



HAL
open science

Mouse, Tactile, and Tangible Input for 3D Manipulation

Lonni Besançon, Paul Issartel, Mehdi Ammi, Tobias Isenberg

► **To cite this version:**

Lonni Besançon, Paul Issartel, Mehdi Ammi, Tobias Isenberg. Mouse, Tactile, and Tangible Input for 3D Manipulation. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), May 2017, Denver, United States. pp.4727-4740, 10.1145/3025453.3025863 . hal-01436206

HAL Id: hal-01436206

<https://inria.hal.science/hal-01436206>

Submitted on 16 Jan 2017

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

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

Mouse, Tactile, and Tangible Input for 3D Manipulation

Lonni Besançon

University Paris Saclay
Orsay, France
lonni.besancon@inria.fr

Paul Issartel

University Paris Saclay
Orsay, France
paul.issartel@u-psud.fr

Mehdi Ammi

Limsi CNRS
Orsay, France
mehdi.ammi@limsi.fr

Tobias Isenberg

Inria
Orsay, France
tobias.isenberg@inria.fr

ABSTRACT

We evaluate the performance and usability of mouse-based, touch-based, and tangible interaction for manipulating objects in a 3D virtual environment. This comparison is a step toward a better understanding of the limitations and benefits of these existing interaction techniques, with the ultimate goal of facilitating an easy transition between the different 3D data exploration environments. For this purpose we analyze participants' performance in 3D manipulation using a docking task. We measured completion times, docking accuracy, as well as subjective criteria such as fatigue, workload, and preference. Our results show that the three input modalities provide similar levels of precision but require different completion times. We also discuss our qualitative observations as well as people's preferences and put our findings into context of the application domain of 3D data analysis environments.

ACM Classification Keywords

H.5.2. User Interfaces: Input devices and strategies

Author Keywords

3D interaction; mouse; tactile interaction; tangible interaction; TUI; usability study.

Many application domains rely on effective, efficient, and intuitive means of interacting with 3D data [37, 48]. Traditionally, this interaction has often relied on mouse and keyboard inputs. Recent developments of interaction technology, however, have led to new input modalities becoming available, in particular tactile input [30, 61, 74]¹ and tangible interaction [31, 59].² Several researchers have thus started to explore their use for interaction with 3D data. Nevertheless, the three input modalities—mouse, touch, and tangibles—are not identical in characteristics such as their capabilities or usability: their advantages and disadvantages depend on the the interaction goal and the given application domain. For example, while one may use a tangible input device intuitively in a game, scientific visualization applications may require a level of accuracy that

¹I. e., interfaces based on finger or pen input on display surfaces.

²I. e., interfaces that follow Ullmer & Ishii's [68] four characteristics.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2017, May 6-11, 2017, Denver, CO, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-4655-9/17/05 ...\$15.00.

<http://dx.doi.org/10.1145/3025453.3025863>

one could expect to better be provided by touch-based or in particular mouse-based input.

Tangibles are often regarded as the best way to interact with 3D data. We question this assumption here with our study that measured several usability factors: participants' accuracy (i. e., rotational difference and Euclidean distance), their perceived fatigue levels, and their perceived workload. We also took into account participants' preferences and their general feedback for each technique. The study consisted of 15 abstract 3D docking tasks—bringing an abstract virtual object to a given target orientation and position—for each of the three modalities. Our study confirmed that mouse, tactile, and tangible input are all valid means to control 3D manipulations. Much to our surprise, however, we found that all three input modalities allow users to achieve the same level of accuracy. Differences only arose with respect to task completion times and preferences. Qualitative observations of the participants during the study provided additional insights on what users tend to do when facing a docking task with these three input techniques which we discuss in detail below.

In summary, we contribute (1) an in-depth analysis of people's understanding and use of mouse-based, tactile, and tangible input for 3D interaction, (2) a study design that compares the three modalities, and (3) in-depth qualitative observations and people's preferences in the context of 3D data analysis environments. We thus shed light on the advantages and disadvantages of the techniques and serves as a basis for their further development and evaluation, in particular for 3D visualization.

RELATED WORK

Much of past work has focused on the comparison of interaction techniques or devices—many academic studies compare novel technique(s) or device(s) to established ones. For instance, many studies were conducted to compare the advantages and limitations of mouse interaction compared to touch interactions for tasks as various as selection, pointing, exploration etc. (e. g., [20, 39, 54]). Our review of the literature, however, revealed a lack of studies that would analyze these modalities for 3D manipulation tasks—only few researchers actually conducted such analyses [11, 28, 67, 76].

Among them, Chen et al. [11] and later Hinckley et al. [28] compared input techniques for 3D manipulation. Both studies, however, narrowly focused on rotation and did not take into account other parameters such as Euclidean distance to the target or usability. Tuddenham et al. [67] compared mouse, tactile, and tangible interaction for a matching task on a tabletop, thus constraining the interaction to two dimensions. They

measured the task completion time, the ease of use, and people’s preference. Yu et al. [76], finally, compared mouse and touch interaction to validate their FI3D widget for 7DOF data navigation. In contrast, we aim to get a holistic and general view of how the three input methods affect the interaction in 3D environments, ultimately to understand how we can better support the analysis of complex 3D datasets.

Moreover, most comparative studies focus on comparing either mouse and tactile interaction or tangible and tactile (and many concentrate on 2D tasks). The literature indeed contains many comparisons of touch and mouse input for a whole variety of tasks and a whole variety of parameters: speed [20, 22, 57], error rate [20, 57], minimum target size [1], etc. Similarly, much research has compared tactile with tangible interaction for tasks as various as puzzle solving [66, 72], layout-creation [44], photo-sorting [66], selecting/pointing [53], and tracking [34]. Most of the work comparing tangible to other interfaces builds on the assumption that physical interfaces are necessarily better because they mimic the real world. However, this assumption was rightfully questioned by Terrenghi et al. [66]. A 2DOF input device (e. g., a mouse) may, in fact, perform well in a 3D manipulation task due to its inherent accuracy or people’s familiarity with it. To better understand advantages and challenges of the three input modalities we thus compare them with each other in a single study.

Esteves and Oakley [19] also emphasize the fact that most studies comparing tangible interaction to other interaction paradigms are hard to generalize due to the highly simplistic tasks assigned to participants. Studies can thus only support very general claims on tangible interaction and its possible benefits. The lack of generalizability of such studies may also be explained by the overly focused participant groups in such studies. Very young participants often seem to be chosen to evaluate tangible interaction: school-aged children, for instance, were asked to evaluate the entertainment of Tangible User Interfaces (TUIs) [75], to solve puzzles [2], or asked to collaborate to understand which paradigm can be used to reduce conflicts in collaboration tasks [45, 50]. Similarly, Lucchi et al. [44] asked college students to recreate layouts using tactile and tangible interfaces. The learning effects of tangible interaction was also tested on non-adult participants in a study conducted by Price et al. [52]. We try to avoid this lack of generalizability by having a variety of participants and by using a task that is highly generalizable to 3D manipulation—3D docking. Such tasks have often been used in the literature to evaluate new 6DOF devices [21, 77], new interaction techniques [24], and for paradigm comparison studies [67] (for the latter, the docking was only conducted in two dimensions). We argue that using a low-level 3D docking task is the key to be able to generalize results from comparative studies.

Related to our work are also remote 3D manipulations through tactile input that benefit from the increasing availability of large displays and the pervasive nature of mobile, tactile-enabled devices. For instance, Liang et al. [43] investigated the use of two back-to-back mobile devices—to facilitate tactile input above and under the mobile device—with a combination of tactile gestures and sensors to support rotation, translation, stretching, slicing, . . . They also conducted an experiment to

examine the use of dedicated regions on the mobile device to control objects or the 3D environment. Similarly, Du et al. [18] investigated the use of a smartphone to navigate within a virtual environment on screen, while Katzakis et al. [36] examined the combination of mobile sensors and tactile input for 3D translation and rotation through a docking task. Coffey et al. [13], however, used ‘indirect’ tactile manipulation to navigate and examine a volumetric dataset to overcome the inherent issues of tactile interaction with stereoscopic rendering [69]. We are interested, in contrast, in a more ‘direct’ interaction³ which also displays the 3D information (e. g., [7])—we do not focus on remote manipulation using separate displays.

Our study mainly builds on the work by Hinckley et al. [28] and Tuddenham et al. [67]. Hinckley et al. [28] conducted comparative 3D docking studies focused on rotation with four different techniques including a 3D ball (our equivalent is a tangible interface) and a mouse. We go beyond their approach in that we consider a full 6DOF manipulation and evaluate more than time and accuracy. We go beyond Tuddenham et al.’s approach [67] in that we, while also comparing mouse, tactile input, and tangible interfaces, use 3D manipulation tasks—including for the tangible input device.


