Still room for improvement in traditional 3D interaction: selecting the fixed axis in the virtual trackball

Daniel Gonzalez-Toledo¹ · Maria Cuevas-Rodriguez¹ · Luis Molina-Tanco¹ · Arcadio Reyes-Lecuona¹

Accepted: 25 December 2021 / Published online: 29 January 2022
© The Author(s) 2022

Abstract

Virtual trackball techniques are widely used when 3D interaction is performed through interfaces with a reduced number of degrees of freedom such as mice and touchscreens. For decades, most implementations fix a vertical axis of rotation, which is a suitable choice when the vertical axis should indeed be fixed, according to some mental model of the user. We conducted an experiment involving the use of a mouse and a touch device to study usability in terms of performance, perceived usability and mental workload when selecting different fixed axes in accordance with the user's mental model. The results we obtained indicate that the consistency between the axis fixed by the technique and the object's intrinsic axis has a positive effect on usability. We believe that implementations that allow to select different fixed axis for each specific object should be considered when designing future reduced-DoF interaction interfaces.

Keywords Usability · 3D interaction · 3D User interfaces · Virtual trackball

1 Introduction

A very common 3D interaction task in several application fields (such as modelling, computer aided design, medical diagnosis or simple browsing of 3D object repositories) is the inspection of a virtual object. Controlling a virtual camera to navigate a 3D scene is, in general, not a trivial problem [26]. A number of interaction techniques have been proposed to map users' actions into the commands needed for navigation, using both standard and specific input devices, tangible props or gesture recognition [4, 19]. In the context of desktop or web-based applications, the use of specific devices for interaction is limited to keyboard and mouse, which provide limited degrees of freedom (DoF). A similar situation is given in applications within portable devices, where interaction is

usually performed through a touch screen. In this paper, we focus on the inspection of 3D objects using standard devices, still widely and commonly used, which constrain the user input to two DoF, provided by mouse or finger gestures over the screen as horizontal (Δx) and vertical (Δy) gestures. Most modern implementations of these interaction techniques choose to fix a vertical axis of reference, which is consistent with the users' mental model for scenes where objects are placed on the ground, such as buildings (Fig. 1a), but not necessarily so when objects are presented as suspended in an empty space, such as mechanical parts (Fig. 1b).

The virtual sphere or virtual trackball was proposed by Chen et al. in an early study [9], for the 3D rotation of objects using 2D devices. Their proposal consisted of simulating a virtual sphere that can be rotated around any arbitrary axis. The 3D object or scene to be rotated could be seen as encased in this sphere. To rotate the object, the user has to roll the sphere by clicking on the surface of the sphere and moving the mouse in the desired direction. Therefore, the axis of rotation is perpendicular to the direction of movement on the surface of the sphere. In the original paper, the virtual sphere was drawn on the screen as a circle and, when the mouse moves outside it, the object is rotated around the z axis (normal to the screen). Alternatively, on touch devices (where the input is finger gestures), such techniques can be easily adapted

✉ Arcadio Reyes-Lecuona
areyes@uma.es

Daniel Gonzalez-Toledo
dgonzalez@uma.es

Maria Cuevas-Rodriguez
mariaquevas@uma.es

Luis Molina-Tanco
lmtanco@uma.es

¹ Departamento de Tecnología Electrónica, Universidad de Málaga, 29071 Málaga, Spain

by treating a touch input as a mouse input (e.g. [3, 23, 42], and [13]). The virtual sphere and the subsequent variations [37, 38] were especially suitable for manipulating objects suspended in an empty space without a ground reference. In fact, the trackball metaphor suggests that user actions make the object rotate rather than making the camera orbit around the object. All are variations of the same idea—the Virtual Trackball—and all are able to rotate the object around three axes.

The Two-Axis Valuator Technique (TAV) is a simpler technique that predates the Virtual Trackball. The TAV directly maps user gestures in X- and Y-coordinates into rotations of the object around vertical and horizontal axes [40]. Despite having more than forty years of age, a variation of the TAV technique is still used in many applications and is one of the preferred techniques nowadays. This variation is known as the TAV with fixed Up-vector [2], or just Fixed Trackball [32]. The TAV with fixed Up-vector fixes a world vertical reference rotation vector, which is consistent with scenes where objects are placed on a horizontal surface. This technique is used in 2021 in 3D modelling tools, such as Blender [5], 3D Max [1] and Sketchup [18] or 3D object stores on the web like 3D Warehouse [30] or Sketchfab [39].

Having the vertical axis as the fixed reference is appropriate when there is a ground, which is either explicitly displayed or not (Fig. 1a). However, there are objects with an intrinsic rotation axis that is not vertical. In the users' mental model, the representations of a wheel or a turbine may have a horizontal rotation axis (Fig. 1b). A representation of the Earth globe may have a diagonal rotation axis representing the tilting of the Earth's rotation away from the direction perpendicular to the orbit around the Sun. Designing an interface that assumes that these axes are fixed could improve interaction, following previous works on mental models and interfaces [29]. A generic implementation should allow for fixing any axis. In this paper, we advance in the direction of this generalization by investigating the importance of selecting which axis to fix depending on the intrinsic characteristics

of the object to be inspected for objects with both vertical and horizontal intrinsic rotation objects.

González-Toledo et al. [16] introduced a tool able to fix either the Up-vector or Right-vector, but no study was made into the importance of this selection. To the best of our knowledge, very few studies have focused on the Two-Axis Valuator with a fixed axis. To the best of our knowledge, only Bade et al. [2] compared the fixed Up-vector with different techniques, but not with a Right-vector. Buda [8] described the two possible fixed axes (Up and Right) and compared them without finding any major differences. However, none of these studies considered any intrinsic characteristics of the object. In this paper, we investigate the importance of selecting which axis to fix, depending on the user's mental model about the intrinsic characteristics of the object to be inspected.

This study was originally motivated by the authors' experience in a research project which involved industries dealing with high investment, low volume products with relatively long service lives. One of the use cases in the project involved creating a 3D collaborative tool to support decision making processes in maintenance inspection of energy production turbines. Some of the tool requirements were selection and manipulation of parts at different levels of the product hierarchy; adaptive transparency and exploded views for visualization of occluded parts; visual representation and spatial location of inspection data; and basic navigation around the product using the mouse. The first stages of the project required modelling the turbine, and to evaluate the early models we used standard visualization tools. Our impression was that, in these tools, navigation around the model of the turbine was awkward, unnatural. In the 3D collaborative tool that we subsequently created, we fixed a horizontal rotation axis for consistency with our mental model of the turbine. This decision was praised by the maintenance inspection operators testing the tool and motivated us to investigate if this result could be generalised.

