. For the selection of specific primitives beyond points, specific techniques exist. For instance, for linear structures in neuroimaging, selections can be performed by manipulating boxes or ellipsoid-shaped regions [ASM*04, SAM*05]. Akers et al. [Ake06] proposed to combine a trackball and a pen to help neuroscientists mark 3D pathways in neural datasets.

For unsegmented data, image segmentation can be used to infer the intended 3D subsets within the data [CSSM06, ONI05, YZNC05]. *Volume Catcher* [ONI05], e. g., avoids the initial phase of manually segmenting the volume. Users trace the target’s contour with a 2D free-form stroke, and a segmentation algorithm run by the system returns a corresponding ROI. In other words, 2D user input is algorithmically extrapolated into 3D. This approach was followed by others [CSSM06, YZNC05] and improved by Yuan et al. [YZNC05] to alleviate the user’s commitment when sketching. Context-aware and structure-aware techniques rely on a similar concept. These techniques have been shown to be efficient for 2D datasets [DS08, DS09] using the perceptual grouping of objects. The third dimension of most spatial datasets, however, increases the complexity of the problem, since most input paradigms have limited degrees of freedom [YEII16]. *CloudLasso* [YEII12] then allowed users to create a 3D selection based on 2D lasso input then extended into 3D based on the density of particles within that lasso. *CloudLasso* was extended by Shan et al. [SXL*14] by analyzing the different clusters created by *CloudLasso* and only selecting the one

with the largest 2D projection. Later, three interactive context-aware techniques (CAST) [YEII16] also selected a single connected component, two of which were based on the shape of the drawn lasso. We also use a drawn lasso with our approach but, instead of algorithmically transforming the 2D shape into a selection volume, we allow the user to specify the 3D selection volume directly by means of tangible input (meaning that the visualization $\mathcal{V}$, manipulation $\mathcal{M}$, interaction $\mathcal{I}$, and user $\mathcal{U}$ spaces are collapsed according to Bruckner et al.'s model of spatial interaction directness [BIRW19]).

2.3. Combination of Touch and Tangible Interaction

The benefits and limitations of touch and tangible input highlighted in the literature [BIAI17a, BRLD17, KSMC05, ZGO13] suggest that they are complementary paradigms that can be combined. Previous systems combining them are found in tangible additions to tabletops [AMR16, JGAK07, PBD*16]. In particular for 3D visualization, Sultanum et al. [SSSS11] created a table-based system for geologic reservoir data. In their setup, props facilitate data read-outs and parametrize a focus+context view, while touch specifies data navigation and dedicated exploration (splitting, peeling, ...). Linked to our spatially-aware-device approach, Olwal and Feiner [OF09] used a spatially-aware small display on a large tabletop. It was tracked in 2D and showed a section of the data on the tabletop at a higher resolution. Touch input was used on both devices. Lapides et al. [LSSS06] used a 2D tablet on a height-adjustable surface to create a 3D drawing board. We build on this idea to use a spatially-aware display with touch input. These tabletop-based systems, however, do not take full advantage of tangible physicality since they remain on the tabletop. Only their 2D positions and orientation are used, thus losing the rich 3D manipulations offered by physical devices.

In VR, some approaches [HG02, HGR03] combined isotonic and isometric 6-DOF input to rotate and translate 3D scenes with a pen-based tablet, to facilitate 2D selections. More related to 3D visualization, Steimle et al. [SJM13] investigated the use of 6 DOF spatially-aware tangible deformable surfaces to visualize slices of 3D volume/2D+time datasets. The use of touch/pen-based and 3D pen-based input with such surfaces also have been investigated for 3D visualization and interactions [STSD10, SBD12, SSFG98]. Lopez et al.'s [LODI16] investigated touch input on a mobile device for 3D visualization, combining the mobile device's monoscopic view and a stereoscopic view of the data and including a tangible interaction mode with the tablet's orientation to control the data. Also, Song et al. [SGF*11] combined a monoscopic view with a mobile device with touch input to place annotation lines into a 3D dataset. The physicality of the interaction has also been exploited for time-dependent data exploration by Sollich et al. [SvZP*16]. Besançon et al. [BIAI17b] then evaluated a system combining touch and tangible input with a spatially-aware tablet for 3D visualization. Their techniques take advantage of both input modalities to improve fluid dynamic researchers' workflow. We base our technique on these approaches to take advantage of both modalities for 3D selection, using touch input to draw the 2D lasso and tangible manipulations to extend this 2D shape into 3D, enabling true 3D spatial selection.

3. Classification of 3D Selection Techniques

Several surveys have previously covered different taxonomies to classify 3D selection techniques. All of them narrowly focus on the selection of dedicated 3D objects or the selection of 3D volumes (or regions of interest). While the techniques used for each of these two purposes are different, the abstract task is still 3D selection. Also, the degree to which the final selection can be precisely influenced or controlled by a user is rarely discussed in existing taxonomies. In contrast, we are interested precisely in this aspect of control that the user has for making the final selection. We thus expand existing taxonomies to include both 3D object and 3D volume selection as well as the user's level of control over the final selection.

3.1. Past Taxonomies of 3D Selection Techniques

Poupyrev and Ichikawa [PI99] focused on 3D object selection. They distinguished techniques based on whether the manipulation was *exocentric* or *egocentric* and then classified them based on the used *metaphor*. This latter classification is useful as different metaphors are needed in different settings or for different targets (raycasting is good for object selection but not for, e.g., unsegmented data). Issartel et al. [IBG*16] showed that exocentric and egocentric manipulations result in different performances depending on the visualized scene (for manipulations of spatially-aware displays). The distinction between egocentric and exocentric techniques can also be used to assess if a technique works in AR [LBGC07]. This taxonomy is important: it distinguishes techniques based both on their employed metaphor and on user's input which we believe is essential.

Later, Bowman et al. [BJH01] considered *feedback*, *confirmation mechanism*, and *object indication* as criteria to classify 3D selection techniques. While these criteria are important, Bowman et al.'s taxonomy does not include a classification of the inherent selection mechanism. Moreover, it introduces some redundancy because the choice of feedback is coupled to the object indication. More recently, Argelaguet and Andujar [AA13] provided a comprehensive survey and thorough classification of 3D selection techniques for virtual environments, basing their classification on the *selection tool*, the *degrees of freedom*, *DOF* of the selection tool manipulation, the employed *disambiguation mechanism*, the *control-display ratio*, and the *relationship between motor and visual space*. Their classification centered around the technique and its characteristics and focused on 3D selection techniques for dedicated objects. In our classification, however, we also include selection techniques that facilitate the specification of 3D selection volumes/regions of interest that thus are also applicable to volumetric data or particle datasets. Such forms of selection are important, in particular, in exploratory data analysis selection techniques such as those discussed in Sect. 2.

Surveys that focus on selecting ROIs—to the best of our knowledge—cover the topic of image/volume segmentation. Fu and Mui [FM81] distinguished between characteristic feature thresholding, edge detection, and region extraction. Haralick and Hapiro [HS85] differentiated between measurement space, space-guided spatial clustering, etc. Later, Khan [Kha13] distinguished between edge-based, PDE-based, region-based, etc. approaches. However, most segmentation techniques do not rely on user input and, hence, differ from the 3D spatial selection scenario on which we focus.

Table 1: Adjusted taxonomy (new parts in bold) based on Argelaguet and Andujar's [AA13] classification of 3D selection techniques.

criterion	description	∈ [AA13]	∈ [PI99]
selection metaphor/tool	ray, cone, brush, lasso, ...	✓	✓
selection control-DOF	refined in <i>selection tool control</i> below	✓	×
disambiguation-mechanism	not linked to control or input	✓	×
motor/visual-space-relation	not linked to control or input	✓	×
target selection type	dedicated objects vs. ROIs	×	×
selection shape creation	parameter-based, semi-automatic, user-controlled	×	×
selection shape adjustments	impossible, automatic, manual	×	×
selection tool control	selection tool DOF, input device DOF, and CD ratio	•	×

3.2. An Extended Taxonomy Focusing on User Control

This classification focuses on the amount of control a user has to define the final selection. While Argelaguet and Andujar [AA13] put some emphasis on input strategies with their taxonomy and, in particular, included the DOF of the selection tool (Table 1), we go beyond this approach as we include criteria that define whether the user or an algorithm is in control of the different steps involved in the selection mechanism. This way we can distinguish between a technique that gives little to no control to the user and another which is essentially built on the user input. It is then easier to determine whether a technique can easily adapt to new kinds of datasets or regions/objects of interest. We thus extend Argelaguet and Andujar's classification [AA13] by subdividing their *selection control DOF* criterion and adding others as shown in Table 1. To simplify, we also removed two further criteria because they had no relation with the user's input and control of the technique but rather deal with the mismatch between visual and motor space or how multiple selections were disambiguated (but these two aspects could again be included). We describe next the criteria that we use in our classification.

Selection Metaphor. Past taxonomies [PI99, AA13] used this criterion to specify the desired selection. A ray metaphor, e.g., computes the intersection with scene objects to highlight the currently selected object [AA13]. Similarly, the intersection of the box with volume data can be computed to obtain a selection of the data (e.g., [BF07, CMB*14]). The metaphor is an important criterion: it affects the control (e.g., w.r.t. the available DOF) of users over the interaction. It also dictates, in most cases, the final selection shape.

Target Selection Type. Previous taxonomies only investigated either dedicated object selection or ROI-based selection (Sect. 3.1). We cover both types with our taxonomy, thus add this criterion as a primary classification criterion. Specifically, we distinguish between single-object, multiple-object, and ROI-based selection.

Selection Shape Creation. ROI-based selections require the specification of a region. Simple approaches adjust pre-defined geometric shapes (e.g., [BF07, ASM*04, SAM*05, ZBM94]) to give quick selections, but with limited control over the result. Hybrid approaches rely on user input (e.g., 2D drawings on a projection of the 3D data) and the data/view context (e.g., lasso-based [CSSM06, LBCW05b, ONI05, YEII12, YEII16, YZNC05]). Such techniques provide more control over the result but require more input to specify the selection. Finally, the highest level of control would rely purely on (3D) user input, without any system

assistance to control parts of the selection process. To the best of our knowledge, no existing technique so far relies solely on user input to define the selection volume—but our new selection technique falls into this category. Even the 3D Lasso technique [ZCL08] only computes which tracts of interests are within the 3D-drawn lasso. Similarly, the 3D live-wire technique [MVN06] rather creates the volume that connects the two user-created 2D shapes.

