

Pointing at Trivariate Targets in 3D Environments

Tovi Grossman, Ravin Balakrishnan

Department of Computer Science

University of Toronto

www.dgp.toronto.edu

tovi | ravin @dgp.toronto.edu

ABSTRACT

We investigate pointing in true 3D environments where the target size varies in three spatial dimensions. We also study the effect of the user's physical movement angle on pointing performance. Results show that target size dimension along the primary axis of movement has a greater impact on performance than the other two dimensions. Movement angle also significantly affects performance, and changes the relative impact of the three target dimensions. Building upon recent results in the modeling of bivariate pointing, we propose and validate a new model that describes pointing at trivariate targets. This model also accounts for movement angle, and outperforms previously published models.

CATEGORIES AND SUBJECT DESCRIPTORS: H.1.2

[**Models and Principles**]: User/Machine Systems – *Human Factors*; *Human Information Processing*; H.5.2

[**Information Interfaces and Presentation**]: User Interfaces – *Theory and methods*.

GENERAL TERMS: Experimentation, Human Factors

KEYWORDS: Pointing, 3D Fitts' Law, volumetric displays.

INTRODUCTION

Advances in three dimensional display technology [6, 13] have reached the point where early models are now available commercially [10, 11]. In order to fully leverage the unique features of these displays, applications must be developed with appropriate user interfaces whose design is grounded in a sound understanding of the human's capabilities when interacting in this new medium. At the most basic level, these interfaces must enable users to easily select and manipulate virtual elements in the 3D display space. While many studies have been conducted on selection and manipulation interfaces for 3D environments (e.g., [8, 20-23]), few have looked at developing and evaluating appropriate predictive theoretical models that describe the underlying human behavior.

In this paper, we study and model user performance in the most fundamental interaction task – pointing – in a 3D

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2004, April 24–29, 2004, Vienna, Austria.

Copyright 2004 ACM 1-58113-702-8/04/0004...\$5.00.

display where the size of the target varies in three spatial dimensions. We review previous related work on pointing and Fitts' law models in 1D and 2D; identify various factors – target dimensions, movement amplitude, movement angle and interactions between them – that could affect pointing performance in 3D; propose several models that could appropriately characterize such 3D pointing behavior; present the results of a controlled experiment that both investigates the effects of the factors identified as well as evaluates the proposed models; and conclude by discussing implications for user interface design.

BACKGROUND

One of the more successful quantitative models in HCI is Fitts' law [7], which is used to model pointing tasks in user interfaces [14]. It predicts the time MT to select a target of width W and distance (or amplitude) A from the cursor as:

$$MT = a + b \log_2 \left(\frac{A}{W} + 1 \right)$$

where a and b are empirically determined constants. The logarithmic term is the index of difficulty (ID) of the task. Fitts' law as originally formulated is a one dimensional model, with the target width W being the only movement constraint. It also assumes that the direction of movement is collinear with this W dimension. Numerous studies have been conducted over the years which amply validate this model for 1D pointing tasks (see [14] for a review). However, pointing tasks in most standard interfaces are typically two-dimensional, with targets having both height and width constraints. While not as extensive as the literature in modeling 1D pointing, there has been some research studying the effects of varying the width and height dimensions of 2D targets, as well as varying the angle at which the cursor approaches the target [1, 5, 9, 15].

In the first study on bivariate pointing in the HCI literature, MacKenzie and Buxton [15] proposed several different formulas for the index of difficulty for a rectangular target, and found two formulations which correlated highly with their experimental data. Their first formulation considers W to be the dimension of the target in the direction of movement (W'). This $ID_{W'}$ model thus reduces the 2D task to a 1D task performed along a line from the start point through the centre of the target. Thus, directional constraints are not captured by this model. Their second formulation, which had

the highest correlation with their experimental data, is expressed by:

$$ID_{min} = \log_2 \left(\frac{A}{\min(W, H)} + 1 \right)$$

where W and H are the width and height of the target. In this case the target is treated as a two dimensional object. The amplitude, however, is still considered to be a one dimensional scalar. This ID_{min} model has since been used in follow-up work [17, 20, 22], and was also proposed apparently independently by Hoffman and Sheikh [9].

A more recent and very thorough study conducted by Accot and Zhai [1] identifies various problems with the ID_w and ID_{min} formulations. The fundamental problem with the ID_w model is that it completely ignores the directional constraints. With the ID_{min} model, the directional constraint (or height) is considered, but not if it becomes greater than the width. Similarly, the width is not considered if it is greater than the height. Thus, this model does not account for data by, for example, Sheikh and Hoffman [19] that showed it is harder to acquire a square than a rectangle. Another problem with the ID_{min} model is that it ignores the angle of approach. Lastly, the ID_{min} model allows the width and height factors to be interchanged without changing the index of difficulty. In Accot and Zhai's [1] study it was shown that this last property does not hold. Figure 1 illustrates all of these limitations.




