

Prediction of user action in moving-target selection tasks

by

Juan Sebastián Casallas

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Major: Human Computer Interaction

Program of Study Committee:
James H. Oliver, Major Professor
Jonathan W. Kelly
Jin Tian
Stephen B. Gilbert
Michael C. Dorneich

Iowa State University
Ames, Iowa
2015

Copyright © Juan Sebastián Casallas, 2015. All rights reserved.

STATEMENT OF INTERNATIONAL COTUTELLE

The work described in this dissertation was jointly supervised by James H. Oliver, Professor at Iowa State University, Frédéric Mérienne, Professor at Arts et Métiers ParisTech, and Samir Garbaya, Maître de Conférences at Arts et Métiers ParisTech. In accordance with French regulations, this dissertation was reviewed by Vincent Hugel, Professor at Université de Toulon, and Samir Otmane, Professor at Université d'Evry Val d'Essonne. Drs. Mérienne, Garbaya, and Hugel were also part of the international examining board for the defense of this doctoral dissertation, together with the Iowa State University Program of Study Committee.

À Catherine.
Para Gaël Elías e Inés Sofía.

...all models are approximations. Essentially all models are wrong, but some are useful. However, the approximate nature of the model must always be borne in mind. [Box and Draper 1987, p. 424]

TABLE OF CONTENTS

LIST OF TABLES	ix
LIST OF FIGURES	x
ACKNOWLEDGMENTS	xiii
ABSTRACT	xiv
CHAPTER 1. INTRODUCTION	1
1.1 Moving-target selection	1
1.1.1 Moving-target selection challenges	1
1.1.2 Moving-target selection enhancements	2
1.1.3 Control strategies in moving-target selection	3
1.2 Research Layout	5
CHAPTER 2. LITERATURE REVIEW	6
2.1 Selection performance	6
2.1.1 The Hick-Hyman Law	6
2.1.2 Fitts' Law	6
2.1.2.1 Two-part formulations of Fitts' Law and gain	7
2.1.2.2 Fitts' Law formulations for 2-D static-target pointing tasks . .	9
2.1.2.3 Fitts' Law formulations for 3-D static-target pointing tasks . .	10
2.1.2.4 Fitts' Law formulations for 1-D moving-target pointing tasks .	11
2.1.2.5 Fitts' Law formulations for 2-D moving-target pointing tasks .	13
2.2 Subjective difficulty	14
2.2.1 The importance of the rating method	14
2.2.2 Subjective difficulty in moving-target selection	14
2.3 Prediction of intention	15
2.3.1 Scoring functions	15
2.3.2 Gaze	16
CHAPTER 3. INTENTION IN UNDIRECTED 3-D MOVING-TARGET SELECTION	18
3.1 Hypotheses	19
3.1.1 Task-specific features for intention prediction in undirected moving- target selection	19

3.1.2	Generalizable features for intention prediction in undirected moving-target selection	19
3.2	Pilot study	19
3.3	Methods	19
3.3.1	Participants	19
3.3.2	Apparatus	20
3.3.3	Procedure	20
3.3.4	Design	22
3.3.5	Data integrity	23
3.3.6	Predictive methods	23
3.3.6.1	Model comparison	24
3.4	Task-specific feature analysis	24
3.4.1	Results	25
3.5	Generalizable feature analysis	26
3.5.1	Relative user-target features	27
3.5.1.1	Distance score feature	27
3.5.2	Time-window selection	28
3.5.3	Evaluation	30
3.5.4	Results and discussion	30
3.5.4.1	Generalizable user-target features	30
3.5.5	Combined task-specific and generalizable features	31
3.6	Discussion	33
3.6.1	Task-specific features for intention prediction in undirected moving-target selection	33
3.6.2	Generalizable features for intention prediction in undirected moving-target selection	34
3.7	Conclusion	35
CHAPTER 4. PROSPECTIVE DIFFICULTY OF 2-D STATIC-TARGET AND MOVING-TARGET SELECTION		36
4.1	Formulations	37
4.1.1	<i>PD</i> formulation for 2-D static-target pointing tasks	37
4.1.2	<i>PD</i> formulation for 1-D moving-target pointing tasks	37
4.1.3	<i>PD</i> formulation for 2-D moving-target pointing tasks	37
4.2	Modeling	38
4.2.1	Choice of <i>ID</i> formulation	38
4.2.2	Formulations of prospective difficulty in 2-D static-target pointing tasks	39
4.2.3	Formulations of prospective difficulty in 1-D moving-target pointing tasks	41
4.2.4	Formulations of prospective difficulty in 2-D moving-target pointing tasks	42

