Bimanual Rotation in Midair Combining Multiple and Single Degree of Freedom Manipulation

Bachelor's Thesis at the
Media Computing Group
Prof. Dr. Jan Borchers
Computer Science Department
RWTH Aachen University

by
Marty Pye

Thesis advisor:
Prof. Dr. Jan Borchers

Second examiner:
Prof. Dr. Torsten Kuhlen

Registration date: 08.06.2012
Submission date: 30.09.2012

I hereby declare that I have created this work completely on my own and used no other sources or tools than the ones listed, and that I have marked any citations accordingly.

Aachen, September 2012
Marty Pye

Contents

List of Figures

List of Tables

Abstract

This thesis introduces a quasi-mode for 1DOF rotation during bimanual midair 3DOF rotation. Several existing interaction techniques for 3D rotation are analysed. 2D input devices, such as mice, lack a direct mapping between the virtual and real world. 3D interaction techniques, such as hand-tracking in midair, alleviate this limitation by mapping object manipulation to real world tasks with the aid of metaphors. The benefits and limitations of existing metaphors, such as the sheet-of-paper and handle bar metaphor are discussed, upon which a new interaction design *BASH* is proposed. The compliance of *BASH* with various guidelines for asymmetric bimanual gesture design is discussed. A detailed mathematical model for each feature in *BASH* is provided with consideration of possible interferences between individual actions.

The advantages of providing constrained rotation are outlined in a quantitative study, which was conducted with a prototype containing the elements relevant to rotation. The results indicate that fixed rotation is faster and more accurate for 1DOF tasks and independent of the users' mental rotation ability. This is evidence which supports the design choice of a quasi-mode for fixed rotation.

Acknowledgements

First and foremost, I would like to thank Chatchavan Wacharamanotham, M. Sc. Without his valuable guidance and support, this thesis would not have been possible.

Secondly, I would like to thank my supervisors Prof. Dr. Jan Borchers and Prof. Dr. Torsten Kuhlen.

Additionally, I would like to thank every participant for their time and collaboration for the tests.

Last, but not least, special thanks goes to my family for always supporting me throughout my studies. To them, I dedicate this thesis.

Thank you,

– Marty Pye.

Conventions

Throughout this thesis, the following conventions are used:

Text conventions

Definitions of technical terms or short excursus are set off in coloured boxes.

EXCURSUS:
Excursus are detailed discussions of a particular point in a book, usually in an appendix, or digressions in a written text.

Definition: *Excursus*

The thesis is written in British English. The first person is written in plural form.

Mathematical conventions

If nothing else is specified, all axes refer to the coordinate system depicted below.

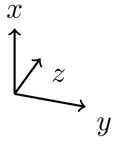

If n is a vector, then $\hat{n}$ is the normalised vector.

All error bars are constructed using a 95% confidence interval of the mean.

Chapter 1

Introduction

Most current Computer Aided Design (CAD) Systems still use two dimensional input and output devices: a mouse for input and a flat screen for output. The user wants to specify shapes, positions and orientations of objects in a virtual 3D scene. These tasks usually require the user to specify three or more variables (e.g. the (x, y, z) components of a translation). However, a 2D mouse does not allow the specification of these three variables in a direct and intuitive manner. It requires the user to switch the planes in which he operates or work with specific constraints.

Usability of 2D input devices in 3D environments is limited.

In order to solve this problem, there are specific 3D input devices on the market, such as 3D mice, which are still not as intuitive and direct as a 3D input method. This is due to the fact that object rotation is still mapped to the rotation of a tool. An even more direct method for interaction in a virtual environment is hand tracking. It is one of the favoured approaches because it integrates the ease with which humans employ their hands in everyday life. Using everyday two-handed actions as metaphors for computer interaction assists in learning this interaction style.

Hand tracking provides the ease with which humans employ their hands in everyday life.

In order for bimanual input to be effective, it must be designed carefully. A classic criticism of two handed input is the "tapping the head and rubbing the stomach" argument. Performing different actions with each hand simultaneously is difficult for most humans without considerable

Design considerations are necessary, due to the complexity of bimanual input.

training. But if used properly, bimanual input can increase the bandwidth of input and therefore results in better performance than one-handed input (Owen et al. [2005]).

Divided attention problem results switching costs.

One of the reasons for the failure of bimanual interaction designs is the divided attention problem (Kahneman [1973]). When there are multiple sources of information (e.g. two cursors), the user has to make a choice where he wants to direct his attention. This results in switching costs. Bimanual interaction works well when, for example, scaling and moving a rectangle by controlling its two opposite corners (Casalta et al. [1999]). The operations of moving and scaling are visually integrated and acan be cognitively chunked as one task.

1.1 Research Questions

Is fixed rotation fast enough?

The bimanual rotation technique presented in this thesis contains a quasi-mode for fixed rotation. Previous studies have shown that free rotation is generally the fastest and most intuitive way to achieve rotation matching tasks, in midair as well as with a mouse (Chen et al. [1988], Jacob et al. [1994]). However, other studies (Veit et al. [2009]) have shown that certain tasks are easier to achieve if the rotation is fixed to one axis. For the fixed rotation to be of use, the axis fixation and subsequent angle specification needs to be achievable with sufficient precision and speed. This is investigated in form of a quantitative study in this thesis.

1.2 Contributions

The main contribution of this thesis is a bimanual quasi-mode rotation technique embedded in the interaction design *BASH*, its partial implementation and an evaluation of certain key factors (speed and accuracy of rotation) in the form of a quantitative study. *BASH* combines two existing metaphors in order to provide an overall more efficient interaction. *BASH* focuses on consistently avoiding the divided attention problem and complying with Guiard's principles for bimanual interaction. The efficiency and benefits of *BASH* are underlined by a quantitative experiment.

Chapter 2

Related work

This chapter focuses on the literature which lead to the evolution of *BASH* and the research questions investigated in chapter 4—"Quantitative Study". It covers existing rotation techniques and contains a discussion of benefits and limitations. It also contains the guidelines with which *BASH* was designed and explains these guidelines with concrete examples.

2.1 3D Rotation Techniques with Input Devices

Several techniques for interactive 3D rotation have evolved, making use of different input devices. The main categories can be separated as follows:

2.1.1 2D Controllers

Virtual Sphere

This 2D interface simulates a physical trackball. The virtual object is shown on the screen, and when the user drags on the object, the movements are applied to the simulated

The virtual sphere simulates a physical trackball.

trackball. To provide the third degree of freedom, a circle is drawn around the object, and when the user clicks outside this circle, the rotation is constrained to the axis perpendicular to the computer screen.

Arcball

Arcball is described as the best known 2D technique for 3D rotation.

This technique is similar to the virtual sphere, but is a mathematically more elegant implementation. Arcball has been described as the best known 2D technique for 3D rotation (Shoemake [1992]). With the virtual sphere, some orientations can only be achieved by performing multiple rotations. Theoretically, Arcball permits the user to rotate an object 360° around any axis with a single mouse drag. In practice however, most users cannot anticipate where to start and end their dragging motion.

2.1.2 3D Controllers

3D input devices come in many shapes and designs. Basically, they can be separated into the following categories:

3D Ball

The 3D ball is an absolute rotation controller.

The 3D ball (Figure 2.1) acts as an absolute rotation controller. The orientation of the object being manipulated always matches the orientation of the 3D ball. With the addition of a clutching mechanism for engaging and disengaging the ball from a virtual object, it is possible to use the 3D ball as a relative rotation controller. Other ball-shaped 3D controllers provide buttons for clutching or other functions integrated into the device.

Figure 2.1: *The 3D ball input device. This input device acts as an absolute controller.*

Trackball

The second type of 3D input controller is the trackball (Figure 2.2). The rotation of the virtual object is mapped to the rotation of the trackball. These trackball devices benefit from the fact that the user can rest his arm on the table as if using a mouse and his movements are stabilised.

The trackball stabilises the users movement while providing up to 6DOF.

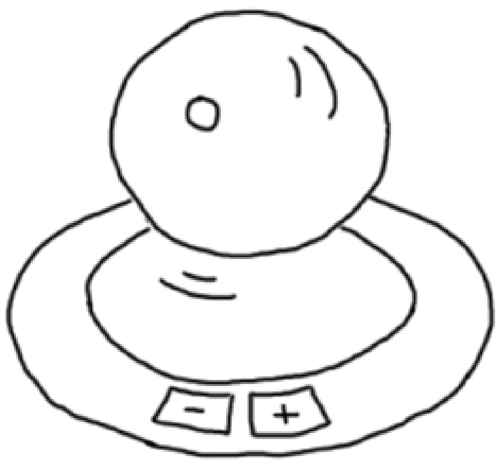

Figure 2.2: *The trackball. The user can rest his arm on the table, which stabilises his movements.*

2.2 Midair

Midair interaction over the keyboard in front of the screen.

Midair rotation techniques provide the ease with which humans employ their hands in everyday life. No additional hardware is needed immediately on the desk, the user can perform rotations at will in midair above traditional I/O devices like mouse and keyboard. Different bimanual rotation techniques are presented in the following sections.

2.2.1 The Sheet-of-Paper Metaphor

An object can be rotated the same way as one rotates a sheet of paper.

Several methods for bimanual rotation in midair have been developed. The first one to be presented is the so called sheet-of-paper metaphor presented by Wang et al. [2011] The user can rotate an object the same way as if he were rotating an imaginary piece of paper with two hands. There are three different gestures for rotating the object around the three axes respectively. These gestures are depicted in Figure 2.3. Rotating the sheet about the y or z axis involves moving the hands in opposite directions along the xz- or xy-plane respectively. To rotate the paper about the x axis, one lifts or lowers the hands while bending the wrists. This method has several benefits and limitations:

Benefits:

- Object translation and object rotation can be performed simultaneously.
- All three axes can be manipulated without having to reposition the hands.

Limitations:

The sheet-of-paper metaphor requires tracking of hand position and orientation.

- Due to occlusion problems, hand position can, in general, be detected more accurately than hand orientation. The "sheet-of-paper" metaphor uses hand position for manipulating the y and z rotation but hand orientation for the x rotation. Therefore, depending on the system setup, rotation around the x axis may

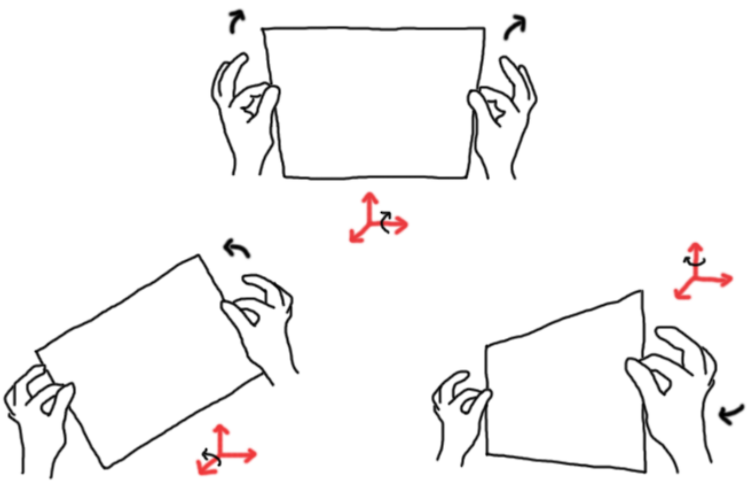

Figure 2.3: *Gestures for rotation around the x, y and z axis. (Adapted from Wang et al. [2011])*

not be as accurate as the other two axes. However, this is a technical concern, which does not necessarily have to influence the best design choice.

- 360° rotation without regrabbing the object can only be achieved around the y and z axis. Due to the physical limitation of the human wrist, the user cannot perform a 360° rotation around the x axis. He is forced to regrab several times. A possibility to relieve this limitation is to amplify the rotated angles to minimise motion.

User has to regrab object due to limitations of the human wrist.

2.2.2 The Handle Bar Metaphor

The next method presented is the handle bar metaphor, which is a bimanual interaction technique designed for the Microsoft Kinect by Song et al. [2012]. With this technique, the user manipulates objects bimanually as if they were skewered on a handle bar (see Figure 2.4).

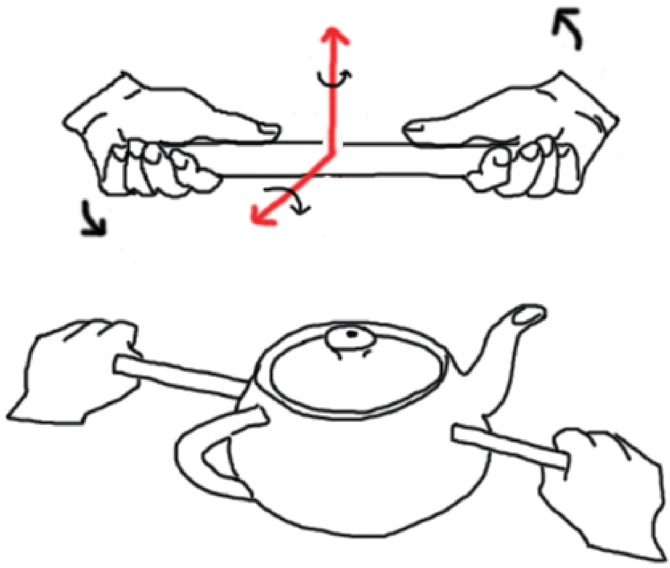

Figure 2.4: *The handle bar metaphor: object (tea pot) can be manipulated as if skewered on a handle bar, which can be moved with both hands. (Adapted from Song et al. [2012])*

Free Rotation

Object is manipulated as if skewered on a handle bar.

When the user closes his hands around the virtual handle bar, he can rotate the object as if it were on the bar. Rotation around the y and z axis works analogously to the sheet of paper metaphor with a fist gesture instead of a pinch gesture. Due to the low resolution of the Kinects image, orientation of a fist cannot be stably detected. Therefore, the authors created a different technique for rotation around the x axis. The user can execute an appropriate concurrent bimanual rotation about the y and z axes simultaneously, later referred to as a "pedaling" motion. This rotates the object incrementally around the x axis (see Figure 2.5).