Figure 1. Limitations of ID_w and ID_{min} models: all target pairs in this figure are considered identical by the respective ID formulations. (a) No difference in ID_w for targets of different heights. (b) Width has no effect on ID_{min} if greater than height. (c) Movement angle has no effect on ID_{min} . (d) Interchanging width and height has no effect on ID_{min} .

To resolve these problems with existing formulations of ID , Accot and Zhai [1] proposed a number of properties a bivariate pointing model should contain, and developed and experimentally validated a weighted Euclidean model, which we dub ID_{WEuc} :

$$ID_{WEuc} = \log_2 \left(\sqrt{\left(\frac{A}{W} \right)^2 + \eta \left(\frac{A}{H} \right)^2} + 1 \right)$$

where η is empirically determined. This model considers $(A/W, A/H)$ to be the “constraint vector”, and they take a weighted norm of this vector and consider it to be the “appropriate distance in a two dimensional space” [1]. As can be seen, it is similar to the Euclidean norm, with the addition of the parameter η , which allows the model to weight the effect of the height differently from the effect of the width. This ID_{WEuc} model is a significant improvement over the ID_{min} model in that it alleviates the problem of the larger of

the two dimensions not being considered. However, different approach angles and different two-dimensional shapes are still not addressed by this model.

While previous work [1, 5, 9, 15] has resulted in a good understanding of bivariate pointing in 2D, the same cannot be said for pointing to trivariate targets in 3D. The only work we are aware of which attempts to model trivariate pointing in 3D is by Ware and colleagues [20, 22] who note that the ID_{min} model can be easily extended to 3D:

$$ID_{min} = \log_2 \left(\frac{A}{\min(W, H, D)} + 1 \right)$$

Unfortunately, this model suffers from the same problems as the equivalent 2D model. In addition, this model has only been applied in the context of studies that were primarily concerned with other factors, such as lag and frame rate [20] and 2D vs. 3D selection [22], and has not been validated in an explicit experiment that manipulated the three target dimensions W , H , and D . Other work on pointing, selection, and manipulation tasks in 3D environments are similarly concerned with other factors such as device resistance and the influence of muscle groups, and not on modeling pointing behavior per se. We refer the reader to [23] for a review of this literature.

GOALS AND DIRECTIONS OF THE CURRENT STUDY

Our work is ultimately motivated by the desire to build a sound foundation of theories and empirical data that can guide the development of user interfaces for emerging 3D displays. The present work is one step in this direction, and our goal here is to obtain a thorough understanding of the factors that govern pointing at trivariate targets in 3D, and to develop a predictive model that accurately characterizes such behavior. Our work extends the prior art in three key directions: a new experimental hardware platform, manipulation of experimental parameters, and modeling.

Experimental Hardware Platform

Most studies on 3D manipulation to date have relied on stereoscopic displays, either immersive VR systems [3], or non-immersive fish-tank VR systems using LCD shutter stereo-glasses [20, 22, 23]. While these displays enable users to infer depth by fusing together the two different images, the display's single image plane does not adequately support accommodation (the ability to focus the eye's lens on objects at different depths in the frontal field of view) (see [18] for a review of human stereoscopic vision). Imperfectly aligned stereo goggles can also confuse cues users obtain from their ocular muscles. Head tracking is also often used to provide motion parallax, but hindered by technical difficulties such as lag and poor accuracy. Lag can also contribute to nausea and dizziness [16]. Due to these limitations, previous studies of 3D manipulation have a potentially uncontrolled confound in that participants' perception of the depth dimension is poorer relative to their perception of the other two dimensions.

Unlike stereoscopic displays, volumetric displays [6] generate true 3D volumetric images by illuminating points in 3D space. As such, viewing such a display is akin to viewing physical objects in a real 3D scene. The human viewer uses their natural physiological mechanisms for depth perception such as true motion parallax and stereopsis through eye convergence and accommodation, resulting in a richer understanding of the virtual 3D scene. In addition, these displays can be viewed from any direction around a 360° hemispherical projection, and the user does not have to wear hardware such as shutter glasses or head-trackers. These displays, however, are in their early stages of development and have lower resolution and refresh rates than 2D displays. Despite these limitations, this is a unique platform for studying 3D pointing, since participants have the ability to reliably perceive all three dimensions. We use a volumetric display for the present study.

Manipulation of Experimental Parameters

Target dimensions (W, H, D). It has been shown in bivariate 2D pointing that H affects MT to a lesser degree than W [1]. We verify if this property holds for trivariate targets, and how the additional parameter D affects performance relative to W and H .

Accot and Zhai [1] found that increasing H reduced MT only when it was smaller than W , but increasing W reduced MT regardless of H . In other words, the W - H ratio was significant. For trivariate targets in 3D, there are three such interactions which must be examined (W - H , W - D , H - D).