4.3	Methods	43
4.3.1	Apparatus	43
4.3.2	Participants	43
4.3.3	Procedure	43
4.3.4	Design	45
4.3.5	Data integrity	46
4.3.6	Statistical methods	46
4.3.6.1	Models for inferential statistics	46
4.3.6.2	Models for Regression Analysis	47
4.3.7	Hypotheses	48
4.3.7.1	Static-target selection	48
4.3.7.2	Moving-target selection	48
4.4	Analysis	48
4.4.1	Static-target block	48
4.4.1.1	Inferential statistics	48
4.4.1.2	Regression analysis	50
4.4.1.3	Participant performance and self assessment	52
4.4.2	Moving-target block	52
4.4.2.1	Inferential statistics	52
4.4.2.2	Regression analysis	55
4.4.2.3	Participant performance and self assessment	62
4.5	Discussion	62
4.5.1	Prospective difficulty of static-target selection	63
4.5.2	Prospective difficulty of moving-target selection	64
4.6	Conclusion	65
CHAPTER 5. PERFORMANCE AND INTENTION IN DIRECTED 3-D MOVING-TARGET SELECTION		67
5.1	Extending the 2-D formulae to 3-D	67
5.2	Subjective difficulty and performance	69
5.3	Predicting user intention in directed tasks	70
5.4	Hypotheses	70
5.4.1	Moving-target selection performance in 3-D	70
5.4.2	Predicting user intention in 3-D moving-target directed-selection tasks .	71
5.5	Methods	71
5.5.1	Participants	71
5.5.2	Apparatus	71
5.5.3	Procedure	71
5.5.4	Design	72
5.6	General results	73

5.6.1	Experimental issues	73
5.6.1.1	Technical issues	74
5.6.1.2	Participant-wellness issues	74
5.6.1.3	Simulator sickness questionnaire	75
5.7	Performance analysis	76
5.7.1	Exploratory data analysis	76
5.7.1.1	Between-block trial performance	76
5.7.1.2	Effects on <i>MT</i> for the successful trials in the 1-sphere block . .	77
5.7.2	Regression analysis	81
5.7.2.1	Change of coefficients per <i>V</i>	82
5.8	Predictive analysis	84
5.8.1	Three-sphere trials	85
5.8.1.1	Wand positions at the highlight frame	85
5.8.1.2	Predictive accuracy	86
5.8.2	Six-sphere trials	86
5.8.2.1	Wand positions at the highlight frame	86
5.8.2.2	Predictive accuracy	86
5.9	Discussion	90
5.9.1	Moving-target selection performance in 3-D	90
5.9.2	Predicting user intention in 3-D moving-target directed-selection tasks .	90
5.10	Conclusion	91
CHAPTER 6.	CONCLUSION	92
6.1	Limitations and future work	93
6.1.1	Intention prediction models	93
6.1.2	<i>PD</i> prediction models	93
6.1.2.1	Two-dimensional static-target tasks	94
6.1.2.2	Moving-target tasks	94
6.1.2.3	Methods for assessing <i>PD</i>	95
6.1.3	Performance prediction models	95
6.1.3.1	Target velocity	95
6.1.3.2	Target distance	95
6.1.3.3	Additional measures of performance	95
6.1.4	General issues	96
6.1.4.1	Target shape	96
6.1.4.2	Target acceleration	96
BIBLIOGRAPHY	97