Constrained Rotation

Constrained Rotation allows rotation around one axis.

The user also has the possibility to perform constrained rotation. Sometimes, the user might want to make a fast

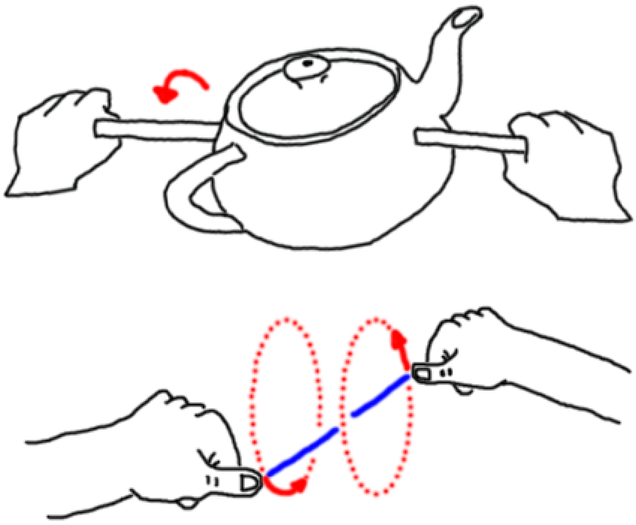

Figure 2.5: *The "pedaling" motion: incremental x axis rotation using continuous rotation around the y and z axes. (Adapted from Song et al. [2012])*

and precise rotation about a specific straight line. In this case, constrained rotation is preferable (see section 2.3—"Separability of Rotation DOF"). In order to begin with a constrained rotation, the user has to draw a crank in midair (see Figure 2.6). With this virtual crank, he can then rotate the object around one axis.

The handle bar metaphor also has several benefits and limitations (Song et al. [2012]):

Benefits:

Handle bar metaphor is stable and flexible.

- All gestures use hand position as opposed to hand orientation. This ensures stable detection by the Kinect.
- The interaction design contains gestures for free and constrained rotation. This allows the user to manipulate individual axes one at a time, if he desires.
- Support for both object and non-object centered manipulation is provided. The user can manipulate the

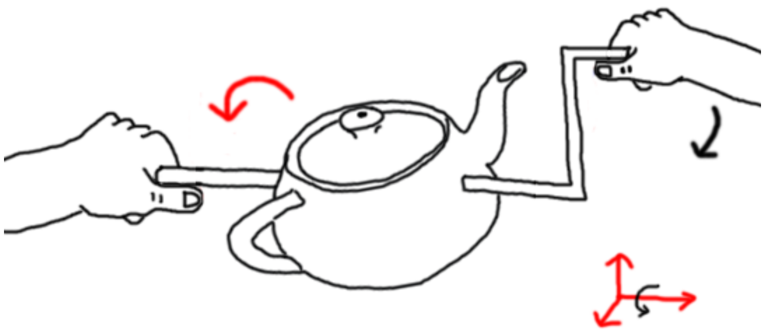

Figure 2.6: *The "cranking" motion: constrained rotation around the handle bar with the aid of a virtual crank. (Adapted from Song et al. [2012])*

position of the handle bar to any location relative to the selected 3D object.

- Unlike the "sheet-of-paper" metaphor, this interaction technique allows a 360° rotation around all three axes at any point in time, without the necessity of repositioning the handle bar several times.

Pedaling motion is unintuitive and tiring.

Limitations:

- The "pedaling" motion for rotation around the x axis is not intuitive for inexperienced users.
- During a user study, some users complained of arm fatigue after 20-30 minutes. This is a general drawback for all midair gestures which provide no physical support to the outstretched arms.
- The bimanual asymmetric gestures caused memory lapse for some of the users.

2.2.3 Grabbing Techniques

Before implementing a rotation technique, a consistent grabbing technique needed to be specified. The act of grabbing and rotating are closely related, because grabbing often precedes rotation. According to Zachmann [2000], grabbing is mostly realised by using one of the following techniques:

- *Single-Step*: The object is attached to the hand at a certain event (e.g. spoken command, gesture)
- *Two-Step*: In the first step, the user switches into a "grabbing mode" (e.g. spoken command, gesture). In the second step, the object is attached to the users hand, usually by collision detection.
- *Natural*: In this technique, the user closes his fingers around the object with a physical grabbing motion.
- *Dwell Time Threshold*: Here the user lingers with his hands on an object, which results in the object being grabbed after a certain amount of time.

Natural grabbing is the "best" technique.

Each technique has its own limitations. Dwell time thresholding induces a constant lag in the interaction. In the single- and two-step technique it is possible for the user to forget the command for grabbing, particularly if the interaction design contains a lot of commands or gestures. Additionally, the two-step grabbing technique is modal, which can always result in mode errors (see Norman [1983]). While natural grabbing is probably the technique with the least flaws for the user, it is also the hardest to implement so that it works consistently. Both the grabbing gesture and the collision with the desired object need to be detected, preferably at the same time.

2.3 Separability of Rotation DOF

Each Task has a perceptual structure, either integral or separable.

Integrating degrees of freedom (DOF) means permitting a user to manipulate several attributes of an object simultaneously, whereas separating DOF constrains the user to manipulation of only certain attributes at once. Jacob et al. [1994] introduced the integrality and separability of the perceptual structure of an input device. Every task has its own perceptual structure. The perceptual structure of a task is integral if the semantic difference between the attributes is low. The perceptual structure of a task is separable, if the semantic distance between the attributes is important. For example, value (lightness) and chroma (saturation) of a colour are perceived integrally, while size and lightness of an object are perceived separably (Handelt and Imai [1972]). Jacob et al. showed that to obtain the best performance, the perceptual structure of the input device needs to correspond to the perceptual structure of the task.

Studies confirm that rotation performances are better with 6DOF devices.

Previous studies (Hinckley et al. [1997], Poupyrev et al. [2000]) have confirmed that six DOF input devices reduce task completion times. For example, Chen et al. [1988] performed a study where users carried out orientation matching tasks with sliders having a separable perceptual structure, and a six DOF input device, which had an integral perceptual structure. While the precision remained similar, achievement times for the six DOF device were significantly lower than for the sliders.

No existing research for separable rotation tasks, i.e. rotations around one axis.

So while integrated degrees of freedom generally permit faster completion times for orientation tasks which require rotation around 3 DOF, none of the above analysed the performance when the orientation tasks structure is separable. For example, if the rotation tasks are only around one axis, an input device with a separated perceptual structure could be superior.

Veit et al. [2009] investigated this influence of the integration and separation of the DOF on the users' performance during 3D orientation tasks. They also measured the number of DOF simultaneously manipulated during an orientation task. The context of this study is presented in the following three sections and the results are presented in section 2.3.4—"Impact on Task Completion Times" and section 2.3.5—"Impact on DOF" and the

2.3.1 Task Conditions

In order to inspect the above, Veit et al. distinguished between the following types of orientation complexity:

- Simple Orientations: for these tasks, the object only needed to be rotated around one axis to reach the target position.

- Complex Orientations: for these tasks, the object needed to be rotated around several axes to reach the target position.

For each task, they distinguished between small and large orientations regarding the minimal amplitude of the rotation:

- Small Orientations: for these tasks, the object needed to be rotated by less than 60 degrees.

- Large Orientations: for these tasks, the object needed to be rotated by more than 60 degrees.

2.3.2 Measurements

They recorded the following variables for all tasks:

- Angular Distance (in degrees): the final angular distance between the manipulated object and the target.

- Achievement Time (in seconds): the time elapsed until the last rotation was made.

- Coarse Achievement Time (in seconds): the time elapsed until the angular distance was lower than 18 degrees for the first time.

- Finer Achievement Time (in seconds): the time elapsed until the angular distance was lower than 9 degrees for the first time.

- Number of DOF (in percent): the amount of time manipulating one, two or three DOF respectively.

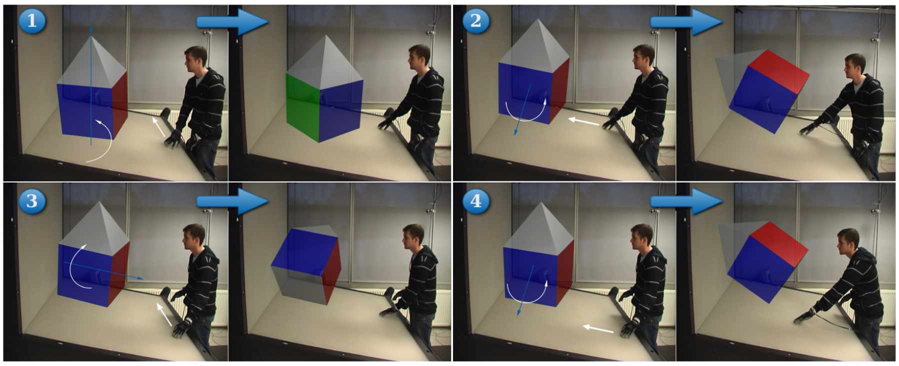

Figure 2.7: *The four mappings between the hand movement and the rotation of the object. The right hand controls two axes (1, 2) while the left hand controls the third (3). The upward movement of the left hand results in the same rotation as the same motion with the right hand (4).*

2.3.3 Compared Techniques

The rotation tasks were performed with an integrated and separated input technique.

The tasks were executed with two different interaction techniques. In the first technique, the users rotated a physical cube in midair and the rotation of the target object was mapped to the rotation performed on the physical cube. This was the integrated DOF technique. In the second technique, the rotation was executed on a plane. This was the separated DOF technique. The dominant hand gave the user access to two specific axes of rotation and the non-dominant hand gave him access to the third axis (see Figure 2.7). This eased the tasks decomposition into DOF, but the user could also manipulate several DOF simultaneously by using both hands at the same time or by performing transverse movements on the screen. There were several conclusions drawn after this study, the most relevant of which are presented below. These results were the building blocks for this thesis.

2.3.4 Impact on Task Completion Times

The input technique had a significant effect on the achievement times of rotation tasks. The details are listed below:

- The mean coarse achievement time is 16.72% faster for the separated input technique than for the integrated input technique. The difference is significant for simple orientations and large orientations.
- The mean finer achievement time was 11.61% faster using the separated input technique. Here also, the difference was only significant for the simple orientations and large orientations.

Separated technique was faster for complex tasks.

An explanation for the results above is that users subset complex tasks into smaller tasks of 1DOF each. In this case, the separated input technique aids the tasks decomposition into individual DOF and therefore results in faster achievement times.

Users subset complex tasks.

It must be kept in mind that Veit et. al used a midair interaction technique for the integrated rotation but used a planar interaction technique for the separated rotation. Therefore it cannot be concluded whether the results can be generalised to pure midair interaction or not. This generalisation to midair interaction is investigated in chapter 4—"Quantitative Study". If the results can be generalised to midair, they will provide evidence that the fixed rotation quasi-mode in *BASH* is a benefit.

Integrated technique was in midair while the separated technique was on a surface.

2.3.5 Impact on DOF

As mentioned previously, Veit et al. measured the number of DOF simultaneously manipulated during each task. Users only manipulated three DOF during 8.75% of the time with the integrated input technique. During 52% of the time, they manipulated two DOF simultaneously. This confirms that mental rotation around three axes is very difficult. Users are usually not able to manipulate all three

Users tend not to manipulate 3DOF simultaneously.

DOF at the same time, even when using an integrated interaction technique. Instead they try to decompose the task into more simple orientations. When using the separated interaction technique, users manipulated one DOF at a time during 80% of the time. This confirms that the separated input technique eases the tasks decomposition into three tasks of one DOF each.

Tasks are decomposed into two phases, the first with an integral structure, the second with a separable structure.

From the results we can conclude that users tend to decompose orientation tasks into two phases. During the first phase, users try to coarsely orientate the object, integrating the DOF. During this phase, an input technique which integrates the DOF is beneficial because it provides the user with the same perceptual structure. During the second phase, the users try to make small adjustments to rotate the object into an accurate position. In this phase, the separated input technique is beneficial, because the users can successively align the object around the three axes of rotation.

2.4 Guiard's Principles for Bimanual Interaction

Different kinds of manual interactions can be classified as follows:

- *Unimanual interactions*: These are realized with one hand, e.g. turning a door handle.
- *Bimanual interactions*: These are realized with two hands which can be divided into
  - *Symmetric bimanual interactions*: e.g. riding a bicycle, steering a car.
  - *Asymmetric bimanual interactions*: e.g. playing the guitar, playing the piano, eating with knife and fork.

Guiard [1987] presented a theoretical model for human bimanual, asymmetric interaction. The relevance of this model is limited for unimanual and symmetric actions due to their lack of complexity. When designing asymmetric bimanual gestures however, it is important to bear Guiard's principles in mind. In the following sections, Guiard's principles for bimanual interaction are presented and explained with the example of driving a nail into a wall with a hammer. The dominant hand is referred to as the right hand and the non-dominant hand is referred to as the left hand. The principles for left-handed people are the same, just mirrored.

Bimanual asymmetric interactions should comply with Guiard's principles.

2.4.1 Right-to-Left Spatial Reference in Manual Motion

According to this principle, motions of the right hand find their spatial references in the results of motions of the left hand. An example is hammering a nail into a wall. The left hand immobilises the nail against the wall while the right hand uses the hammer (see Figure 2.8)

Right hand operates in space prepared by the left hand.