Movement and approach angles. We use the following definition of the three primary axes: the X-axis is the left-right axis, the Y-axis is the up-down axis, and the Z-axis is the forward-backward axis. For the present study, we limit our exploration to physical movements in the XZ plane in 3D space (i.e., $Y=0$), and explore *movement angles* θ of 0° (i.e., along the X-axis), 90° (Z-axis), and 45° (XZ-axis). While the movement angle is the human user's axis of movement, the *approach angle* is the angle between the movement vector (defined by θ) and the axis parallel to the width of the target, as shown in Figure 2a,b. In [15], the approach angle was adjusted so that target width was not collinear with the user's movement vector. It is quite probable that this will affect MT , and the ID_W model was proposed to compensate for this. However, in altering the approach angle, the physical movement angle θ was also simultaneously altered. Evidence that θ affects MT has been seen in [2] and is partially due to differences in the muscle groups required to affect the different movements.

Because both θ and approach angle were simultaneously altered in [15], there was no way to positively determine which variable was primarily responsible for the changes in MT for the different conditions tested. To eliminate this confound, our study uses a constant approach angle of 0°.

In a trivariate target 3D pointing task, as in 2D [2], it is likely that the changes in θ will affect MT because of the different

physical movement requirements. It is also likely that the changes in θ in the XZ plane will effect the user's perception of the targets due to the targets being at different distances in the visual field. This is likely independent of the quality of the 3D display, since such an effect would occur in the "perceptually ideal" physical 3D world as well. Thus, altering θ may affect MT due to both physical and perceptual differences. We compensate for this possible perceptual difference by counterbalancing the location of targets to both ends of each movement vector.




Figure 2. (a) Both targets have approach angle of 0° but movement angles of 0° and 45° respectively. (b) Both targets have approach angle 45°. (c) Stimulus used in experiment. Targets are constrained to XZ plane and have movement angles of 0°, 45°, and 90°. W is measured along movement axis, H along the Y-axis, and D is perpendicular to both.

Modeling

We wish to develop a single model that will accurately capture the various parameters that define user performance when pointing to trivariate targets in 3D.

Despite the deficiencies inherent in the ID_{min} model, for the sake of completeness and continuity with previous work, we include it as the first baseline model in our work:

$$ID_{min} = \log_2 \left(\frac{A}{\min(W, H, D)} + 1 \right)$$

Following [1], we extend this model to accommodate the possible relative effects of different target dimensions, by assigning weights α and β to these parameters, resulting in:

$$ID_{W\tau min} = \log_2 \left(\frac{A}{\min(W, \alpha H, \beta D)} + 1 \right)$$

We compare this model to a 3D version of Accot and Zhai's [1] weighted Euclidian model:

$$ID_{WtEucl} = \log_2 \left(\sqrt{\left(\frac{A}{W}\right)^2 + \alpha \left(\frac{A}{H}\right)^2 + \beta \left(\frac{A}{D}\right)^2} + 1 \right)$$

The problem with these weighted extensions is that, like their corresponding 2D versions, they do not account for possible differences in performance due to varying movement angles. To accommodate movement angles, all components should also be weighted by an additional parameter $f_{W,H,D}(\theta)$ which takes on different empirically determined values dependent on movement angle θ . Incorporating $f(\theta)$ into the $ID_{W\tau min}$ model results in:

$$ID_{W_{\min}\theta} = \log_2 \left(\frac{A}{\min(f_W(\theta)W, f_H(\theta)H, f_D(\theta)D)} + 1 \right)$$

Note that the α and β parameters in $ID_{W_{\min}}$ simply get absorbed into the new $f_H(\theta)$ and $f_D(\theta)$ parameters in $ID_{W_{\min}\theta}$. Incorporating $f(\theta)$ in a similar manner into the $ID_{W_{\text{Euc}}}$ model results in:

$$ID_{W_{\text{Euc}}\theta} = \log_2 \left(\sqrt{f_W(\theta) \left(\frac{A}{W} \right)^2 + f_H(\theta) \left(\frac{A}{H} \right)^2 + f_D(\theta) \left(\frac{A}{D} \right)^2} + 1 \right)$$

EXPERIMENT

Apparatus

We used a 3D volumetric display developed by Actuality Systems (www.actuality-systems.com), which generates a 10" spherical 3D volumetric image by sweeping a semi-transparent 2D image plane around the Y-axis (Figure 3). A total of 198 2D images (slices) of 768x768 pixels each are displayed uniformly around the Y-axis, resulting in a total of 116 million voxels. The display's refresh rate is 24Hz. The display was driven by a 2 GHz Pentium4 computer on which the experimental software ran. The input device used was an Ascension Flock-of-birds electromagnetic 6-dof tracker equipped with a single button. The tracker controlled a 3D cursor with a direct one-to-one mapping and a control-display gain of one. The origin was defined to be the center of the volumetric display, with positive X-axis pointing right, Y pointing up, and Z pointing away from the user. Graphical objects were modeled in units relative to the radius (5") of the display volume (1 unit = 5").