Selection Shape Adjustments. Once the selection has been specified, users may still want to adjust it. As part of the core selection, there is either no post-creation control at all, automatic system-based control, or user-based control. For ray-casting or cone-casting, e.g., there is generally no control over the ray itself once the selection has been made. Some techniques such as the *flexible pointer* [OF03], however, employ user-controlled ray-bending to assist with the disambiguation of targets. Similarly, *Aperture* [FHZ96] uses a cone-casting selection whose size (apex angle) can be manually adjusted. We can also make this distinction for other techniques. The *Silk Cursor* [ZBM94], e.g., uses a pre-defined shape whose size cannot be adjusted. The *Bubble cursor* [VGC07] relies on a spherical selection tool that is dynamically and automatically extended to reach objects close to its center. *Aperture* [FHZ96], in contrast, allows users to dynamically control the size of the selection cone. Similarly, context-aware techniques [YEII12, YEII16] allow users to adjust an initially system-derived selection threshold.

Selection Tool Control. In their taxonomy, Argelaguet and Andujar [AA13] consider *tool control* as the way the “user is able to control” the selection shape. They distinguish two aspects: the selection tool DOF and the control-display ratio. We add another factor: the input device DOF (that they briefly mention as affecting the performance)—it is essential to understand the correspondence between physical input and virtual selection operator.

Selection tool DOF: To alleviate the lack of control given to the user on the volume selection creation, most techniques provide the users with mechanisms to adjust the selection volume's origin and orientation. Techniques lacking such controls which do not allow the users to create their own selection volume would have a negative impact on the variety of possible selections. Similarly, techniques that provide control over the origin of the selection tool only [ZBM94] are necessarily more limiting than techniques also providing orientation control [AAT08, FHZ96].

Input Device DOF: While the number of available DOF for the control of selection tool is important, it is similarly important to consider the DOF provided by the used input device. Having a mismatch between these two variables could lead to a bad performance. On the one hand, having more input DOF than necessary could be confusing for users as they do not see the changes provided by the additional unused DOF [WTW*05]. On the other hand, having less DOF than necessary is also a problem as users can report being confused by integrated manipulations [BIAI17a] or the interaction requires excessive mode switching.

Control-display (CD) ratio: It determines the conversion of manipulations of the input devices into manipulations in the virtual world. It is usually isomorphic (one-to-one mapping) but can take other static values. Because 3D selection is based on 3D input, a down-scaling CD ratio can lead to many user manipulations and increase fatigue. Similarly, having an up-scaling CD ratio can lead to

Table 2: Our taxonomy applied to a selection of techniques to illustrate the design space exploration of control.

technique	metaphor	target	shape creation	shape adjustment	selection DOF	input DOF	CD control
Balloon Selection [DFK12]	sphere	object	no control	no control	3 DOF	2 DOF	none
Go-Go [PBW196]	hand avatar	object	no control	no control	3 DOF	6 DOF	none
RayCasting [Min95]	ray	object	no control	no control	5 DOF (trans.: 3, rot.: 2)	up to 6 DOF	none
Framing Hands [PFC*97]	ray	object(s)	no control	no control	5 DOF	6 DOF	none
iSith [WBB06]	ray	object	no control	no control	5 DOF	12 DOF (2 hands)	none
Flexible Pointer [OF03]	curved ray	object	no control	no control	10 DOF	12 DOF (2 hands)	none
Silk Cursor [ZBM94]	box	object	no control	no control	3 DOF	6 DOF	none
Bubble Cursor [VGC07]	sphere	object	no control	automated adjust.	6 DOF	6 DOF	none
Aperture [FHZ96]	cone	object	no control	manual size adjust.	5 DOF (trans.: 2, rot.: 3)	6 DOF	none
SQUAD [KBB11]	sphere + ray	object	no control	automated adjust.	5 DOF	6 DOF	none
Senseshapes [OBF03]	cone/sphere/box	object	no control	no control	5 DOF	6 DOF	none
Cylinder Selection [LBCW05a]	lasso	ROI	semi-autom., 2D input	no control	2 DOF	2 DOF	none
CloudLasso [YEII12]	lasso	ROI	semi-autom., 2D input	threshold adjust.	2 DOF	2 DOF	none
CAST [YEII16]	lasso	ROI	semi-autom., 2D input	threshold adjust.	2 DOF	2 DOF	none
Tangible Brush	lasso + extrusion	ROI / objects	manual	none	up to 6 DOF	up to 6 DOF	manual

imprecise selections of small targets [KBB11]. Dynamic CD ratios address both issues. Argelaguet and Andujar [AA13] based their classification on König et al.'s [KGDR09] distinction between CD manual switching, target oriented CD, and velocity-oriented CD. We argue here that a velocity-oriented technique still provides a manual control of the CD ratio. We thus gather manual switching and velocity-oriented techniques under the manual control of the CD ratio. A further distinction can be made by distinguishing between explicit manual control (setting a value with a slider, buttons, or using specific gestures [VB05]) and implicit manual control (e. g., velocity-based [KGDR09] or pressure-based [BAII7]).

3.3. Taxonomy Classification of Existing Approaches

We can now classify classic 3D selection techniques with our taxonomy (Table 2). Ray Casting [Min95] uses the ray metaphor to select objects. This metaphor does not involve any shape creation so users don't need to control or adjust it. Classical implementations of Ray Casting rely on hand or eye tracking and uses up to 6 DOF but the users controls 5 DOF for the selection (3 for translations and 2 for rotations) with no need for a CD ratio control. Cylinder Selection [LBCW05a] uses a lasso to specify ROIs. Users partially control the volume creation as the computer generates a volume based on the user's 2D lasso. This technique usually does not let users modify the initial shape or the generated volume. User input relies on 2 DOF and so does the selection. The list of techniques described by our extended taxonomy in Table 2 is by no means exhaustive. However, we have carefully chosen techniques that represent each specific selection metaphor and their commonly envisioned alternatives.

We note that our taxonomy does not suggest that more control over the shape/volume creation leads to better technique. Indeed, one can quickly notice that context-dependent selection with semi-automatic shape creation would not fare well in VR where techniques usually target whole objects, while Ray Casting would perform poorly for particle data selection. Our taxonomy highlights two things, however. First, different tasks and environments call for different levels of control over the shape creation. Second, a manual shape creation technique has not yet been explored, while it could be useful in scientific visualization which sometimes requires specific yet arbitrary (i. e., no automatically identifiable features or

properties) 3D selections. We propose to fill this gap with our new manual selection technique, Tangible Brush (last row in Table 2 and Sect. 4) in order to address this need of scientific visualization to select specific yet arbitrarily shaped volumes.

4. Tangible Selection Brush

We are inspired by the *Tablet Freehand Lasso* [Luc05, LBCW05a] which allows users to draw a lasso on a 2D data projection, which is extruded as a generalized cone into 3D data space. We rely, however, on both touch and tangible inputs, thus use a hybrid tactile/tangible paradigm. While such hybrid techniques were shown to be efficient for data exploration, manipulation, and annotation [SGF*11, BIAII7b], we focus on the more sophisticated task of 3D spatial selection. We hence use a similar setup but design the interaction to serve the needs of 3D selection. Specifically, we ask users to draw a 2D brush on the tangible device that they then extend into 3D space by physically moving the device. In contrast to existing techniques for 3D selection, we thus give users full interactive control of the final selection volume, without algorithmic extrapolation.

4.1. Overall Interaction Design

Similar to the mentioned previous approaches, we start by asking users to draw a closed shape on the mobile device with respect to the displayed 2D projection of the data. To create closed strokes, we compute the Euclidean distance between start and end point and—if this distance is smaller than 0.2 (device screen units whose length is 2)—we connect the start and end points with a straight line. If this is not possible we remove the drawn path so that users can start anew.

In contrast to sketch-based modeling (e. g., [IMT99]) or context-aware selection [SXL*14, YEII12, YEII16], we facilitate the 3D extension of the 2D shape in an entirely user-controlled fashion by linking the drawn 2D brush (lasso) to the tablet's physical location. Users move the brush through the volumetric data by physically manipulating the tablet in 3D. During the user-controlled selection phase, we thus record all positions/orientations of the brush at each sample to compute a selection volume and to apply it to the data.

While users can control all 6 DOF of their tangible tablet (3 transl.,




Figure 2: Tablet interface to control the selection.




Figure 3: The tablet being used together with the synchronized split-screen view that is shown on a separate, large display.

3 rot.) to extrude the 2D shape, this fully-integrated manipulation may be difficult to achieve for some users or may want to consider manipulations only along a single axis. We thus allow users to constrain their (virtual) movements when moving the brush to along the tablet's normal axis (by clicking a button on the tablet; Fig. 2).

4.2. Selection Computation

We apply a regular grid to the 2D space of the tablet's surface with an adjustable resolution (we used 200×200). We then use the drawn lasso shape to mark cells as inside or outside the selection brush. For each sampled position/orientation of the tablet during the selection interaction, we derive the position of the Boolean grid within the data and connect two consecutive grid arrangements. We then use the resulting 3D grid slice to carry out the selection operation: either we select/deselect 3D particles directly or we manipulate a 3D regular selection data structure (i. e., a discretized 3D selection volume).

While most 3D selection techniques rely on simple selection shapes, some make use of progressive refinement strategies to improve the precision of the final 3D volume (e. g., [BKB13, EDF08, KBB11, YEII12, YEII16]). We use a similar approach based on Boolean operations, either adding to (OR), intersecting with (AND), or subtracting from the previous selection (AND NOT). We then use the resulting selection volume to process both volumetric data and particle clouds. We update the selection dynamically with each new sample to give users immediate feedback on their input.

4.3. Setup and User Interface

As in previous work [BIAI17b, LODI16, SvZP*16, SGF*11], we combine (a) a spatially-aware mobile device with touch input and (b)

a stationary large display (Fig. 3). All interactions are performed on the mobile device, including the drawing of the 2D selection brushes and the tangible movement of the selection brush through the 3D data space. The stationary device shows a high-resolution view of the dataset as a frame of reference for the interaction. Both devices display 2D projections of the 3D data—the larger stationary display from a static vantage point and the mobile device based on its 3D location and orientation. The more expensive computations i. e., the 3D selection volume creation and its application to the dataset are carried out on the PC that also drives the stationary display.