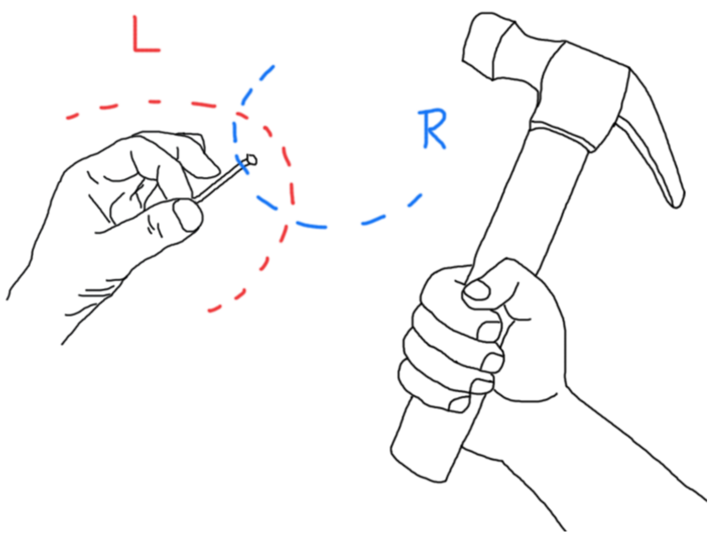

Figure 2.8: *Right-to-Left Spatial Reference in Manual Motion: The right hand operates in the space which the left hand has prepared for action.*

2.4.2 Left-Right Contrast on the Spatial-Temporal Scale of Motion

The right and left hand are on opposite sides of the spatial-temporal scale of motion (see Figure 2.9).

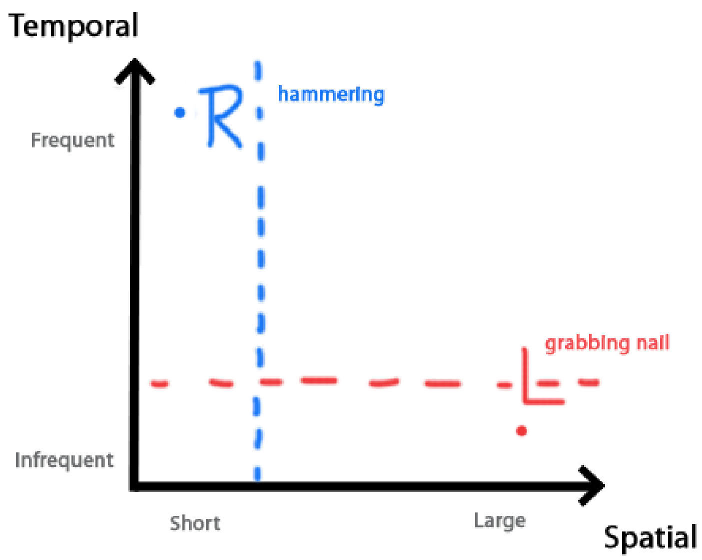

Figure 2.9: *The Spatial-Temporal Scale: The right hand makes short movements frequently while the left hand makes large movements less often.*

Right hand performs short motions frequently. Left hand performs large motions infrequently.

On the spatial axis, the right hand performs shorter and more accurate motions (moving the hammer up and down on the nail) while the left hand performs larger motions (grabbing a new nail from the box). On the temporal axis, the right hand performs more frequent actions (hammer has to be lifted and dropped many times on the nail head) while the left hand performs motions with less frequency (holding the nail in place for the hammer to hit).

2.4.3 Left Hand Precedence in Action

Left hand operates before right hand.

According to this principle, the left hand's contribution to the current action starts before that of the right hand. For the hammer-nail example, this is the case when the left hand grabs a nail and places it against the wall before the right hand starts hammering (see Figure 2.10).

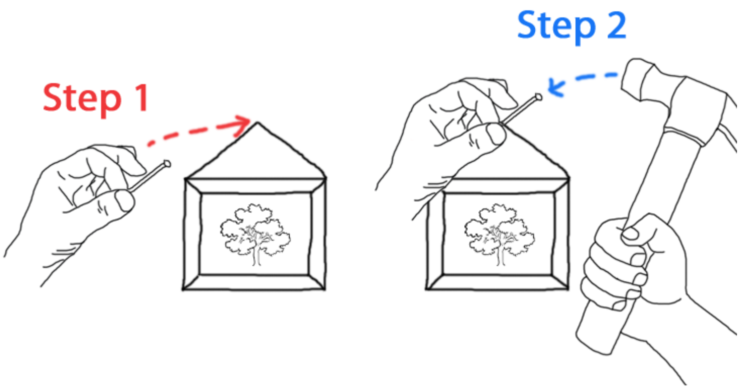

Figure 2.10: *Left Hand Precedence in Action. The left hand's contribution starts before that of the right hand.*

The compliance of *BASH* with Guiard's principles is demonstrated in section 3.4—"Compliance with Guiard's Principles". This chapter presented a discussion of different interaction techniques. The following chapter presents the interaction design *BASH* as a result of the conclusions drawn out of this chapter.

Chapter 3

The Interaction Design – BASH

The final interaction design contributing to this thesis is a Bimanual Asymmetric interaction design combining the Sheet-of-paper metaphor and Handle bar metaphor, further referred to as *BASH*. It provides free rotation with a quasi-mode for fixed rotation. It attempts to combine the benefits of both metaphors to provide a richer interaction. Even though it solves some of the problems of each interaction technique, it also exhibits some new limitations. *BASH* is presented in detail in this chapter.

3.1 Free Rotation

Free rotation is based on sheet-of-paper metaphor.

Free Rotation works similarly to the sheet-of-paper metaphor for rotation around the x and z axis. The user can grab the virtual object with both hands using a pinching gesture as if to grab a sheet of paper at two ends and then rotate the object by moving his hands in the xy or zy plane. Rotation around the y axis however, is different to the sheet-of-paper metaphor. Instead of tracking hand orientation and the bending of the wrists, the object must be grabbed at the top and at the bottom and then the hands moved in the xz plane (see Figure 3.1).

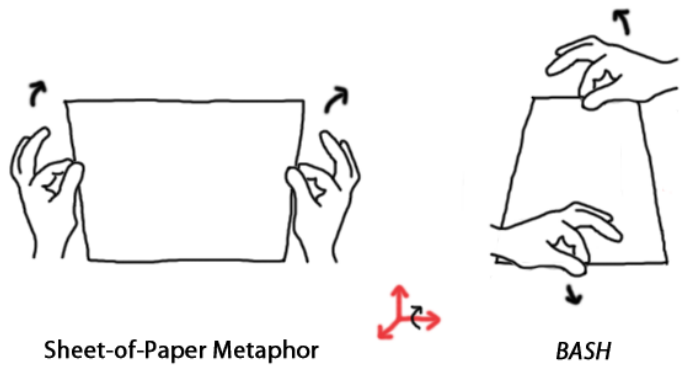

Figure 3.1: *Rotation around y axis.* ***Left****: Sheet-of-Paper metaphor.* ***Right****: BASH*

3.1.1 Mathematical Model

The free rotation technique in *BASH* is built according to the following mathematical model :

p_h^i := Position of the hand $h \in \{L, R\}$ in frame i.

$\hat{n}^i$:= Normalised vector between left and right hand in frame i.

$$\hat{n}^i = \frac{p_R^i - p_L^i}{\|p_R^i - p_L^i\|}$$

θ^i := Axis of rotation (through object centre) in frame i.

$$\theta^i = \hat{n}^{i-1} \times \hat{n}^i$$

α^i := Angle of rotation in frame i.

$$\alpha^i = arctan2(\|\theta^i\|, \hat{n}^{i-1} \cdot \hat{n}^i)$$

R_{Free}^i := Free Rotation around θ^i of angle α^i.

To summarise the above, as soon as the object is grabbed, θ^i is computed as the vector perpendicular to the plane defined by the previous and current position of both hands.

Subsequently, α^i is computed, which together with θ^i forms R^i_{Free} and is then applied to the object as a transformation matrix. This computation of R^i_{Free} takes place incrementally in every frame, until the object is released.

3.1.2 Benefits and Limitations

The benefits and limitations of this gesture for rotating an object around the y axis in comparison to the sheet-of-paper and handle bar metaphor are listed below.

Benefits:

Rotation around y axis is more intuitive in *BASH*.

- The handle bar metaphor rotates an object around the y axis with the aid of the pedaling motion, which is not intuitive. However, grabbing an object at the top and bottom and flipping it is an everyday action, and therefore intuitive. Also pedaling is a relative positioning technique which could be slower if one has to perform a large rotation.

BASH only relies on tracking of hand position.

- The sheet-of-paper metaphor uses hand orientation to detect rotation around the y axis. Hand orientation is harder to track and more prone to error than hand position. However, *BASH* uses hand position to track rotation around the y axis. This can be more stably tracked.
- This gesture also removes the limitation of the human wrist. The object can be rotated 360° without the need to regrab.
- The movement direction is consistent. The movement of the hands is always in the plane orthogonal to the axis of rotation.

Limitations:

Regrabbing is necessary.

- The user is not constantly able to manipulate all three DOF without regrabbing in between. For example, if he has grabbed the object on the left and right side, he can not rotate the object around the y axis without either rotating the object 90° around the x axis first or alternatively releasing the object first and then repositioning his hands.

3.2 Fixed Rotation

Fixed rotation quasi-mode is based on the handle bar metaphor.

The fixed rotation quasi-mode provides the user with the option of performing constrained rotation at any point in time. The reason why constrained rotation can be of benefit is explained in section 2.3—"Separability of Rotation DOF". The fixed rotation technique is inspired by the handle bar metaphor. The user can switch into the fixed rotation quasi-mode by releasing the object with one hand and forming a fist with the other, as if grabbing a crank. He can then rotate this hand, rotating the virtual crank (see Figure 3.2). The axis of rotation is the line between the two hands just before the user released the object with one hand.

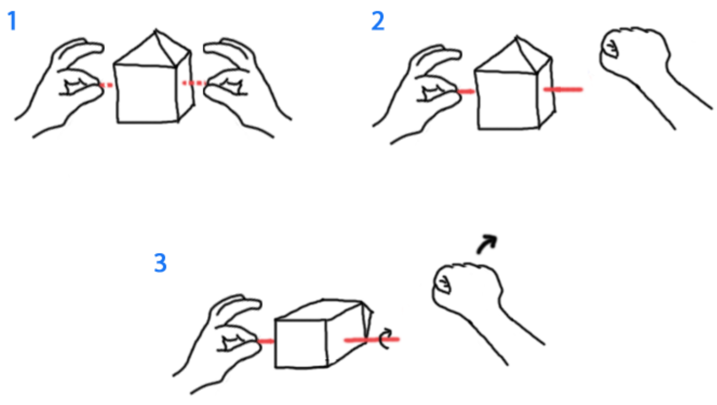

Figure 3.2: *Gesture for fixed rotation.* ***Left:*** *Object is grabbed in free rotation mode.* ***Centre:*** *One hand forms a fist, which triggers the fixation of the axis.* ***Right:*** *Fist performs cranking motion to rotate the object around the fixed axis.*

3.2.1 Mathematical Model

The mathematical model of the fixed rotation technique in *BASH* is presented below:

As soon as the object is released with the right hand, the axis of rotation is determined:

$$\theta^i = \frac{p_R^{i-1} - p_L^{i-1}}{\|p_R^{i-1} - p_L^{i-1}\|}$$

Then, when the right hand forms a fist, the angle of rotation is computed:

$$\alpha^i = arctan2(\|\theta^i\|, p_R^i \cdot p_R^{i-1})$$

R_{Fixed}^i := Fixed rotation around θ^i with amplitude α^i.

To summarise the above, as soon as the object is released by the right hand, the axis of rotation θ^i is determined. As soon as the right hand forms a fist, α^i is computed and R_{Fixed}^i is applied to the object. Until the object is released by either the left or right hand, R_{Fixed}^i is computed and applied in every frame. During this time, θ^i remains constant. As soon as the object is released by either hand, the fixed rotation is terminated. The whole mechanism works analogously, if at the beginning, the object is released by the left hand instead of the right hand.

3.2.2 Benefits and Limitations

Again, *BASH* attempts to use the benefits of the handle bar metaphor and remove some of the limitations:

Benefits:

- Entering the cranking mode is less complicated than with the handle bar metaphor, where the user has to open his fist, draw a crank, and then reform a fist to grab the crank again. Here the user just has to form a fist and can start cranking.

Limitations:

- The user needs to have grabbed the object with both hands, before he can enter the fixed rotation quasi-mode. Otherwise the rotation axis is ambiguous.

3.3 Further Object Manipulation Techniques

BASH contains translation and scaling.

Further object manipulation like translation and scaling can also be integrated into *BASH*. Translation can be achieved by grabbing an object and moving both hands while maintaining a constant distance between both hands. Scaling can be achieved by grabbing an object and then decreasing or increasing the distance between the two hands. These features were not implemented in the prototype as they were not needed for the experiment. They are however, embedded in the interaction design in order to show that *BASH* can be extended to further interactions. In future work, *BASH* could be provided with these features, by implementing the following mathematical models for scaling and translating.

3.3.1 Scaling

The model resolves accidental scaling during rotation of an object.

Scaling can be achieved by increasing or decreasing the distance between the two hands while grabbing an object. Unfortunately, this technique would lead to an accidental scaling when rotating or translating an object, because it is difficult for the user to keep his hands at a constant distance to each other when rotating or translating an object. A suitable mechanism is needed for determining when the objects scale should be coupled to the hands' movement or not.