COMPARATIVE STUDY

As we aim to understand the use of mouse, tactile input, and tangibles for the manipulation of 3D scenes or datasets, our study investigates a task representative of 3D manipulation, in a realistic scenario, using a wide range of participants. Beyond time and error metrics, we observed people’s actions, learnt about their realistic preferences, and their subjective ratings of the techniques. We aimed to understand four of Nielsen’s five factors of usability [49]: effectiveness, efficiency, subjective satisfaction, error tolerance, and ease of learning. Error tolerance, was not within the scope of our study. The effectiveness is reflected by an accuracy score (in both angular and Euclidean distance), the efficiency by means of the time to complete the task, the subjective satisfaction by looking at participants’ answers to our questions, and the ease of learning by looking at the evolution of task completion times.

Task. The docking task we employ comprises translation in 3 DOF, re-orientation in 3 DOF, and precise final positioning of 3D shapes—actions representative of interactive 3D data exploration. A docking task⁴ consists of bringing a virtual object to a target position and orientation. The docking target is shown on the screen as a wire-frame version of the object, without the users having any control over the target’s position or orientation. Such a docking interaction thus mimics many aspects of typical 3D interaction, even though an actual docking target may only implicitly exist in real-life scenarios.

³The terms ‘direct’ and ‘indirect’ interaction have to be used carefully. While mouse input is arguably indirect, tangible and tactile input have both direct and indirect properties. Tactile input, in our case, occurs directly on the displayed data (albeit on a projection of the 3D shape) and is thus typically considered to be a direct interaction [41, 42, 47, 51, 56, 57, 63]. Tangible input directly manipulates a 3D shape (tangible) where the virtual shape is thought to be, but our visuals are projected onto the separate display. We thus argue that tactile and tangible interaction are more direct than mouse interaction.

⁴Other examples of docking task studies: [11, 21, 22, 24, 28, 71, 77].



(a) Screenshot of the task. (b) Person interacting in the tangible condition. (c) Tracking using cameras.

Figure 1. Study setup. Participants were asked to move and orient the shaded object such that it matches the target.

In practice, we used the Utah teapot as the 3D object to manipulate. It is a generic shape most people understand and does not present any orientation ambiguity. Other objects could have been used (Hinckley et al. [28] and Chen et al. [11] used a house with different colors on each side, Zhai et al. [77] used a tetrahedra with colored edges). Our pilot studies confirmed that there was no ambiguity in the orientation of the teapot. We randomly generated and validated the target positions beforehand (to ensure that all targets are reachable by all input modalities), yielding a pool of 15 valid target positions (see example in Fig. 1(a)). Our pilot studies confirmed that the use of perspective and relative size were enough to allow depth perception on a void background. Per input modality, we asked our participants to carry out 15 repetitions. For each of them we randomly selected the positions from the remaining positions in the pool. We used the same pool of positions for all modalities. We counter-balanced the order of input modalities each participant saw to reduce the bias from learning effects. Our within-participants design thus comprised of 3 input modalities \times 1 task \times 15 trials = 45 trials in total for each participant.

Each trial was started and validated on a key press by the participant (similar to Chen et al. [11] or Hinckley et al. [28]). We considered using a pedal for validation (e. g., [28]) but our pilots showed its triggering precision to be inferior to a key press. We asked participants to balance accuracy and speed, and intentionally did not reveal their achieved accuracy after each trial (as done by others [11, 28]) to avoid a bias toward accuracy [28]. In addition, to avoid participant response bias [15], we explicitly told them before the experiments that none of the techniques was developed by us.

Apparatus. For all three input modalities we used the same touch-enabled 21" LCD screen with a resolution of 1920×1080 pixels and a 60 Hz refresh-rate. Participants were seated in front of the screen which was slightly tilted (approx. 15°) to provide a comfortable tactile input setting (see Fig. 1). We decided against using a stereoscopic display as this causes a parallax issue [23, 69], as well as ‘touch-through’ issues [10]—users touch through the 3D objects to reach the touch-enabled screen. The mouse condition used a classical computer mouse: a Logitech m100 at 1000 dpi with a 125 Hz polling rate. The tangible condition was based on an optically tracked hand-held cardboard-based cuboctahedron (see Fig. 1(b)), each edge measuring 65 mm. The lack of embedded electronic parts make the tangible prop weigh only 26g. Mark-

ers on each face facilitated its 3D tracking with 6 DOF. Each marker was as big as the cuboctahedron face it was placed on to ensure an optimal tracking. The optical tracking system comprised two Project Tango tablets.⁵ Since camera refresh rates depends on lighting conditions (the darker the room, the lower the refresh rate), we set up a room with only artificial lighting.⁶ The lighting was then improved by using two 220W lightbulbs—each one producing 3300 lumen—reflected by photography umbrellas to avoid a direct over-lighting of the tangible prop which would hinder the optical tracking. Ultimately, our setup yielded camera framerates of 30 fps at a resolution of 800×600 . We adjusted the tablet positions according to a previous pilot study. In the final setup, the two cameras were located as shown in Fig. 1(c): one above to see both the screen and the tangible probe from above, and one on the participants’ left side (at approx. head level) so that the space in front of the screen was visible. Together, they allowed us to avoid dead angles: participants could comfortably hold the cuboctahedron without blocking the camera’s view. Programmatically, the optical tracking was realized thanks to a combination of the Vuforia⁷ and ARToolkit⁸ frameworks and stabilized by using the 1 € filter [9].⁹ The tactile input, finally, was captured using capacitive touch sensing built into the screen. This touch sensor provided up to 10 points—captured via TUIO [35].¹⁰ The overall setup (distance to screen, camera placement) also allowed users to rest their arms/wrists (mouse+keyboard condition) as well as to rest their arms, elbows, and shoulders (tactile/tangible conditions) on the table.

Interaction Mappings. As much as possible, we chose established mappings for the evaluated input modalities as follows.

(a) Mouse+Keyboard. Inspired by the mappings used by Blender,¹¹ Autodesk MDT,¹² or Catia and software tools based on VTK such as Paraview,¹³ we used the following mappings:

- right button: translation along the x - y -axes,

⁵See <https://www.google.com/atap/project-tango/>.

⁶In practice, the setting of refresh rates for cameras is not fully reliable on Android systems—even with our precautions. Nevertheless, our setup reduced the refresh rate variability as much as possible.

⁷See <https://www.vuforia.com/>.

⁸See <https://www.hitl.washington.edu/artoolkit/>.

⁹See <http://www.lifl.fr/casiez/1euro/>.

¹⁰See <http://www.tuio.org/>.

¹¹See <https://www.blender.org/>.

¹²See <http://www.autodesk.fr/products/autocad-mechanical/overview>.

¹³See <http://www.vtk.org/> and <http://www.paraview.org/>.




Figure 2. Tactile mappings for mobile 3D interaction.

- left button: Virtual Trackball rotation for the x -/ y -axes,
- keyboard modifier + right button: z -axis translation,
- keyboard modifier + left button: rotation around the z -axis (leftward mouse motion = clockwise rotation), and
- the use of the scroll wheel was disabled since zooming needed to be inaccessible for the docking task.

While several rotation techniques have been implemented (see the surveys by Chen et al. [11] and Bade et al. [3]), Bell [6]’s Virtual Trackball (VT) and Shoemake [62]’s Arcball seem to be the ones most frequently used in available softwares. Yet, they are often seen as frustrating by users because they violate a number of principles for intuitive interaction [3]. Based on our pilot studies we decided to use an improved version of Bell’s VT; one that respects the third principle mentioned by Bade et al. [3] and provides a transitive 3D rotation.

(b) Tactile Input. In contrast to mouse+keyboard and tangible input, no single established standard or quasi-standard for touch-based interaction with 3D data exists. Based on our survey of 36 commercial and academic mobile applications on Android and iOS (see Fig. 2), we found that most interaction mappings do not provide the 6 DOF we need. From those which do, most used the mapping that relies on either one or two fingers, with the latter providing rotation round the z -axis, uniform scaling, and translation along the x -/ y -axes using pinching (RST—Rotation, Scale, Translation). Some systems provide a RST technique with a system-controlled moding: once the user’s intention is captured by the system the control mode is locked. However, we decided not to use system-control moding because this could hinder the way users understand the interaction mapping. While studies have shown that it is possible to outperform the classical RST technique by separating the degrees of freedom [46], we believe that the intuitiveness of the pinching mapping can be of advantage in our case, so we decided to use the following mappings:

- 1 finger motion: virtual trackball rotation for the x -/ y -axes,
- 2 fingers—RST:
 - translation: translation along the x -/ y -axes,
 - rotation: rotation around the z -axis, and
 - pinching: z -axis translation (cf. Hancock et al. [25]).