Figure 3. Volumetric display

Participants

Five female and seven male paid volunteers participated in the experiment. Participants were screened for adequate stereo vision using the StereoOptical RANDDOT stereopsis test. Participants ranged in ages from 20 to 25, were all right-handed, and controlled the tracking device and consequently the cursor with their right hand.

Procedure

The task was reciprocal 3D target acquisition, which required participants to point to two targets back and forth in

succession. The targets were rendered as wireframe cuboids, equidistant from the centre of the display in opposite directions along the given axis of movement. This effectively counterbalanced any differences in perception of targets in different parts of the display. All targets were constrained to the XZ plane (i.e., $Y = 0$). The target to be selected was yellow, and the other red. When participants clicked on the yellow target, the targets would swap colors, as an indication that the participant had to now move to and select the other target. The 3D cursor that the participant controlled with the handheld tracker was displayed as a crosshair with short line segments along the X, Y, and Z directions all intersecting at its origin.

Design

A repeated measures within-participant factorial design was used. The independent variables were amplitude A (0.44, 0.88, 1.28 units), target height H , width W , depth D (all 0.04, 0.08, 0.16, 0.32 units), and movement angle θ (0° , 90° , 45°). A fully crossed design resulted in a total of 576 combinations of A , H , W , D , and θ .

We use a target-centric definition of W , H , and D as illustrated in Figure 2c, where H is always measured along the Y-axis, W is the dimension of the target along the direction of movement, and D is perpendicular to W and H .

The experiment was performed in three sessions, each occurring on separate days. In each session, participants would complete trials for all 192 H, W, D, A permutations for one of the three movement angles. For each of the H, W, D, A permutations, presented in random order in the session, participants performed a trial set consisting of seven yellow target selections (i.e., six reciprocal movements between the two targets). Because the first selection was used to signal the beginning of the trial set, it required an accurate selection – the 3D cursor needed to be positioned inside the target. Subsequent selections did not enforce this accuracy, but an audible buzzing sound provided feedback when an error was made. After each trial set, the display informed participants of their error rate in the immediately completed set, their cumulative error rate for that session, and the number of trial sets remaining. Participants were asked to minimize their errors. Participants could take breaks between trial sets, but not within each trial set.

Before each session participants were given two practice trial sets to familiarize themselves with selection for the given movement angle.

Participants were randomly divided into 6 groups of 2 each. Assignment of movement angle to groups on each day was counterbalanced using a balanced latin square. Each session lasted approximately one hour.

Performance Measures

The dependent variables were movement time MT – defined as the time between clicks in a trial set, and error rate – defined as the average number of errors per trial. Errors

occurred when participants clicked when the cursor was outside the target.

Results

Outliers were removed based on *MT* and accuracy – defined as distance between the click and the target center. Any data point further than 2 standard deviations away from its condition's mean (by *MT*, or by accuracy) was removed. 7.6% of the data were removed as outliers.

Movement Time Analysis

Main effects:

Analysis of variance showed that the independent variables *W* ($F_{3,105} = 2527, p < .0001$), *H* ($F_{3,105} = 715, p < .0001$), *D* ($F_{3,105} = 1577, p < .0001$), and Θ ($F_{2,22} = 892, p < .0001$) all had a significant main effect on *MT*. Of particular interest is the effect of Θ . Multiple means comparisons showed no significant difference in *MT* for $\Theta = 0^\circ$ and 45° , but a significantly higher *MT* for $\Theta = 90^\circ$. This effect is most likely due to the different muscle groups utilized, as movement for $\Theta = 0^\circ$ and 45° requires limited use of the shoulder muscles, while movement for $\Theta = 90^\circ$ is achieved primarily with the shoulder. This is consistent with the findings of previous studies [2, 4, 12] which suggest that the use of smaller muscle groups (hand, wrist) will result in better performance in a pointing task than larger muscle groups (upper arm, shoulder). This effect clearly indicates that movement angle must be accounted for in any attempt to model pointing behavior in 3D.

Effect of movement direction:

Recall that for each Θ we tested ($0^\circ, 45^\circ, 90^\circ$), two targets were placed equidistant from the origin at opposite directions along the movement axis defined by Θ on the XZ plane. Our reciprocal pointing task required back-and-forth movements in both directions along these three angles. Ideally, one would expect that two equally sized targets on either end of any of these vectors would be equally difficult to select. However, there is a possibility that targets displayed further away from the user in the depth axis are harder to perceive and therefore select. Analysis of variance, however, showed that movement direction did not have a significant effect on *MT* ($F_{1,35} = 0.51, p = .4764$), as Figure 4a illustrates. This implies that perception of targets was equally good in all locations.

Interactions between Θ and target dimensions:

The effect of Θ on *MT* was dependent on *W* and *D* as indicated by the interactions $W \times \Theta$ ($F_{6,121} = 17.4, p < .0001$), and $D \times \Theta$ ($F_{6,121} = 22.5, p < .0001$). However, there was no significant $H \times \Theta$ interaction. This is an important result that likely occurs because targets were located in the XZ plane and the perception of the height of the target remains constant as Θ changes in this plane. In contrast, the perception of *W* is likely best with $\Theta = 0^\circ$ and worst at $\Theta =$

90° , while the perception of *D* is best with $\Theta = 90^\circ$ and worst at $\Theta = 0^\circ$. Figure 4b,c shows these interactions.