Fig. 1 3D objects with different reference axes [16]. **(a)** A reconstruction of the Parthenon with a vertical reference axis. **(b)** A gas turbine with a horizontal reference axis

2 Hypothesis

When virtual objects are presented as suspended in an empty space, the natural rotating axis of the object depends on what the virtual model is representing. The presented study aims to demonstrate that having the object's intrinsic axis consistent with the axis fixed in the interaction technique will improve interaction since it will match the user's mental model and, therefore, will be more natural. In the same way, when there is an inconsistency between both axes, the interaction, and therefore the user experience, will be hindered. In addition, this should happen regardless of the device we use (mouse or touch screen) since it is related with the consistency between the user's mental model and the interaction technique, not with the specific characteristics of the interaction devices.

In the light of all the above, our hypothesis is the following: having consistency between both axes, either vertical or horizontal, will benefit the overall usability of the system, which will be reflected in the following ways:

1. it will lead to higher performance,
2. it will lead to subjects perceiving increased usability, and
3. it will lead to a reduction in the workload of the subject when using the system.

3 Material and methods

3.1 Experiment overview

The participants in the experiment were asked to use a piece of software to complete a series of inspection tasks. The virtual object inspected was always a textured sphere with equidistant holes on its surface (Fig. 2). The users had to interact with the sphere in order to find a small target hidden in one of the holes. The texture on the sphere represents meridians and parallels, so that an intrinsic rotation axis is suggested to the participant. Rather than choosing objects of different shapes, we deliberately chose the spheres to avoid suggesting any intrinsic rotation axis through the geometry of the object—only through texture.

Fig. 2 (a) 3D model of the textured sphere with twenty holes used in the inspection tasks. (b) Vertices of a dodecahedron inscribed in the sphere. The vertices of the dodecahedron are used to evenly distribute holes on the sphere

To interact with the sphere, the software implemented a Two-Axis Valuator technique that allows fixing either a vertical, up-vector or a horizontal, right-vector, independently of the axis suggested by the meridian and parallel texture. The details of both variants of the technique are illustrated in the Appendix at the end of this paper. The participants solved the inspection task with different combinations of suggested intrinsic rotation axis of the object, and rotation axis fixed by the interaction technique. The time to complete the inspection tasks was measured by the software. Each combination was used several times in a row, after which participants had to fill-in questionnaires on perceived usability, workload and fatigue.

3.2 Inspection task design

Each participant in the experiment had to complete an inspection task a number of times. The object inspected was a sphere that presented 20 cylindrical holes on its surface, as shown in Fig. 2a. The holes were distributed evenly on the sphere, in alignment with the vertices of a dodecahedron inscribed in the sphere, as in Fig. 2b. The goal of the inspection task was to find small target—a yellow sphere—appearing inside only one of the 20 holes (see the enlarged detail in Fig. 2a). In order to do so, participants had to rotate the sphere, looking for the target, and place it at the centre of the screen inside a circular viewfinder, which turned yellow when the target was inside. After one second, the target disappeared, a 'success' sound was played back through the speakers marking the completion of one trial, and the viewfinder returned to its initial colour. Then, a new trial started, presenting the participant with a new target to find. This was repeated 20 times in such a way that the target appeared once in every available hole. During this phase, in the upper right corner of the screen a countdown counter was displayed showing the number of targets remaining. Full details on the arrangement of targets are given in Sect. 3.6 *Experimental Apparatus*.

3.3 Experimental procedure

The experiment was designed with four within-subject conditions, which account for the vertical and horizontal orientations of both the technique's fixed rotation axis (technique reference axis) and the object's intrinsic rotation axis (object's intrinsic axis). This defines two within-subject independent variables: the first one is the consistency between the object's intrinsic axis and the technique reference axis, and the second independent variable is the technique axis orientation. These combine into the four conditions of Table 1, illustrated in Fig. 3. We refer to Conditions A and D as *consistent* because the technique reference axis and the object's intrinsic axis are in line. Conditions B and C are *inconsistent* since the technique reference axis and the object's intrinsic axis do not match. In A and B, the technique reference axis is *vertical*, while in C and D it is *horizontal*.

Every participant performed the whole study with the four conditions. Half of the participants used a tablet, while the other half used a notebook computer with a mouse. The application was designed to prevent differences in performance between both devices (full details are given in Sect. 3.6 on the experimental apparatus). Therefore, the used platform is a between-subject independent variable.

At the beginning of the experiment, the participants were assigned to one between-subject condition (mouse or tablet) and given a set of instructions about the procedure and how the interaction is performed with the device. No explanations or demonstrations of the rotation technique were presented

Fig. 3 Sphere representation for the four experimental conditions. The rows indicate the object's intrinsic axis. The columns show the technique reference axis, which is also indicated with a white axis that crosses the sphere. A and D are consistent, and B and C are inconsistent

Fig. 4 Experiment procedure. (a) shows the complete procedure of the experiment. (b) shows the procedure of a block. The conditions and sequence order for participant #3 are shown as an example

Table 1 Technique reference axis and the object's intrinsic axes for the four experimental conditions

Conditions	Technique reference axis	Object's intrinsic axis
A	Vertical	Vertical
B	Vertical	Horizontal
C	Horizontal	Vertical
D	Horizontal	Horizontal

to the participants. The subjects were told that they were going to try four variants of a technique for manipulating 3D objects. They were also told that each time the target would roughly come out of the antipodes of the previous point.

Figure 4a outlines the structure of the whole experiment for each participant. Once the informed consent and the questionnaires are filled in, the participants had to complete four blocks, one for each within-subject condition. The condition of each block (Table 1: A, B, C or D) is selected according to a counterbalanced order using Latin squares. Finally, once these four blocks were completed, they were asked to fill out the fatigue questionnaire a second time.

The procedure of a block is illustrated in Fig. 4b. Each block started with a training phase, where participants could practice the inspection task as long as they needed in order to become familiar with the rotation technique. During this phase, the word 'Training' was displayed on the screen and the countdown counter, showing the number of remaining targets, was not shown. Once the participants felt that they were ready, they click or touch a start button located in the lower right corner of the screen. The scenario shown in Fig. 2a is loaded, and the 20 inspection trials described in Sect. 3.2

begin. The software internally balances the location of the hidden target for each trial, following a strategy described later in Sect. 3.6. The participants were asked to complete each trial as quickly as possible. At the end of the block and after finishing the 20 trials, participants filled in usability and workload questionnaires regarding the rotation technique used in this block.

3.4 Data collection and analysis

The experiment goals were to measure the task performance, perceived usability, perceived workload and perceived fatigue and compare it under the different experimental conditions illustrated in Fig. 3.

For each trial, we recorded the time that participants took to align the target inside the circular viewfinder. For each participant and each condition, we computed three average times: one for the 10 first trials, one for the last 10 trial and the overall average for the 20 trials. For this, we compute the geometric means of the trial times by applying a logarithmic transformation to the time measurements, computing the arithmetic mean and then exponentiating back to get the geometric mean. This is done to correct the positive asymmetry due to some users taking unusually excessively long times to complete the trials [33].