On the tablet interface (Fig. 2), menus (top) allow users to load datasets and change general view settings. A toggle button (right) controls whether selection extrusion uses all 6 DOF of the tangible interaction or whether it is constrained to the tablet's normal axis. A slider (right) reflects the current CD factor obtained from pressure sensors on the tablet's back, similar to Besançon et al.'s [BAI17] design. In addition, we added a classical slider-based control of the CD factor. In addition to isomorphic mappings, users can thus also be either more precise in their interactions or can cover larger space ranges using a single motion, with the slider always providing the visual feedback recommended for pressure-based input [CIS07, HW78, WSB10]. Users can switch between view manipulation and selection mode using another toggle button (bottom left). A final set of toggle buttons (left) control the Boolean operations. In addition to these system-controlled states, a user-controlled button (red, left) starts or stops the tangible manipulations of either the view or the selection extrusion during the interaction. The tangible manipulation of the tablet is interpreted in such a way that, during normal handling of the tablet, the data retains its location with respect to the tablet and thus appears to be attached to it. Only during active manipulation, the tablet is relocated with respect to the dataset.

The projection of the data in the center of the tablet provides both visual feedback and a canvas for the drawing of the 2D selection brush. An orthographic projection of the data allows users to relate the drawn 2D shape to the visible features of the data. While this projection lacks perspective depth clues, data features do not shrink or grow w.r.t. the drawn brush due to perspective foreshortening during the selection extrusion when the tangible device is moved through the dataset. This physical motion is also reflected in the data view (while either the data manipulation or the selection is actively engaged by the user) because the tablet is interpreted as a cutting plane, thus adding perceptual depth cues through the interaction.

Moreover, the separate display shows two synchronized data views (Fig. 3) using perspective projection. In this split-screen arrangement, the left view represents the data as it is shown on the tablet—the display acts as a physical representation of the tablet position w.r.t. the projected data. On the right, we show a view rotated by 45° around the vertical y-axis (based on pilot tests) from the left view. It also shows a representation of the tablet (which remains at the same screen-space location all the time) and the 2D selection brush within the data space. We also interpret this tablet-specific plane as a cutting plane such that the data between tablet and viewer is removed. This arrangement gives a good impression of the location of the data w.r.t. the tablet to steer the selection input.

4.4. Implementation and Performance

Our system is modular: the tablet runs a native Android app and the PC a separate Linux-based software. Both use VTK (VTK 6.0 on tablet, VTK 6.3 on PC), and render the data using OpenGL (OpenGL ES 2.0 on the tablet, OpenGL 3.0 on the PC) and shaders (GLSL 1.0 on the tablet, GLSL 1.3 on Linux). The tablet communicates with the PC over UDP (via Wifi). It sends status updates that allow the PC application to adjust its view and compute the selection. The communication lag was largely negligible in all our experiments.

We compute the selection on the PC and highlight the selection result depending on the data type. For particle data, we render the particles in a lighter shade of their initial color. For other data we show a semi-transparent selection volume. While not implemented, we envision to directly adjust the transfer function of volumetric datasets to show the selected voxels. Our rendering performance on both the tablet and the desktop was approx. 60 fps, independent of the size and characteristics of the different datasets. Only for large selections (brushes that take approx. 3/4 of the actual tablet size), the selection processing on the PC led to the view on the stationary display to somewhat lag behind during extrusion. A multi-threaded implementation of the PC application would address this issue.

4.5. Classification

Let us now revisit our taxonomy from Sect. 3 with respect to Tangible Brush. As a *selection metaphor/tool* we obviously use a combination of brush-based selection and 2D lasso. Our *target selection type* is based on ROIs because we extrude the 2D brush into 3D space. With only a few modifications, however, we could also easily apply the Tangible Brush to dedicated object selection (i. e., the objects touched during the extrusion interaction). In contrast to existing techniques that always rely on some form of automated computation for their selection shape, Tangible Brush solely relies on user input to define the selection volume. Since users already have manual control over the created final selection volume with Tangible Brush, the technique does not need any post-interaction *selection shape adjustments* that are required for most techniques that rely on an automated process. We could, however, add such an adjustment later-on as suggested by the participants of our study (see Sect. 6). We also do not need any technique to adjust the selection volume's origin and orientation that most techniques provide because our technique is already flexible in its selection volume creation.

To provide *selection tool control*, our technique relies on 2 DOF to draw the lasso using regular single-finger touch input as well as either movements only along the tablet's normal axis (1 DOF) or the full 6 DOF of the spatially-aware device, depending on the user's preference or the requirements of the dataset. Finally, we provide manual CD ratio control based on 1 DOF pressure input. With Tangible Brush we thus provide a **manual and structure-independent selection** that has not yet, to the best of our knowledge, been explored and which can be used in a variety of application scenarios. It works with virtually any type of 3D data, without relying on internal structures to produce a good selection.

5. Controlled Experiment

To understand how this manual control compares to structure-dependent selection techniques, we compared our approach with an existing state-of-the-art technique in a controlled experiment.

5.1. Structure-Dependent Selection Technique

Among the related work covered in Sect. 2, a number of techniques drew our attention. More specifically, all techniques that are based on a drawn 2D lasso were potential candidates for a comparison (i. e., [CSSM06, MVN06, ONI05, SXL*14, YEII12, YEII16]). We chose the SpaceCast approach from the CAST selection techniques [YEII16] because SpaceCast is a semi-automatic structure-dependent technique. Like all CAST techniques SpaceCast selects only a single connected component, but specifically constrains the selection to the drawn lasso shape (similar to our technique), and had already been compared to the other CAST selections, to cylinder/cone selection [Luc05, LBCW05a], and to CloudLasso [YEII12]. As such, using our results, a reader can extrapolate how our technique performs w.r.t. CloudLasso, or to other CAST techniques or even selections based on cylinders/cones which were found to be outperformed by CAST techniques but can be considered as a baseline for studies in other contexts. A final reason for our choice was that both SpaceCast and our technique do not require complex setups (e. g., a CAVE) and are thus more likely to be adopted by domain experts than ones that rely on a complex setup [Bes17]. For a fair comparison, we allowed participants to adjust the threshold used by SpaceCast. While our general prototype allows users to adjust the CD ratio with a slider or pressure sensors (Sect. 4.3), we deactivated this input in our controlled experiment to remove possible confounds. Moreover, the datasets were not heterogeneous enough that participants would have needed this control.

5.2. Hypotheses

Based on our general experience with spatial selection and our pilot studies, we formulated the following hypotheses:

- H1** Since SpaceCast was explicitly designed with density-based selection in particle datasets in mind, its accuracy will be better than that of Tangible Brush for datasets where the target can be identified based on density. If the target cannot be identified based on particle density, Tangible Brush will be more accurate.
- H2** SpaceCast is partially automated and only requires a 2D lasso as an input, so it will be generally faster to use than our Tangible Brush that also needs the tangible motion of the tablet in space.
- H3** Technique preference will depend on the dataset characteristics: for simple shapes that can easily be identified based on particle density participants will prefer SpaceCast, for other datasets where this identification of selection target based on particle density is not possible participants will prefer Tangible Brush.

5.3. Apparatus and Datasets

We used a 7" Google Tango tablet ("Yellowstone," 1920 × 1200 resolution, 323 ppi, 370 g) and a 55" (139.7 cm diagonal, 3840 × 2160 resolution, 79 ppi) static vertical screen, mounted at shoulder height. The latter used a PQLabs overlay to capture touch input. It served as




Figure 4: Datasets from our study; particles to be selected are red.

the primary input/output screen for the structure-dependent SpaceCast and as a secondary synchronized screen for Tangible Brush. For the latter, we captured input on the tablet, and restricted elaborate computations to the main PC (to save battery power and increase performance). We used the four particle datasets shown in Fig. 4:

- D1** Simple shapes (Fig. 4(a)): two simple shapes, a cylinder and a cube, having a higher density than the noise around them.
- D2** Shell & ring (Fig. 4(b)): ring, sphere, half-ball surrounded by a semi-spherical shell of interfering particles (all same density).
- D3** Pipe (Fig. 4(c)): target cylinder encapsulated in a larger cylinder (both same density).
- D4** Simulation (Fig. 4(d)): simulation of two colliding galaxies with varying densities, the target was one of the galaxies.

We designed D1 and D2 to be easy for SpaceCast. We made D2 to be more problematic for Tangible Brush by closely placing the ball and distractor shell, thus requiring a high precision. We designed D3 to be difficult for SpaceCast by using the same density for target and distractors, requiring several Boolean operations. Last, we chose D4 as it is difficult for both techniques and representative of possible real-world use cases that will make our results more generalizable.

5.4. Participants and Procedure

We recruited 16 unpaid participants (5 female; ages 21–53, mean = 25.8, med = 24.5, SD = 7.4, normal or corrected-to-normal vision). 11 had at least a university degree, while the other five had at most an A-level equivalent. Nine were experienced with 3D manipulations through video games and all had extensive experience with touch interaction, through the daily use of their smartphones. Three participants had previously been exposed to tangible interaction through their job, one of them on a daily basis. Other had no or very little exposition to it. Three participants were left-handed.

We first gathered the participants' demographics and explained

the study's purpose. Each participant performed selections using both techniques on D1–D4. We introduced each technique and then provided a training phase on four simple 3D shapes. We explained to participants that they could adjust the threshold of the selection of SpaceCast and that Tangible Brush was initially programmed to be depth-constrained but could that they could also perform unconstrained extrusions. Then, we asked them to achieve selection on datasets D1–D4. We repeated each task three times, each time from a different starting angle. Before each task, participants could freely explore the dataset to understand its structure and the targets. After that, we asked them to select as fast and accurately as possible. We did not stress any factor more than the other (mirroring Yu et al.'s [YEII16] study). We also mentioned that a perfect selection was generally impossible to avoid overly long interaction times.

Overall, our within-subjects study design thus used 2 selection techniques $\times$ 4 datasets $\times$ 3 repetitions = 24 trials per person. We counter-balanced the technique and dataset sequences to avoid order effects, the latter using a Williams design Latin square [Wil49] to control how often a condition precedes and follows another. After each technique, we asked participants to assess workload (using NASA's Task Load Index [TLX]) and fatigue. After all trials, we asked them, for each dataset, what technique they preferred using and why. We also recorded our participants' observations/comments based on thinking-aloud. One study session took approx. 75 min.