Relative effect of target dimensions:

One goal of our study was to investigate if any particular target affected *MT* more than the others. Multiple means comparisons indicated that regardless of the values of *H* and *D*, *MT* significantly increased when *W* is increased from 0.04 to 0.08, from 0.08 to 0.16, and from 0.16 to 0.32. For *H* and *D*, however, this was only true until the parameter reached a size of 0.16; the increase from 0.16 to 0.32 did not significantly affect *MT*. This is illustrated by the slopes of the lines in Figure 4e, where it is clear that changes in *W* has the most impact on *MT*. This result is consistent with the same effect found in 2D pointing [1].

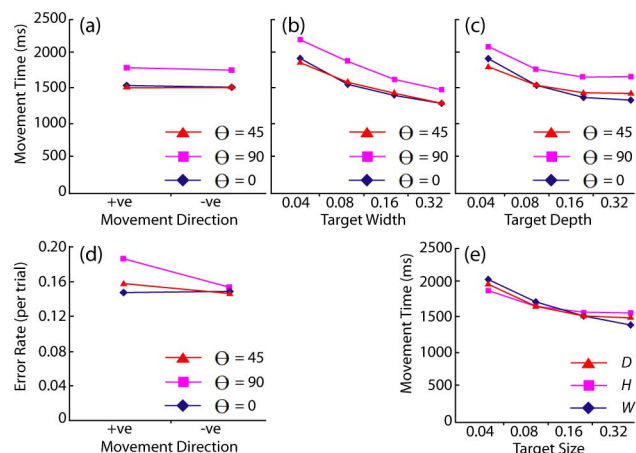


Figure 4. (a) Effect of movement direction on *MT*, by movement angle. (b) Interaction between *W* and Θ . (c) Interaction between *D* and Θ . (d) Effect of movement direction on error rate, by movement angle. (e) Relative effect of target dimensions on *MT*.

Interactions between target dimensions:

Also of interest is whether changes in any one of *W*, *H*, and *D* would have differing effects on *MT* depending on the size of the other dimensions. Indeed, our analysis showed significant interaction effects for all combinations of dimensions: $H \times W$ ($F_{9,525} = 38.6, p < .0001$), $H \times D$ ($F_{9,525} = 17.4, p < .0001$), $W \times D$ ($F_{9,525} = 42.7, p < .0001$).

An alternative way of looking at this issue is to consider the effect on performance of the ratios between these target dimensions. Figure 5 shows how the *W/H*, *W/D*, and *H/D* ratios affect *MT*, broken down by movement angle and the three different amplitudes tested. Note that the ratios 0.125 and 8 represent only single data pairs (e.g., *W/H* ratio of 8 is generated only when *W*=8 and *H*=1), whereas the ratios 0.25, 0.5, 1, 2, 4 represent double data pairs (e.g., *W/H* ratio of 4 is generated for *W*=8, *H*=2 and *W*=4, *H*=1). As such, to keep the analysis balanced by number of data points, we exclude the 0.125 and 8 ratios from this analysis.

W/H ratio: There was a significant effect on *MT* ($F_{4,140} = 167, p < .0001$), and significant interactions with amplitude *A*

($F_{8,490} = 7.1, p < .0001$), movement angle Θ ($F_{8,154} = 5.9, p < .0001$), and depth D ($F_{12,665} = 2.4, p < .0001$).

For $\Theta = 0^\circ$ (Figure 5a), increasing W/H from 1 to 2 had a significant effect on MT for all distances. However, increasing it from 2 to 4 significantly affects MT only for the shortest amplitude $A = 0.48$. Looking at the left half of the graph, the only significant effects are when W/H decreased from 1 to 0.5 for the furthest amplitude $A = 1.28$, and when it is decreased from 1 to 0.25 at $A = 0.88$ (all $p < .01$). Examining the symmetry about unity shows that at $A = 0.48$, MT is significantly lower for $W/H = 2, 4$ than $W/H = 0.5, 0.25$. This corresponds to the effect observed by Accot and Zhai [1]. However, when $A = 0.88$ and 1.28 , MT for corresponding ratios (0.25 and 4, 0.5 and 2) are not significantly different. Thus, we see that at high amplitudes, increasing H , even when W is the constraining factor, can still significantly reduce MT . While this is in contrast to Accot and Zhai [1]'s results, they did not break down their results by amplitude, and it is therefore possible that their data exhibited a similar trend for large amplitudes but was simply not reported. As Figures 5b,c illustrate, the results for $\Theta = 90^\circ$, and $\Theta = 45^\circ$ show similar trends to $\Theta = 0^\circ$.