Perceived usability is measured with the ten-item SUS questionnaire [7], which was translated into Spanish by Devin [14]. The subjects scored each question using a 5-point Likert scale ranging from (strongly disagree) to (strongly agree). The answers to these ten items were assessed to obtain an SUS score [34] which summarizes the subject's perception of the usability. In addition, as described by Lewis and Sauro [25] and Borsci et al. [6], the SUS score was divided in two factors or subscales, *learnability* and *usability*.

The workload is measured through the Raw TLX (RTLX) questionnaire, a modification of the NASA Task Load Index (NASA-TLX) that has been found to be comparatively sensitive while being much simpler [20]. We used the Spanish translation by Arquer and Nogareda [12]. The questionnaire consists of six questions that are evaluated using a 21-point scale ranging from very low to very high. These data were averaged in order to obtain an RTLX score for each participant and condition. Along with this summary score, the average of each of the six sub-scales that compose the TLX has also been obtained for each condition and each participant. The sub-scales are Mental, Physical and Temporal Demand, Frustration, Effort and Performance.

Finally, to know whether the different platforms caused similar fatigue and if they were therefore comparable, a fatigue questionnaire based on Shaw [36] was filled out. Participants evaluated their fatigue using a 10-point Likert scale (from no fatigue to extreme fatigue) of six body areas (eyes, neck, hands, wrists, arms and shoulders).

We used estimation techniques based on the effect size and confidence intervals (CI) to analyse data and to report the results in detail, as recommended by Cumming [11] and APA [41], among others. In addition, a mixed design ANOVA, with one between-subject factor (the platform) and two within-subject factors (consistency and technique reference axis), was carried out. These analyses have been done for the task completion time, the usability, and the workload. In addition, an independent sample T-test was carried out to analyse the fatigue caused by the platform.

3.5 Participants

All participants read and signed an informed consent document before the experiment was conducted. They were subsequently asked to complete a demographic questionnaire, indicating their gender, age, and experience with 3D object manipulations. A second questionnaire was used to determine their fatigue level at the beginning of the experiment. 32 subjects, including 19 males and 13 females, participated in the experiment, with ages between 18 and 59 ($M = 28.8$, $SD = 7.9$). 20 participants were 18–29 years old, 9 were 30–39, two were 40–49, and one was 50–59. None of the participants reported any uncorrected visual acuities or impairments. No monetary incentive was offered. No experience in the use of 3D modelling software or similar software was required. However, all participants were very familiar with the use of mouse and touchscreens, and 60% of them stated that they had some or a lot of experience using 3D modelling software. All procedures were reviewed and approved by the Ethical Committee for Experimentation of the University of Malaga.

3.6 Experimental apparatus

The experiment was set up in a closed room. The laptop computer was an Asus K541UJ-GO319T running Windows 10, and the tablet was a Samsung Galaxy Tab S4 running Android 9.

The software used in the experiment was developed in Unity3D to integrate HOM3R, a 3D viewer for complex hierarchical product models [17]. The software was run on a full screen on a 15.6-inch display with a resolution of 1366x768 pixels for the PC and on a 10.50-inch touchscreen display with a resolution of 2560x1600 pixels for the tablet. The control-to-display ratio was chosen according to the device resolution in order to rotate the 3D model 180° when the mouse or the fingers are moved from side to side on the screen. This makes the application resolution independent. The application is non-cpu- nor gpu-intensive, to ensure no performance differences between the devices. The computer was operated with a mouse and the tablet with fingers. Both the computer and the mouse rested on a desk. Participants

Table 2 The four dodecahedron vertex sequences used in the experiment. The vertex IDs are shown in Fig. 2b. For each block, the targets are hidden in holes corresponding to one of this vertex sequences

Sequence Id	Vertices sequence
1	16, 05, 09, 13, 02, 18, 06, 12, 08, 14, 10, 20, 04, 17, 01, 19, 03, 15, 11, 07
2	13, 07, 11, 15, 03, 19, 01, 17, 04, 20, 02, 18, 06, 10, 14, 08, 12, 16, 05, 09
3	18, 02, 20, 04, 07, 13, 09, 05, 16, 03, 15, 11, 01, 17, 14, 10, 06, 12, 08, 19
4	20, 02, 18, 06, 12, 16, 05, 09, 13, 07, 04, 17, 01, 11, 15, 03, 19, 08, 14, 10

using the tablet were able to decide whether to rest it on the desk or to hold it on their lap, but then they were asked to keep the tablet in the same position during the whole experiment. For the mouse, the x-/y-axis rotation controller is operated by holding down the left mouse button and tracking the mouse movement. For the touchscreen, one-fingered gestures control the rotation for the x- and y-axis.

The viewfinder and the target's radius were, respectively, 15% and 5.6% of the sphere's radius. Targets were arranged across subsequent trials to follow a sequence that ensures that the next target to find is always located on the hidden side of the sphere. In this way, we ensured that the participant always had to rotate the sphere to reach the next target. For this, we pre-computed several sequences that place each target in one of the holes around the antipodal point from the previous one via an offline algorithm that generates sequences of vertices of a dodecahedron inscribed in the sphere. Out of these sequences, we selected four in which the first target is hidden (Table 2). This is a 3D extension of the ring of targets proposed by ISO9241-9 [22], which ensures that the distance between consecutive targets is always the same. During each block, one of the four sequences (Table 2: 1, 2, 3 or 4) is used to set the order of the target's appearance. The order in which the sequences are selected is counterbalanced using Latin squares.

4 Results

The experiment goals were to measure the task performance, perceived usability, perceived workload and perceived fatigue and compare it under the four experimental conditions illustrated in Fig. 3. We report the results in detail in the following sections, where we consider three factors: consistent vs. inconsistent conditions is the main factor, but we also wanted to examine if there were differences between a fixed vertical axis, which is commonly implemented by software tools that the participants might have used before, vs. a fixed horizontal axis. Finally, we also examine if there were differences between participants that used the mouse vs. those that used finger gestures on the tablet.

4.1 Task performance

To estimate the performance, the software measured the time that participants took complete the task, i.e. to align the tar-

get inside the circular viewfinder, as detailed in the previous section.

We found a significant effect of consistency in task performance ($F(1, 30) = 13.077$, $p = 0.001$). For consistent configurations, i.e. those in which the fixed axes and the object's intrinsic axis coincide, the mean¹ was $7.24^{+0.51}_{-0.48}$ s, while for the inconsistent configurations the mean was $7.9^{+0.7}_{-0.64}$ s. The estimated difference between non-consistent and consistent was $0.67^{+0.45}_{-0.41}$ s.

The axis fixed by the interaction technique has a marginally significant effect ($F(1, 30) = 3.681$, $p = 0.065$). When the vertical axis was fixed, the participants took $7.4^{+0.6}_{-0.57}$ s. When the horizontal axis was fixed, they took $7.73^{+0.62}_{-0.57}$ s. The estimated difference was $0.34^{+0.4}_{-0.37}$ s.