LIST OF TABLES

Table 3.1	Accuracy and 95% confidence intervals for the evaluated feature-sets .	25
Table 3.2	Tree size, number of leaves, accuracy and 95% confidence intervals for the evaluated generalizable moving-target feature-sets	30
Table 3.3	Accuracy difference and 95% confidence intervals for the evaluated generalizable moving-target feature-sets. Asterisks (*) denote a significant difference ($\alpha = 0.05$).	31
Table 3.4	Tree size, number of leaves, accuracy and 95% confidence intervals for the evaluated target-based feature-sets.	33
Table 3.5	Accuracy difference and 95% confidence intervals for the target-based feature-sets. Asterisks (*) denote a significant difference ($\alpha = 0.05$), dots (.) denote a marginal difference ($\alpha = 0.1$).	33
Table 4.1	Regression estimates for $\overline{PD} = a + b ID$	51
Table 4.2	Regression estimates for $\overline{PD} = a_\theta + b ID$	52
Table 4.3	Regression estimates for $\overline{PD} = a_V + b_V \sqrt{D} + c_V ID$	57
Table 4.4	Least-squares estimates for the regressions of \hat{a}_V , \hat{b}_V , and \hat{c}_V , from Table 4.3, on V	57
Table 4.5	Hierarchical regression $\overline{PD}_i \sim \mathcal{N}(a_{V[i]} + b_{V[i]} \sqrt{D_i} + c_{V[i]} ID_i, \sigma^2)$	58
Table 4.6	Regression estimates for $\overline{PD} = a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m$	59
Table 4.7	Regression estimates for $\overline{PD} = a_V + b \sqrt{D_s} + c_V \sqrt{D_m} + d ID_m$	60
Table 4.8	Least-squares estimates for the regressions of \hat{a}_V and \hat{c}_V , from Table 4.7, on V	62
Table 4.9	Hierarchical regression $\overline{PD}_i \sim \mathcal{N}(a_{V[i]} + b \sqrt{D_{si}} + c_{V[i]} \sqrt{D_{mi}} + d ID_{mi}, \sigma^2)$	63
Table 5.1	Regression estimates for $\overline{PD} = a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m$	83
Table 5.2	Least-squares estimates for the regressions of the point estimates from Table 5.1— $\hat{a}_V = Mdn(a_V)$, $\hat{b}_V = Mdn(b_V)$, $\hat{c}_V = Mdn(c_V)$, and $\hat{d}_V = Mdn(d_V)$ —on target speed V	83
Table 5.3	Hierarchical regression $\overline{MT}_i \sim \mathcal{N}(a_{V[i]} + b_{V[i]} \sqrt{D_{si}} + c_{V[i]} \sqrt{D_{mi}} + d_{V[i]} ID_{mi}, \sigma^2)$	84

LIST OF FIGURES

Figure 2.1	Two-dimensional task with a bivariate target. The red disc indicates the cursor; the green rectangle represents a target with width W , and height H . D is the cursor–target distance, and θ is the target angle.	9
Figure 2.2	Three-dimensional trivariate target. The red disc indicates the cursor; the green rectangular cuboid represents a target with width W , height H , and depth D . D is the cursor–target distance, α is the z – x azimuth angle, and θ is the altitude angle.	11
Figure 3.1	Experimental setup of the pilot study.	20
Figure 3.2	Experimental setup with an array two spheres	21
Figure 3.3	Possible row positions— <i>left</i> , <i>center</i> and <i>right</i> —with respect to the user in the two-sphere block	22
Figure 3.4	Decision tree for feature-set 4, suggesting that participants based their decisions only on sphere size, with a preference for the right sphere. Leaves represent prediction outcomes (sph_1 is the left sphere, and sph_2 is the right sphere), while the other nodes represent tested attributes (r_1 or r_2). The numbers in parenthesis within the leaves represent the total number of instances that fall into that leaf, over the number of incorrectly predicted instances among these instances.	26
Figure 3.5	Δdot vs. time. Each line corresponds to a trial, colored according to the selected sphere. The plot has been trimmed to the 5 th percentile of the selection times (2.35 s).	29
Figure 3.6	Generated decision tree for feature-set $\{\overline{\Delta dot}, \overline{\Delta D}\}$. The numbers in parenthesis within the leaves represent the total number of instances that fall into that leaf, over the number of incorrectly predicted instances among these instances.	32
Figure 4.1	Measurements relevant to the target-selection tasks considered in this chapter	36
Figure 4.2	Two-dimensional moving-target model	42
Figure 4.3	Screenshot of a moving-target question with $\phi = 45^\circ$	44