(c) Tangible Input. Tangible input is not yet widely established outside academic research so we could not draw from established mappings in software tools. We thus decided to use the intuitive isomorphic position control: a one-to-one mapping that moves and rotates the virtual object similar to the motions of the tangible object in real life. While such an interaction could be classified as a minimal TUI, it fulfills the four characteristics of TUIs as defined by Ullmer and Ishii [68]—similar to other comparable tangible input devices in the literature [27, 64]—and is thus well suited for our study.

(d) Input Range. The input range of each modality was adjusted so that translations would not exceed the cameras’ Field of View (FoV) in the tangible condition. In other words, it was possible to achieve all 3D docking tasks without clutching for translations. Rotations, however, were not constrained by the cameras’ FoV and ranged from 19° to 228° . Clutching could be used for each modality by releasing the finger-pressure on mouse button, removing fingers from the tactile screen, or briefly using a second-hand grasp with the tangible object.

Participants. 36 unpaid participants (10 females) took part in our comparative study. Their ages ranged from 19 to 52 years (mean = 30.2, SD = 8.7; median = 26). Three were left-handed, the remaining 33 right-handed. With respect to their expertise with 3D manipulation on a computer, 12 participants ranked themselves as skilled due to frequent use of video-games or 3D softwares, while 24 participants stated they had no significant prior experience. Furthermore, 22 of the participants had a university degree, while 14 had a high school degree. They all had either normal or corrected-to-normal vision.

Procedure. Participants were guided through the study by means of a study controller software that presented the different task blocks in turn. Before starting the trials of a new input modality, participants were introduced to the interaction technique. They were intentionally given minimal instruction on using each device, they were only informed that they could

- use the mouse’s left and right buttons and the keyboard’s shift key in the mouse+keyboard condition,
- use multiple fingers on the tactile screen in front of them for the tactile condition, and
- use the tangible object for the tangible condition.

Further, the space in which the tangible object could be used was pointed out because participants had to keep within the field of vision of the cameras. An evaluator was present to answer potential questions during the experiment as well as take notes about the usage of each of the three input modalities.

Throughout the study, we asked participants to fill in several questionnaires. A first questionnaire captured their demographics and their level of fatigue before the experiment. After each condition, participants filled a questionnaire to assess their workload and fatigue level. For the former we used NASA’s Task Load Index,¹⁴ the latter was based on Shaw’s approach [60]. A final questionnaire assessed the subjective ratings for the different techniques. We go beyond the usual Likert-scale or ranking approach suggested by Nielsen [49] undergone in most studies: to confirm this last self-assessment, we informed participants that they would have to do a final set of 15 docking tasks, for which they could pick their favorite technique. Only after they had voiced their choice, we informed them that, in fact, the study was over and that the last question was only used to understand their true preferences. We used this procedure to better understand their preferences and to avoid a bias toward the technological advantages of tangible input. Because the experiment already took approx. more than an hour, we conjectured that, if asked to perform an additional set of trials, participants would have a strong incentive to

¹⁴See <http://humansystems.arc.nasa.gov/groups/tlx/downloads/TLXScale.pdf>

pick the solution they really preferred to use. We finally asked whether, if given the free choice, they would have carried the additional batch of 15 tasks—to better understand people’s eagerness to interact with the chosen technique. Indeed, Nielsen [49] explains that “data showing voluntary usage is really the ultimate subjective satisfaction rating,” which is what we assessed by this last question. **Variables.** In our comparative study we thus analyze one independent variable—the interaction modality—and five dependent variables—completion time, accuracy, fatigue, workload, and preferences. We took two different types of accuracy into account: the Euclidean distance to the target in 3D space as well as the rotational difference (in degrees) to the target.

Hypotheses. Based on our previous experience with the three input modalities, we hypothesized that:

- H1** The time spent on trials would be shorter in the tangible condition than in the tactile condition due to the inherent and fully integrated [33] structure. Tactile-based interaction would also be faster than mouse-based input due to its higher directness and partially integrated structure.
- H2** The accuracy for both the rotation and the Euclidean distance to the target would be better for the mouse than the tactile condition due to the better support of the hand when using a mouse. The accuracy of the tactile input, in turn, would be better than the tangible condition due to the lack of support for the hand when using tangibles.
- H3** The workload for the tangible condition would be low overall due to its intuitive mapping and fast interaction times—yet the need to have to hold the object and fine-position it would have a negative impact. The higher mental demand necessary to understand the mapping of tactile and mouse interaction balanced by the reduced physical demand of these techniques would produce a slightly higher workload than for the tangible.
- H4** The resulting fatigue would be highest for tangible input due to having to hold the physical object, lower for tactile input due to the added rest on the surface, and minimal for mouse input due to the arm resting on the table.
- H5** People prefer both tangible and tactile inputs over mouse input: tactile for its “intuitive” mappings and reasonable accuracy, tangible because it benefits from the similarity to real-world interaction (but lacks a bit of accuracy). Mouse-based input is not preferred because it forces the separation of input DOF, while the others provide means of controlling several DOF in an integrated fashion.

RESULTS

We collected a total of 1620 docking trials from 36 participants, i. e., 540 trials for each input modality. To compare the three conditions, we measured the task completion times as well as an accuracy score for each condition and each participant based on their results in each of the trials for a given condition.

While HCI experiment data is traditionally analyzed by applying null-hypothesis significance testing (NHST), this form of analysis of experimental data has come under increasing criticism within the statistics [5, 14] and HCI communities [17, 16]. We thus report our results using estimation techniques




Figure 3. Task completion times: (a) absolute values in seconds and (b) pairwise comparison ratios (left-side technique divided by right one, 1 means similar performances). Error bars: 95% confidence intervals.




Figure 4. Task completion times in seconds: (a) mouse condition, (b) tactile condition, and (c) tangible condition. Error bars: 95% CIs.

with effect sizes¹⁵ and confidence intervals (instead of *p*-value statistics), consistent with recent APA recommendations [70].

Task Completion Time. We analyze log-transformed time measurements to correct for positive skewness and present our results anti-logged, as it is standard in such cases [55]. Consequently, we arrive at geometric means.¹⁶ They dampen the effect of potential extreme trial completion times which could otherwise have biased an arithmetic mean.

We present the completion time results in Fig. 3(a). It shows that it took participants 61 s to complete the task in the mouse condition, 47 s in the touch condition, and 26 s in the tangible condition. While the confidence intervals reveal a difference in favor of the tangible condition over the mouse and touch conditions, they do not allow us to say anything more with confidence. We thus computed a pairwise comparison between the different conditions, see Fig. 3(b). The differences in these pairwise comparisons were also anti-logged and thus present ratios between each of the geometric means. These ratios all being clearly $\neq 1$ allows us to interpret the time differences of completing the task. Fig. 3(b) shows that there is strong evidence for the tangible condition to clearly outperform the mouse condition: it is more than twice as fast as the mouse condition. The difference between the tangible condition and the touch condition is also quite strong: the tangible condition is almost twice as fast as the touch condition. The difference between mouse and touch is not as strong; yet, the touch condition can still be considered faster than the mouse condition.

¹⁵The term *effect size* here refers to the different means we measured. We do not refer to standardized effect sizes [12] because reporting them is not always recommended [4], but rather to simple effect size.

¹⁶While an arithmetic mean uses the sum of a set of values to obtain the mean, a geometric mean uses the product of the set’s values.




Figure 5. Euclidean distances: (a) absolute values in space units and (b) pairwise comparison ratios. Error bars: 95% CIs.




Figure 6. Rotational distances: (a) absolute values in $^{\circ}$ and (b) pairwise comparison ratios. Error bars: 95% CIs.

We also checked for learning effects by dividing the 15 trials of each condition into three subsets of 5. We thus analysed the completion times for the three thirds of trials in Fig. 4. As shown in Fig. 4(a), the completion time in the mouse condition drops from 75 seconds in the first set of 5 trials to approximately 55 seconds in the second and third subsets of trials. In the tactile condition, we can observe a strong evidence of a reduction of the completion time between the first subset of trials and second subset and less evidence for a decrease from the second to the last subset. In the tangible condition, however, we did not find any evidence of a difference in completion time between each subsets of trials.

Accuracy. An inspection of Q-Q plots on the Euclidean and angular distance showed that the data did not follow a normal distribution but instead approximately followed a log-normal distribution. Thus, we also log-transformed both measurements for the analysis and we present the results anti-logged.

(a) Euclidean Distance. We report the Euclidean distance to the target in Fig. 5(a). It is computed as the distance between the target’s 3D center to the movable teapot’s 3D center. Fig. 5(a) shows that all three techniques lead to similar accuracies, with means of 5 mm for the mouse condition and the tangible condition, and 6 mm for tactile input. Pairwise comparison between the conditions (Fig. 5(b)) suggest that the tangible and the mouse input may have a slight advantage over tactile interaction, while both mouse and tangible inputs are very similar in accuracy to each other for our chosen task.