Finally, we could not measure any effect of the platform, i.e. mouse v. finger gestures on the touchscreen ($F(1, 30) = 2.33$, $p = 0.137$). Participants using the mouse took $7.2^{+0.97}_{-0.86}$ s, while participants using the touchscreen took $7.95^{+0.7}_{-0.65}$ s. The estimated difference is $0.75^{+1.2}_{-1.07}$ s, a large CI including the non-difference as a plausible case.

Figure 5a plots the average task completion time of each experimental condition, suggesting that the consistent conditions, A and D, involve lower completion times. Additionally, when conditions are consistent, the vertical axis condition seems to have better performance than the horizontal one. However, when conditions are inconsistent, the times seem to be similar for both orientation axes.

4.1.1 Interaction between factors for task performance

We did not find a significant interaction effect between the platform and the within-subjects factors (The results of the ANOVA are $F(1, 30) = 2.076$, $p = 0.16$ for the interaction between the platform and the consistency, $F(1, 30) = 0.359$, $p = 0.554$ for the interaction between the axis and the platform, and $F(1, 30) = 1.17$, $p = 0.28$ for the interaction between the consistency and the axis. Thus, we consider that these interaction effects can be neglected and will base our analysis on the main effects.

¹ It is conventional to use the superscript–subscript notation for plus or minus one standard deviation. Instead, we use it across this section to report 95% confidence intervals on the measured times [10].

Fig. 5 Estimated means and CIs for the task completion time separated by each experimental condition. These are consistent vertical, consistent horizontal, inconsistent vertical and inconsistent horizontal. The conditions of the same orientation axis have been joined with lines. To

the left, in (a), we show the results for all the twenty trials. To the right, in (b), we show the results separately for the first ten and the last ten trials to illustrate possible learning effects (see text). All results use $n=32$. All error bars are 95% CIs

Fig. 6 Lineal regression model calculated to analyse the learning effects

4.1.2 Analysis of learning effects on task performance

We analysed the possible learning effects along the experiment using two strategies. We observed the overall task performance of subjects as a function of the number of trials they accumulated. Also, we divided the twenty trials into two halves, analysing the first ten trials and the last ten separately for the factors for which we have observed some effect: consistency and axis.

We computed a lineal regression model which attempts to predict the performance time as a linear function of the number of trials. All the trial completion times were used as samples for this analysis and are represented in Fig. 6. The graph shows descending time along the trials. The effect of the number of trials on the completion time was found to be significant with $F(1, 2558) = 31.0276$, $p < 0.001$. Nevertheless, the effect of learning is likely small since only 1.2% of the time variance was explained by the number of trials (R^2 of 0.012).

We then proceeded to analyse the factors for the first ten trials and the second ten trials separately. Consistent conditions seem to help participants to learn faster: In the first ten

trials, for consistent conditions the average is only slightly shorter: $7.56^{+0.54}_{-0.51}$ s vs. $8.08^{+0.7}_{-0.64}$ s for the inconsistent ones, an estimated difference of $0.52^{+0.56}_{-0.52}$ s. However, in the last ten trials, we measured a stronger effect of the consistency: $6.92^{+0.50}_{-0.49}$ s, vs. $7.73^{+0.55}_{-0.51}$ for the inconsistent configurations, an estimated difference of $0.81^{+0.45}_{-0.42}$ s. This is a 55.8% larger difference for the last trials, with a smaller CI.

Figure 5b displays the average task completion times for each condition, in combination with the axis fixed by the technique. Consistent conditions show an average performance improvement of 8.45% while inconsistent conditions show only a 4.2% improvement. We observed a small interaction between the consistency and the orientation axis only when the first ten trials are considered; however, the ANOVA showed no significant differences, yielding $F(1, 30) = 1.748$, $p = 0.196$.

We observed a likely small difference in learning depending on the orientation axis. At the first ten trials, the performance time was very similar: $7.68^{+0.76}_{-0.69}$ s for fixed vertical axis vs. $7.95^{+0.74}_{-0.68}$ s for fixed horizontal, a difference of $0.27^{+0.59}_{-0.54}$ s which does not suggest an advantage for any of

Fig. 7 Estimated means and CIs for the SUS scores according to each experimental condition. These are the following: consistent vertical, consistent horizontal, inconsistent vertical and inconsistent horizontal. Conditions with the same axis have been joined by lines. All error bars are 95% CIs

the axes. For the second half of the trials, the mean for fixed vertical axis was $7.12^{+0.63}_{-0.58}$, while the mean for trials with a fixed horizontal axis was $7.52^{+0.69}_{-0.63}$ s, an estimated difference of $0.4^{+0.49}_{-0.45}$ s.

4.2 Perceived usability

The usability of the different conditions perceived by the participants was measured with the ten-item SUS questionnaire, as detailed in Sect. 3.4.

We measured a significant effect on usability again in the within-subject consistency factor ($F(1, 30) = 29.183$, $p < 0.001$)—the analysis provides good evidence that consistency outperforms inconsistency in terms of participant usability scores: On average, consistent conditions scored $72^{+5}_{-5.1}$ points while inconsistent conditions scored $53^{+6.1}_{-6.2}$ points, an estimated difference of $18.98^{+7.22}_{-7.18}$, a large CI but far from zero.

Regarding the second within-subject factor, there is also good evidence that the fixed vertical axis scores better at perceived usability than the fixed horizontal axis. In the case of the vertical axis, participants scored an average of $69^{+5.2}_{-6.1}$ points, while for horizontal one, they scored an average of $56^{+6.3}_{-5.7}$ points. The estimated difference in this case was $12.19^{+7.78}_{-7.79}$ points. The ANOVA signals a significant effect ($F(1, 30) = 10.232$, $p = 0.003$).

Similarly to the analysis of task performance, we could not measure any effect of the platform—mouse vs. touchscreen, between-subjects—on perceived usability measured by the SUS score. The ANOVA resulted in $F(1, 30) = 0.164$, $p = 0.688$.

Fig. 8 Estimated means and CIs for the workload RTLX score questionnaires, separated by each experimental condition. These are consistent vertical, consistent horizontal, inconsistent vertical and inconsistent horizontal. The conditions with the same axis have been joined by lines. All error bars are 95% CIs

4.2.1 Interaction between factors for perceived usability

Studying the four conditions separately, we can see that condition A (vertical consistent) yielded higher scores than all the others while condition C (horizontal inconsistent) produced the lowest scores (see Fig. 7). We could not find a significant interaction effect between consistency an axis: the result of the ANOVA was $F(1, 30) = 2.494$, $p = 0.125$. Neither could we measure any interaction between the platform and the within-subject factors. We computed $F(1, 30) = 0.084$, $p = 0.775$ for the interaction between the consistency and the platform, and $F(1, 30) = 0.203$, $p = 0.655$ for the interaction between the axis and the platform.