6. Results

We gathered a total of 384 trials (24 trials $\times$ 16 participants) for analysis. Recent criticism of NHST to analyze experimental data [Bak16, Cum14, DCH14, Dra16] and recent APA recommendations [Van09], led us to report our results using estimation techniques with effect sizes[‡] and confidence intervals instead of p -values.[§] We interpret them as providing different strength of evidence about the population mean [BD17, BD19, Cum14, Dra16, Gig04, Goo99, SH97].

Completion Time. We analyzed log-transformed measurements to correct for positive skewness and present anti-logged results.[¶] We show completion times in Fig. 5.^{||} The confidence intervals in Fig. 5(a) are not overlapping, providing strong evidence that SpaceCast is about twice as fast as Tangible Brush. There is also strong evidence of SpaceCast being faster for all datasets (Fig. 5(b)), with the difference being strongest for the second dataset for which it is more than 2.5 times as fast. The effect is smaller for D3. We also present the total number of selection steps performed per trial (Fig. 5(c)). The small overlap of both confidence intervals suggests that users created more selections with Tangible Brush.

Accuracy. Similar to Yu et al. [YEII12, YEII16], we computed two accuracy scores, F1 and MCC. Both are based on three factors: the number of true positives (rightfully selected particles, TP), false positives (incorrectly selected particles, FP), and false negatives (number of particles that should be selected but are not, FN). F1 is

[‡] Effect size refers to the means we measured. We do not use standardized effect sizes [Coe02]: reporting them is not always recommended [Bag09].

[§] A p -value-approach reading of our results can still be inferred [KA13].

[¶] This approach is standard for data analysis processes such as ours [SL10].

^{||} We present graphs individually per dataset in the additional material.




Figure 5: Task completion time in sec. (a), direct comparisons by dataset/overall (b), number of selections (c). Error bars: 95% CIs.



Figure 6: MCC (a) and F1 (b) accuracy. Error bars: 95% CIs.




Figure 7: Workload measurement in TLX units ($\in [0, 100]$, 0 is best) for (a) mental, (b) physical, and (c) temporal demand, (d) effort, (e) performance, (f) frustration, and (g) total. Error bars: 95% CIs.

computed as $F1 = 2 \cdot (P \cdot R) / (P + R)$ with $P = TP / (TP + FP)$ and $R = TP / (TP + FN)$. MCC also considers the number true negatives (particles correctly omitted, TN) as

$$MCC = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

We present these results in Fig. 6. For F1, 1 indicates a perfect performance and 0 the worst possible performance, while MCC results range from -1 (worst performance) to 1 (perfect performance). Both techniques obtained good MCC (Tangible Brush: 0.95, SpaceCast: 0.92) and F1 (Tangible Brush: 0.91, SpaceCast: 0.86) scores. Both scores (Fig. 6(a) and (b), resp.) show strong evidence for Tangible Brush being more accurate. Yet, both figures also show that only a small effect for the accuracy difference between both techniques.

Workload. We gathered workload measurements with a Raw-TLX (equally well suited as a regular TLX [Har06]), presented in Fig. 7. Our data shows no evidence of a difference between SpaceCast and Tangible Brush for physical demand (Fig. 7(b)), perceived performance (Fig. 7(e)), or frustration (Fig. 7(f)). There is clear evidence, however, that the mental demand is higher for Tangible Brush (Fig. 7(a)) and weak evidence that effort is higher for Tangible Brush Fig. 7(d)). Consequently, Fig. 7(g) displays a higher total workload for Tangible Brush.

Fatigue. We present the overall fatigue measurements and their sub-aspects in Fig. 8. The finger fatigue shown in Fig. 8(a) seems to be higher for SpaceCast. The overlapping confidence intervals




Figure 8: Fatigue measurements (scale 0–10) for (a) fingers, (b) hands, (c) arms, (d) shoulders, and (e) total. Error bars: 95% CIs.




Figure 9: Technique preferences shown by dataset.

for hand fatigue (Fig. 8(b)), arm fatigue (Fig. 8(c)), shoulder fatigue (Fig. 8(d)), and the total fatigue (Fig. 8(e)) suggest that we could not find evidence of a difference between the two techniques.

Preferences and Qualitative Results. After each study, we asked participants for their preferred technique per dataset (Fig. 9). For the first dataset (D1), a slight preference exists for SpaceCast (7 $\times$) over Tangible Brush (5 $\times$). For D2 which we had designed to be easier for SpaceCast, SpaceCast was largely preferred by participants (10 $\times$) over Tangible Brush (3 $\times$). For D3 which we had designed to be easier for Tangible Brush, we expectedly saw a strong advantage for Tangible Brush (11 $\times$) over SpaceCast (2 $\times$). Finally, the preferences for D4, the galaxy simulation, exhibit a slight preference for Tangible Brush (8 $\times$) over SpaceCast (6 $\times$).

Many participants (11 $\times$) stated that they wish they could draw the lasso on the tablet more precisely. Five suggested to use a stylus, five to allow zooming for drawing only, and one suggested both. Six participants said that SpaceCast was “way easier for simple shapes” or shapes with a homogeneous density. However, five others reported that they did not understand why or how SpaceCast derived the final selection volume for shapes along a varying density (with D4 mainly). One even reported that he considered the technique to be plainly “a pain” for D4. Two participants suggested to combine SpaceCast and Tangible Brush to “use SpaceCast for simple shapes or continuous densities,” and then to “adjust the selection for more complicated cases with Tangible Brush.” Finally, while participants were given the option to use the unconstrained 6DOF manipulations for Tangible Brush, all used the initial depth-only manipulations.

7. Discussion

We discuss our most relevant and interesting findings next.

Completion Time. Our data provided strong evidence to support hypothesis H2: SpaceCast was faster than Tangible Brush, for all datasets. Of course, participants were generally less familiar with tangible interaction than with touch input, causing at least a part of the increased interaction times. Users are likely to get faster with the

Tangible Brush as they would get more used to tangible interaction, in the same way they are now used to touch interaction. Furthermore, we noticed that most participants wasted time redrawing the lasso shape several times for Tangible Brush because they were not satisfied with their first few tries, mirroring qualitative feedback we reported earlier (though we could not quantify this behavior). We hypothesize that with zooming capabilities for the lasso drawing or by providing stylus-based input, the completion time with Tangible Brush would be shorter. We also observed more selection steps with Tangible Brush which also explains the higher interaction times. Reducing the number of selections with a greater expertise with the technique could probably also yield better time performances. For this reason, we believe that the time difference between the two techniques could be, with training, much smaller. We note that no participant reported that the techniques take too long to use or that it became frustrating such that prevented them from using them.

Accuracy. Our results seem to provide strong evidence against H1: Tangible Brush was *more accurate* than SpaceCast, overall and for each dataset. Yet, while there is clear evidence for the better performance of Tangible Brush, the effect is rather small. Both techniques exhibit a good performance, and are thus both well suited for spatial selection. We note, however, that only D4 is representative of a complicated real-world scenario. Based on our study we hypothesize that more complicated datasets would further decrease SpaceCast's accuracy (it only gives control over the x -/ y -axes). Moreover, completion time differences between both techniques would also shrink because users would spend more time with SpaceCast to remove unwanted regions from their initial selection.

Trade-off between Accuracy and Completion Time. We saw that Tangible Brush is more accurate than SpaceCast—one of the most accurate state-of-the-art techniques. This comes at the cost of completion time, although we discussed that the completion time of Tangible Brush could be reduced. By combining a semi-automated and a manual input, however, experts could have a tradeoff between a short completion time and a high precision: they could use the fast SpaceCast to select regions that are easy to derive from 2D primitives, while using Tangible Brush for less obvious cases. Such a solution can be useful for complex datasets and selections that cannot solely be based on one technique. We envision multiple combinations. First, simple shapes or data-dependent selection could be achieved by a CAST technique, while Tangible Brush would be used to handle complex and data-independent selections. We also envision that Tangible Brush can be modified to include a data-dependent strategy to assist the drawing of the lasso (e. g., based on active contours [KWT88]), keeping the manual extrusion. Such a combination would reduce the need for accurately drawing the lasso.

Extrusion in Constrained Mode. During the experiment, we observed that all participants used the tablet in the constrained mode (i. e., using only 1 DOF). They mainly rotated and placed the data in the *correct extrusion axis*, drew the lasso, and then extruded the selection shape. The constrained mode appears to be obviously effective for datasets D1, D2, and D3 since participants can see all the particles of interest using one axis. An exception may be D3 with the sphere, but participants approximated this case to a cylinder. We did not expect constrained mode to be used for D4; yet participants still did not use unconstrained mode and instead use a lot of time

to place the tablet w.r.t. the dataset, before starting any selection. Even if our interface gives manual control regarding the selection, future research may focus on where and when unconstrained mode is useful by, e. g., starting with bend or curving features.

Workload. Our workload results are surprising: we expected Tangible Brush not to require much mental demand, but our data seem to highlight that it requires more than SpaceCast. This can be explained by the fact that tangible interaction is not yet, contrary to touch input, widely adopted, as we noted before. Moreover, the mismatch between input and output space for Tangible Brush may similarly contribute to a high mental demand: while Tangible Brush is moved through 3D space during the selection extrusion, participants could only observe the effects on the projected 2D view of the dataset. We saw that some participants struggled when attempting even simple dataset rotations with the tangible tablet, further supporting this interpretation.

Fatigue. The fatigue results were, overall, similar for both techniques. While arm/shoulder pain is surprising for SpaceCast, it was caused by the position that our setup enforced as nine participants remarked: standing and lifting up the arms to interact. Even though the fatigue results are high, no participant reported that it was unbearable or that they had to take a break during the experiment. Both techniques can thus be used without worrying about fatigue.

Comparison to Other Techniques. We compared Tangible Brush to SpaceCAST because it is also a lasso-based technique (Sect. 5.1). It is structure-dependent and has been compared to other techniques that are considered as baselines for other studies (such as cylinder selections or CloudLasso) [YEIII6]. SpaceCAST was shown to be approximately three times as fast as Cylinder-based selection and approximately twice as fast as CloudLasso [YEIII2] (for some datasets) while providing at least similar accuracy. Given our previously discussed results, we assume that Tangible Brush would be faster than Cylinder selection, the baseline for many selection studies. Such inferences cannot be made, however, with respect to CloudLasso as completion time for this technique are dataset-dependent according to past studies [YEIII6]. Concerning accuracy, SpaceCast is at least equally accurate as baseline approaches [YEIII6]. Our technique being slightly more precise than SpaceCast and we thus deduct that Tangible Brush is very likely to be more precise than Cylinder selection or CloudLasso.