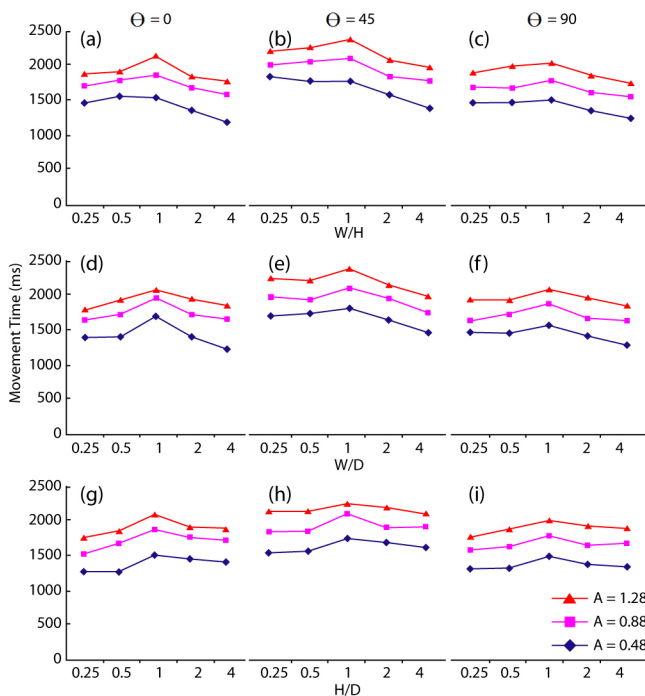


Figure 5. (a-c) Effect of W/H ratio on MT , by Θ and amplitude. (d-f) W/D ratio. (g-i) H/D ratio.

W/D ratio: There was a significant effect on MT ($F_{4,140} = 167, p < .0001$), and significant interactions with A ($F_{8,490} = 2.9, p < .0001$) and Θ ($F_{8,154} = 8.3, p < .0001$). There was no significant interaction with H ($F_{12,665} = 1.2, p = .2774$). Figures 5d-f illustrate the trends. The most interesting trend is when $\Theta = 0^\circ$ (Figure 5d), where the curves are very symmetric about unity. For the two larger amplitudes ($A = 0.88, 1.28$), the symmetric pairs ($W/D = 0.5$ and 2, 0.25 and 4) do not result in significantly different MT , indicating that

increasing D had the same effect as increasing W . However, for $A = 0.48, W/D = 4$ results in a significant lower MT than its counterpart ratio $W/D = 0.25$. ($p < .001$ in all cases). This result is not observed for the other two movement angles, as might be expected from the interactions between Θ and D which indicated that increasing D reduces MT most for $\Theta = 0^\circ$ and less for $\Theta = 45^\circ$ and 90° (Figure 4c).

H/D ratio: There was a significant effect on MT ($F_{4,140} = 111, p < .01$), and significant interactions with A ($F_{8,490} = 2.0, p < .01$), Θ ($F_{8,154} = 4.0, p < .0001$), and W ($F_{12,665} = 3.2, p < .01$) (Figures 5g-i). Note that symmetry about unity were observed in both $\Theta = 90^\circ$ and 45° : for all A , MT for the symmetric pairs ($H/D = 0.5$ and 2, 0.25 and 4) are not significantly different. However, this isn't true for $\Theta = 0^\circ$, where increasing D reduces MT more than increasing H due to Θ and target dimension interactions (Figure 4c).

Fit of the models

We fit the MT data to the five candidate models described earlier, using a least-squares method (Table 1). Where appropriate, models were fitted by Θ , allowing us to analyze results independent of how Θ effects the time. The table shows parameter estimates where applicable and the corresponding standard errors for those estimates. The last column provides the R^2 values for the regression.

The poor correlation of the ID_{min} model supports our earlier argument as to its shortcomings. Adding weights to the target dimensions improves the fit, as does including $f(\Theta)$. The ID_{WtEuc} and $ID_{WtEuc\Theta}$ models, however, outperforms all the ID_{min} model variants. The best fit overall, with a correlation of 0.912, was with the $ID_{WtEuc\Theta}$ model where $f(\Theta)$ was calculated independently for each target dimension (note that $f(\Theta)$ was not calculated for the H parameter due to the lack of a significant $H \times \Theta$ interaction as discussed earlier), and is hence our preferred model:

$$MT \approx 56 + 508 \log_2 \left(\left(\sqrt{f_w(\theta) \left(\frac{A}{W}\right)^2 + \frac{1}{9.2} \left(\frac{A}{H}\right)^2} + f_D(\theta) \left(\frac{A}{D}\right)^2 + 1 \right) \right)$$

with $f_w(0^\circ) = 0.211, f_w(90^\circ) = 0.717, f_w(45^\circ) = 0.242, f_D(0^\circ) = 0.194, f_D(90^\circ) = 0.312, \text{ and } f_D(45^\circ) = 0.147$.

Figure 6 illustrates the differences between $ID_{WtEuc\Theta}$ and $ID_{WtMin\Theta}$. Note how the points in the Euclidian model are less spread out, especially as ID increases.