(b) Rotational Distance. Fig. 6(a) reports the rotational distance to the target. The results are 3.4° for mouse input and 3.7° for both tactile and tangible input. Fig. 6(b) shows the pairwise comparison between the conditions. Similar to the Euclidean distance, these comparisons indicate that all techniques are similar. There is weak evidence that the mouse may yield slightly more rotationally-precise results than tactile or tangible. However we did not find evidence for a performance difference between tactile and tangible for the rotation.

Our analysis of both types of accuracy did not yield evidence for a large difference in accuracy between the different input modalities. This result did not change if we—to account for




Figure 7. Total workload in overall NASA TLX units ($\in [0, 100]$). Error bars are 95% CIs for the total workloads.

learning effects—only analyzed the latter 2/3 or even the last 1/3 of the trials of each participant in the different conditions.

Measuring Workload. When collecting workload measurements using NASA’s TLX we noticed that the pilot-study participants were often confused by its second part—weighing each of the different sub-aspects (i. e., mental, physical, and temporal demand, performance evaluation, effort, and frustration) for a given task. To avoid the seemingly random choices which would lead to inconclusive or even incorrect results we decided not to consider this second part of the TLX. We were thus left with what is called a *Raw TLX* (RTLX). According to Hart’s [26] survey, the RTLX may be equally well suited as the regular TLX. We thus compute the workload for each task as the average of the RTLX ratings by participants.

The results of this analysis are shown in Fig. 7. Here, we show the total workload for each condition as well as the specific sub-aspects rated by participants. The non-overlapping confidence intervals between the tactile and the tangible condition show that there the tangible condition requires a lower workload than the tactile condition, yet for differences between the tangible and the mouse condition and even more so between the mouse and the tactile condition there is much less evidence.

The individual sub-aspects of the workload differs somewhat between the different conditions, but we did not observe many striking differences between the three input modalities. Fig. 8 shows a detailed analysis of the differences of the sub-aspects. We can observe that there are only clear differences in the rating of mental demand between the mouse and tangible condition (Fig. 8(a)), for the physical demand between the mouse and the other two (Fig. 8(b)), as well as for the temporal demand between tactile and tangible condition (Fig. 8(c)). The other comparisons between conditions for the sub-aspects only show gradual differences (also evident in the respective lengths of the colored patches in Fig. 7). Yet, we can observe a slight advantage of mouse over tactile for performance evaluation (Fig. 8(d)), a small advantage of tangible over the other two for effort (Fig. 8(e)), as well as a lower frustration in the tangible condition (Fig. 8(f)). The difference in temporal demand between mouse and tangible (Fig. 8(c)) matches the differences observed in overall interaction times between them (Fig. 3). In contrast, there was no difference between the mouse and tactile condition even though we observed a clear difference in the completion time between them.

Measuring Fatigue. We present the analysis of the fatigue measurement in Fig. 9. Interestingly, none of the conditions exhibits a particularly high level of fatigue with the means all being lower than 4 on the scale of 0 to 10. While the mean of our measurements is highest for the tactile condition, based on the confidence intervals there is no evidence that there would be an important difference between any of the conditions.




Figure 8. Workload sub-aspects of Fig. 7’s data in individual TLX units ($\in [0, 100]$): (a) mental, (b) physical, and (c) temporal demand, (d) performance (0 is best), (e) effort, and (f) frustration. Error bars: 95% CIs.




Figure 9. Total fatigue on a scale from 0 to 10. Error bars are 95% CIs for the total fatigue ratings.




Figure 10. Participant preferences: (a) self-reported preferred technique, (b) technique chosen for the additional (but hypothetical) set of 15 trials, and (c) technique chosen by those participants who would have voluntarily stayed to complete the additional set of 15 trials.

Measuring Preferences. In addition to the measured values we asked for participants’ preferences. As described above, we asked for both a normal preference rating and the technique they would choose if faced with another set of 15 trials, as well as if they would want to actually stay for these additional 15 trials. Fig. 10 reports these self-ratings.

Interestingly, the tangible condition was chosen most often for the stated preference (24 \times). Among those, however, 5 participants hesitated between touch and tangible, all ultimately picking the tangible as their favorite. The remaining 12 participants stated that they preferred tactile over mouse (tactile: 8 \times ; mouse: 4 \times). When faced with an additional set of trials, a majority still preferred the tangible condition (22 \times). The tactile vs. mouse preference, however, changed with the mouse now being rated higher than the tactile (tactile: 5 \times ; mouse: 9 \times). Of the 16 participants who freely decided to do the tasks again ((c) in Fig. 10), 11 preferred the tangible condition, 4 favored the mouse, and 1 picked tactile.

The Impact of Experience. Based on the demographics of the participants as well as their experience in 3D manipulation we also analyzed the difference between experienced and non-experienced participants. Fig. 11 shows the Euclidean and rotational distances as well as the completion times for each condition, for different levels of experience. The confidence intervals seem to always suggest a more accurate task completion of experienced participants for each input modal-

ity. For tactile input we can even observe strong evidence for this difference, both for Euclidean and angular distances. For task completion times there is strong evidence of a better performance of experienced user only for the mouse condition.

QUALITATIVE OBSERVATIONS

In addition to the quantitative analysis based on the captured data we also provide a summary of qualitative, observational data that was captured by the experimenter during the study.

Mouse. Among our 36 participants, we observed that 21 had issues with the mouse’s mapping and moding. They did not intuitively try to combine the key and the mouse to perform the 3D operations they wanted to carry out. The subset of eight participants stating to be experienced in 3D manipulation through video games/software did not encounter any problem to find the mapping. The seven remaining participants—without prior experience of 3D manipulation with a mouse—did not exhibit any issue with our mapping. Among the 21 participants who were observed to have issues, only 12 reported that the mapping was too difficult to find alone or remember. 22 participants praised the accuracy the mouse offered, nine stated that it was nice because they were used to using a mouse, and four reported that the lack of physical demand of the device is one of its assets. In the completion time measures however, we can observe a clear evidence for a better performance of experienced participants in the mouse condition.

Tactile. We observed that 20 participants had difficulties with the two-finger RST interaction. Our impression was that these difficulties arose from the RST technique integrating rotation and scaling into a single interaction, as opposed to only affecting a single DOF at a time. Participants had troubles moving the two touching fingers without modifying their distance to each other, resulting in unwanted zoom-in or -out actions while translating the object. Eleven participants stated that they would have preferred a mapping that would allow them to translate the object along the x - y -axes without affecting its distance or its orientation around the z -axis. Still, 13 participants assessed the tactile interaction as an intuitive input.

We also observed that 20 participants used fingers from different hands for the pinch interaction, while 16 used two fingers from the same hand for the same interaction. This important difference in providing the input for the same type of interaction mapping likely had a large impact on people’s accuracy and speed during the tasks as well as their preference. In both cases, however, participants reported the tactile-based interaction and the corresponding mapping to be “intuitive” and “more natural than the mouse”—13 participants made such statements when asked to assess the different conditions. Of them, eight specifically praised the tactile input for its perceived accuracy, while five reported that “they were faster with it” than with a mouse. Two participants stated that they felt in control of the data they were manipulating, mirroring previous statements in other studies [73, 76]. Three participants stated that they resented the fact that the removal of their fingers from the tactile screen led to little or even big transformations being issued inadvertently—Tuddenham et al.’s [67] *exit error*.

Tangible. The observation of the tangible condition showed that most participants were indeed not familiar with this type




Figure 11. Impact of experience on (a) Euclidean distance (in mm), (b) rotational distance (in °), and (c) completion time (in ms). Error bars: 95% CIs.

of input and manipulated the tangible object in interesting ways. For rotations $>90^\circ$, for instance, 29 participants used two hands, while seven used just a single hand. The second hand, however, was used only briefly for clutching in large rotations. Additional adjustments with the manipulating hand after clutching then ensured fine positioning, freeing the other hand for the trial validation (in contrast to other setups [71]). Bi-manual clutching thus did not affect our measurement of the task completion time. We also observed that 12 participants completed the docking task by sequentially manipulating the different types of transformations. In contrast to an integrated interaction used by the other participants, they first translated the tangible object along the x -/ y -axes, then rotated it to match the orientation of the target, and finally translated it again to obtain the correct z -location. It is unclear, however, if these 12 participants did not take advantage of the integrated interaction due to being used to the separated interaction offered by traditional 3D user interfaces, due to being afraid of losing the optical tracking, or due to not feeling comfortable with the DOF-integrated manipulation offered by the tangible interaction. Noteworthy, 17 participants (i. e., about half of them) used their non-dominant hand to interact if the docking target happened to be on the non-dominant side of the participant. Finally, it is interesting to note that 20 participants reported a lack of accuracy with the tangible condition. Among them, only four thought it was solely due to the technology while the remaining 16 believed this was due to the inaccuracy of their hand movements. Overall, 25 participants reported that the tangible condition was simple (11 \times) or intuitive (14 \times).