Preferences and Suggested Improvements. Our results tend to support hypothesis H3: participants tend to prefer SpaceCast for datasets in which the selection can be easily obtained with simple shapes or particle density (D1, D2). For complicated cases, they prefer Tangible Brush strongly (D3) or slightly (D4). Three participants said it was “nice to be able to control the selection based on spatial input” and that it was “helpful particularly for the galaxy dataset [D4] and the cylinder dataset [D3].” Moreover, no participant favored a single technique for all datasets, reinforcing the suggestion by two participants to combine both techniques. While Tangible Brush is already a hybrid technique, a combination of a partially automated with a manual approach would allow users to pick the technique depending on the dataset and the selection target.

Participants also suggested several improvements. Currently, users extrude the 2D shape with any tablet motion. One participant suggested to consider backward moves to cancel previous selections.

He considered going backward was a way to deselect, as if doing “Ctrl+Z.” This is an interesting idea that would enable physical undo operations and that could be explored within tangible systems. When a dataset is loaded, we also currently compute a dataset-dependent scale factor on the tablet to show the whole dataset as large as possible without touching the borders of the screen. This may negatively affect one’s ability to select small sections of the dataset, and five participants suggested to use zooming operations (for drawing purposes). We could combine this zoom with stylus input to further improve accuracy when drawing by avoiding the fat-finger issue.

Limitations. Despite our careful experimental design, some limitations should be mentioned. First, the tablet screen was much smaller than the screen on which we recorded SpaceCast’s input, leading to precision issues when drawing the lasso with Tangible Brush as highlighted by participants. A larger screen would probably lead to better results especially for the task completion time but would have made the tablet handling more difficult. The zoom suggested by participants may be a better solution. We also did not recruit experts who are used to selecting ROIs—they may have used different strategies for both interaction techniques. Indeed, our participants sometimes failed to see possible Boolean strategies that would have helped them, which most experts would probably have used. Participant preferences may also have been biased by the novelty effect of tangible interaction. Yet, to avoid this issue we were careful to ask for participants’ preferences per dataset, trying to get them to focus on the benefits of each technique rather than their entertaining values. We also asked them to justify their preferences and did not find any justification based on the enjoyment or novelty of tangible interaction. We thus believe that the novelty effect of tangible interaction had only little impact on our preference results. Finally, the mentioned mismatch between output and input space could be an issue for Tangible Brush. For SpaceCast and similar techniques, in contrast, users see a 2D projection of the data and provide matching 2D input—even if the data arguably lives in 3D. To avoid the space mismatch for Tangible Brush one could envision an AR setup (e.g., Microsoft’s HoloLens [Hol]) or a stereoscopic data projection on the large display: the stereoscopic data view would provide an ideal context for the tangible operations, thus linking output and tangible input spaces. While the 2D lasso that provides the brush would probably still require the tablet’s display to show a 2D projection (to avoid parallax issues [BS13, CHS*15, VSB*10, VSBH11]), we want to explore such setups in the future to further examine the potential of tangible interaction for exploring 3D visualizations.

8. Conclusion

We have further explored the design space of *spatial selection for 3D datasets*, adding a **structure-independent and manual selection** to the interaction spectrum for exploratory visualization as exemplified by our technical contribution, Tangible Brush, which can be applied to various types of 3D data (e.g., particle, volume, surface-based, and linear). Our **taxonomy** of interactive data selection techniques showed that new classification criteria were needed that describe the control offered to users by several techniques. Filling a hole that existed in the scope of data selection techniques, our Tangible Brush—in contrast to past approaches—gives **manual control to the user** performing a selection. Our experiment showed the poten-

tial of such a technique for 3D selection, highlighting its **excellent accuracy**, in particular, in complicated cases where it was also the preferred technique. Tangible brush provides a welcome **trade-off between accuracy and speed**, that can be chosen depending on the specific application case. Our study also highlighted the need for several selection techniques/strategies in any practical visualization application, depending at least on the dataset and the selection targets. Ultimately, it seems ideal that **an application should provide several different selection approaches from the entire spectrum** that cover different aspects of our taxonomy as we previously mentioned in Sect. 3.3 and Sect. 7. Ultimately, we envision it to be **used in tandem with context-dependent technique**, with users selecting the technique that is best suited for a given task at hand.

References

- [AA09] ARGELAGUET F., ANDUJAR C.: Efficient 3D pointing selection in cluttered virtual environments. *IEEE Computer Graphics and Applications* 29, 6 (Nov./Dec. 2009), 34–43. doi: 10.1109/MCG.2009.117
- [AA13] ARGELAGUET F., ANDUJAR C.: A survey of 3D object selection techniques for virtual environments. *Computers & Graphics* 37, 3 (May 2013), 121–136. doi: 10.1016/j.cag.2012.12.003
- [AAT08] ARGELAGUET F., ANDUJAR C., TRUEBA R.: Overcoming eye-hand visibility mismatch in 3D pointing selection. In *Proc. VRST* (2008), ACM, New York, pp. 43–46. doi: 10.1145/1450579.1450588
- [Ake06] AKERS D.: CINCH: A cooperatively designed marking interface for 3D pathway selection. In *Proc. UIST* (2006), ACM, New York, pp. 33–42. doi: 10.1145/1166253.1166260
- [AMR16] AL-MEGREN S., RUDDLE R. A.: Comparing tangible and multi-touch interaction for interactive data visualization tasks. In *Proc. TEI* (2016), ACM, New York, pp. 279–286. doi: 10.1145/2839462.2839464
- [ASM*04] AKERS D., SHERBONDY A., MACKENZIE R., DOUGHERTY R., WANDELL B.: Exploration of the brain’s white matter pathways with dynamic queries. In *Proc. Visualization* (2004), IEEE Comp. Soc., Los Alamitos, pp. 377–384. doi: 10.1109/VSUAL.2004.30
- [Bag09] BAGULEY T.: Standardized or simple effect size: What should be reported? *British Journal of Psychology* 100, 3 (Aug. 2009), 603–617. doi: 10.1348/000712608X377117
- [BAI17] BESANÇON L., AMMI M., ISENBERG T.: Pressure-based gain factor control for mobile 3D interaction using locally-coupled devices. In *Proc. CHI* (2017), ACM, New York, pp. 1831–1842. doi: 10.1145/3025453.3025890
- [Bak16] BAKER M.: Statisticians issue warning over misuse of *P* values. *Nature* 531, 7593 (Mar. 2016), 151. doi: 10.1038/nature.2016.19503
- [BCD*18] BÜSCHEL W., CHEN J., DACHSELT R., DRUCKER S., DWYER T., GÖRG C., ISENBERG T., KERREN A., NORTH C., STUEZLINGER W.: Interaction for immersive analytics. In *Immersive Analytics*. Springer, Berlin/Heidelberg, 2018, ch. 4, pp. 95–138. doi: 10.1007/978-3-030-01388-2_4
- [BD17] BESANÇON L., DRAGICEVIC P.: La différence significative entre valeurs *p* et intervalles de confiance (The significant difference between *p*-values and confidence intervals). In *Conférence Francophone sur l’Interaction Homme-Machine* (Aug. 2017), AFIHM, Poitiers, France. In French, English translation in [Bes17].
- [BD19] BESANÇON L., DRAGICEVIC P.: The Continued Prevalence of Dichotomous Inferences at CHI. In *CHI Extended Abstracts* (2019), ACM, New York. doi: 10.1145/3290607.3310432
- [Bes17] BESANÇON L.: *An Interaction Continuum for 3D Dataset Visualization*. Theses, Université Paris-Saclay, Dec. 2017.
- [BF07] BENKO H., FEINER S.: Balloon selection: A multi-finger technique for accurate low-fatigue 3D selection. In *Proc. 3DUI* (2007), IEEE Comp. Soc., Los Alamitos, pp. 79–86. doi: 10.1109/3DUI.2007.340778