Figure 4.4	Static-target task	45
Figure 4.5	Moving-target task	45
Figure 4.6	Main-effect coefficient plot for the fitted ordered probit model on static-target PD ratings	49
Figure 4.7	Two-way interaction coefficient plot for the fitted ordered probit model on static-target PD ratings	50
Figure 4.8	Regression for $\overline{PD} = a + b ID$	51
Figure 4.9	Regression for $\overline{PD} = a_\theta + b ID$	52
Figure 4.10	Main-effect coefficient plot for the fitted ordered probit model on moving-target PD ratings	53
Figure 4.11	Two-way interaction coefficient plot for the fitted ordered probit model on moving-target PD ratings	54
Figure 4.12	Regression for $\overline{PD} = a_V + b_V \sqrt{D} + c_V ID$	56
Figure 4.13	Regression for $\overline{PD} = a_V + b \sqrt{D_s} + c_V \sqrt{D_m} + d ID_m$	61
Figure 5.1	Three-dimensional moving-target model	68
Figure 5.2	Identical model parameters V , D_m , and D_s for three different γ rotations of \mathbf{V} around \mathbf{D}	68
Figure 5.3	Experimental setup with six spheres. Left, the sphere starting positions, middle the spheres approximately at their controlled positions, right the spheres after the goal sphere gets highlighted. The green sphere represents sph_0	74
Figure 5.4	Distributions of the pre- and post-experiment assessment scores of the ssq by scale. Lower and upper lines of the boxes represent the first and third quantiles, their distance called the inter-quantile range (IQR), thick box lines represent the median values, upper and lower whiskers represent values that are within 1.5 IQR of the box hinges. Points represent values that are greater than the third quantile plus 1.5 IQR.	75
Figure 5.5	Distribution of participant success rates (left), and successful trial completion times (right) per experimental block	76
Figure 5.6	Distributions of movement times for the successful trials per experimental factor in the 1-sphere block	78
Figure 5.7	Movement times per γ angle for the successful trials in the 1-sphere block.	79
Figure 5.8	Front and right views of the distribution of wand positions at the frame where the sphere was highlighted in the 1-sphere trials	80

Figure 5.9	Movement times for each of the distances between the target sphere and the wand at the highlight frame $ P_w - P $, for the successful trials in the 1-sphere block. The top bars show the distribution of $ P_w - P $ binned every 0.025 m. The orange lines represent the distance between the target sphere and $P_{w,0} \pm 0.05$ m.	81
Figure 5.10	Residuals $(\overline{MT} - \widehat{MT})$ vs. fitted values (\widehat{MT}) for the $\overline{MT} \sim \mathcal{N}(a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m, \sigma^2)$ regression. The dotted lines represent $\pm \hat{\sigma}$	82
Figure 5.11	Front and right views of the distribution of wand positions at the frame where the goal sphere was highlighted in the 3-sphere trials	85
Figure 5.12	Three-sphere accuracy comparison for the $\widehat{T}Score$, and $dScore$ predictors for trial percentages $0.1 T, 0.2 T, \dots, 0.9 T, T$, using different N , and $decay$ parameters. The dotted lines represent the accuracy given by chance, i.e. $1/3$	87
Figure 5.13	Front and right views of the distribution of wand positions at the frame where the goal sphere was highlighted in the 6-sphere trials	88
Figure 5.14	Three-sphere accuracy comparison for the $\widehat{T}Score$, and $dScore$ predictors for trial percentages $0.1 T, 0.2 T, \dots, 0.9 T, T$, using different N , and $decay$ parameters. The dotted lines represent the accuracy given by chance, i.e. $1/3$	89