This mechanism in *BASH* is inspired by the grabbing mechanism from Schlattmann and Klein [2009] and is presented below:

t^i := Time stamp of the frame i.

v_S^i := Signed scaling velocity with which the hands are moved together (positive) or apart (negative) between previous frame $i-1$ and current frame i.

$$v_S^i = \frac{\|p_L^{i-1} - p_R^{i-1}\| - \|p_L^i - p_R^i\|}{t^i - t^{i-1}}$$

v_M^i := Manipulation velocity, which is the sum of the translational velocities of both hands minus the scaling velocity.

$$v_M^i = \frac{\|p_L^i - p_L^{i-1}\|}{t^i - t^{i-1}} + \frac{\|p_R^i - p_R^{i-1}\|}{t^i - t^{i-1}} - |v_S^i|$$

With the help of manipulation and scaling velocity, the scaling factor can be determined:

$$f_S^i = \begin{cases} sgn(v_S^i) \cdot (|v_S^i| - v_M^i) & \text{, if } |v_S^i| > v_M^i \\ 0 & \text{, else} \end{cases}$$

The scaling factor f_S^i is zero while the manipulation velocity is dominant. This ensures that the user can rotate and translate an object without accidentally scaling it. When the scaling velocity is dominant, the size of the currently grabbed object can be incremented by f_S^i. If this mechanism is too sensitive, an additional threshold can be introduced, so that scaling is only triggered when the hands move together or apart at a certain speed. This threshold could be determined in a short test scenario, where several users perform a scaling motion in midair while the velocities of their gestures are recorded.

3.3.2 Translation

Model resolves ambiguity of translation by splitting it into a symmetric and asymmetric part.

Translation can be achieved by bimanually moving the object from one position to another. This is simple when both hands travel the same distance. It becomes more complicated, when the user performs a translation and rotation motion in one gesture. Therefore the translation needs to be split up into a symmetric and an asymmetric part. The symmetric part is used for translation and the asymmetric part is processed by the model for free rotation. The mechanism in *BASH* for computing translation is presented below:

First, the total amount of translation is computed:

t_h^i := Translation of the hand $h \in \{L, R\}$ between frame $i-1$ and i.

$$t_h^i = p_h^i - p_h^{i-1}$$

t^i := Total amount of translation.

$$t^i = \frac{1}{2}(t_L^i + t_R^i)$$

Then, the amount of symmetric translation in percent is computed:

μ^i := Ratio of symmetric translation.

$$\mu^i = \begin{cases} 0 & \text{, if } t_L^i = t_R^i = 0 \\ \frac{min(\|t_L^i\|, \|t_R^i\|)}{max(\|t_L^i\|, \|t_R^i\|)} & \text{, else} \end{cases}$$

s := Arbitrary factor to alter translation speed.

T^i := Translation in frame i.

$$T^i = s \cdot \mu^i \cdot t^i$$

To summarise the above, the translation is split into a symmetric and asymmetric part. The asymmetric part is processed by the mechanism for free rotation. The symmetric part ($\mu^i \cdot t^i$), together with the factor s, forms T^i, which is the translation to be applied to the object in the current frame. If an object is grabbed, T^i is computed and applied to the object in every frame until the object is released.

3.3.3 Integration with One-Handed Interaction

Symmetric two-handed object manipulation is suitable as an extension of one-handed manipulation (see Owen et al.

[2005]). Only if high precision and control of the object are needed, does using both hands improve the current task. In *BASH*, the user can grab and move an object with one hand.

3.4 Compliance with Guiard's Principles

All the gestures in *BASH* are compliant with Guiard's principles for bimanual interaction. Guiard's model is primarily relevant for bimanual, asymmetric gestures. The only gesture in *BASH* which fits into this category is the gesture for fixed rotation. This gesture complies with all three of Guiard's principles:

BASH complies with all three of Guiard's principles.

- *Right-to-Left Spatial Reference in Manual Motion*: The motions of the dominant hand find their spatial references in the results of the motions of the non-dominant hand. The non-dominant hand holds the axis, while the dominant hand performs the cranking motion around this pinned axis.
- *Left-Right Contrast on the Spatial-Temporal Scale of Motion*: On the spatial scale, the dominant hand performs the accurate rotation necessary for acquiring the desired angle, while the non-dominant hand just holds the axis in place. On the temporal scale, the dominant hand performs more motion, while the non-dominant hand plays a passive, infrequent role.
- *Left Hand Precedence in Action*: The non-dominant hand's contribution to the task starts before that of the dominant hand. The non-dominant hand first holds the axis, only then does the dominant hand start performing the cranking gesture.

3.5 Gesture Transition

State transition graph is not complete.

The transitions between individual actions in *BASH* are summarised in a state machine as shown in Figure 3.3. The

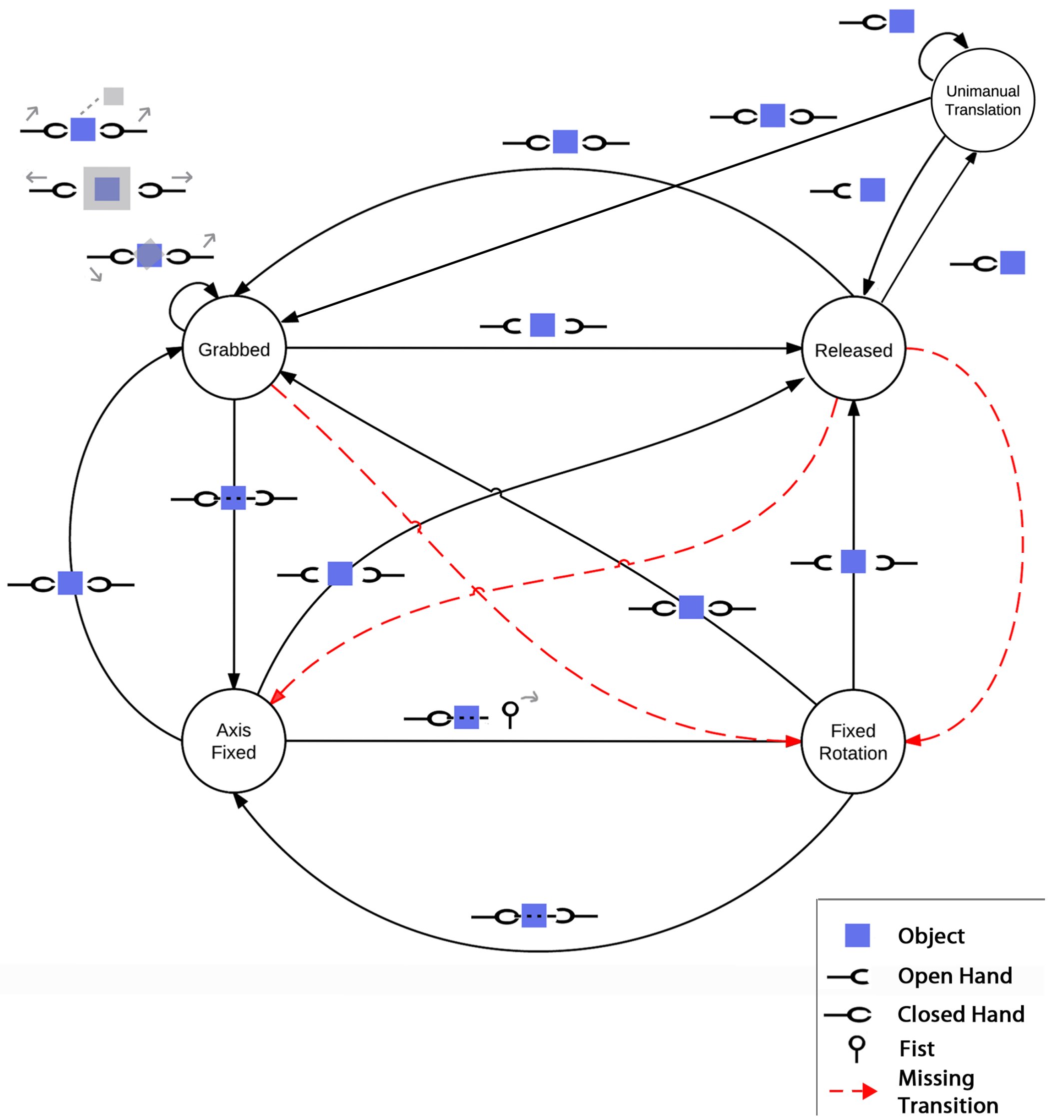

Figure 3.3: *State machine for gesture transition. Red transitions are missing from the interaction design and represent limitations of BASH.*

red transitions represent transitions which are missing from *BASH*. These represent one of the limitations of *BASH*: it is not possible to to perform fixed rotation without fixing the axis first.

Chapter 4

Quantitative Study

> *"Facts, and facts alone, are the foundation of science. . . When one devotes oneself to experimental research it is in order to augment the sum of known facts, or to discover their mutual relations."*
>
> *—François Magendie 1944*

This study was conducted in order to determine whether axis and subsequent angle specification in midair can be performed with sufficient accuracy and speed. The efficiency of the fixed rotation quasi-mode depends on these two factors. The detailed setup of the study is described in this chapter.

4.1 Bimanual Rotation Prototype

The prototype provides the user with the possibility to grab and rotate an object using a bimanual pinching technique. While *BASH* contains a natural grabbing technique, this prototype implements the single-step grabbing technique (section 2.2.3—"Grabbing Techniques"). The object is grabbed as soon as the distance between index finger and thumb of the right hand falls below $1cm$. For this study, a single-step grabbing technique sufficed, due to the fact that the participants only needed to manipulate one object. So

Single-step grabbing technique sufficed for this study.

As additional feedback, the icons change colour when the object is grabbed.

the prototype only needed a binary state for *grabbed* and *released*. When *BASH* is extended to multiple objects, a natural grabbing technique should be used. The icons representing the left and right hand on the screen turn green when the object is grabbed in order to provide the user with active feedback about the grabbing state (see Figure 4.1).

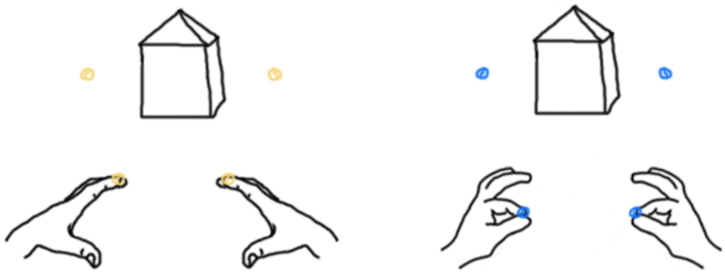

Figure 4.1: ***Left:*** *Object is released.* ***Right:*** *Object is grabbed. Markers change colour to signify the grabbed state.*

Hand position could be tracked regardless of hand orientation.

More functionality like translating or scaling were not added, because the main purpose was to implement the prototype in order to test the users' rotation performances. The prototype was implemented for use with the *Vicon*[1]. Reflective Markers are placed on the users' hands and the *Vicon* cameras can then track the markers in 3D. The cameras are arranged in a cube around the user (see Figure 4.2), allowing the markers visibility for basically all hand orientations. The continuous tracking of the markers was a key attribute to the interaction. If the markers disappeared for any length of time, inconsistencies occurred in the interaction. The cameras were fixed using an aluminium profile construction instead of the standard tripods to ensure that the cameras were not moved accidentally, which would result in a loss of tracking accuracy.

[1]http://www.vicon.com/

Figure 4.2: *Camera constellation as used during the study. The hands were consistently tracked regardless of the orientation.*

4.2 Background and Aims

An elaboration of the background for this study is presented in chapter 2—"Related work". However, a brief summary is provided below.

Axis fixation and angle specification need to be achievable with sufficient accuracy and speed.

Previous studies have shown that free rotation is generally the fastest and most intuitive way to achieve rotation matching tasks, in midair as well as with a mouse (Chen et al. [1988], Jacob et al. [1994]). A freehand interaction technique has the advantage that it allows very direct manipulation of the object in 3D. However, other studies (e.g., Veit et al. [2009]) have shown that certain tasks are easier to achieve if the rotation is fixed to one axis. In order to have the best of both worlds in midair, *BASH* is proposed as an interaction technique, which allows a fluent transition between free and fixed rotation. Fixed Rotation consists of two subtasks: First, specifying the axis and second, specifying the angle. These two subtasks are investigated separately in this study. For the fixed rotation to be of use, the axis fixation and subsequent angle specification needs to be achievable with sufficient precision and speed.

The general aim is to find out whether rotation axis and angle can be specified quickly and accurately. If this is the case, we can focus on the fluent transition between the "free rotation mode" and the "fixed rotation mode".

4.2.1 Hypotheses

For this study, two hypotheses were made:

H1: Users are able to specify the axis of rotation regardless of its direction.

H2: Users are able to rotate the object faster and more accurately in fixed rotation mode (given that the axis of rotation has already been fixed) than in free rotation mode.

4.3 Method of Investigation

4.3.1 Study Design

The participants were students between the age of 20 and 28. The total number of participants was 15 (see section 4.3.5—"Sample Size Calculation"), not counting pilot studies. The tasks were "within-subject", which resulted in 45 trials (3 interaction techniques $\times$ 3 axes $\times$ 5 angles) per subject. The order of the tasks was counterbalanced with a Latin Square.

4.3.2 Tasks

House had to be rotated into a given orientation.

The users were asked to perform a series of orientation matching tasks in free and fixed rotation mode. The object used for the matching task was always a house, similar to the study performed by Chen et al. [1988]. On the left-hand side of the screen, subjects were shown a solid-rendered, coloured house. The subjects were asked to match the orientation of this house to a tilted house on the right-hand side of the screen. The house was coloured differently on all of its faces so as to aid the user in identifying these faces. The tasks were always 1DOF rotations around either the x, y or z axis and the target angles were 45°, −60°, 90°, 135° or −180°.

The other tasks were axis placement tasks. Again, they were given a house on the left-hand side and a tilted house on the right-hand side. The users had to place the appropriate axis for the given rotation. For the pilot study, the axis was rendered as a line but for the study, the axis was rendered as a very thin pipe, which improved perspectival issues.

Axis of rotation had to be placed.

4.3.3 Procedure

The purpose of the study was outlined to all subjects. When the signed consent form had been obtained, the subject participated in the study, which had a duration of approximately 30 minutes. Reflective markers were attached to index finger and thumb of both hands.

The study included three tasks:

1. The user performed 15 rotation matching trials with the free rotation technique. The rotation trials were only around one axis at a time, so this task consisted of 5 trials around the x, y and z axis respectively.

   Free rotation.