DISCUSSION

With our ultimate goal of better understanding the different input modalities that are available for spatial manipulation in the context of the exploration of 3D scientific data, we now discuss those aspects of our results that are most surprising and/or most relevant for our target application domain.

Efficiency. In line with our hypothesis H1, we found that the tangible interaction was faster than the tactile input which, in turn, was faster than mouse control. The reason for this difference in completion times is likely the inherent and straightforward integration of DOF control in the tangible condition, whereas the tactile and mouse condition need to switch interaction modes—with all the negative implications arising from user- or even system-controlled interaction modes (e. g., [8, 58]). While tactile input still facilitates some degree of direct manipulation and DOF integration (4 DOF in the RST mode), the mouse only controls 2 DOF at any given time and is also the most indirect input device.

We conjecture that, despite the established benefits of the RST mapping, participants encountered difficulties with it that may have impacted their performance, in particular the completion

time. We also hypothesize that the tangible condition’s fast completion time may be a reason for its high accuracy: an approximate docking is achieved much faster than in the other conditions, giving participants time to fine-tune their docking.

Learnability. According to Nielsen [49], learnability is one of the most important factors of usability. We noticed during the experiment that not a single user decided to give up on reaching the level of accuracy he/she wanted to achieve in a given trial with a given technique. In other words, they were all able to complete the tasks successfully. Looking at Fig. 4 we can clearly see that learning happens in the mouse and tactile conditions. In the mouse condition, a third of trials (i. e., 5 trials) were enough to achieve significantly better results and master the mouse interaction. In the tactile condition, after a first subset of trials, the completion time required for a trial was also visibly decreased. From the evolution of the completion time in the tactile condition, we could however wonder if results would have gotten any better if participants were given an additional set of trials. In the tangible condition, we cannot find any evidence of a learning in Fig. 4. Even when comparing the completion time of the first trial to the others, we could not detect signs of an improvement. These results thus support previous statements concerning the affordances of TUIs: they do not require learning as people are used to performing physical manipulation in the real world.

Effectiveness. The effectiveness was measured in form of an accuracy score of each modality. We initially thought that the different input modalities provided different degrees of accuracy. A mouse has a high-dpi sensor and a well-rested grasp configuration, while tactile relies on the finger as a rather blunt instrument with less support. The tangible condition, finally, needs optical tracking with the arm operating in empty space. Yet, surprisingly, our data does not provide evidence for any of the three techniques being more efficient (not providing evidence for H2). These results are even more surprising since they contradict as well the results obtained in the 2D docking task studied by Tuddenham et al. [67] who found that the tangible condition exhibits an easier and more accurate manipulation than the tactile condition. Similarly, they contradict results by Vuibert et al. [71] who found that a constrained desktop device—such as the PHANTOM—leads to a better accuracy than unconstrained interaction. The results extend previous finding from Hinckley et al. [28] who found no difference of rotational accuracy between a tangible-like interface (3D ball and 3D tracker) and the mouse condition. However, many participants still reported that they *perceived* that they had precise control over their actions in the mouse (22 \times) and tactile conditions (8 \times). In the tangible condition, however, they felt that they had uncontrollable and involuntary hand movements and 20 of them reported the lack of accuracy they

experienced. We believe that this *perceived* level of accuracy should not be disregarded in a decision of which interaction device to use or to offer for tasks that require a high accuracy. A possible explanation is that, overall, tangible and tactile interaction are less accurate than mouse+keyboard interaction but all inputs allow users to achieve a similar final accuracy. We believe that this *perceived* level of accuracy should not be disregarded in a decision of which interaction device to use or to offer for tasks that require a high accuracy. A possible explanation is that, overall, tangible and tactile interaction are less accurate than mouse+keyboard interaction but all inputs allow users to achieve a similar final accuracy.

Workload. With our data we cannot confirm hypothesis H3, but the overall measurements show—for our task and participant group—the same tendency as argued in H3: The perceived workload for the tangible interaction is lower than for the tactile condition as well as slightly lower than for the mouse condition. We believe, however, that the tactile input (as well as mouse input) can be improved. We saw that many participants kept their arms in the air while interacting in the tactile condition which contributed to the workload. This issue could be improved upon using a better (tactile-only) setup and a better interaction mapping. For the latter we noticed that many participants had problems with the sensitivity of the *z*-translation—caused by them starting the interaction with their fingers very close together as they are used to interact that way on smart phones and tablets. Tactile interaction—even or in particular if it uses the same interaction mappings—may require people to re-learn some of their familiar interaction techniques as they transition from small to larger screens.

Similarly, we also observed some frustration with tangible input. Some participants who felt at ease with tangible input tried to manipulate it fast with one or two hands. Our optical tracking system, however, was only good enough for slow to medium movements but could not follow relatively fast manipulations, leading to participant frustration. Similarly, participants occasionally occluded both cameras of our tracking system, leading them to report frustration due to the interrupted tracking—maybe even focusing on such issues when rating the frustration and not concentrating on other interaction issues.

Fatigue. Based on fatigue measurements we cannot confirm our hypothesis H4. The study setup was created such that—to facilitate a fair comparison—there was both enough space for mouse-based and tangible input as well as an equivalent view on the screen for all conditions. This arrangement, however had an implication on the self-assessed fatigue values. Indeed, many participants did not rest their elbows in the tactile condition, potentially resulting in shoulder and arm fatigue that would probably not have been perceived on a setup created specifically for tactile interaction. Such a setup would have also reduced the physical demand of the workload for tactile interaction. This arrangement would only reduce the arm and shoulder fatigue but would not impact the finger fatigue that we observe in Fig. 9. Nevertheless, the fatigue ratings for all techniques are quite similar, so that at least the fatigue measurement seems to have little impact on the choice of interaction modality. Because our tangible prop was comparatively light (26 g) it probably had no influence on the overall fatigue

of the users, and we thus cannot generalize these results to other types of props relying on self-tracking which are heavier.

We would also like to emphasize that tangible interaction lacks the possibility to easily maintain the virtual object in a given position and orientation as people release it. This was reported by four participants when asked what they liked about each condition. We can thus conjecture that an extended use of the tangible could drastically impact fatigue if it is impossible to release the tangible object without causing exit errors.

Subjective Preferences. Our data shows an overwhelming preference for tangible input, thus contradicting our hypothesis H5. We believe, however, that this result should be taken with a grain of salt. Our participants' preference for tangible interaction is likely biased by them being used to mouse-based and tactile interfaces, while tangible input is new to most of them. Indeed, some of the participants who selected tangible input as their favorite explained that they would use this technique for the forced and free choice (i. e., (b) and (c) in Fig. 10) because they do not have the opportunity to “play” with such technology at home, while they have easy access to tactile screens and mice. The novelty effect thus clearly made a difference at least for 5 out of the 11 participants who picked the tangible option for the last preference choice (i. e., (c) in Fig. 10). We also believe the use of the word “play” by the participants is noteworthy. While usually subjective satisfaction measures focus on aspects such as simplicity, safety, completeness, and irritation/frustration, TUIs introduce the concept of fun. This may further bias subjective preference studies. We can thus conclude that, thanks to its entertaining dimension and the novelty effect, the tangible interaction is the preferred mean of interaction. While the novelty effect may fade, the entertaining property of tangible interaction will probably remain, making tangibles perfectly suitable, for instance, for children—as studied, e. g., by Horn et al. [29].

Experience. The faster completion times in the mouse condition for experts is not surprising: most of tools available for 3D manipulation use the classical mouse+keyboard interface and these results were predictable. It is interesting to notice; however, that experience had less influence in the mouse condition over the accuracy achieved by the two groups of participants. Similarly, since tangible interaction is still largely a focus of research activities as of today, experience had likely not a big influence on the results we obtained. All participants were equally prepared for this type of interaction due to their general experience manipulating objects directly in 3D space. We have no clear explanation, however, for our observation of a small improvement in accuracy for experienced participants for tactile input. While some of them may have tried one of the few 3D exploration or modification applications on mobile environments, the lack of a standard way of interacting with 3D data in mobile apps (Fig. 2) leads us to believe that is probably not the reason for the observed difference.