Figure 6. Raw data for (a) $ID_{WtEuc\Theta}$ and (b) $ID_{WtMin\Theta}$

Model	Parameter Estimates (with std error indicated below)											R^2
	θ	a (ms)	b (ms/bit)	α	β	$f_W(0^\circ)$	$f_D(0^\circ)$	$f_W(90^\circ)$	$f_D(90^\circ)$	$f_W(45^\circ)$	$f_D(45^\circ)$	
ID_{min}	0°	-187 69.6	454 17.4									0.781
	90°	78 68.9	451 17.2									0.783
	45°	-16 53.2	411 13.3									0.833
ID_{Wmin}	0°	-163 53.4	472 14.1	1.656 0.087	1.162 0.052							0.855
	90°	118 45.7	477 12.3	1.751 0.081	1.394 0.056							0.887
	45°	11 40.5	428 10.7	1.515 0.065	1.278 0.048							0.894
ID_{WEuc}	0°	-395 48.7	496 11.9	0.372 0.051	0.905 0.096							0.901
	90°	-23 43.4	494 11.2	0.259 0.033	0.421 0.045							0.910
	45°	-149 36.8	445 9.21	0.412 0.044	0.582 0.056							0.924
$ID_{Wmin\theta}$	All	79 25.5	467 7.26	2.128 0.757		1.440 0.517	1.670 0.597	0.851 0.309	1.236 0.445	1.442 0.518	1.821 0.651	0.878
$ID_{WEuc\theta}$	All	56 21.6	508 6.58	0.109 0.063		0.211 0.123	0.194 0.114	0.717 0.426	0.312 0.148	0.242 0.142	0.147 0.044	0.912

Table 1. Summary of model fitting results. $f_H(\theta)$ was not calculated since there was no significant $H \times \theta$ interaction. α and β are calculated instead of $f_W(\theta)$ and $f_D(\theta)$ for ID_{Wmin} and ID_{WEuc} when presented by $\theta = 0^\circ, 45^\circ, 90^\circ$, because the breakdown of data by θ makes it unnecessary to compute $f(\theta)$ for each θ as is done when data is aggregated across all θ in $ID_{Wmin\theta}$ and $ID_{WEuc\theta}$.

Error Analysis

Error rate was significantly affected by H ($f_{3,105} = 206.2, p < .0001$), W ($f_{3,105} = 467.9, p < .0001$), D ($f_{3,105} = 246.9, p < .0001$), and θ ($f_{2,70} = 13.47, p < .0001$). There was no significant effect for A ($f_{2,70} = 2.46, p = .086$). The overall mean error rate was 15.7%. While this is higher than the error rate of 4% seen in typical Fitts' Law experiments, our observations of the participants rules out the possibility of subjects "racing through the experiment". Rather, our analysis indicates that when any of the target dimensions W , H , or D takes on the smallest value of 0.04 units, the error rate goes up significantly ($p < .0001$). When this smallest dimension is not considered, the error rate decreases to an average of 5.4%, which is in line with results reported in the literature. We also believe that secondary sources of error were the greater perceptual and motor difficulty of the 3D task, and noise in our 3D tracker.

Finally, we looked at how the direction of movement affected error rate. Recall that for all three movement angles tested, the direction of movement did not have a significant effect on MT . With error rate, however, there was a significant interaction between movement direction and θ ($f_{2,175} = 7.12, p < .0001$), and multiple means comparisons indicated that the error rate for -ve direction movements when $\theta = 90^\circ$ was significantly higher than for $\theta = 0^\circ$ or $\theta = 45^\circ$ (Figure 4d). This effect is likely due slight difficulty in perceiving targets at the far back of the display. However, recall that we did not find a similar effect on MT (Figure 4a). This indicates that users may not actually think they have perception difficulties with these targets, believing they are in the target and thus completing the task without compromising MT but unwittingly making an erroneous selection.

IMPLICATIONS for USER INTERFACE DESIGN

A primary motivation of our work was to provide empirical results and practical models that could guide the development of user interfaces for 3D displays. Our study investigated how target dimensions and movement angles affect selection performance, and the results can provide us with significant guidelines about the layout of selectable targets, i.e. how items such as widgets, menus, and other objects should be sized and positioned in 3D displays.

Target sizing: As in any interface, widgets should be large enough so that they can be easily selected and manipulated, but their size should be minimized to reduce the amount of visual space which they take up. An observation from our experiment is that as error rates significantly spiked after they were reduced to a certain size (1/5 of an inch), so clearly targets should not be made this small. More interesting is that designers should avoid targets with grossly disproportionate dimensions, since our results show that having one dimension more than four times as large as another does not make it any easier for the user to select the target. Another important finding was that the dimension parallel to the line of approach (defined as W in our experiment) was always the most critical one. Thus if a target was at the back of the display, it would be beneficial for it to have a larger size in the Z-axis, since the user would most likely be moving forward to select it. Similarly if were at the side of the display, a larger size along the X-axis would be beneficial.