4.3 Perceived workload

We measured workload through the Raw TLX questionnaires, as detailed in Sect. 3.4.

Both consistency and axis have a significant effect on workload perceived by the participants. For consistent conditions, participants reported an average workload of $37^{+6.2}_{-5.6}$ points (note that less points is better), but $45^{+6.7}_{-5.9}$ for inconsistent ones—an estimated difference of $8.10^{+8.5}_{-4.5}$. The ANOVA also supports an effect of consistency since $F(1, 30) = 13.61$, $p = 0.001$. As for the axis orientation, the average workload score for the trials with vertical fixed axis was $38^{+5.5}_{-6.1}$, and a larger $45^{+6.1}_{-5.1}$ for the horizontal ones. The estimated difference was $7.31^{+3.77}_{-3.78}$. Consistent with those results, the ANOVA outcome is $F(1, 30) = 15.63$, $p < 0.001$. These results seem indicate that fixing the vertical axis, rather than the horizontal, has a positive effect on the workload perceived by the participants. Once again we did not find an effect of the platform—mouse or touchscreen.

Fig. 9 Fatigue reported by participants divided into the two groups: those who used the mouse and those who used the touchscreen. (a) overall results by platform. (b) results separated by platform and body part. All error bars are 95% CIs, and $n = 16$

4.3.1 Interaction between factors in workload perception

Figure 8 shows the detailed results when all the experimental conditions are analysed separately. In case A (consistent-vertical), the average RTLX score is relatively different than those in the other cases, but the confidence intervals are similar for all four conditions. For Case A (consistent-vertical), we measured the lowest RTLX score, while the other conditions yielded scores that suggest a similar effect. However, we could not measure a significant interaction effect between these factors—consistency and axis. The ANOVA resulted in $F(1, 30) = 2.343$, $p = 0.136$.

We were not able to find a significant interaction between the platform and the within-subjects factors. The ANOVA resulted in $F(1, 30) = 0.366$, $p = 0.55$ for the interaction between the consistency and the platform and $F(1, 30) = 0.122$, $p = 0.73$ for the interaction between the axis and the platform.

4.4 Fatigue results

The level of fatigue reported was in general very low (always less or around 1 on a scale from 0 to 10). We could not detect any significant difference between both platforms: $0.38^{+0.23}_{-0.32}$, and a T-test result of $t(30) = 0.963$, $p = 0.343$.

Figure 9a shows the results of the fatigue questionnaires for both platforms. These are averaged results for different parts of the body. If we look at the results for the different parts of the body separately Fig. 9b. The hands and the wrists are the ones that most contribute to mouse fatigue, while, for the touchscreen, the neck is the most relevant. These results do fit in with the participants' behaviour when using the tablet. Since we did not impose any restrictions, most of them decided to leave the tablet horizontal on the table without holding it. It seems plausible that this causes more fatigue in the neck due to its angled posture when looking down than it would cause in the hands, wrist or shoulders.

5 Discussion and conclusion

5.1 Discussion of results

5.1.1 On the importance of consistency between rotation axis fixed in the interaction and intrinsic axis of an inspected object

We have found evidence that the consistency between the axis fixed by the interaction technique and the inspected object's intrinsic rotation axis positively influences usability in terms of the task performance, perceived usability, and mental workload. In consistent situations, participants took on average less time to perform the tasks, rated the system as more usable, and indicated lower workloads. All this would support our hypothesis. As we mentioned before, a previous study by Buda [8] compared both the TAV with fixed Up-vector and that with fixed Right-vector, obtaining that the first outperformed the second. However, his experiments were performed using objects with an intrinsic vertical axis, like a head. This is actually in accordance with our results since those objects are consistent only with the TAV with fixed Up-vector.

Considering our study in the framework of the schema theory proposed by Schmidt [35], the proposed inspection task requires a combination of different discrete tasks, like a finger slide on the touchscreen. When the user wants to rotate the object in a specific direction, an action is performed, which produces a feedback that is compared with the expected behaviour of the system. This allows the user to correct the programmed action. If the rotation produced by the user's action is similar to the one expected, less corrections take place and the interaction is more efficient. We interpret our findings in this framework. Consistent conditions produce more expected feedback and, therefore, higher performance, better perceived usability and a lower mental workload. In other words, when the fixed axis is consistent with the intrinsic rotation axis of the object, the technique better fits the user's mental model, and interaction is more fluid.

However, this represents only a small portion of the whole interaction and, hence, the differences in the task completion times are relatively small, but enough to reveal the effect. Nevertheless, these differences are large in absolute terms for usability and the workload perceived by participants, which supports our interpretation.

This problem is also similar to the classical Stimulus-Response Compatibility effect [31] and the more general theory of event coding [21]. These theories are proposed to explain the response time for a stimulus, depending on whether certain characteristics of the stimulus are compatible or not with the response. These theories assume that both the stimulus and response share a common representational structure in the cognitive system, thus explaining how action may influence the perception of the stimulus. Although our task has a different structure, similar principles can be used here to explain our results. In this sense, some authors have studied spatial compatibility effects, which are different from the relationship between the stimulus location and response location, such as when manipulating tools. For instance, Kunde et al. [24] studied task performance when hand movements are transformed into direct or inverse movements of a tool, showing that responding was delayed when the hand and the tool moved in non-corresponding rather than corresponding directions. The theory tells us that users execute a movement by activating the anticipatory codes of the movement's sensory effects [28]. Our results can be explained in the framework of these models. When the object has an intrinsic rotation axis, and the interaction technique fixes that axis, the action codes activated to navigate around the object and the perception of that object share a common cognitive representation, which is consistent. Therefore, performance and user experience is improved.

5.1.2 On the difference between fixing the vertical and fixing the horizontal axis

Regarding the orientation axis factor, we have observed that the vertical axis has also a positive effect on usability. We have found an effect on all three measured parameters, although in different ways. We measured a greater effect in the subjective parameters, the SUS and RTLX questionnaires, than in the task completion time where we have not been able to measure a clear difference. We conjecture that participants find the vertical axis more comfortable, either because they are more used to it, or because it is more natural to them. Indeed, navigation systems with a fixed vertical axis already exist and are widely used in 3D modelling applications or other applications offering a 3D view of objects or scenes. The use of a fixed horizontal axis, following the intrinsic axis of the object, is a novel idea.

When the two factors are analysed together, we observed that consistency has a greater effect on the three measured

parameters than the orientation axis. The consistent condition with the vertical axis obtained the best results, followed by those consistent with a horizontal axis (see Figs. 5a, 7 and 8). This would be also explained in the light of the schema theory since the motor program taken in action would benefit from previous experience and would therefore be more adjusted to the desired behaviour and require fewer corrections.