Realistic Application Scenarios. While our study scenario and tasks were chosen to be representative of generic 3D interaction as needed for visual data exploration, for realistic scenarios we likely face different requirements. We envision, for example, that longer interaction periods will be needed

with different types of tasks and more complex interaction techniques. The longer interaction periods will have an effect on fatigue and workload, in particular for tangible and tactile input. Realistic tasks, moreover, require more than 6 DOF interaction: uniform or non-uniform scaling are needed as well as interactions constrained to specific DOF should at least be included. In addition, many other interaction modalities are needed for practical applications such as cutting plane interaction, parameter specification, view or data selection, etc. (e. g., [13, 38, 76]). All these are likely to favor mouse-based and tactile input, as tangible interaction will likely be more difficult to use for generic interaction—unless multiple tangible input devices are used. Tangible input, however, may have some benefits for specialized input (e. g., [32, 65]), while tactile input may be better for integrated approaches (e. g., [13, 40, 65]). A final aspect to consider for realistic application scenarios is that, unlike the participant population we tested, we would be faced with experts in 3D interaction as they carry out such tasks on an everyday basis. Even though the learning effects we saw did not affect the results of our study overall, we may see other preference ratings among domain experts after longer periods of use than the ones voiced by our participants.

Summary of Limitations. The discussion so far has, in fact, mentioned many of the limitations of this work already, so we only provide a brief summary here. We strove to conduct a study that would avoid the numerous pitfalls of such a comparison study by having a population of users that was more representative than in other HCI studies, facilitating a fair comparison of each technique, and limiting the impact of biases. Yet, our study was limited by the need for a setup that would accommodate all three input modalities, while in practice dedicated setups better suited to a given modality would lead to better individual results. Moreover, practical applications will require more complex interaction scenarios, for which mouse and tactile-based input are likely better suited than tangible interaction. In addition, the chosen participant population for a quantitative experiment such as ours is different for the ultimate target audience, and the novelty factor of tangible interaction also introduced a bias—in particular for the self-reported preferences. Another influence of the chosen participants is that we faced learning effects, that would disappear if the techniques would be used in practice for a longer time. Finally, the chosen mapping for, in particular, tactile interaction may be successful in one type of application, but other applications and combinations with additional interface elements may require other mappings that may better be suited for visual exploration of 3D data. We believe that this mapping question should be the focus of future research.

CONCLUSION

We have compared mouse, tactile and tangible interaction in the context of 3D manipulation with a 3D docking task. We have provided a study design that limited the biases involved in this kind of study—participant response bias [15], or learning effect. We set reliable and comparable methods in a setup that was not in the advantage of any of the techniques. We also imagined a technique to better assess the subjective preference of participants by tricking them in thinking that they had an additional set of trials to perform.

	advantages	disadvantages
mouse	<ul style="list-style-type: none"> • availability, familiarity • perceived accuracy • DOF separation • low physical fatigue • moding for complex tasks 	<ul style="list-style-type: none"> • difficult mapping • slowest interaction • moding required
tactile	<ul style="list-style-type: none"> • availability, familiarity • perceived precision • increased directness • faster than mouse • easier mapping • multiple mapping options 	<ul style="list-style-type: none"> • unclear suitability of given mappings • slower than tangible • physical fatigue, exit error
tangible	<ul style="list-style-type: none"> • fastest interaction • intuitive mapping • impression of control • novelty factor 	<ul style="list-style-type: none"> • complex tasks unsupported • relies on 3D tracking • physical fatigue, exit error • separate object needed • rigid interaction mapping • always on, extra moding needed to stop interacting

Table 1. Advantages/limitations of each input modality.

Despite the limitations mentioned, our study has provided valuable insights on the potential of the three input modalities—mouse, tactile, and tangible—for the use in 3D interaction in general and, specifically, for the visual exploration of 3D data. In particular, we found that they are all equally well suited for precise 3D positioning tasks—contrary to what is generally assumed about tactile and tangible as input modalities. Our analysis of task completion time showed that tangible interaction was fastest, tactile slower, and mouse slowest. However, we did observe learning effects that may play out for longer-term usage, even though our data still showed the same advantage for tangible interaction if only the last third of trials was examined. Moreover, we discussed several additional considerations that need to be taken into account when designing practical interaction scenarios that put the observed advantages of tangible interaction into perspective. Researchers can now build on our findings by knowing that there is not a single input modality that would be a clear favorite for controlling 3D data during visual exploration, but that all three have their respective advantages and disadvantages that which be considered and which are summarized in Table 1.

Our findings also facilitates further studies that can now focus on other aspects of the different input modalities. In particular, the interaction mapping for tactile input will remain a focus of future research. In addition, the issue of the exit error will have to be addressed for both tactile and tangible inputs. The presence or the lack of spatial multiplexing of DOF control for tactile (which some participants did not use despite this being possible) is another aspect that should be investigated. A closer investigation of people’s use of dominant and non-dominant hands during interaction for both the tangible and the tactile conditions also would be an interesting path to follow. Ultimately, however, we want to continue our examination of how to best create an interaction continuum that allows one to fluidly switch between different interaction scenarios and interaction environments—picking the best one for a given task or situation. This direction of work will be facilitated by the insights we gained with this study.

REFERENCES

1. Pär-Anders Albinsson and Shumin Zhai. 2003. High precision touch screen interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 105–112. DOI: <http://dx.doi.org/10.1145/642611.642631>
2. Alissa N. Antle, Milena Droumeva, and Daniel Ha. 2009. Hands on what? Comparing children’s mouse-based and tangible-based interaction. In *Proceedings of the International Conference on Interaction Design and Children*. ACM, New York, 80–88. DOI: <http://dx.doi.org/10.1145/1551788.1551803>
3. Ragnar Bade, Felix Ritter, and Bernhard Preim. 2005. Usability comparison of mouse-based interaction techniques for predictable 3D rotation. In *Proceedings of the International Symposium on Smart Graphics*. Springer, Berlin/Heidelberg, 138–150. DOI: http://dx.doi.org/10.1007/11536482_12
4. Thom Baguley. 2009. Standardized or simple effect size: What should be reported? *British Journal of Psychology* 100, 3 (Aug. 2009), 603–617. DOI: <http://dx.doi.org/10.1348/000712608X377117>
5. Monya Baker. 2015. Statisticians issue warning over misuse of P values. *Nature* 531, 7593 (March 2015), 151. DOI: <http://dx.doi.org/10.1038/nature.2016.19503>
6. Gavin Bell. 1988. Bell’s trackball. Source code, formerly at <http://www.dispersoid.net/code/trackball.c>, available through the Internet Archive’s copy of July 17, 2007. (1988). <http://web.archive.org/web/20070717090837/http://www.dispersoid.net/code/trackball.c>
7. Lonni Besançon, Paul Issartel, Mehdi Ammi, and Tobias Isenberg. 2017. Hybrid tactile/tangible interaction for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 881–890. DOI: <http://dx.doi.org/10.1109/TVCG.2016.2599217>
8. William Buxton. 1986. Chunking and phrasing and the design of human-computer Dialogues. In *Proceedings of the IFIP World Computer Congress*. 475–480. <http://www.dgp.toronto.edu/OTP/papers/bill.buxton/chunking.html>
9. Géry Casiez, Nicolas Roussel, and Daniel Vogel. 2012. 1 € filter: A simple speed-based low-pass filter for noisy input in interactive systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 2527–2530. DOI: <http://dx.doi.org/10.1145/2207676.2208639>
10. Li-Wei Chan, Hui-Shan Kao, Mike Y Chen, Ming-Sui Lee, Jane Hsu, and Yi-Ping Hung. 2010. Touching the void: Direct-touch interaction for intangible displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 2625–2634. DOI: <http://dx.doi.org/10.1145/1753326.1753725>
11. Michael Chen, S. Joy Mountford, and Abigail Sellen. 1988. A study in interactive 3-D rotation using 2-D control devices. *ACM SIGGRAPH Computer Graphics* 22, 4 (Aug. 1988), 121–129. DOI: <http://dx.doi.org/10.1145/378456.378497>
12. Robert Coe. 2002. It’s the effect size, stupid: What effect size is and why it is important. In *Proceedings of the Annual Conference of the British Educational Research Association*. <http://www.leeds.ac.uk/educol/documents/00002182.htm>
13. Dane Coffey, Nicholas Malbraaten, Trung Le, Iman Borazjani, Fotis Sotiropoulos, Arthur G. Erdman, and Daniel F. Keefe. 2012. Interactive slice WIM: Navigating and interrogating volume datasets using a multi-surface, multi-touch VR interface. *IEEE Transactions on Visualization and Computer Graphics* 18, 10 (Oct. 2012), 1614–1626. DOI: <http://dx.doi.org/10.1109/TVCG.2011.283>
14. Geoff Cumming. 2014. The new statistics: Why and how. *Psychological Science* 25, 1 (Jan. 2014), 7–29. DOI: <http://dx.doi.org/10.1177/0956797613504966>
15. Nicola Dell, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies. 2012. “Yours is better!” Participant response bias in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 1321–1330. DOI: <http://dx.doi.org/10.1145/2207676.2208589>
16. Pierre Dragicevic. 2016. Fair statistical communication in HCI. In *Modern Statistical Methods for HCI*, Judy Robertson and Maurits Kaptein (Eds.). Springer International Publishing, Cham, Switzerland, Chapter 13, 291–330. DOI: http://dx.doi.org/10.1007/978-3-319-26633-6_13
17. Pierre Dragicevic, Fanny Chevalier, and Stéphane Huot. 2014. Running an HCI experiment in multiple parallel universes. In *Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, 607–618. DOI: <http://dx.doi.org/10.1145/2559206.2578881>
18. Yuan Du, Haoyi Ren, Gang Pan, and Shjian Li. 2011. Tilt & touch: Mobile phone for 3D interaction. In *Proceedings of the International Conference on Ubiquitous Computing*. ACM, New York, 485–486. DOI: <http://dx.doi.org/10.1145/2030112.2030183>
19. Augusto Esteves and Ian Oakley. 2011. Informing design by recording tangible interaction. In *Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, 2077–2082. DOI: <http://dx.doi.org/10.1145/1979742.1979893>
20. Clifton Forlines, Daniel Wigdor, Chia Shen, and Ravin Balakrishnan. 2007. Direct-touch vs. mouse input for tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 647–656. DOI: <http://dx.doi.org/10.1145/1240624.1240726>
21. Bernd Froehlich, Jan Hochstrate, Verena Skuk, and Anke Huckauf. 2006. The GlobeFish and the GlobeMouse. In

- Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 191–199. DOI: <http://dx.doi.org/10.1145/1124772.1124802>
22. David Glesser, François Bérard, and Jeremy R. Cooperstock. 2013. Overcoming limitations of the trackpad for 3D docking operations. In *Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, 1239–1244. DOI: <http://dx.doi.org/10.1145/2468356.2468578>
 23. Morris Goodman and Arthur Teicher. 1988. To touch or not to touch. *Psychotherapy: Theory, Research, Practice, Training* 25, 4 (1988), 44–64. DOI: <http://dx.doi.org/10.1080/07351692009348874>
 24. Mark Hancock, Sheelagh Carpendale, and Andy Cockburn. 2007. Shallow-depth 3D interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 1147–1156. DOI: <http://dx.doi.org/10.1145/1240624.1240798>
 25. Mark Hancock, Thomas ten Cate, and Sheelagh Carpendale. 2009. Sticky tools: Full 6DOF force-based interaction for multi-touch tables. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*. ACM, New York, 133–140. DOI: <http://dx.doi.org/10.1145/1731903.1731930>
 26. Sandra G. Hart. 2006. 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: <http://dx.doi.org/10.1177/154193120605000909>
 27. Ken Hinckley, Randy Pausch, John C Goble, and Neal F Kassell. 1994. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 452–458. DOI: <http://dx.doi.org/10.1145/191666.191821>
 28. Ken Hinckley, Joe Tullio, Randy Pausch, Dennis Proffitt, and Neal Kassell. 1997. Usability analysis of 3D rotation techniques. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*. ACM, New York, 1–10. DOI: <http://dx.doi.org/10.1145/263407.263408>
 29. Michael S Horn, R Jordan Crouser, and Marina U Bers. 2012. Tangible interaction and learning: The case for a hybrid approach. *Personal and Ubiquitous Computing* 16, 4 (2012), 379–389. DOI: <http://dx.doi.org/10.1007/s00779-011-0404-2>
 30. Tobias Isenberg. 2016. Interactive exploration of 3D scientific visualizations on large display surfaces. In *Collaboration Meets Interactive Surfaces and Spaces: Theory and Practice*, Craig Anslow, Pedro Campos, and Joaquim Jorge (Eds.). Springer, Berlin. To appear. DOI: <http://dx.doi.org/10.1007/978-3-319-45853-3>
 31. Hiroshi Ishii. 2008. The tangible user interface and its evolution. *Communications of the ACM* 51, 6 (June 2008), 32–36. DOI: <http://dx.doi.org/10.1145/1349026.1349034>
 32. Bret Jackson, Tung Yuen Lau, David Schroeder, Kimani C. Toussaint Jr., and Daniel F. Keefe. 2013. 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: <http://dx.doi.org/10.1109/TVCG.2013.121>
 33. Robert JK Jacob, Linda E Sibert, Daniel C McFarlane, and M Preston Mullen Jr. 1994. Integrality and separability of input devices. *ACM Transactions on Computer-Human Interaction* 1, 1 (1994), 3–26. DOI: <http://dx.doi.org/10.1145/174630.174631>
 34. Yvonne Jansen, Pierre Dragicevic, and Jean-Daniel Fekete. 2012. Tangible remote controllers for wall-size displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 2865–2874. DOI: <http://dx.doi.org/10.1145/2207676.2208691>
 35. Martin Kaltenbrunner, Till Bovermann, Ross Bencina, and Enrico Costanza. 2005. TUIO: A protocol for table-top tangible user interfaces. In *Proceedings of the International Workshop on Gesture in Human-Computer Interaction and Simulation*. Vannes, France. <http://opensourcecontrol.org/tuio-protocol-table-top-tangible-user-interfaces>
 36. Nicholas Katakis, Kiyoshi Kiyokawa, and Haruo Takemura. 2013. Plane-casting: 3D cursor control with a smartphone. In *Proceedings of the Asia Pacific Conference on Computer Human Interaction*. ACM, New York, 199–200. DOI: <http://dx.doi.org/10.1145/2525194.2525275>
 37. Daniel F. Keefe. 2010. Integrating visualization and interaction research to improve scientific workflows. *IEEE Computer Graphics and Applications* 30, 2 (March/April 2010), 8–13. DOI: <http://dx.doi.org/10.1109/MCG.2010.30>
 38. Daniel F. Keefe and Tobias Isenberg. 2013. Reimagining the scientific visualization interaction paradigm. *IEEE Computer* 46, 5 (May 2013), 51–57. DOI: <http://dx.doi.org/10.1109/MC.2013.178>
 39. Kenrick Kin, Maneesh Agrawala, and Tony DeRose. 2009. Determining the benefits of direct-touch, bimanual, and multifinger input on a multitouch workstation. In *Proceedings of Graphics Interface*. CIPS, Toronto, 119–124. DOI: <http://dx.doi.org/10.20380/GI2009.16>
 40. Tijmen Klein, Florimond Guéniat, Luc Pastur, Frédéric Vernier, and Tobias Isenberg. 2012. A design study of direct-touch interaction for exploratory 3D scientific visualization. *Computer Graphics Forum* 31, 3 (June 2012), 1225–1234. DOI: <http://dx.doi.org/10.1111/j.1467-8659.2012.03115.x>