ACKNOWLEDGMENTS

I wish to express my gratitude to my advisers, James Oliver, Frédéric Mérienne, and Samir Garbaya, for their trust and guidance throughout this joint PhD program. I am especially thankful to Dr. Oliver who met with me well-nigh every week during these years, and was always supportive and enthusiastic about my research, writing, and personal accomplishments.

I am also grateful to my international committee for having shared their expertise in their domains throughout my doctoral studies. In particular, I want to thank Michael Dorneich who mentored me in academic life in general, and Jonathan Kelly who also actively collaborated with me and mentored me in research. I also want to thank Vincent Hugel and Samir Otmame for reviewing this dissertation.

A venture like this one would not have been possible without all my family, friends, colleagues, and personnel in France, the USA, and Colombia. Thanks to you, these five years have been an incredibly enriching experience. I am afraid that no proper acknowledgment can be given in such a short space.

I must also acknowledge the profound role that my parents have played in this achievement, by preparing me, perhaps unknowingly, for this accomplishment. I also want to thank my sister, Laura, future *real* doctor, for her love, and her active and tacit encouragement.

Last, but never least, I am obliged to my wife, Catherine, who blissfully followed me in this endeavor, and gave me the most significant result of these years—*Gaël* and *Inés*.

ABSTRACT

Selection of moving targets is a common task in human–computer interaction (HCI), and more specifically in virtual reality (VR). In spite of the increased number of applications involving moving–target selection, HCI and VR studies have largely focused on static-target selection. Compared to its static-target counterpart, however, moving-target selection poses special challenges, including the need to continuously and simultaneously track the target and plan to reach for it, which may be difficult depending on the user’s reactivity and the target’s movement. Action prediction has proven to be the most comprehensive enhancement to address moving-target selection challenges. Current predictive techniques, however, heavily rely on continuous tracking of user actions, without considering the possibility that target-reaching actions may have a dominant pre-programmed component—this theory is known as the pre-programmed control theory.

Thus, based on the pre-programmed control theory, this research explores the possibility of predicting moving-target selection prior to action execution. Specifically, three levels of action prediction are investigated: action performance, prospective action difficulty, and intention. The proposed performance models predict the movement time (*MT*) required to reach for a moving target in 2-D and 3-D space, and are useful to compare users and interfaces objectively. The prospective difficulty (*PD*) models predict the subjective effort required to reach for a moving target, without actually executing the action, and can therefore be measured when performance can not. Finally, the intention models predict the target that the user plans to select, and can therefore be used to facilitate the selection of the intended target.

Intention prediction models are developed using decision trees and scoring functions, and evaluated in two VR studies: the first investigates undirected selection (i.e., tasks in which the users are free to select an object among multiple others), and the second directed selection (i.e., the more common experimental task in which users are instructed to select a specific object). *PD* models for 1-D, and 2-D moving-target selection tasks are developed based on Fitts’ Law, and evaluated in an online experiment. Finally, *MT* models with the same structural form of the aforementioned *PD* models are evaluated in a 3-D moving-target selection experiment deployed in VR. Aside from intention predictions on directed selection, all of the explored models yield relatively high accuracies—up to $\sim 78\%$ predicting intended targets in undirected tasks, $\hat{R}^2 = .97$ predicting *PD*, and $\hat{R}^2 = .93$ predicting *MT*.