2. The user had to perform 15 axis matching trials. An object was given before and after a rotation around one axis. The user had to place this axis. The axes to be specified were always either the x, y or z axis.

   Axis placement.

3. The user had to perform 15 rotation matching trials with the fixed rotation technique. The object was already fixed to the appropriate axis, the user only had to specify the angle of rotation. The appropriate axis was highlighted in order to avoid confusion about which axis the object was to be rotated around.

   Fixed rotation.

The order of these tasks was counterbalanced to cancel out any learning or fatigue effects.

4.3.4 Measurement

Definition:
Euler Angles

EULER ANGLES:
The Euler angles are three angles in the xy, xz and yz plane to describe the orientation of a rigid body in 3D Euclidean space, introduced by Leonhard Euler.[2]

Throughout the study, time in ms and accuracy as three Euler angles was measured and logged for each task. The accuracy of the axis placement was measured as the angle between target axis and the placed axis in the plane defined by those two axes. The rotation accuracy of the object was recorded when the participant indicated that he had finished rotating the object. The course of each task can be

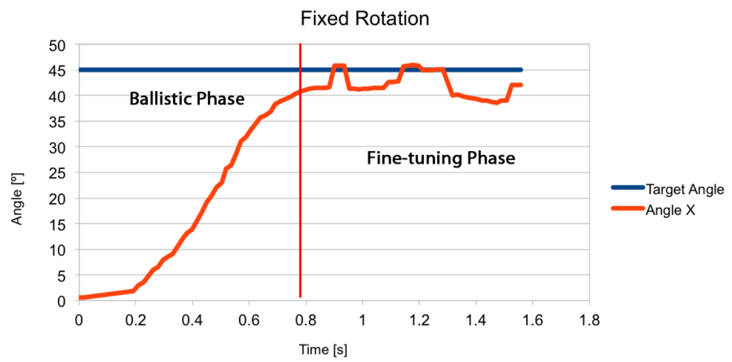

Figure 4.3: *Exemplary course of a rotation task fixed to the x axis. The objects angle of rotation around the x axis is tracked over time. The target angle is $45°$.*

plotted as a graph of the current angle of the object over time. As one can see in Figure 4.3, the course of each task can be separated into two phases:

[2]Novi Commentarii academiae scientiarum Petropolitanae 20, 1776, pp. 189–207 (E478)

1. *Ballistic*: In this phase, the object is roughly aligned according to the target orientation. The curve approaches the target angle comparatively quickly.

2. *Fine-tuning*: In this phase the user performs the final adjustments. He generally under- and overshoots the target orientation before reaching his desired level of precision.

Criteria from Liu et al. [2009] was used in order to determine whether the rotation is currently in the ballistic phase or fine-tuning phase. The analysis of this study was to show whether the technique ("fixed" or "free") had any effect on the ballistic or fine-tuning phase respectively. Although there were three tasks, only task 1 and 3 are comparable. Task 2 is reported with the aid of descriptive statistics, e.g. distribution of the angular error as a histogram.

Each task was separated into ballistic and fine-tuning phase.

4.3.5 Sample Size Calculation

In order to calculate the sample size, power analysis was performed on the data of the first pilot user. Unfortunately with purely that data, the power analysis reported that even with 30 participants there would still be a 41.23% chance that type II errors would occur. Therefore, another three pilot studies were performed and their data analysed. After this, the power analysis reported that a significant effect (at $\alpha = 0.05$) was very likely to occur with 10–12 participants. To be on the safe side, the study was conducted with 15 participants.

Study was performed with 15 participants.

4.3.6 Statistical Methods

The variables used in this experiment were the following:

Dependent Variables: completion time in ms, accuracy in degrees

Independent Variables: task (free rotation, axis specification or fixed rotation), rotation axes (x, y or z axis) and target angles.

The completion time and accuracy were compared using mixed-model ANOVA. Users were treated as a random effect. (2 techniques $\times$ 3 axes). If the data, especially completion time, was not normally distributed, a log transformation was applied. If the sphericity assumption of the data did not hold, the result was adjusted using the Greenhouse-Geisser correction.

Chapter 5

Results and Analysis

5.1 Achievement Time Analysis

5.1.1 Descriptive Statistics

To provide an overview on general tendencies, the mean achievement times for the ballistic and fine-tuning phase are plotted in Figure 5.1. The individual plots are sorted by interaction technique and axis of rotation. The error bars are constructed using a 95% confidence interval of the mean. Three general tendencies can be observed from Figure 5.1.

- The rotation tasks were completed faster with the fixed rotation technique than with the free rotation technique in both the ballistic phase and the fine-tuning phase.

Fixed rotation faster.

- Variance in achievement times among target angles was more pronounced with the free rotation technique.

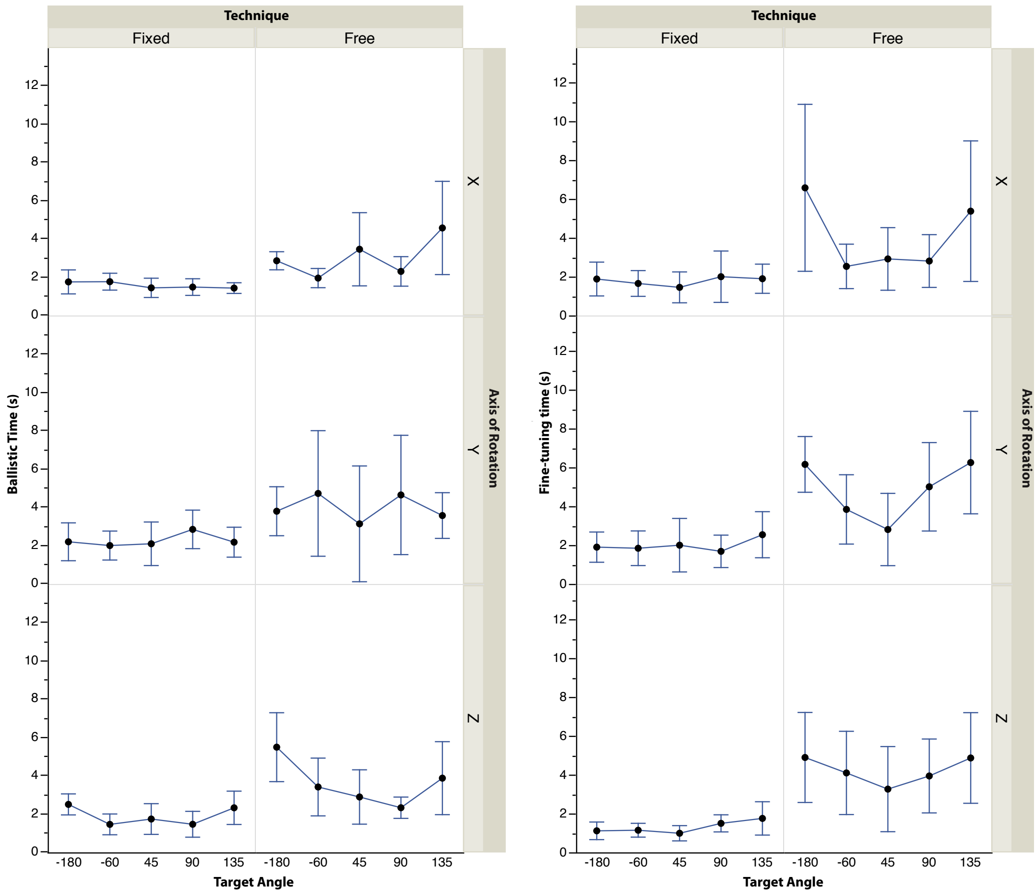

Figure 5.1: *Mean achievement times over target angle sorted by interaction technique and axis of rotation.* ***Left:*** *ballistic phase* ***Right:*** *fine-tuning phase*

5.1.2 Fine-tuning Phase Analysis

The mixed-model analysis of variance for the time spent in the fine-tuning phase delivered the results presented in Table 5.1.

Technique has significant effect.

The axis of rotation, the target angle and the interaction technique (free or fixed) have a significant effect on the time spent in the fine-tuning phase. The main point of interest is the significance of the technique, which confirms that fixed rotation is significantly faster than free rotation for the fine-tuning phase. However, there is an interaction between the

Mixed-Model ANOVA Fine-Tuning Phase				
Source	DF	DFDen	F Ratio	p-Value
Axis of Rotation	2	405	5.1093	0.0064*
Target Angle	4	405	11.9353	< 0.0001*
Axis of Rotation * Target Angle	8	405	1.0224	0.4182
Technique	1	405	92.1455	< 0.0001*
Axis of Rotation * Technique	2	405	0.5006	0.6066
Target Angle * Technique	4	405	3.9214	0.0039*
Axis of Rotation * Target Angle * Technique	8	405	0.6821	0.7074

Table 5.1: *Results of the mixed-model ANOVA for the time spent in the fine-tuning phase.*

target angle and the interaction technique, which needs further investigation. This interaction is visible in Figure 5.2

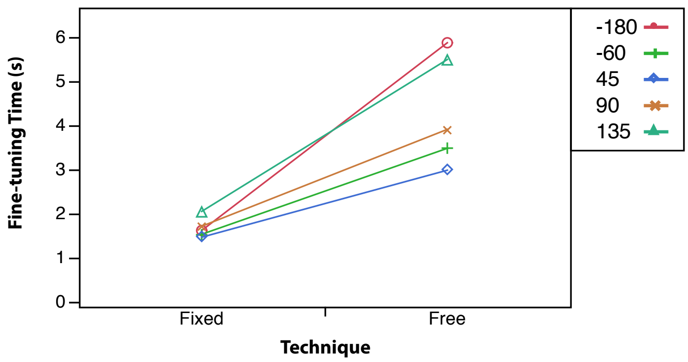

Figure 5.2: *Least square means of time in fine-tuning phase plotted over interaction technique for each target angle. A unidirectional interaction with different magnitudes between technique and target angle is visible.*

Figure 5.2 shows that the interaction between target angle and method is unidirectional. The mean times spent in the fine tuning phase for each angle are always lower with the fixed rotation rotation technique and higher with the free rotation technique. However, there are different magnitudes of interaction effects. While the magnitude of effect seems to remain relatively constant for the angles -60°, 45° and 90°, the magnitude of effect for the target angles 135°

Magnitude of effect higher for large angles.

and -180° is higher. This observation also corresponds with observations made by Veit et al. [2009], where the advantage of the separated interaction technique was only significant for larger angle amplitudes.

Simple Effect of Target Angle

Fixed rotation faster, regardless of the target angle.

The analysis of each individual target angle showed that the interaction technique had a significant main effect for every target angle (p-value < 0.0001 for all 5 target angles). This confirms that the observations made from Figure 5.2 are in fact significant. The fixed rotation technique was always significantly faster in the fine-tuning phase, regardless of the target angle.

Finally, in order to investigate the magnitude of effect, the effect size was calculated (see Table 5.2). The Cohen's d obtained from the effect size calculation underline the observations made in Figure 5.2. The magnitude of effect grows with the amplitude of the target angle, 45° having the lowest ($d = 0.5716$) and -180° having the highest ($d = 1.7204$).

Effect Size Calculation	
Target Angle	Cohen's d
-180°	1.72
-60°	0.82
45°	0.57
90°	0.91
135°	1.02

Table 5.2: *Effect size calculation of the interaction technique on the fine-tuning achievement times per target angle.*

Simple Effect of Method

Now that the interaction in Figure 5.2 has been analysed by target angle, this section analyses the interaction by method. The results of this analysis showed that for fixed rotation, the axis of rotation had a significant effect ($F_{2,196} = 4.78, p = 0.0094$) on the time spent in the fine-tuning phase. The effect is visible in Figure 5.3. The time spent in the fine-tuning phase around the z axis seems to be shorter than for the other two axes. Tukey's honest significance test confirms precisely this. The reason for this cannot be explained with any of the qualitative data gathered from the experiment.

Fine-tuning time around z axis was significantly shorter.

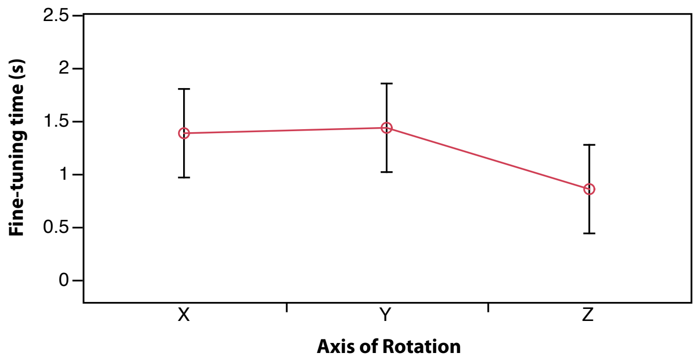

Figure 5.3: *Least square means of time spent in fine-tuning phase for fixed rotation plotted over axis of rotation. Time for rotations around the z axis is significantly lower.*

5.1.3 Ballistic Phase Analysis

The mixed-model analysis of variance for the time spent in the ballistic phase delivered the results presented in Table 5.3. The axis of rotation, the target angle, and the interaction technique have a significant effect on the time spent in the ballistic phase. However, there is a significant interaction between the axis of rotation and the target angle. This interaction is clearly visible in Figure 5.4. This does not directly affect our hypothesis, which is investigating the effect of interaction technique, but it is helpful to analyse this interaction as well. It is not obvious which direction the interaction effect has. The three lines have different slopes for different angles, and cross each other several times. There-

Technique has significant effect.

Significant interaction Axis of Rotation – Target Angle.

Mixed-Model ANOVA Ballistic Phase				
Source	DF	DFDen	F Ratio	p-Value
Axis of Rotation	2	406	4.9033	0.0079*
Target Angle	4	406	7.3225	< 0.0001*
Axis of Rotation * Target Angle	8	406	2.2652	0.0223*
Method	1	406	22.5027	< 0.0001*
Axis of Rotation * Method	2	406	0.0461	0.9549
Target Angle * Method	4	406	0.7529	0.5565
Axis of Rotation * Target Angle * Method	8	406	0.9690	0.4597

Table 5.3: *Results of the mixed-model ANOVA for the time spent in the ballistic phase.*

fore, main effect analysis was performed first by target angle, then by axis of rotation. The results of these tests are presented in the following two sections.