41. Sebastian Knoedel and Martin Hachet. 2011. Multi-touch RST in 2D and 3D spaces: Studying the impact of directness on user performance. In *Proceedings of the IEEE Symposium on 3D User Interfaces*. IEEE Computer Society, Los Alamitos, 75–78. DOI: <http://dx.doi.org/10.1109/3DUI.2011.5759220>
42. Vincent Levesque, Louise Oram, Karon MacLean, Andy Cockburn, Nicholas D. Marchuk, Dan Johnson, J. Edward Colgate, and Michael A. Peshkin. 2011. Enhancing physicality in touch interaction with programmable friction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 2481–2490. DOI: <http://dx.doi.org/10.1145/1978942.1979306>
43. Hai-Ning Liang, Cary Williams, Myron Semegen, Wolfgang Stuerzlinger, and Pourang Irani. 2013. An investigation of suitable interactions for 3D manipulation of distant objects through a mobile device. *International Journal of Innovative Computing, Information and Control* 9, 12 (2013), 4737–4752. <http://www.ijicic.org/apchi12-291.pdf>
44. Aurélien Lucchi, Patrick Jermann, Guillaume Zufferey, and Pierre Dillenbourg. 2010. An empirical evaluation of touch and tangible interfaces for tabletop displays. In *Proceedings of the International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, New York, 177–184. DOI: <http://dx.doi.org/10.1145/1709886.1709917>
45. Paul Marshall, Rowanne Fleck, Amanda Harris, Jochen Rick, Eva Hornecker, Yvonne Rogers, Nicola Yuill, and Nick Sheep Dalton. 2009. Fighting for control: Children’s embodied interactions when using physical and digital representations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 2149–2152. DOI: <http://dx.doi.org/10.1145/1518701.1519027>
46. Anthony Martinet, Géry Casiez, and Laurent Grisoni. 2010. The effect of DOF separation in 3D manipulation tasks with multi-touch displays. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*. ACM, New York, 111–118. DOI: <http://dx.doi.org/10.1145/1889863.1889888>
47. Shawna Meyer, Oryx Cohen, and Erik Nilssen. 1994. Device comparisons for goal-directed drawing tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 251–252. DOI: <http://dx.doi.org/10.1145/259963.260468>
48. Tamara Munzner, Chris Johnson, Robert Moorhead, Hanspeter Pfister, Penny Rheingans, and Terry S. Yoo. 2006. NIH-NSF visualization research challenges report summary. *Computer Graphics and Applications, IEEE* 26, 2 (March 2006), 20–24. DOI: <http://dx.doi.org/10.1109/MCG.2006.44>
49. Jakob Nielsen. 1993. *Usability Engineering*. Morgan Kaufmann, San Francisco. DOI: <http://dx.doi.org/10.1016/B978-0-08-052029-2.50001-2>
50. Izabel C. Olson, Zeina Atrash Leong, Uri Wilensky, and Mike S. Horn. 2011. “It’s just a toolbar!” Using tangibles to help children manage conflict around a multi-touch tabletop. In *Proceedings of the International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, New York, 29–36. DOI: <http://dx.doi.org/10.1145/1935701.1935709>
51. Ivan Poupyrev and Shigeaki Maruyama. 2003. Tactile interfaces for small touch screens. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*. ACM, New York, 217–220. DOI: <http://dx.doi.org/10.1145/964696.964721>
52. Sara Price, Yvonne Rogers, Mike Scaife, Danae Stanton, and Helen Neale. 2003. Using ‘tangibles’ to promote novel forms of playful learning. *Interacting with Computers* 15, 2 (April 2003), 169–185. DOI: [http://dx.doi.org/10.1016/S0953-5438\(03\)00006-7](http://dx.doi.org/10.1016/S0953-5438(03)00006-7)
53. Mathieu Raynal, Guillaume Gauffre, Cédric Bach, Bénédicte Schmitt, and Emmanuel Dubois. 2010. Tactile camera vs. tangible camera: Taking advantage of small physical artefacts to navigate into large data collection. In *Proceedings of the Nordic Conference on Human-Computer Interaction*. ACM, New York, 373–382. DOI: <http://dx.doi.org/10.1145/1868914.1868958>
54. Farzan Sasangohar, I. Scott MacKenzie, and Stacey Scott. 2009. Evaluation of mouse and touch input for a tabletop display using Fitts’ reciprocal tapping task. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 53, 12 (Oct. 2009), 839–843. DOI: <http://dx.doi.org/10.1177/154193120905301216>
55. Jeff Sauro and James R. Lewis. 2010. Average task times in usability tests: What to report?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 2347–2350. DOI: <http://dx.doi.org/10.1145/1753326.1753679>
56. Dominik Schmidt, Florian Block, and Hans Gellersen. 2009. A comparison of direct and indirect multi-touch input for large surfaces. In *Proceedings of Human-Computer Interaction – INTERACT*. Springer, Berlin, Heidelberg, 582–594. DOI: http://dx.doi.org/10.1007/978-3-642-03655-2_65
57. Andrew Sears and Ben Shneiderman. 1991. High precision touchscreens: Design strategies and comparisons with a mouse. *International Journal of Man-Machine Studies* 34, 4 (1991), 593–613. DOI: [http://dx.doi.org/10.1016/0020-7373\(91\)90037-8](http://dx.doi.org/10.1016/0020-7373(91)90037-8)
58. Abigail J. Sellen, Gordon P. Kurtenbach, and William A. S. Buxton. 1992. The prevention of mode errors through sensory feedback. *Human Computer Interaction* 7, 2 (June 1992), 141–164. DOI: http://dx.doi.org/10.1207/s15327051hci0702_1
59. Orit Shaer and Eva Hornecker. 2010. Tangible user interfaces: Past, present, and future directions. *Foundations and Trends in Human-Computer Interaction*

- 3, 1–2 (2010), 1–137. DOI:
<http://dx.doi.org/10.1561/11000000026>
60. Christopher D. Shaw. 1998. Pain and fatigue in desktop VR: Initial results. In *Proceedings of Graphics Interface*. CIPS, Toronto, 185–192. DOI:
<http://dx.doi.org/10.20380/GI1998.23>
61. Ben Shneiderman. 1991. Touch screens now offer compelling uses. *IEEE Software* 8, 2 (March 1991), 93–94. DOI:<http://dx.doi.org/10.1109/52.73754>
62. Ken Shoemake. 1992. ARCBALL: A user interface for specifying three-dimensional orientation using a mouse. In *Proceedings of Graphics Interface*. CIPS, Toronto, 151–156. DOI:<http://dx.doi.org/10.20380/GI1992.18>
63. Adalberto L. Simeone and Hans Gellerseny. 2015. Comparing indirect and direct touch in a stereoscopic interaction task. In *Proceedings of the IEEE Symposium on 3D User Interfaces*. IEEE Computer Society, Los Alamitos, 105–108. DOI:
<http://dx.doi.org/10.1109/3DUI.2015.7131733>
64. Peng Song, Wooi Boon Goh, Chi-Wing Fu, Qiang Meng, and Pheng-Ann Heng. 2011. WYSIWYF: Exploring and annotating volume data with a tangible handheld device. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 1333–1342. DOI:
<http://dx.doi.org/10.1145/1978942.1979140>
65. Nicole Sultanum, Sowmya Somanath, Ehud Sharlin, and Mario Costa Sousa. 2011. “Point it, split it, peel it, view it”: Techniques for interactive reservoir visualization on tabletops. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*. ACM, New York, 192–201. DOI:
<http://dx.doi.org/10.1145/2076354.2076390>
66. Lucia Terrenghi, David Kirk, Abigail Sellen, and Shahram Izadi. 2007. Affordances for manipulation of physical versus digital media on interactive surfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 1157–1166. DOI:
<http://dx.doi.org/10.1145/1240624.1240799>
67. Philip Tuddenham, David Kirk, and Shahram Izadi. 2010. Graspables revisited: Multi-touch vs. tangible input for tabletop displays in acquisition and manipulation tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 2223–2232. DOI:
<http://dx.doi.org/10.1145/1753326.1753662>
68. Brygg Ullmer and Hiroshi Ishii. 2000. Emerging frameworks for tangible user interfaces. *IBM Systems Journal* 39, 3.4 (2000), 915–931. DOI:
<http://dx.doi.org/10.1147/sj.393.0915>
69. Dimitar Valkov, Frank Steinicke, Gerd Bruder, and Klaus Hinrichs. 2011. 2D touching of 3D stereoscopic objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 1353–1362. DOI:
<http://dx.doi.org/10.1145/1978942.1979142>
70. Gary R. VandenBos (Ed.). 2009. *Publication Manual of the American Psychological Association* (6th ed.). American Psychological Association, Washington, DC.
<http://www.apastyle.org/manual/>
71. Vanessa Vuibert, Wolfgang Stuerzlinger, and Jeremy R Cooperstock. 2015. Evaluation of docking task performance using mid-air interaction techniques. In *Proceedings of the ACM Symposium on Spatial User Interaction*. ACM, New York, 44–52. DOI:
<http://dx.doi.org/10.1145/2788940.2788950>
72. Sijie Wang. 2010. *Comparing tangible and multi-touch interfaces for a spatial problem solving task*. Master’s thesis. Simon Fraser University, Canada.
<http://summit.sfu.ca/item/11500>
73. Diane Watson, Mark Hancock, Regan L. Mandryk, and Max Birk. 2013. Deconstructing the touch experience. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*. ACM, New York, 199–208. DOI:
<http://dx.doi.org/10.1145/2512349.2512819>
74. Daniel Wigdor and Dennis Wixon. 2011. *Brave NUI World: Designing Natural User Interfaces for Touch and Gesture*. Morgan Kaufmann Publishers Inc., San Francisco. DOI:
<http://dx.doi.org/10.1016/B978-0-12-382231-4.00037-X>
75. Lesley Xie, Alissa N. Antle, and Nima Motamedi. 2008. Are tangibles more fun? Comparing children’s enjoyment and engagement using physical, graphical and tangible user interfaces. In *Proceedings of the International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, New York, 191–198. DOI:
<http://dx.doi.org/10.1145/1347390.1347433>
76. Lingyun Yu, Pjotr Svetachov, Petra Isenberg, Maarten H. Everts, and Tobias Isenberg. 2010. FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (Nov./Dec. 2010), 1613–1622. DOI:
<http://dx.doi.org/10.1109/TVCG.2010.157>
77. Shumin Zhai and Paul Milgram. 1998. Quantifying coordination in multiple DOF movement and its application to evaluating 6 DOF input devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, 320–327. DOI:<http://dx.doi.org/10.1145/274644.274689>