- [BIA117a] BESANÇON L., ISSARTEL P., AMMI M., ISENBERG T.: Mouse, tactile, and tangible input for 3D manipulation. In *Proc. CHI* (2017), ACM, New York, pp. 4727–4740. doi: [10.1145/3025453.3025863](https://doi.org/10.1145/3025453.3025863)
- [BIA117b] BESANÇON L., ISSARTEL P., AMMI M., ISENBERG T.: Hybrid tactile/tangible interaction for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 881–890. doi: [10.1109/TVCG.2016.2599217](https://doi.org/10.1109/TVCG.2016.2599217)
- [BIRW19] BRUCKNER S., ISENBERG T., ROPINSKI T., WIEBEL A.: A model of spatial directness in interactive visualization. *IEEE Transactions on Visualization and Computer Graphics* 25 (2019). To appear. doi: [10.1109/TVCG.2018.2848906](https://doi.org/10.1109/TVCG.2018.2848906)
- [BJH01] BOWMAN D. A., JOHNSON D. B., HODGES L. F.: Testbed evaluation of virtual environment interaction techniques. *Presence: Teleoperators and Virtual Environments* 10, 1 (Feb. 2001), 75–95. doi: [10.1162/105474601750182333](https://doi.org/10.1162/105474601750182333)
- [BKB13] BACIM F., KOPPER R., BOWMAN D. A.: Design and evaluation of 3D selection techniques based on progressive refinement. *International Journal of Human-Computer Studies* 71, 7 (Jul/Aug 2013), 785–802. doi: [10.1016/j.ijhcs.2013.03.003](https://doi.org/10.1016/j.ijhcs.2013.03.003)
- [BRLD17] BÜSCHEL W., REIPSCHLÄGER P., LANGNER R., DACHSELT R.: Investigating the use of spatial interaction for 3D data visualization on mobile devices. In *Proc. ISS* (2017), ACM, New York, pp. 62–71. doi: [10.1145/3132272.3134125](https://doi.org/10.1145/3132272.3134125)
- [BS13] BRUDER G., STEINICKE F.: 2.5D touch interaction on stereoscopic tabletop surfaces. In *Proc. ISIS3D* (2013), pp. 1–4.
- [CHS*15] COLLEY A., HÄKKILÄ 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*, vol. 8433 of LNCS. Springer, Berlin/Heidelberg, 2015, pp. 47–67. doi: [10.1007/978-3-319-16940-8_3](https://doi.org/10.1007/978-3-319-16940-8_3)
- [CIS07] CECHANOWICZ J., IRANI P., SUBRAMANIAN S.: Augmenting the mouse with pressure sensitive input. In *Proc. CHI* (2007), ACM, New York, pp. 1385–1394. doi: [10.1145/1240624.1240835](https://doi.org/10.1145/1240624.1240835)
- [CMB*14] CABRAL M., MONTES A., BELLOC O., FERRAZ R., TEUBL F., DORETO F., LOPES R., ZUFFO M.: Bi-manual gesture interaction for 3D cloud point selection and annotation using COTS. In *Proc. 3DUI* (2014), IEEE Comp. Soc., Los Alamitos, pp. 187–188. doi: [10.1109/3DUI.2014.6798883](https://doi.org/10.1109/3DUI.2014.6798883)
- [Coe02] COE R.: It's the effect size, stupid: What effect size is and why it is important. In *Proc. Annual Conf. British Educational Research Assoc.* (2002).
- [CSD03] COURNIA N., SMITH J. D., DUCHOWSKI A. T.: Gaze-vs. hand-based pointing in virtual environments. In *CHI Extended Abstracts* (2003), ACM, New York, pp. 772–773. doi: [10.1145/765891.765982](https://doi.org/10.1145/765891.765982)
- [CSSM06] CHEN H.-L. J., SAMAVATI F. F., SOUSA M. C., MITCHELL J. R.: Sketch-based volumetric seeded region growing. In *Proc. SBIM* (2006), EG Assoc., Goslar, Germany, pp. 123–130. doi: [10.2312/SBM/SBM06/123-129](https://doi.org/10.2312/SBM/SBM06/123-129)
- [Cum14] CUMMING G.: The new statistics: Why and how. *Psychological Science* 25, 1 (Jan. 2014), 7–29. doi: [10.1177/0956797613504966](https://doi.org/10.1177/0956797613504966)
- [DCH14] DRAGICEVIC P., CHEVALIER F., HUOT S.: Running an HCI experiment in multiple parallel universes. In *CHI Extended Abstracts* (2014), ACM, New York, pp. 607–618. doi: [10.1145/2559206.2578881](https://doi.org/10.1145/2559206.2578881)
- [DFK12] DAIBER F., FALK E., KRÜGER A.: Balloon selection revisited: Multi-touch selection techniques for stereoscopic data. In *Proc. AVI* (2012), ACM, New York, pp. 441–444. doi: [10.1145/2254556.2254641](https://doi.org/10.1145/2254556.2254641)
- [DHKP05] DE HAAN G., KOUTEK M., POST F. H.: IntenSelect: Using dynamic object rating for assisting 3D object selection. In *Proc. IPT/EGVE* (2005), EG Assoc., Goslar, Germany, pp. 201–209. doi: [10.2312/EGVE/IPT_EGVE2005/201-209](https://doi.org/10.2312/EGVE/IPT_EGVE2005/201-209)
- [Dra16] DRAGICEVIC P.: Fair statistical communication in HCI. In *Modern Statistical Methods for HCI*, Robertson J., Kaptein M., (Eds.). Springer International Publishing, Cham, Switzerland, 2016, ch. 13, pp. 291–330. doi: [10.1007/978-3-319-26633-6_13](https://doi.org/10.1007/978-3-319-26633-6_13)
- [DS08] DEHMESHI H., STUERZLINGER W.: Intelligent mouse-based object group selection. In *Proc. Smart Graphics* (2008), Springer, Berlin/Heidelberg, pp. 33–44. doi: [10.1007/978-3-540-85412-8_4](https://doi.org/10.1007/978-3-540-85412-8_4)
- [DS09] DEHMESHI H., STUERZLINGER W.: GPSel: A gestural perceptual-based path selection technique. In *Proc. Smart Graphics* (2009), Springer, Berlin/Heidelberg, pp. 243–253. doi: [10.1007/978-3-642-02115-2_21](https://doi.org/10.1007/978-3-642-02115-2_21)
- [ÉAC*16] ÉVAIN A., ARGELAGUET F., CASIEZ G., ROUSSEL N., LÉCUEY A.: Design and Evaluation of Fusion Approach for Combining Brain and Gaze Inputs for Target Selection. *Frontiers in Neuroscience* 10, 454 (Oct. 2016), 14. doi: [10.3389/fnins.2016.00454](https://doi.org/10.3389/fnins.2016.00454)
- [EDF08] ELMQVIST N., DRAGICEVIC P., FEKETE J.-D.: Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (Nov./Dec. 2008), 1141–1148. doi: [10.1109/TVCG.2008.153](https://doi.org/10.1109/TVCG.2008.153)
- [FHZ96] FORSBERG A., HERNDON K., ZELEZNIK R.: Aperture based selection for immersive virtual environments. In *Proc. UIST* (1996), ACM, New York, pp. 95–96. doi: [10.1145/237091.237105](https://doi.org/10.1145/237091.237105)
- [FM81] FU K.-S., MUI J.: A survey on image segmentation. *Pattern Recognition* 13, 1 (1981), 3–16. doi: [10.1016/0031-3203\(81\)90028-5](https://doi.org/10.1016/0031-3203(81)90028-5)
- [GB06] GROSSMAN T., BALAKRISHNAN R.: The design and evaluation of selection techniques for 3D volumetric displays. In *Proc. UIST* (2006), ACM, New York, pp. 3–12. doi: [10.1145/1166253.1166257](https://doi.org/10.1145/1166253.1166257)
- [Gig04] GIGERENZER G.: Mindless statistics. *The Journal of Socio-Economics* 33, 5 (2004), 587–606. doi: [10.1016/j.socsc.2004.09.033](https://doi.org/10.1016/j.socsc.2004.09.033)
- [Goo99] GOODMAN S. N.: Toward evidence-based medical statistics. 1: The P value fallacy. *Annals of Internal Medicine* 130, 12 (1999), 995–1004. doi: [10.7326/0003-4819-130-12-199906150-00008](https://doi.org/10.7326/0003-4819-130-12-199906150-00008)
- [Han97] HAND C.: A survey of 3D interaction techniques. *Computer Graphics Forum* 16, 5 (Dec. 1997), 269–281. doi: [10.1111/1467-8659.00194](https://doi.org/10.1111/1467-8659.00194)
- [Har06] HART S. G.: NASA-task load index (NASA-TLX); 20 years later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50, 9 (Oct. 2006), 904–908. doi: [10.1177/154193120605000909](https://doi.org/10.1177/154193120605000909)
- [HG02] HACHET M., GUITTON P.: The interaction table – A new input device designed for interaction in immersive large display environments. In *Proc. EGVE* (2002), Eurographics Assoc., Goslar, Germany, pp. 189–196. doi: [10.2312/EGVE/EGVE02/189-196](https://doi.org/10.2312/EGVE/EGVE02/189-196)
- [HGR03] HACHET M., GUITTON P., REUTER P.: The CAT for efficient 2D and 3D interaction as an alternative to mouse adaptations. In *Proc. VRST* (2003), ACM, New York, pp. 205–212. doi: [10.1145/1008653.1008689](https://doi.org/10.1145/1008653.1008689)
- [Hib99] HIBBARD B.: Top ten visualization problems. *ACM SIGGRAPH Computer Graphics* 33, 2 (May 1999), 21–22. doi: [10.1145/326460.326485](https://doi.org/10.1145/326460.326485)
- [Hol] Hololens website. URL: <https://www.microsoft.com/microsoft-hololens/>. Visited Dec. 2018.
- [HS85] HARALICK R. M., SHAPIRO L. G.: Image segmentation techniques. *Computer Vision, Graphics, and Image Processing* 29, 1 (Jan. 1985), 100–132. doi: [10.1016/S0734-189X\(85\)90153-7](https://doi.org/10.1016/S0734-189X(85)90153-7)
- [HW78] HEROT C. F., WEINZAPFEL G.: One-point touch input of vector information for computer displays. *ACM SIGGRAPH Computer Graphics* 12, 3 (Aug. 1978), 210–216. doi: [10.1145/965139.807392](https://doi.org/10.1145/965139.807392)
- [HWS02] HARDERS M., WILDERMUTH S., SZÉKELY G.: New paradigms for interactive 3D volume segmentation. *Computer Vision, Graphics, and Image Processing* 13, 1 (Feb. 2002), 85–95. doi: [10.1002/vis.277](https://doi.org/10.1002/vis.277)
- [IBG*16] ISSARTEL P., BESANÇON L., GUÉNIAT F., ISENBERG T., AMMI M.: Preference between allocentric and egocentric 3D manipulation in a locally coupled configuration. In *Proc. SUI* (2016), ACM, New York, pp. 79–88. doi: [10.1145/2983310.2985750](https://doi.org/10.1145/2983310.2985750)
- [IMT99] IGARASHI T., MATSUOKA S., TANAKA H.: Teddy: A sketching interface for 3D freeform design. In *Proc. SIGGRAPH* (1999), ACM, New York, pp. 409–416. doi: [10.1145/311535.311602](https://doi.org/10.1145/311535.311602)