5.1.3 On the independence of results with respect to the interaction platform

We studied the differences between the two platforms (mouse and touchscreen), finding no clear differences between them. The results obtained do not show relevant differences in terms of performance, usability or workload for our task. More importantly, the ANOVA did not signal any interaction effect between the platform and the other two factors for any of the three parameters evaluated. When studying the fatigue perceived by the participants, we could not find any difference between the platforms either. We did find differences in how different parts of the body contribute to perceived overall fatigue, but we did not find any important differences in the overall fatigue. For the touchscreen—a tablet device, the contribution to the overall fatigue of body parts differs from that of previous studies [3]. This might be justified by how participants have used this platform since they were not given specific instructions on how to hold the device.

This analysis allows us to assume that, in the context of this experiment, both paradigms, mouse and touchscreen, are equivalent in terms of 3D interaction. Consistently with this result, Martinez et al. [27] show that, when interaction techniques are well adapted, the particular presentation device will have a very small effect in usability. In contrast, previous work has reported small differences in the time needed to perform simpler tasks, such as selection or 2D docking [15]. Since our inspection task could be regarded as a type of 3D docking, it involves predicting the 3D rotation movement that takes place in the sphere with respect to the movement of the device. We hypothesize that this task involves higher level cognitive skills that seem to be shared by both paradigms and that the low level motor task of controlling a mouse and interacting with the touch screen with one finger are less important. For this reason, we have integrated the results of the two platforms to test our hypotheses.

5.1.4 On the learning effects along the experiment

We carried out analysis of the possible learning effects on our results. Participants reported that they found it to be a little difficult each time they started a new condition, and that they gradually became more comfortable with it. This happened even though they had been told that they could stay in the training phase for as long as they needed until

they had mastered the rotation technique. Using a linear regression model, we have found that there is indeed a small learning effect throughout each of the conditions (see Fig. 10). Observing Fig. 5b, two aspects of this learning effect can be highlighted. The first is that learning is more important in consistent situations. Secondly, it seems that there is a small interaction, although it is not significant in the ANOVA, between the technique axis and consistency when the participants are inexperienced in the first half of each block. Nevertheless, this clearly disappears in the second half since experience is gained. It seems that at the beginning of the block, only the condition with consistency and a vertical fixed axis is significantly better, while at the end of the block, performance improves in all conditions but one: inconsistent with a horizontal fixed axis. It could be hypothesized that, at first, participants might have a previous schema based on their experience with computers and 3D, which is based on a fixed vertical axis but requires feedback consistent with the expected movement of the object. However, as soon as they gain some experience, they manage to build new schemata improving performance in all conditions but in the horizontal-inconsistent, which has neither the advantage of action-feedback consistency nor any already built schema based on experience. The subjective evaluation through the SUS and RTLX questionnaires seems to validate these results since the questionnaires reveal that the consistent conditions were perceived by the participants as the ones that had more learning capacity and as those that generate the least frustration.

5.2 Conclusions and recommendations: still room for improvement in traditional 3D interaction

We aimed to understand whether the right selection of a fixed axis when designing an interaction technique can improve inspection tasks by making navigation more natural. To this aim, we carried out a study on two different platforms, a mouse and a touch screen. Their limited number of DoF make 3D interaction challenging. Three different variables have been analysed to assess the usability of the techniques: performance—in terms of time to complete the tasks—and the SUS and RTLX questionnaire scales. The fatigue caused by each of the platforms was also investigated.

We found evidence that confirms our hypothesis. When the fixed axis is consistent with the intrinsic axis of an object being inspected, the interaction is more usable, more efficient, easier to learn, produces a lower mental workload and generates less frustration. These results reflect the importance of the fixed axis in a TAV technique and how this axis should be fixed according to the object and the mental model that the user has of it. Specifically, when the object appears to be on a surface, like the ground or the floor, the vertical axis is an intrinsic axis that should be kept fixed since

gravity is always vertical. However, in those cases where the object is suspended in an empty space and there is an intrinsic axis, as with a rotating axis of a wheel or a turbine rotor, it is worthwhile to adapt the interaction technique that has that intrinsic axis fixed. This is not considered in currently used 3D viewers and would improve usability in those cases. Likewise, current 3d model authoring tools do not allow the annotation of a model with an intrinsic rotation axis' attribute, which could be automatically read by 3D viewers, allowing application of our findings to improve the usability of the interaction. Another option to implement consistency is via interactive mechanisms such as visualization widgets. Visualization widgets are virtual objects (e.g. cutting planes, lenses or particle emitters) that help users to interactively manipulate 3D spatial objects and data. A recent survey on 3D Interaction techniques for visualization [4] suggests that this is a clear area for future work. This might be a useful mechanism to interactively fix an axis, in inspection of 3D objects of which the shape is not known in advance, or to allow applications to adapt navigation to mental models which can vary across different contexts or users.

In summary, we found significant differences between consistency and inconsistency across performance time, perceived usability, and perceived fatigue. Still, a reader of this paper, perhaps involved in developing a 3D authoring or visualization tool, might wonder if the effort of implementing this consistency is worth the time. This question is not for the authors of this paper to answer, as it depends on a large number of factors. Perhaps it would not be worth for a general purpose 3D viewer in a 3D object online shop. But for specialised applications, which operators use for long periods of time to carefully and systematically inspect or annotate complex objects, in domains where not making mistakes due to fatigue is important, we firmly believe that implementing such mechanisms would increase the value of the software.

Funding This work was supported by the PLUGGY project (<https://www.pluggypjproject.eu/>), within the European Union's Horizon 2020 research and innovation programme, under grant agreement No 726765. The data and the application used to conduct the experiment are available upon request. Open Access funding provided by Universidad de Málaga/CBUA in the framework of the CRUE-CSIC agreement with Springer Nature.

Declarations

Conflict of interest The authors declare that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence,

unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix: Illustration of the TAV with fixed Up- and fixed right-vectors

The software used by participants in the experiment described in this paper implements two variants of a 3D interaction technique: The Two-Axis Valuator (TAV) with fixed Up-vector and the TAV with fixed Right-vector. Both techniques can be explained using the concept of a virtual sphere that completely encases the 3D scene or object being inspected. The displacements on each of the axes on the screen are mapped onto rotations in the 3D world. For example, in Fig. 10 a gesture from a point on the screen p_1 to another p_2 pro-

duces a rotation of the virtual sphere, that would take a point on its surface A_1 to a new position A_2 . The gesture on the screen plane from p_1 to p_2 is the sum of two orthogonal displacements: Δy (in blue color in Fig. 10), along the vertical, y-axis on the screen; and Δx along the horizontal, x-axis on the screen (in green in Fig. 10).

The TAV with fixed Up-vector maps Δy to rotations $\Delta\phi$ (in blue) around a global X-axis, and Δx to rotations $\Delta\theta$ (in green) around the object's local Y-axis. This implies that the local Y-axis of the object is not generally vertical in world co-ordinates. This is illustrated in Fig. 10a.