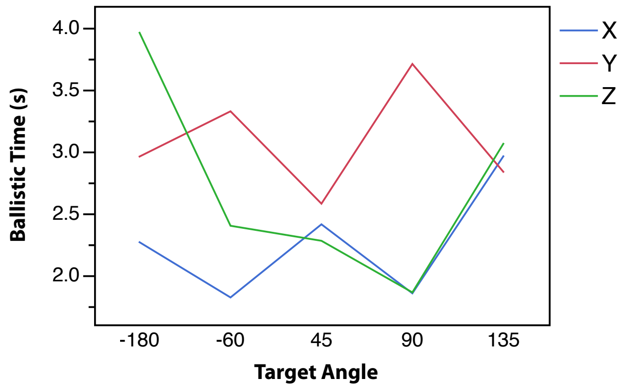

Figure 5.4: *Means of time in ballistic phase over target angle sorted by axis of rotation. The direction of the interaction is not visible from this graph.*

Simple Effect of Target Angle

Investigating all target angles individually, the following observations were made:

- The effect of the interaction technique was significant for every target angle. This confirms that the fixed rotation technique leads to faster achievement times in the ballistic phase, regardless of the target angle.

- The axis of rotation was only significant for the target angles -180° and 90°. The reason for this significance is explained in the following two paragraphs.

Regrabbing costs time.

Regrabbing effect for -180° around z axis: The effect of the axis of rotation on the ballistic phase for a target angle of -180° is shown in Figure 5.5. The time spent in the ballistic phase seems to be larger for the y axis and larger again for the z axis. However, Tukey's honest significance test states that the level of the ballistic times for the x and y axis are not significantly different, but the level of the ballistic times for the z axis is significantly different to those of the other two axes. These results correspond with an observation made during the study. In order to perform a 180° rotation around the z axis without regrabbing, the users had to cross their arms. Many users preferred to regrab several times, rather than crossing their arms. The effect did not occur around the other two axes, because *BASH* does not require crossing of arms for large rotations around the x and y axis.

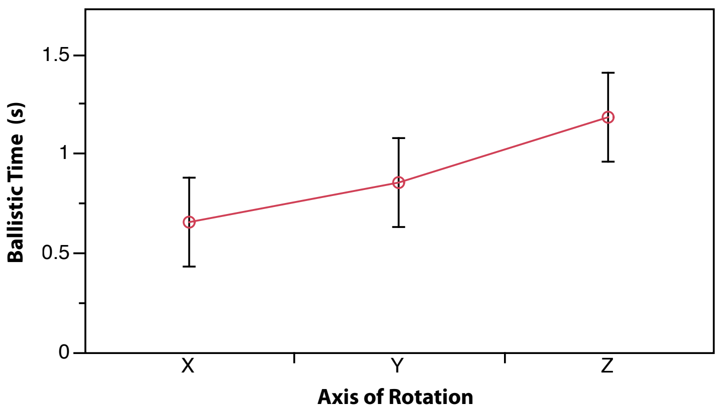

Figure 5.5: *Least square means of time spent in ballistic phase plotted over axis of rotation for the target angle of -180°.*

Hand repositioning for 90° around y axis: The effect of the axis of rotation on the ballistic phase for a target angle of 90° is shown in Figure 5.6. The time spent in the ballistic phase seems to be substantially larger for the y axis than for the other two axes. Tukey's honest significance test confirms this observation. This result for the 90° task reflects one of the flaws of *BASH*. The user cannot manipulate all 3 DOF at all times. In this specific case, after having rotated the object 90° downwards, he cannot correct errors around the x axis without repositioning his hands. This is not the only reason for the higher ballistic times spent around the y axis, for the effect still occurs with less magnitude for the fixed rotation technique, where per definition, there is no error around other axes. So this could imply that rotation around y has a more complicated cognitive process involved, because the rotation gesture does not have the same starting position for the hands as the other two axes. The users have to remember to reposition their hands, before being able to rotate the object around the desired axis.

User has to reposition hands to correct error.

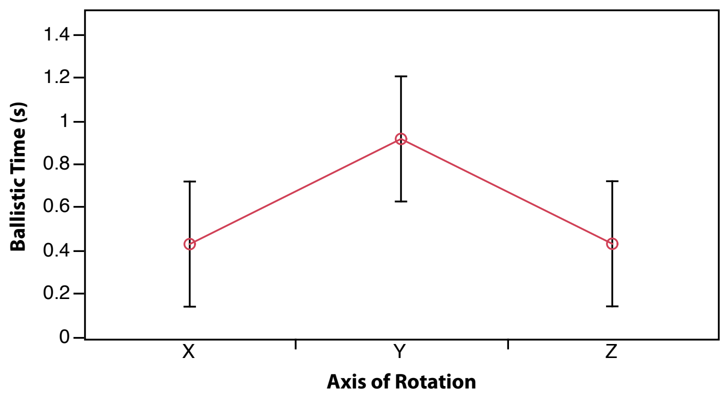

Figure 5.6: *Least square means of time spent in ballistic phase plotted over axis of rotation for the target angle of 90°. The time around the y axis is significantly higher.*

The Cohen's d obtained from the effect size calculation ($d_{-180} = 1.0133$ and $d_{90} = 0.4915$) imply, that even though the effects of both the regrabbing for z axis rotation and the repositioning of the hands for y axis rotation are significant, the effect of regrabbing is quite a bit stronger than the effect of having to reposition hands. One way of alleviating the regrabbing dilemma would be to magnify gestures so they are not mapped 1:1 to the object manipulation.

Simple Effect of Axis of Rotation

Now that the interaction in Figure 5.4 has been analysed by target angle in the previous section, the focus in this section is the analysis by axis of rotation. Observing all axes individually, the following results were obtained:

Fixed rotation is faster, regardless of axis of rotation.

- The interaction technique was significant for every axis of rotation. This implies that the fixed rotation technique is faster in the ballistic phase, regardless of the rotation axis. Together with this exact same observation for the target angle, it can be concluded that fixed rotation is faster in the ballistic phase, independent upon axis of rotation and target angle.
- The target angle had a significant effect on the time spent in the ballistic phase for the z axis. The reason for this significance is explained in the following paragraph.

Ballistic time is longer for large orientations.

Significance for z Axis: The effect of the target angle on the ballistic phase for the z axis is shown in Figure 5.7. The time spent in the ballistic phase seems to be higher for the larger orientations (135° and 180°) than for the other orientations. However, Tukey's honest significance test stated, that only the level of the 180° orientations is significantly different to the three smallest orientations (45°, -60° and 90°). The reason for this is again the regrabbing dilemma.

5.2 Accuracy Analysis

5.2.1 Descriptive Analysis

Observing the data in a descriptive form of angular deviation plotted over target angle, separated by interaction

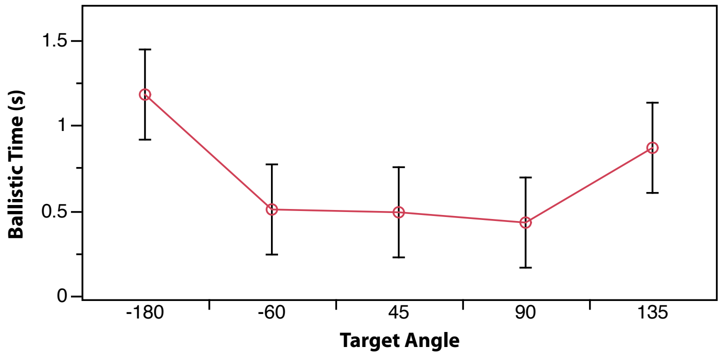

Figure 5.7: *Least square means of time spent in the ballistic phase plotted over target angle for the rotations around the z axis. Time for the 135° and 180° tasks is higher.*

technique and axis of rotation, the following observations were made:

No accumulative error for fixed rotation.

- For the fixed rotation technique, there was no error around the non-primary axes. This was obviously going to be the case, because the rotation is fixed to the required (primary) axis and the user cannot manipulate the other (non-primary) axes.
- For fixed rotation, the error around the x and y axis seems to remain relatively constant with small confidence intervals regardless of the target angle, but around the z axis, there seems to be more variance of error among target angles. This fact needs further investigation.
- There is a lot more variance of error with the free rotation technique, which needs further investigation.

The mixed-model ANOVA showed that the overall accuracy for the fixed rotation technique was indeed significantly better than for the free rotation technique, regardless of target angle and axis of rotation. There were however, interactions between interaction technique and target angles, so the error analysis was also performed separated by axis of rotation.

5.2.2 Simple Effect of Axis of Rotation

The data was subset by axis of rotation. For each axis of rotation, the data was analysed by error around this specific axis. The results in Table 5.4 showed that the interaction technique did not have a significant effect for any of the axes. So, even though the users made a larger sum total of error with the free rotation technique, the error around the relevant axis was not significantly different.

Error around relevant axis in free rotation was not significantly higher.

The target angle did however have a significant effect on the accuracy. Certain target angles were more accurately achieved than others.

Mixed-Model ANOVA for Error around x Axis				
Source	DF	DFDen	F Ratio	p-Value
Target Angle	4	126	6.5958	$< 0.0001^*$
Technique	1	126	0.6596	0.4182
Target Angle * Technique	4	126	0.2368	0.9171
Mixed-Model ANOVA for Error around y Axis				
Source	DF	DFDen	F Ratio	p-Value
Target Angle	4	126	4.1266	0.0036^*
Technique	1	126	1.1110	0.2939
Target Angle * Technique	4	126	1.6143	0.1747
Mixed-Model ANOVA for Error around z Axis				
Source	DF	DFDen	F Ratio	p-Value
Target Angle	4	126	7.7976	$< 0.0001^*$
Technique	1	126	0.3784	0.4182
Target Angle * Technique	4	126	1.0121	0.9171

Table 5.4: *Results of the mixed-model ANOVA for the error around the individual axes.*

Simple Effect of x Axis

For the x axis, the -60° and 135° tasks were significantly less accurate than the others. This is visible in Figure 5.8. The increased inaccuracy for the target angles -60° and 135° is due to perspectival issues. These perspectival issues are shown in Figure 5.9. The 90° task for example can be

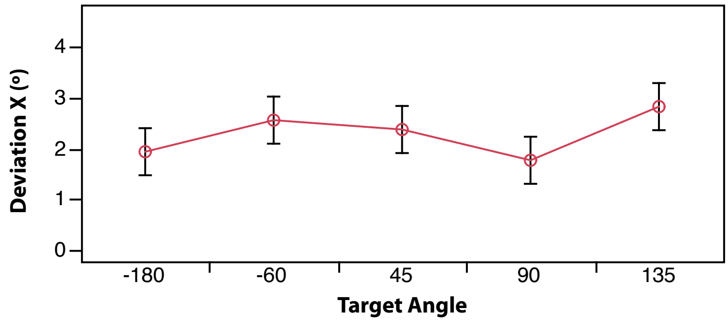

Figure 5.8: *Least square means of error around x axis for tasks around the x axis plotted over target angle.*

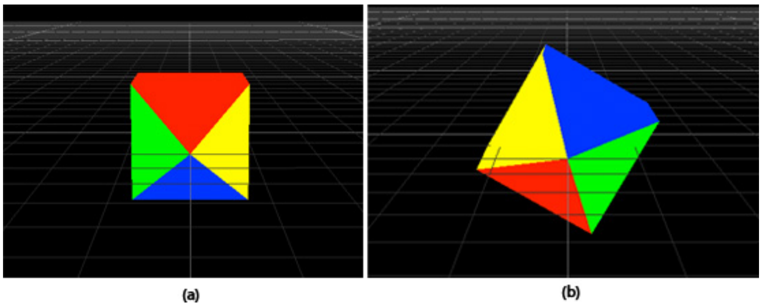

Figure 5.9: *Perspectival issues:* ***(a)*** *90° task, house can be perfectly aligned to the grid.* ***(b)*** *-60° task, house cannot be aligned to the grid as well.*

Possibility of alignment was crucial for accuracy.

aligned accurately to the grid so that the bottom edge of the red face is parallel to the horizontal grid lines. For the -60° task, none of the faces can be aligned accurately to the grid. This caused the users to be less accurate. The same problem occurred for the 135° task.

Simple Effect of y and z Axis

For the y axis, the -60° and 135° tasks were significantly less accurate as well. This is due to similar perspectival issues as for the x axis rotation. The same principle applied to tasks around the z axis, but here it was even more pronounced, because there was no grid in the vertical direction. This meant that the user could not align the object to any grid lines, and therefore, there was more error around the z axis.

At this point one can conclude, that the error accumulates around all three axes with the free rotation technique. The error around the axis relevant to the current rotation task is however, not effected by the interaction technique. The accuracy is effected by the target angle, depending on whether this task aligns accurately with the grid. This confirms previous research which stated that providing a grid for 3D manipulation tasks improves the accuracy.

Previous research confirms that grid improves accuracy.

5.3 Axis Placement Analysis

5.3.1 Descriptive Analysis

The mean angular deviation of the placed axis to the target axis is represented in Figure 5.10. Additionally, the angular deviation is also plotted over individual tasks in Figure 5.11.

The following tendencies can be obtained from Figure 5.10 and Figure 5.11:

- The error around x and z axis for task 14 seems to be substantially larger than for other tasks.
- The error around the y axis seems to be particularly low for tasks 6 to 10.
- The error around the y axis seems to be the lowest (Figure 5.10)

These observations need further analysis, which is presented in the following section.