Target positioning: In typical 2D interfaces, the user's workspace is in the center of the screen taking up most of the display space, while the interface widgets border the screen. Extending this layout to a 3D display could have the workspace in the center volume of the display, and the interface elements on its outside surface. However, our

experiment showed that moving forwards and backwards to select targets was significantly slower than moving left and right. If targets were to be positioned on the outside surface, they should be limited to the left and right sides of the surface, maximizing left and right selections.

Recall that our results, which can guide both target sizing and positioning, were dependant on the location of the user's viewpoint. This presents a challenge for interface designers, since one benefit of volumetric displays is that users can walk around it for full 360° viewing. It would therefore be quite useful if a 3D volumetric display system were able to track the position of its users and then adjust the layout of the interface elements accordingly.

CONCLUSION

We have presented experimental work that investigated how target dimensions, movement angles, and their interactions effect selection performance in a 3D environment. Unlike most experiments involving 3D tasks, ours was performed on a true 3D volumetric display which provided users with reliable perception of all three spatial dimensions. We observed that moving forwards and backwards in depth is slower than moving left and right for selecting targets; that the target width was more critical than the height and depth of the target; and that the effect of the width and depth were dependant on the movement angle, while the effect of the height was constant regardless of the movement angle. Building upon previous work on univariate and bivariate pointing tasks, we introduced and validated a variant of Fitts' law that models pointing to trivariate targets in 3D. Unlike the previous models, our model also accounts for varying movement angles. We have shown that performance time can be predicted from this model to a greater extent than alternative models which have been previously suggested. Finally, we have discussed the implications of our results for the sizing and positioning of user interface widgets in three dimensional displays.

ACKNOWLEDGEMENTS

We thank members of the Dynamic Graphics Project lab at the University of Toronto, Actuality Systems, and all those who participated in our experiment.

REFERENCES

- Accot, J., & Zhai, S. (2003). Refining Fitts' law models for bivariate pointing. *ACM CHI*. p. 193-200.
- Boritz, J., *et al.* (1991). Fitts' law studies of directional mouse movement. *Graphics Interface*. p. 216-223.
- Buxton, W., & Fitzmaurice, G. (1998). HMD's, Caves, & Chameleon: A human-centric analysis of interaction in virtual space. *Computer Graphics*. p. 64-68.
- Card, S., *et al.* (1991). A morphological analysis of the design space of input devices. *ACM TOIS*. 9. p. 99-122.
- Crossman, E. (1956). The measurement of perceptual load in manual operations. PhD Thesis. U. Birmingham.
- Ebert, D., *et al.* (1999). Realizing 3D visualization using crossed-beam volumetric displays. *Communications of the ACM*. 42(8). p. 101-107.
- Fitts, P. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*. 47. p. 381-391.
- Hinckley, K., *et al.* (1997). Usability analysis of 3D rotation techniques. *ACM UIST*. p. 1-10.
- Hoffmann, E., & Sheikh, I. (1994). Effect of varying target height in a Fitts movement task. *Ergonomics*. 37(6). p. 1071-1088.
- Actuality Systems Inc. www.actuality-systems.com
- Dimension Technologies Inc. www.dti3d.com
- Langolf, G., *et al.* (1976). An investigation of Fitts' law using a wide range of movement amplitudes. *Journal of Motor Behaviour*. 8. p. 113-128.
- Lucente, M. (1997). Interactive three-dimensional holographic displays: seeing the future in depth. *Computer Graphics*. May 1997.
- MacKenzie, S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction*. 7. p. 91-139.
- MacKenzie, S., & Buxton, W. (1992). Extending Fitts' law to two-dimensional tasks. *ACM CHI*. p. 219-226.
- McCauley, M., & Sharkey, T. (1992). Cybersickness: Perception of self-motion in virtual environments. *Presence*. 1(3). p. 311-318.
- Murata, A. (1999). Extending effective target width in Fitts' law to a two-dimensional pointing task. *International Journal of HCI*. 11(2). p. 137-152.
- Patterson, R., & Martin, W. Human Stereopsis. *Human Factors*. 34(6). p. 669-692.
- Sheikh, I., & Hoffmann, E. (1994). Effect of target shape on movement time in a Fitts task. *Ergonomics*. 37(9). p. 1533-1548.
- Ware, C., & Balakrishnan, R. (1994). Reaching for objects in VR displays: Lag and frame rate. *ACM TOCHI*. 1(4). p. 331-356.
- Ware, C., & Jessome, D. (1988). Using the bat: a six dimensional mouse for object placement. *IEEE Computer Graphics and Applications*. 8. p. 65-70.
- Ware, C., & Lowther, K. (1997). Selection using a one-eyed cursor in a fish-tank VR environment. *ACM TOCHI*. 4(4). p. 309-322.
- Zhai, S. (1995). Human performance in six-dof input control. PhD Thesis. University of Toronto.