The TAV with fixed Right-vector maps the vertical component Δy of the screen gesture onto rotations around the object's local X-axis ($\Delta\phi$ in blue), while the horizontal component Δx will cause a rotation around the world Y-axis ($\Delta\theta$ in green). This is illustrated in Fig. 10b.

The paper advances in the direction of a generalization of the vector to fix, by way of investigating the importance of selecting which axis to fix depending on the intrinsic characteristics of the inspected object.

Fig. 10 Both techniques illustrated side to side. A gesture on the screen plane of the device (left) is transformed by the interaction techniques into a 3D rotation in the virtual world. (a) is the TAV with fixed Up-vector, where rotations are around a global X axis but a local Y axis; (b) is the TAV with fixed Right-vector, where rotations are around a local X axis, but a global Y axis

References

- Autodesk: 3ds Max - 3d Modelling, Animation and Rendering Software (2012). <https://www.autodesk.com/products/3ds-max/overview>, <http://usa.autodesk.com/3ds-max/>
- Bade, R., Ritter, F., Preim, B.: Usability comparison of mouse-based interaction techniques for predictable 3d rotation. In: Smart Graphics, vol. **3638**, pp. 138–150. Springer, Berlin, Heidelberg (2005). doi: <https://doi.org/10.1007/11536482>. <http://www.springerlink.com/index/5egjuvehfq4lgat0.pdf>
- Besaçon, L., Issartel, P., Ammi, M., Isenberg, T.: Mouse, Tactile, and Tangible Input for 3D Manipulation. Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17 pp. 4727–4740 (2017). doi: <https://doi.org/10.1145/3025453.3025863>. <http://dl.acm.org/citation.cfm?doid=3025453.3025863>
- Besaçon, L., Ynnerman, A., Keefe, D.F., Yu, L., Isenberg, T.: The state of the art of spatial interfaces for 3D visualization. Comput. Graph. For. **40**(1), 293–326 (2021). <https://doi.org/10.1111/CGF.14189>
- Blender Foundation: blender.org - Home of the Blender project - Free and Open 3D Creation Software (2015). <https://www.blender.org/>
- Borsci, S., Federici, S., Lauriola, M.: On the dimensionality of the system usability scale: a test of alternative measurement models. Cogn. Process. **10**(3), 193–197 (2009). <https://doi.org/10.1007/s10339-009-0268-9>
- Brooke, J.: SUS: A quick and dirty usability scale. In: Jordan, P.W., Thomas, B., Weerdmeester, B.A., McClelland, A.L. (eds.) Usability Evaluation in Industry. Taylor and Francis, London (1996)
- Buda, V.: Rotation techniques for 3d object interaction on mobile devices. Tech. rep., Utrecht University, Utrecht (2012). <https://dspace.library.uu.nl/handle/1874/255390>
- Chen, M., Mountford, S.J., Sellen, A., Chen, M., Mountford, S.J., Sellen, A.: A study in interactive 3-D rotation using 2-D control devices. ACM SIGGRAPH Comput. Graph. **22**(4), 121–129 (1988)
- Cowan, G.: Statistical Data Analysis, 1st edition. Clarendon Press, Oxford, New York (1998)
- Cumming, G.: The new statistics: why and how. Psychol. Sci. **25**(1), 7–29 (2014). <https://doi.org/10.1177/0956797613504966>
- de Arquer, I., Nogareda, C.: Estimación de la carga mental de trabajo: el método NASA TLX. Notas técnicas de prevención. Instituto Nacional de Seguridad y Salud en el Trabajo. Gobierno de España NTP **544**, (2001)
- Decle, F., Hachet, M.: A study of direct versus planned 3d camera manipulation on touch-based mobile phones. In: MobileHCI09 - The 11th international conference on human-computer interaction with mobile devices and services (2009). <https://doi.org/10.1145/1613858.1613899>
- Devin, F.: Sistema de Escalas de Usabilidad: qué es y para qué sirve? | UXpañol. <http://uxpanol.com/teoria/sistema-de-escalas-de-usabilidad-que-es-y-para-que-sirve/>
- Forlines, C., Wigdor, D., Shen, C., Balakrishnan, R.: Direct-touch vs. mouse input for tabletop displays. In: Conference on Human Factors in Computing Systems - Proceedings, pp. 647–656. ACM Press, New York, New York, USA (2007). doi: <https://doi.org/10.1145/1240624.1240726>. <http://dl.acm.org/citation.cfm?doid=1240624.1240726>
- González-Toledo, D., Cuevas-Rodríguez, M., Molina-Tanco, L., Reyes-Lecuona, A.: 3D Object Rotation Using Virtual Trackball with Fixed Reference Axis. In: Proceeding of EuroVR2018, 1, pp. 3–5. London, UK (2018). doi: <https://doi.org/10.5281/zenodo.2593170>
- Gonzalez-Toledo, D., Cuevas-Rodríguez, M., Garre-Del-Olmo, C., Molina-Tanco, L., Reyes-Lecuona, A.: HOM3R: a 3D viewer for complex hierarchical product models. J. Virt. Real. Broadcast. (2017). <https://doi.org/10.20385/1860-2037/14.2017.3>
- Google SketchUp: 3D Design Software 3D Modeling on the Web SketchUp (2019). <https://www.sketchup.com/>, <https://www.sketchup.com/page/homepage>
- Hand, C.: A survey of 3D interaction techniques. Comput. Graph. For. **16**(5), 269–281 (1997). <https://doi.org/10.1111/1467-8659.00194>
- Hart, S.G.: Nasa-Task Load Index (NASA-TLX); 20 Years Later. Proceedings of the Human Factors and Ergonomics Society Annual Meeting **50**(9), 904–908 (2006). <https://doi.org/10.1177/154193120605000909>
- Hommel, B.: The theory of event coding (TEC) as embodied-cognition framework. Front. Psychol. **6**, 1318 (2015). <https://doi.org/10.3389/fpsyg.2015.01318>
- ISO/TC 159/SC, T.C.: Ergonomic requirements for office work with visual display terminals (VDTs). Part 9: Requirements for non-keyboard input devices (2000). <https://www.iso.org/standard/30030.html>
- Kratz, S., Rohs, M.: Extending the virtual trackball metaphor to rear touch input. In: 3DUI 2010 - IEEE Symposium on 3D User Interfaces 2010, Proceedings, pp. 111–114 (2010). doi: <https://doi.org/10.1109/3DUI.2010.5444712>
- Kunde, W., Müsseler, J., Heuer, H.: Spatial compatibility effects with tool use. Hum. Fact. **49**(4), 661–670 (2007). <https://doi.org/10.1518/001872007X215737>
- Lewis, J.R., Sauro, J.: The factor structure of the system usability scale. In: International conference on human centered design, pp. 94–103. Springer, Berlin, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02806-9_12
- Marchand, E., Courty, N.: Controlling a camera in a virtual environment. The Vis. Comput. **18**(1), 1–19 (2002)
- Martinez, D., Kieffer, S., Martinez, J., Molina, J.P., Macq, B., Gonzalez, P.: Usability evaluation of virtual reality interaction techniques for positioning and manoeuvring in reduced, manipulation-oriented environments. The Vis. Comput. **26**(6), 619–628 (2010)
- Müsseler, J., Skottke, E.M.: Compatibility relationships with simple lever tools. Hum. Fact. **53**(4), 383–390 (2011). <https://doi.org/10.1177/0018720811408599>
- Norman, D.: Some Observations on Mental Models. In: D. Gentner, A.L. Stevens (eds.) Mental Models, Cognitive Science, chap. 1, pp. 7–14. Lawrence Erlbaum Associates, Inc., New Jersey (1983)
- Official Google, B.: 3D Warehouse (2016). <https://3dwarehouse.sketchup.com/>
- Proctor, R.W., Van Zandt, T.: Human Factors in Simple and Complex Systems, Second Edition, 2nd editio edn. CRC Press, Boca Raton, Florida (2008). <https://learning.oreilly.com/library/view/human-factors-in/9780805841190/chapter-01.html>
- Rybicki, S., DeRenzi, B., Gain, J.: Usability and performance of mouse-based rotation controllers. In: Proceedings - Graphics Interface, pp. 93–100. ACM, Victoria, British Columbia, Canada (2016). <https://people.cs.uct.ac.za/~jgain/publications/rotationcontrollers.pdf>
- Sauro, J., Lewis, J.R.: Average task times in usability tests. In: Proceedings of the 28th international conference on Human factors in computing systems - CHI '10, p. 2347. ACM Press, New York, New York, USA (2010). doi: <https://doi.org/10.1145/1753326.1753679>. <http://portal.acm.org/citation.cfm?doid=1753326.1753679>
- Sauro, J.: MeasuringU: measuring usability with the system usability scale (SUS) (2011). <https://measuringu.com/sus/>
- Schmidt, R.A.: A schema theory of discrete motor skill learning. Psychol. Rev. **82**(4), 225–260 (1975). <https://doi.org/10.1037/h0076770>
- Shaw, C.D.: Pain and Fatigue in Desktop VR: Initial Results. Proc. Graphics Interface pp. 185–192 (1998). doi: <https://doi.org/10.20380/GI1998.23>