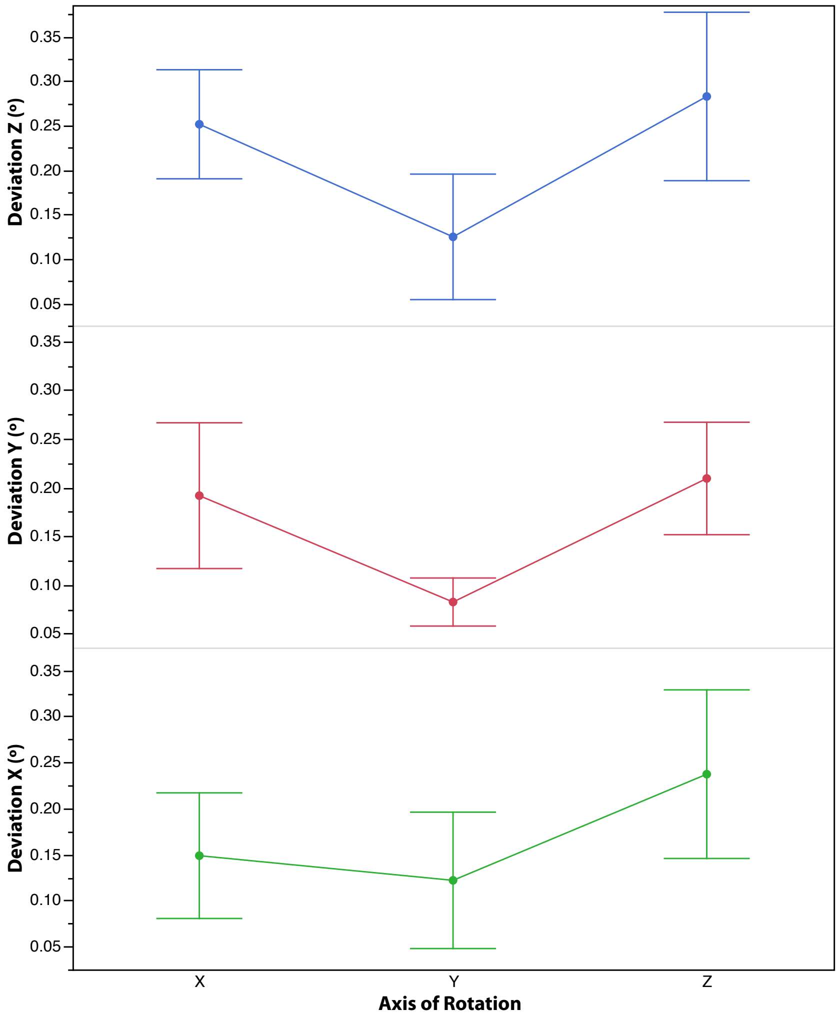

Figure 5.10: *Mean angular deviation for axis placement over axis of rotation. Error around the y axis was the lowest.*

5.3.2 Analysis by Angular Deviation

Axis of rotation had significant effect.

For this analysis, the data was subset by the three angular deviations (angular error around the x, y and z axis respectively) and the effect of the axis of rotation was investigated. The axis of rotation had a significant effect on all three angular errors. The influence of the axis of rotation on the angular deviation around the respective axes is described in the following paragraph.

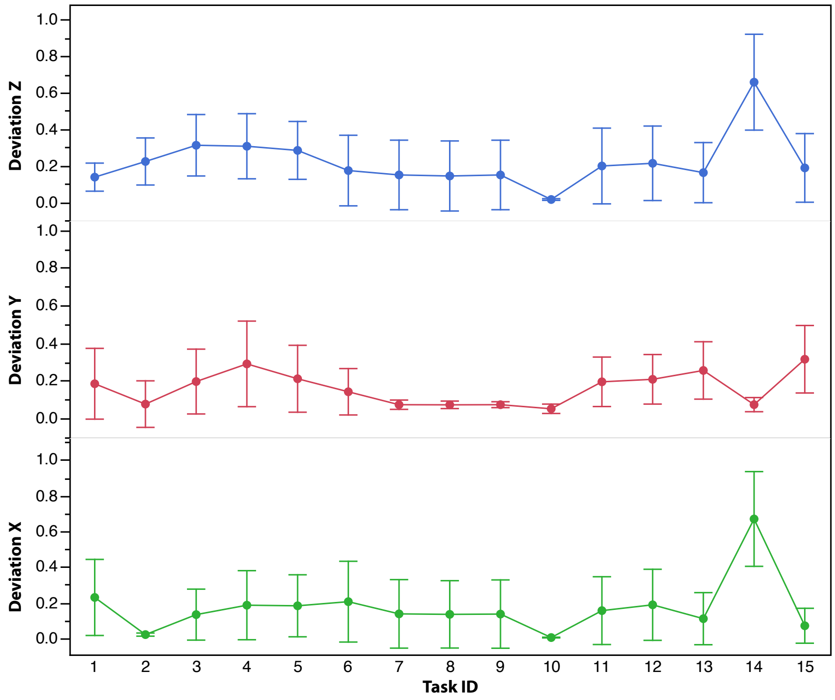

Figure 5.11: *Mean angular deviation for axis placement plotted over task ID. Error for task 14 is particularly high.*

The effect of the axis of rotation on the deviation around the z axis is presented in Figure 5.12. Tukey's honest sig-

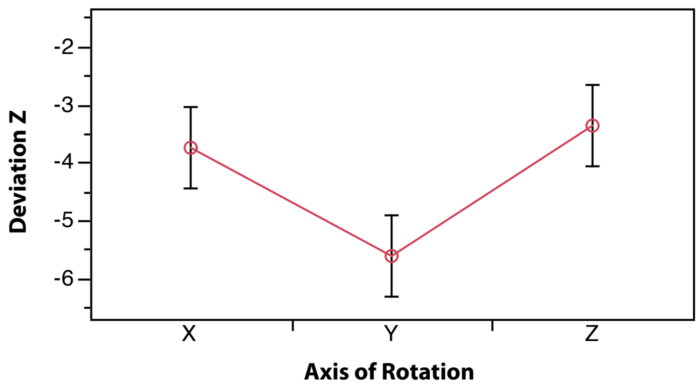

Figure 5.12: *Effect of the axis of rotation on the angular deviation around the z axis.*

Axis perpendicular to user was hard to align accurately.

nificance test showed that there was substantially less error for tasks around the y axis. This is due to perspectival issues. In order to specify y correctly, the users had to place the axis from left to right. They could align the axis with the grid lines going left to right. For the tasks where they had to specify the x or z axis, it was harder to be precise due to depth perception. Without rotating the view, it is impossible to tell whether an axis is perfectly vertical or slightly tilted backwards. This dilemma was visible whenever depth perception was necessary in order to place the axis accurately and explains second and third point observed in the descriptive analysis.

High error for task 14 was due to a cognitive issue.

The high peak for task 14 in Figure 5.11 is due to a participant error which occurred throughout the study. Instead of specifying the z axis, the users often specified the y axis (see Figure 5.13). So the error for task 14 is due to a cognitive issue during mental rotation.

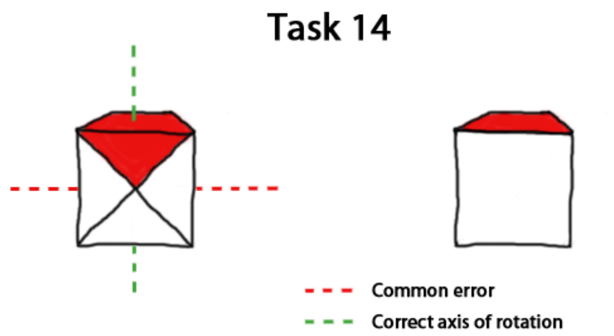

Figure 5.13: *Task 14. Common error was to specify the y axis instead of the z axis.*

At this point one can conclude that the axis could not be specified with the same accuracy regardless of the orientation. Certain perspectives had a negative effect on the axis placement. This has to be kept in mind when implementing the fixed rotation technique.

Chapter 6

Discussion and Future Work

The main goal of this thesis was to determine whether the advantage of a fixed rotation quasi-mode could be generalised to midair interaction and thus provide evidence for the benefits of *BASH*. We therefore conducted a study in order to investigate the speed and accuracy with which axis and subsequent angle specification can be performed. The details of the study results were presented in the previous chapter (see chapter 5—"Results and Analysis"). This chapter provides a discussion of the implications regarding our research questions.

6.1 Discussion

After the thorough analysis in chapter 5—"Results and Analysis", the following conclusions regarding the two hypotheses of the quantitative study can be made:

Axis fixation accuracy depends on perspective.

Axis fixation accuracy depends on viewport perspective: The first hypothesis stated that users would be able to specify the axis of rotation regardless of the direction. This hypothesis is rejected. The accuracy of the axis placement depended too strongly on perspective. During the study, the users could not rotate the view, which would have helped placing the axis accurately. The inaccuracy of the axis specification is a factor which might have a key influence on the actual usability of the fixed rotation technique in *BASH*. Improvement possibilities for these perspectival issues are discussed in section 6.2—"Future Work".

Fixed rotation is fast and accurate.

Object rotation for 1DOF tasks was faster and more accurate with the fixed-rotation technique: The second hypothesis stated that users would be able to rotate the object faster and more accurately using the fixed rotation technique than the free rotation technique. This hypothesis was indeed true. The overall accuracy of the fixed rotation technique was better, because there was no accumulated error around several axes. The accuracy of the task was significantly better, if the orientation task could be aligned to the grid. This confirms that a 3D object manipulation environment should provide such a grid. The fixed rotation technique was faster than the free rotation technique in both the ballistic and fine-tuning phase. From this one can conclude that fixed rotation does in fact present a substantial advantage for rotations around 1DOF in midair. Therefore, providing a midair interaction design with capability of performing constrained rotation is an improvement to that interaction design.

Mental rotation ability has no effect.

Mental rotation ability did not have a significant effect for 1DOF tasks: The preliminary study described in appendix A—"Preliminary Study" showed that the mental rotation ability of the users did not have an effect on the completion times. This suggests that for 1DOF tasks, *BASH* is independent of the users ability to mentally rotate objects.

Results cannot be generalised to all tasks.

Generalisation of the studies is limited: The preliminary and quantitative studies conducted for this thesis contain some limitations. The preliminary study was only performed with eight users. The result regarding the mental rotation ability might not be generalisable to larger populations. The prototype for the quantitative study only contained the elements relevant to rotation and the camera perspective was fixed. The results therefore cannot be generalised to the whole interaction design *BASH*. The results regarding completion time and accuracy cannot be generalised to all rotation tasks because the study was intentionally only conducted with 1DOF tasks. Further studies as suggested in section 6.2—"Future Work" could help in generalising the results of this thesis.

6.2 Future Work

This section provides a short outlook on future work and what issues still need addressing.

The main reason why the axis could not be placed accurately was perspective. It was difficult for the user to see, where exactly the markers representing his hands were and whether they were correctly aligned or not. These perspectival issues will persist in causing problems, particularly when *BASH* is extended to multiple objects. Therefore one needs a consistent technique of providing the user with additional feedback about the 3D position of objects. This will aid the user in actually being able to grab objects naturally. Without additional feedback, users might think their hands are immediately next to the objects even when they are not. One way to realise this would be to provide position pegs, as presented in Glueck et al. [2009]. These position pegs are shown in Figure 6.1. The position of the object in the yz plane is shown by the position of the base circle in the 2D reference grid. Feedback on the x coordinate (height) is provided by the inner radius. The thicker the inner radius is, the closer the object is to the grid. The stalk connects each object to it's respective base circle. A small cone connected to the stalk and base circle provides additional feedback on whether the object is above or below the reference grid.

Perspective is an issue.

Provide position pegs.

When the position pegs have been implemented, a study could be performed in order to investigate whether these position pegs are an improvement and reduce perspective induced errors.

Do position pegs reduce perspective induced error?

As a next step, *BASH* needs to be implemented as a whole according to the mathematical models presented in chapter 3—"The Interaction Design – BASH". Only then can further studies be conducted in order to evaluate the usability of *BASH*. For example, a study could investigate how much time the users spend in the fixed rotation mode, and for what purpose.

What use is made of fixed rotation quasi-mode?

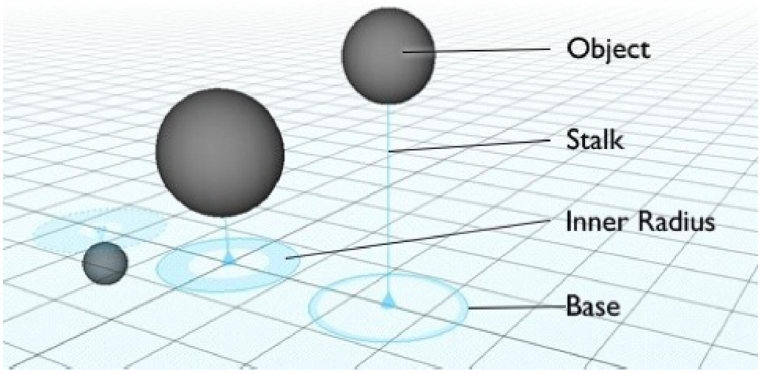

Figure 6.1: *Composition of position pegs. The pegs provide feedback about object position in relation to the grid. (Image from Glueck et al. [2009])*

Currently, the prototype is implemented for use with Vicon. Therefore markers still need to be placed on the users' hands. In future work, the prototype could be extended to markerless tracking with the Kinect.

Appendix A

Preliminary Study

A.1 Prototype

The preliminary study contained two prototypes:

- The first prototype was an arc ball rotation technique for the mouse. The users could rotate an object (a house) with the arc ball rotation technique.
- The second prototype allowed the users to perform fixed rotation around a given axis. The angle of rotation was mapped in a linear fashion, i.e. if the user moved his hand up or down, the angle of rotation was mapped to the movement.

A.2 Background and Aims

The aim of this study was to investigate whether fixed rotation in midair is faster than arc ball rotation with a mouse for 1DOF tasks. The study also investigated whether mental rotation ability had any effect on the achievement times of 1DOF tasks.