- [Ise16] ISENBERG T.: Interactive exploration of three-dimensional scientific visualizations on large display surfaces. In *Collaboration Meets Interactive Spaces*, Anslow C., Campos P., Jorge J., (Eds.). Springer, Berlin/Heidelberg, 2016, ch. 6, pp. 97–123. doi: [10.1007/978-3-319-45853-3_6](https://doi.org/10.1007/978-3-319-45853-3_6)
- [JGAK07] JORDÀ S., GEIGER G., ALONSO M., KALTENBRUNNER M.: The reacTable: Exploring the synergy between live music performance and tabletop tangible interfaces. In *Proc. TEI* (2007), ACM, New York, pp. 139–146. doi: [10.1145/1226969.1226998](https://doi.org/10.1145/1226969.1226998)
- [JH13] JANKOWSKI J., HACHET M.: A survey of interaction techniques for interactive 3D environments. In *Eurographics STARs* (2013), EG Assoc., Goslar, Germany, pp. 65–93. doi: [10.2312/conf/EG2013/stars/065-093](https://doi.org/10.2312/conf/EG2013/stars/065-093)
- [JLS*13] JACKSON B., LAU T. Y., SCHROEDER D., TOUSSAINT 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 (Dec. 2013), 2802–2809. doi: [10.1109/TVCG.2013.121](https://doi.org/10.1109/TVCG.2013.121)
- [Joh04] JOHNSON C. R.: Top scientific visualization research problems. *IEEE Computer Graphics and Applications* 24, 4 (July/Aug. 2004), 13–17. doi: [10.1109/MCG.2004.20](https://doi.org/10.1109/MCG.2004.20)
- [KA13] KRZYWINSKI M., ALTMAN N.: Points of significance: Error bars. *Nature Methods* 10, 10 (Oct. 2013), 921–922. doi: [10.1038/nmeth.2659](https://doi.org/10.1038/nmeth.2659)
- [KBB11] KOPPER R., BACIM F., BOWMAN D. A.: Rapid and accurate 3D selection by progressive refinement. In *Proc. 3DUI* (2011), IEEE Comp. Soc., Los Alamitos, pp. 67–74. doi: [10.1109/3DUI.2011.5759219](https://doi.org/10.1109/3DUI.2011.5759219)
- [Kee10] KEEFE D. F.: Integrating visualization and interaction research to improve scientific workflows. *IEEE Computer Graphics and Applications* 30, 2 (Mar./Apr. 2010), 8–13. doi: [10.1109/MCG.2010.30](https://doi.org/10.1109/MCG.2010.30)
- [KGDR09] KÖNIG W. A., GERKEN J., DIERDORF S., REITERER H.: Adaptive pointing – Design and evaluation of a precision enhancing technique for absolute pointing devices. In *Proc. INTERACT* (2009), Springer, Berlin/Heidelberg, pp. 658–671. doi: [10.1007/978-3-642-03655-2_73](https://doi.org/10.1007/978-3-642-03655-2_73)
- [Kha13] KHAN W.: Image segmentation techniques: A survey. *Journal of Image and Graphics* 1, 4 (Dec. 2013), 166–170. doi: [10.12720/joig.14.166-170](https://doi.org/10.12720/joig.14.166-170)
- [KI13] KEEFE D. F., ISENBERG T.: Reimagining the scientific visualization interaction paradigm. *IEEE Computer* 46, 5 (May 2013), 51–57. doi: [10.1109/MC.2013.178](https://doi.org/10.1109/MC.2013.178)
- [KOM*03] KAISER E., OLWAL A., MCGEE D., BENKO H., CORRADINI A., LI X., COHEN P., FEINER S.: Mutual disambiguation of 3D multimodal interaction in augmented and virtual reality. In *Proc. ICMI* (2003), ACM, New York, pp. 12–19. doi: [10.1145/958432.958438](https://doi.org/10.1145/958432.958438)
- [KSMC05] KONIECZNY J., SHIMIZU C., MEYER G., COLUCCI D.: A handheld flexible display system. In *Proc. Visualization* (2005), IEEE Comp. Soc., Los Alamitos, pp. 591–597. doi: [10.1109/VISUAL.2005.1532846](https://doi.org/10.1109/VISUAL.2005.1532846)
- [KWT88] KASS M., WITKIN A., TERZOPOULOS D.: Snakes: Active contour models. *International Journal of Computer Vision* 1, 4 (Jan. 1988), 321–331. doi: [10.1007/BF00133570](https://doi.org/10.1007/BF00133570)
- [LBCW05a] LUCAS J. F., BOWMAN D. A., CHEN J., WINGRAVE C. A.: Design and evaluation of 3D multiple object selection techniques. In *ACM I3D Poster Compendium* (2005).
- [LBCW05b] LUCAS J. F., BOWMAN D. A., CHEN J., WINGRAVE C. A.: *Design and Evaluation of 3D Multiple Object Selection Techniques*. Tech. rep., Virginia Tech, USA, 2005.
- [LBGC07] LOOSER J., BILLINGHURST M., GRASSET R., COCKBURN A.: An evaluation of virtual lenses for object selection in augmented reality. In *Proc. GRAPHITE* (2007), ACM, New York, pp. 203–210. doi: [10.1145/1321261.1321297](https://doi.org/10.1145/1321261.1321297)
- [LG94] LIANG J., GREEN M.: JDCAD: A highly interactive 3D modeling system. *Computers & Graphics* 18, 4 (July/Aug. 1994), 217–222. doi: [10.1016/0097-8493\(94\)90062-0](https://doi.org/10.1016/0097-8493(94)90062-0)
- [LKM*17] LAVIOLA JR. J. J., KRUIJFF E., MCMAHAN R. P., BOWMAN D. A., POUPYREV I.: *3D User Interfaces: Theory and Practice*, 2nd ed. Addison-Wesley, Boston, 2017.
- [LODI16] 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 (May 2016), 1616–1629. doi: [10.1109/TVCG.2015.2440233](https://doi.org/10.1109/TVCG.2015.2440233)
- [LSS06] LAPIDES P., SHARLIN E., SOUSA M. C., STREIT L.: The 3D tractus: A three-dimensional drawing board. In *Proc. Tabletop* (2006), IEEE Comp. Soc., Los Alamitos, pp. 169–176. doi: [10.1109/TABLETOP.2006.33](https://doi.org/10.1109/TABLETOP.2006.33)
- [Luc05] LUCAS J. F.: *Design and Evaluation of 3D Multiple Object Selection Techniques*. Master's thesis, Virginia Tech, USA, Jan. 2005.
- [MBJS97] MINE M. R., BROOKS JR F. P., SEQUIN C. H.: Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Proc. SIGGRAPH* (1997), ACM, New York, pp. 19–26. doi: [10.1145/258734.258747](https://doi.org/10.1145/258734.258747)
- [Min95] MINE M.: *Virtual Environment Interaction Techniques*. Tech. Rep. TR95-018, CS Dep., Univ. North Carolina at Chapel Hill, 1995.
- [Mun14] MUNZNER T.: *Visualization Analysis and Design*. CRC Press, Boca Raton, FL, USA, 2014. doi: [10.1201/b17511](https://doi.org/10.1201/b17511)
- [MVN06] MALMBERG F., VIDHOLM E., NYSTRÖM I.: A 3D live-wire segmentation method for volume images using haptic interaction. In *Proc. DGCI* (2006), Springer, Berlin/Heidelberg, pp. 663–673. doi: [10.1007/11907350_56](https://doi.org/10.1007/11907350_56)
- [NMVB09] NYSTRÖM I., MALMBERG F., VIDHOLM E., BENGTSSON E.: Segmentation and visualization of 3D medical images through haptic rendering. In *Proc. PRIP* (2009), Belarusian State University Publishing Center, Minsk, Belarus, pp. 43–48. URN: [nbn:se:uu:diva-107657](https://nbn-resolving.org/urn:nbn:se:uu:diva-107657)
- [OBF03] OLWAL A., BENKO H., FEINER S.: Senseshapes: Using statistical geometry for object selection in a multimodal augmented reality. In *Proc. ISMAR* (2003), IEEE Comp. Soc., Los Alamitos, pp. 300–301. doi: [10.1109/ISMAR.2003.1240730](https://doi.org/10.1109/ISMAR.2003.1240730)
- [OF03] OLWAL A., FEINER S.: The flexible pointer: An interaction technique for selection in augmented and virtual reality. In *ACM UIST Poster Compendium* (2003), pp. 81–82.
- [OF09] OLWAL A., FEINER S.: Spatially aware handhelds for high-precision tangible interaction with large displays. In *Proc. TEI* (2009), ACM, New York, pp. 181–188. doi: [10.1145/1517664.1517705](https://doi.org/10.1145/1517664.1517705)
- [ONI05] OWADA S., NIELSEN F., IGARASHI T.: Volume catcher. In *Proc. I3D* (New York 2005), ACM, pp. 111–116. doi: [10.1145/1053427.1053445](https://doi.org/10.1145/1053427.1053445)
- [PBD*16] PLIMMER B., BROWN B., DIPROSE J., DU PREEZ S., LUXTON-REILLY A.: Tabletop 3D object manipulation with touch and tangibles. In *Collaboration Meets Interactive Spaces*, Anslow C., Campos P., Jorge J., (Eds.). Springer International, Cham, Switzerland, 2016, ch. 2, pp. 11–32. doi: [10.1007/978-3-319-45853-3_2](https://doi.org/10.1007/978-3-319-45853-3_2)
- [PBW196] POUPYREV I., BILLINGHURST M., WEGHORST S., ICHIKAWA T.: The Go-Go interaction technique: Non-linear mapping for direct manipulation in VR. In *Proc. UIST* (1996), ACM, New York, pp. 79–80. doi: [10.1145/237091.237102](https://doi.org/10.1145/237091.237102)
- [PFC*97] PIERCE J. S., FORSBERG A. S., CONWAY M. J., HONG S., ZELENIN R. C., MINE M. R.: Image plane interaction techniques in 3d immersive environments. In *Proceedings of the 1997 symposium on Interactive 3D graphics* (1997), ACM, New York, pp. 39–44. doi: [10.1145/253284.253303](https://doi.org/10.1145/253284.253303)
- [PI99] POUPYREV I., ICHIKAWA T.: Manipulating objects in virtual worlds: Categorization and empirical evaluation of interaction techniques. *Journal of Visual Languages & Computing* 10, 1 (Feb. 1999), 19–35. doi: [10.1006/jvlc.1998.0112](https://doi.org/10.1006/jvlc.1998.0112)
- [PM03] POUPYREV I., MARUYAMA S.: Tactile interfaces for small touch screens. In *Proc. UIST* (2003), ACM, New York, pp. 217–220. doi: [10.1145/964696.964721](https://doi.org/10.1145/964696.964721)
- [Rhe02] RHEINGANS P.: Are we there yet? Exploring with dynamic visualization. *IEEE Computer Graphics and Applications* 22, 1 (Jan. 2002), 6–10. doi: [10.1109/38.974511](https://doi.org/10.1109/38.974511)