37. Shoemake, K.: Arcball rotation control. In: P.S. Heckbert (ed.) Graphics gems IV, chap. III.1, pp. 175–192. AP Professional, San Diego (1994). <https://dl.acm.org/citation.cfm?id=180910>
38. Shoemake, K.: ARCBALL: a user interface for specifying three-dimensional orientation using a mouse. In: Proceedings of the conference on Graphics interface '92, pp. 151–156. Canadian Information Processing Society, Vancouver (1992). <https://dl.acm.org/citation.cfm?id=155312>
39. Sketchfab: Sketchfab - Publish & find 3D models online (2019). <https://sketchfab.com/>
40. Thornton, R.W.: The Number Wheel. In: Proceedings of the 6th annual conference on Computer graphics and interactive techniques - SIGGRAPH '79, vol. 13, pp. 102–107. ACM Press, New York, New York, USA (1979). doi: <https://doi.org/10.1145/800249.807430>. <http://portal.acm.org/citation.cfm?doid=800249.807430>
41. VandenBos, G.R.: Publication Manual of the American Psychological Association (6th ed.). American Psychological Association (2009). <http://www.apastyle.org/manual/>
42. Yu, L., Svetachov, P., Isenberg, P., Everts, M.H., Isenberg, T.: FI3D: direct-touch interaction for the exploration of 3D scientific visualization spaces. IEEE Trans. Vis. Comput. Graph. **16**(6), 1613–1622 (2010). <https://doi.org/10.1109/TVCG.2010.157>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Daniel Gonzalez-Toledo earned his M.Sc. degree in Telecommunications engineering in 2013 and M.Sc degree in Telecommunications engineering in 2015 at the University of Malaga (Spain). He started his career as an engineer at IBM Global Services in 2012, before that he was working in different companies as IT staff to pay for his university studies. At the end of 2013, he started his academic career in the department of Electronic Technology at the University of Malaga. Since then, he

has worked as a researcher on four different European projects. He is currently working on a European Project (SONICOM - <https://www.sonicom.eu/>) focused on 3D audio spatialisation while pursuing his Ph.D. in the field of 3D interaction techniques. During his Ph.D. studies, he has done research stays at Imperial College London (United Kingdom) and at microwave and fibre optics laboratory of the National Technical University of Athens.

Maria Cuevas-Rodriguez earned her M.Sc. degree in Telecommunications engineering in 2011 and M.Sc degree in Mechatronic in 2013 at the University of Malaga (Spain). She moved for a year to Tatung University (Taiwan) to improve her skills in 2009/2010. In 2011, she started her career in the department of Electronic Technology of University of Malaga. Since then, she is working as a researcher in different European Projects, on topics related to 3D interaction with

reduced DoF and 3D binaural audio. She is currently working on a European Project (SONICOM - <https://www.sonicom.eu/>) focused on 3D audio spatialisation while pursuing her Ph.D. degree in the same topic. During her Ph.D., she has been working in the Imperial College London (United Kingdom) for 3 months as visiting student and in Facebook Reality Lab (Seattle, EEUU) for 6 months with an internship.

Luis Molina-Tanco is a Lecturer at the University of Malaga and a researcher in DIANA group. He studied Telecommunication Engineering (BSc+MSc) at Universidad Politécnica de Madrid, worked in the Telecommunications Industry and then went back to Academia to undertake a Ph.D. in Computer Science at the University of Surrey (UK), supervised by Professor Adrian Hilton. In 2003, he joined the Department of Electronic Technology (DTE) of the University of Málaga, as

Researcher and Lecturing Assistant, and as a Lecturer since 2007. Since 1996, he has worked in national and EU funded projects in the fields of HCI and pattern recognition, spanning applications in telemedicine, computer animation, robotics and virtual reality.

Arcadio Reyes-Lecuona is an Associate Professor and the head of DIANA research group at the University of Málaga. He obtained his PhD degree in Telecommunication Engineering in 2001. He is a Telecommunication Engineer since 1995 (BSc + MSc equivalent) and Psychologist since 2017 (BSc+MSc equivalent). During the last 20 years, he has been working in VR, haptics, 3D interaction and 3D audio. He has participated in several EU and national projects and has been PI in three national

projects and five European projects. His research interests are in HCI in VR, including 3D interaction with reduced DoF and 3D binaural audio, being one of the coordinators of the 3D Tune-In audio toolkit, an open source library for 3D audio rendering. He is a member of the Steering Committees of EuroXR (European association of Extended Reality) and AIPO (Spanish-speaking association of HCI).