A.2.1 Hypotheses

H1: Mental Rotation has no effect on completion times for 1DOF tasks.

H2: Fixed rotation in midair is faster than free rotation with a mouse.

A.3 Method of Investigation

The tasks where the same as in section 4.3.2—“Tasks”. The study was performed with a “within-subject” design. The mental rotation test was performed at the end of the study so as not to put the users under any unnecessary pressure before the test.

A.4 Results

Both **H1** and **H2** were confirmed. Completion time was significantly faster with the midair rotation technique than with the mouse. With this result, the prototype was extended to investigate the same thing purely in midair in chapter 4—“Quantitative Study”.

Mental Rotation did not have a significant effect on the completion time. There was however an interaction between technique and mental rotation. The mental rotation ability affected the mouse completion times more than the midair completion times, but the effect on the completion times were not significant enough to draw any conclusions.

Appendix B

Mental Rotation Test

Appendix B contains the contains the mental rotation test performed with all users during the study

Explanation and Examples for Cube Test

You will be given 5 cubes to choose from (a,b,c,d,e). There are 6 different symbols on each cube. You can see three of them.

In each of the following tasks you will see one of the given cubes in a different position. You need to find out which cube it is. The cube are rotated, turned, or both. Consequently, this could mean that a new symbol becomes visible.

Hint: The given cubes (a,b,c,d,e) are different cubes. They feature the same symbols but on different sides.

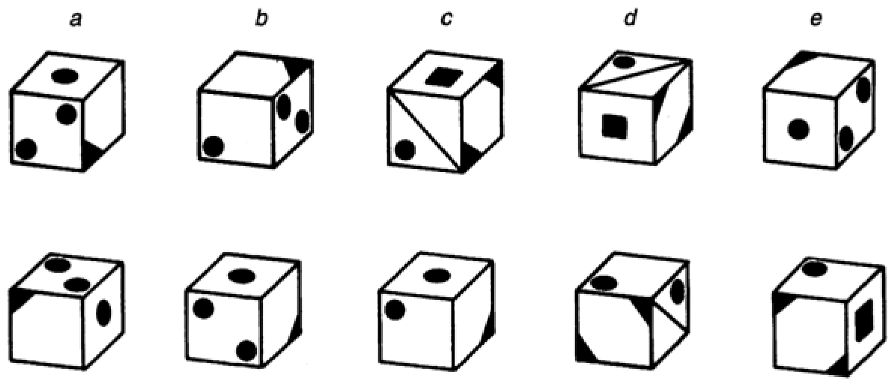

The first cube in this example is cube (a) in a different position. Therefore you should mark (a) on your solution sheet. Cube (a) has been once turned to the right then turned up. The second example matches cube (e), the third matches cube (b). The fourth example is cube (c) and the fifth cube (d).

Stop! Please wait for a signal before you begin.
Do not turn your pages before that signal.

Figure B.1: *Mental rotation test – page 1*

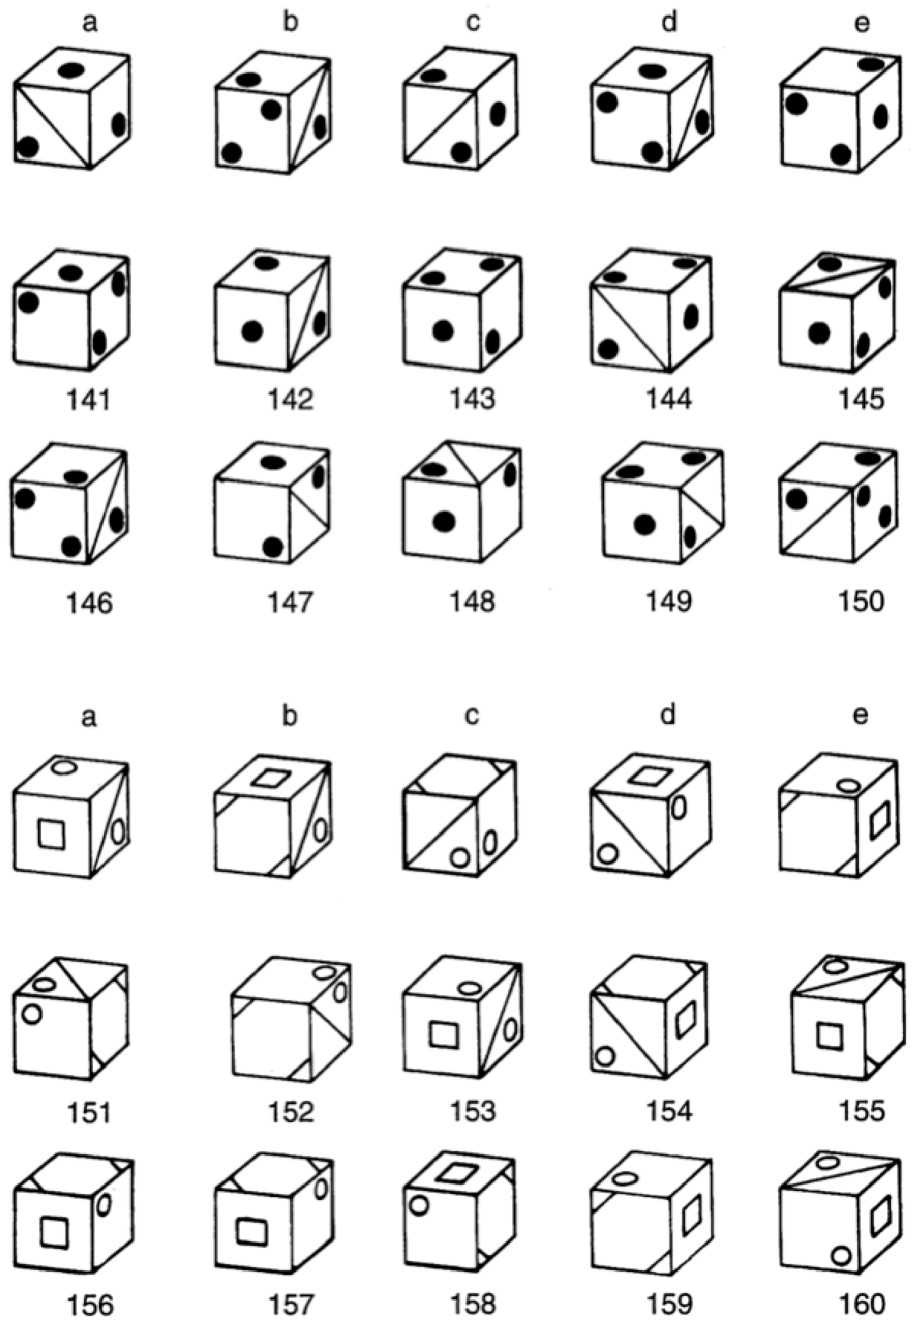

Figure B.2: *Mental rotation test – page 2*

Cube Test: Answer Sheet

Age:__________

	a	b	c	d	e
141					
142					
143					
144					
145					
146					
147					
148					
149					
150					

	a	b	c	d	e
151					
152					
153					
154					
155					
156					
157					
158					
159					
160					

Figure B.3: *Mental rotation test – page 3*

Cube Test: Correction Sheet

	a	b	c	d	e
141					■
142	■				
143					■
144				■	
145				■	
146			■		
147			■		
148			■		
149		■			
150				■	

	a	b	c	d	e
151			■		
152			■		
153				■	
154		■			
155		■			
156					■
157	■				
158				■	
159	■				
160					■

Moderator Instruction

1. Time: 9 minutes, starting when the participant flip to the second page.
2. Each of the correct answers is 1 point.
3. Sum the score up (max = 20)
4. Look up for the standard score from the table on the right
5. The standard score can be treated as an interval variable, i.e., you may use t-test or ANOVA to analyze them

Reference

Amthauer, R., Brocke, B., Liepmann, D., & Beauducel, A. (2001). Intelligenz-Struktur-Test 2000 R (I-S-T 2000 R) Manual. Göttingen: Hogrefe.

Raw Score	Standard Score (by age)			
	21–25	26–30	31–40	From 41
20	125	125	128	131
19	122	122	125	128
18	119	119	123	125
17	116	116	120	123
16	113	113	117	119
15	110	110	114	116
14	107	107	111	113
13	105	105	108	110
12	103	103	106	108
11	100	100	103	105
10	98	98	101	102
9	96	95	98	100
8	93	93	96	97
7	91	91	93	95
6	88	89	91	92
5	86	86	88	90
4	84	84	85	87
3	81	81	82	84
2	78	78	79	81
1	75	76	76	78
0	73	73	73	75

Figure B.4: *Mental rotation test – page 4*

Bibliography

Didier Casalta, Yves Guiard, and Michel Beaudouin-Lafon. Evaluating two-handed input techniques: rectangle editing and navigation. In *CHI '99 extended abstracts on Human factors in computing systems*, CHI EA '99, pages 236–237, New York, NY, USA, 1999. ACM. ISBN 1-58113-158-5. doi: 10.1145/632716.632862. URL http://doi.acm.org/10.1145/632716.632862.

Michael Chen, S. Joy Mountford, and Abigail Sellen. A study in interactive 3-d rotation using 2-d control devices. In *Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '88, pages 121–129, New York, NY, USA, 1988. ACM. ISBN 0-89791-275-6. doi: 10.1145/54852.378497. URL http://doi.acm.org/10.1145/54852.378497.

Michael Glueck, Keenan Crane, Sean Anderson, Andres Rutnik, and Azam Khan. Multiscale 3d reference visualization. In *Proceedings of the 2009 symposium on Interactive 3D graphics and games*, I3D '09, pages 225–232, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-429-4. doi: 10.1145/1507149.1507186. URL http://doi.acm.org/10.1145/1507149.1507186.

Yves Guiard. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model, 1987. URL http://cogprints.org/625/.

Stephen Handelt and Shiro Imai. The free classification of analyzable and unanalyzable stimuli. *Perception & Psychophysics*, 12:108–116, 1972. ISSN 0031-5117. doi: 10.3758/BF03212854. URL http://dx.doi.org/10.3758/BF03212854.

Ken Hinckley, Joe Tullio, Randy Pausch, Dennis Proffitt, and Neal Kassell. Usability analysis of 3d rotation techniques. In *Proceedings of the 10th annual ACM symposium on User interface software and technology*, UIST '97, pages 1–10, New York, NY, USA, 1997. ACM. ISBN 0-89791-881-9. doi: 10.1145/263407.263408. URL http://doi.acm.org/10.1145/263407.263408.

Robert J. K. Jacob, Linda E. Sibert, Daniel C. McFarlane, and M. Preston Mullen, Jr. Integrality and separability of input devices. *ACM Transactions on Computer-Human Interaction*, 1(1):3–26, March 1994. ISSN 1073-0516. doi: 10.1145/174630.174631. URL http://doi.acm.org/10.1145/174630.174631.

Daniel Kahneman. *Attention and Effort*. Prentice-Hall series in experimental psychology. Prentice-Hall, 1973. ISBN 9780130505187. URL http://books.google.co.in/books?id=7kvuAAAAMAAJ.

L. Liu, R. van Liere, C. Nieuwenhuizen, and J.-B. Martens. Comparing aimed movements in the real world and in virtual reality. In *Virtual Reality Conference, 2009. VR 2009. IEEE*, pages 219 –222, March 2009. doi: 10.1109/VR.2009.4811026.

Donald A. Norman. Design rules based on analyses of human error. *Communications of the ACM*, 26(4):254–258, April 1983. ISSN 0001-0782. doi: 10.1145/2163.358092. URL http://doi.acm.org/10.1145/2163.358092.

Russell Owen, Gordon Kurtenbach, George Fitzmaurice, Thomas Baudel, and Bill Buxton. When it gets more difficult, use both hands: exploring bimanual curve manipulation. In *Proceedings of Graphics Interface 2005*, GI '05, pages 17–24, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2005. Canadian Human-Computer Communications Society. ISBN 1-56881-265-5. URL http://dl.acm.org/citation.cfm?id=1089508.1089512.

Ivan Poupyrev, Suzanne Weghorst, and Sidney Fels. Non-isomorphic 3d rotational techniques. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '00, pages 540–547, New York, NY, USA, 2000. ACM.

ISBN 1-58113-216-6. doi: 10.1145/332040.332497. URL http://doi.acm.org/10.1145/332040.332497.

Markus Schlattmann and Reinhard Klein. Efficient bimanual symmetric 3D manipulation for markerless hand-tracking. In *Virtual Reality International Conference (VRIC)*, April 2009.

Ken Shoemake. Arcball: A user interface for specifying three-dimensional orientation using a mouse. In *Proceedings of Graphics Interface 1992*, volume 92, pages 151–156, 1992.

Peng Song, Wooi Boon Goh, William Hutama, Chi-Wing Fu, and Xiaopei Liu. A handle bar metaphor for virtual object manipulation with mid-air interaction. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, CHI '12, pages 1297–1306, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1015-4. doi: 10.1145/2207676.2208585. URL http://doi.acm.org/10.1145/2207676.2208585.

Manuel Veit, Antonio Capobianco, and Dominique Bechmann. Influence of degrees of freedom's manipulation on performances during orientation tasks in virtual reality environments. In *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology*, VRST '09, pages 51–58, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-869-8. doi: 10.1145/1643928.1643942. URL http://doi.acm.org/10.1145/1643928.1643942.

Robert Wang, Sylvain Paris, and Jovan Popović. 6D hands: arkerless hand-tracking for computer aided design. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, UIST '11, pages 549–558, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0716-1. doi: 10.1145/2047196.2047269. URL http://doi.acm.org/10.1145/2047196.2047269.

G. Zachmann. *Virtual Reality in Assembly Simulation - Collision Detection, Simulation Algorithms, and Interaction Techniques*. PhD thesis, Technische Universität Darmstadt, 2000.

Index

Typeset September 26, 2012