- [SAM*05] SHERBONDY A., AKERS D., MACKENZIE R., DOUGHERTY R., WANDELL B.: Exploring connectivity of the brain's white matter with dynamic queries. *IEEE Transactions on Visualization and Computer Graphics* 11, 4 (July/Aug. 2005), 419–430. doi: [10.1109/TVCG.2005.59](https://doi.org/10.1109/TVCG.2005.59)
- [SBB*06] SCHMIDT G., BAILLOT Y., BROWN D. G., TOMLIN E. B., SWAN J.: Toward disambiguating multiple selections for frustum-based pointing. In *Proc. 3DUI* (2006), IEEE Comp. Soc., Los Alamitos, pp. 87–94. doi: [10.1109/VR.2006.133](https://doi.org/10.1109/VR.2006.133)
- [SBD12] SPINDLER M., BÜSCHEL W., DACHSELT R.: Use your head: Tangible windows for 3d information spaces in a tabletop environment. In *Proc. ITS* (2012), ITS '12, ACM, New York, pp. 245–254. doi: [10.1145/2396636.2396674](https://doi.org/10.1145/2396636.2396674)
- [SGF*11] 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* (2011), ACM, New York, pp. 1333–1342. doi: [10.1145/1978942.1979140](https://doi.org/10.1145/1978942.1979140)
- [SH97] SCHMIDT F. L., HUNTER J. E.: Eight common but false objections to the discontinuation of significance testing in the analysis of research data. In *What If There Were No Significance Tests?* Lawrence Erlbaum Associates Publishers, 1997, ch. 3, pp. 37–64. doi: [10.4324/9781315629049](https://doi.org/10.4324/9781315629049)
- [SJM13] STEIMLE J., JORDT A., MAES P.: Flexpad: Highly flexible bending interactions for projected handheld displays. In *Proc. CHI* (2013), ACM, New York, pp. 237–246. doi: [10.1145/2470654.2470688](https://doi.org/10.1145/2470654.2470688)
- [SL10] SAURO J., LEWIS J. R.: Average task times in usability tests: What to report? In *Proc. CHI* (2010), ACM, New York, pp. 2347–2350. doi: [10.1145/1753326.1753679](https://doi.org/10.1145/1753326.1753679)
- [SP04] STEED A., PARKER C.: 3D selection strategies for head tracked and non-head tracked operation of spatially immersive displays. In *Proc. 8th Intern. Immersive Projection Technology Workshop* (2004), pp. 163–170.
- [SSFG98] SZALAVÁRI Z., SCHMALSTEIG D., FUHRMANN A., GERVAUTZ M.: “Studierstube”: An environment for collaboration in augmented reality. *Virtual Reality* 3, 1 (Mar. 1998), 37–48. doi: [10.1007/BF01409796](https://doi.org/10.1007/BF01409796)
- [SSSS11] 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* (2011), ACM, New York, pp. 192–201. doi: [10.1145/2076354.2076390](https://doi.org/10.1145/2076354.2076390)
- [Ste06] STEED A.: Towards a general model for selection in virtual environments. In *Proc. 3DUI* (2006), IEEE Comp. Soc., Los Alamitos, pp. 103–110. doi: [10.1109/VR.2006.134](https://doi.org/10.1109/VR.2006.134)
- [STSD10] SPINDLER M., TOMINSKI C., SCHUMANN H., DACHSELT R.: Tangible views for information visualization. In *Proc. ITS* (2010), ACM, New York, pp. 157–166. doi: [10.1145/1936652.1936684](https://doi.org/10.1145/1936652.1936684)
- [Sut66] SUTHERLAND I. E.: Ten unsolved problems in computer graphics. *Datamation* 12, 5 (1966), 22.
- [SvZP*16] SOLLICH H., VON ZADOW U., PIETZSCH T., TOMANCAK P., DACHSELT R.: Exploring time-dependent scientific data using spatially aware mobiles and large displays. In *Proc. ISS* (2016), ACM, New York, pp. 349–354. doi: [10.1145/2992154.2996779](https://doi.org/10.1145/2992154.2996779)
- [SXL*14] 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 (Aug. 2014), 145–156. doi: [10.1007/s12650-014-0206-5](https://doi.org/10.1007/s12650-014-0206-5)
- [Tan] Project Tango on Wikipedia. URL: [https://en.wikipedia.org/wiki/Tango_\(platform\)](https://en.wikipedia.org/wiki/Tango_(platform)). Visited Dec. 2018.
- [THA10] TOP A., HAMARNEH G., ABUGHARBIEH R.: Spotlight: Automated confidence-based user guidance for increasing efficiency in interactive 3D image segmentation. In *Proc. MCV* (2010), Springer, Berlin/Heidelberg, pp. 204–213. doi: [10.1007/978-3-642-18421-5_20](https://doi.org/10.1007/978-3-642-18421-5_20)
- [TJ00] TANRIVERDI V., JACOB R. J.: Interacting with eye movements in virtual environments. In *Proc. CHI* (2000), ACM, New York, pp. 265–272. doi: [10.1145/332040.332443](https://doi.org/10.1145/332040.332443)
- [TLX] Nasa TLX website. URL: <https://humansystemsarcnasagov/groups/tlx/downloads/TLXScale.pdf>. Visited Dec. 2018.
- [TM04] TORY M., MÖLLER T.: Human factors in visualization research. *IEEE Transactions on Visualization and Computer Graphics* 10, 1 (Jan./Feb. 2004), 72–84. doi: [10.1109/TVCG.2004.1260759](https://doi.org/10.1109/TVCG.2004.1260759)
- [Tuk77] TUKEY J. W.: *Exploratory Data Analysis*. Pearson, Reading, MA, USA, 1977.
- [Van09] VANDENBOS G. R. (Ed.): *Publication Manual of the American Psychological Association*, 6th ed. APA, Washington, DC, 2009.
- [VB05] VOGEL D., BALAKRISHNAN R.: Distant freehand pointing and clicking on very large, high resolution displays. In *Proc. UIST* (2005), ACM, New York, pp. 33–42. doi: [10.1145/1095034.1095041](https://doi.org/10.1145/1095034.1095041)
- [VGC07] VANACKEN L., GROSSMAN T., CONINX K.: Exploring the effects of environment density and target visibility on object selection in 3D virtual environments. In *Proc. 3DUI* (2007), IEEE Comp. Soc., Los Alamitos. doi: [10.1109/3DUI.2007.340783](https://doi.org/10.1109/3DUI.2007.340783)
- [VSB*10] 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* (2010), EG Assoc., Goslar, Germany, pp. 17–24. doi: [10.2312/EGVE/MRC10/017-024](https://doi.org/10.2312/EGVE/MRC10/017-024)
- [VSBH11] VALKOV D., STEINICKE F., BRUDER G., HINRICHS K.: 2D touching of 3D stereoscopic objects. In *Proc. CHI* (2011), ACM, New York, pp. 1353–1362. doi: [10.1145/1978942.1979142](https://doi.org/10.1145/1978942.1979142)
- [WBB06] WYSS H. P., BLACH R., BUES M.: isith - intersection-based spatial interaction for two hands. In *3D User Interfaces (3DUI'06)* (March 2006), pp. 59–61. doi: [10.1109/VR.2006.93](https://doi.org/10.1109/VR.2006.93)
- [WBG*19] WANG X., BESANÇON L., GUÉNIAT F., SERENO M., AMMI M., ISENBERG T.: A Vision of Bringing Immersive Visualization to Scientific Workflows. In *CHI-IA 2019 - Workshop on Immersive Analytics at ACM* (May 2019), Glasgow, United Kingdom.
- [WHB06] WINGRAVE C. A., HACIAHMETOGLU Y., BOWMAN D. A.: Overcoming world in miniature limitations by scaling and scrolling. In *Proc. 3DUI* (2006), IEEE Comp. Soc., Los Alamitos, pp. 11–16. doi: [10.1109/VR.2006.106](https://doi.org/10.1109/VR.2006.106)
- [Wil49] WILLIAMS E.: Experimental designs balanced for the estimation of residual effects of treatments. *Australian Journal of Chemistry* 2, 2 (1949), 149–168. doi: [10.1071/CH9490149](https://doi.org/10.1071/CH9490149)
- [Wil96] WILLS G. J.: Selection: 524,288 ways to say “this is interesting”. In *Proc. InfoVis* (1996), IEEE Comp. Soc., Los Alamitos, pp. 54–60. doi: [10.1109/INFVIS.1996.559216](https://doi.org/10.1109/INFVIS.1996.559216)
- [WSB10] WILSON G., STEWART C., BREWSTER S. A.: Pressure-based menu selection for mobile devices. In *Proc. MobileCHI* (2010), ACM, New York, pp. 181–190. doi: [10.1145/1851600.1851631](https://doi.org/10.1145/1851600.1851631)
- [WTW*05] WINGRAVE C. A., TINTNER R., WALKER B. N., BOWMAN D. A., HODGES L. F.: Exploring individual differences in raybased selection: Strategies and traits. In *Proc. VR* (2005), IEEE Comp. Soc., Los Alamitos, pp. 163–170. doi: [10.1109/VR.2005.1492770](https://doi.org/10.1109/VR.2005.1492770)
- [WVFH12] 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 (Dec. 2012), 2236–2244. doi: [10.1109/TVCG.2012.292](https://doi.org/10.1109/TVCG.2012.292)
- [YEII12] 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 (Dec. 2012), 2245–2254. doi: [10.1109/TVCG.2012.217](https://doi.org/10.1109/TVCG.2012.217)
- [YEII16] 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 (Jan. 2016), 886–895. doi: [10.1109/TVCG.2015.2467202](https://doi.org/10.1109/TVCG.2015.2467202)
- [YZNC05] YUAN X., ZHANG N., NGUYEN M. X., CHEN B.: Volume cutout. *The Visual Computer* 21, 8 (Sept. 2005), 745–754. doi: [10.1007/s00371-005-0330-2](https://doi.org/10.1007/s00371-005-0330-2)

- [ZBM94] ZHAI S., BUXTON W., MILGRAM P.: The “Silk Cursor”: Investigating transparency for 3D target acquisition. In *Proc. CHI* (1994), ACM, New York, pp. 459–464. doi: [10.1145/191666.191822](https://doi.org/10.1145/191666.191822)
- [ZCL08] ZHOU W., CORREIA S., LAIDLAW D. H.: Haptics-assisted 3D lasso drawing for tracts-of-interest selection in DTI visualization. In *IEEE Visualization Poster Compendium* (2008).
- [ZGO13] ZUCKERMAN O., GAL-OZ A.: 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 (July/Aug. 2013), 803–820. doi: [10.1016/j.ijhcs.2013.04.003](https://doi.org/10.1016/j.ijhcs.2013.04.003)