CHAPTER 1

INTRODUCTION

1.1 Moving-target selection

Selection of moving targets is a common task in human–computer interaction (HCI) and more specifically in virtual reality (VR). Targets may move independently from user input, as in interactive, or “clickable,” video [Ilich 2009; Hasan et al. 2011; Silva et al. 2012], and air traffic control displays [Mould and Gutwin 2004; Hasan et al. 2011]. Targets may also move relative to the user, a case in point being navigation in VR [Mould and Gutwin 2004] and augmented reality (AR) [You et al. 2012]. In some applications, including video games [Mould and Gutwin 2004; Pavlovych and Gutwin 2012] and interactive 3-D simulations [Mould and Gutwin 2004; Hasan et al. 2011], both kinds of movements are present.

In spite of the increased number of applications involving moving-target selection, HCI and VR studies have largely focused on static-target selection. In the taxonomy for VR manipulation tasks presented by Poupyrev et al. [1997], selection parameters such as target size and distance are given some level of detail, including measurement variables, whereas target movement is classified under “other parameters,” without any further description. In the more recent taxonomies of 3-D interaction techniques by Bowman et al. [2004], and 3-D selection techniques by Argelaguet and Andujar [2013], target motion is not included as a parameter for target selection. This example is, perhaps, reflective of the numerous HCI studies on static-target selection based on Fitts’ Law [Fitts 1954], whose inputs are target size and distance only.

1.1.1 Moving-target selection challenges

Moving-target selection poses special challenges compared to its static-target counterpart. Unlike static-target selection, which can be executed without ongoing visual control under certain conditions [Hoffmann and Chan 2012], the nature of moving-target selection requires the user to continuously and simultaneously track targets and plan to reach for them [Hasan et al. 2011], which may be difficult considering the inherent speed and precision limitations of the human sensory-motor system [Shadmehr et al. 2010]. In some applications targets may move along unpredictable paths [Ilich 2009; Hasan et al. 2011; Pavlovych and Gutwin 2012;

You et al. 2012; Ortega 2013], with changing speeds [Pavlovych and Gutwin 2012; You et al. 2012], making the task more challenging. Additionally, certain combinations of target and viewport motions may cause targets to fall out of the user’s field of view, become occluded, or change in visible size, this problem is illustrated in interactive sport-videos in which players can be selected while moving [Ilich 2009]. Even in cases in which target visibility, size, and velocity are constant, there is a critical speed beyond which selecting the target becomes impossible [Hoffmann 1991].

Finally, common HCI challenges are exacerbated in VR moving-target selection. End-to-end latency (the delay between user input and system output), for example, is often revealed to users when tracking a moving target [Reddy 1994]. More importantly, the presence of latency sharply affects moving-target reaching accuracy and, to a lesser degree, moving-target tracking [Pavlovych and Gutwin 2012]. Similarly, VR induced symptoms such as sensory conflict, simulator sickness, confusion, and frustration [Cobb et al. 1999] are likely worsened in moving-target selection.

1.1.2 Moving-target selection enhancements

Ilich [2009] identifies two major categories for moving-target selection enhancement: pointer enhancement, and target enhancement. Ilich also suggests the general “task simplification” category, but the strategies classified in this category are not described since they can also be considered as pointer and target enhancements.

According to Ilich, pointer enhancement is possible through an increase in the speed or area of the pointer, as well as an added affinity to certain targets. The sole increase in pointer size extends the time window for selection when targets are to be intercepted [Tresilian 2005], and may reduce the effective target distance on target-chase tasks [Ilich 2009], but it may also affect target visibility and acquisition accuracy, especially when targets are small—this issue is known in touch screens as the “fat finger” problem [Holz and Baudisch 2010]. The increase in pointer speed in the form of velocity control has been shown to increase the selection accuracy in tasks with a single moving-target [Jagacinski et al. 1980], but such control techniques, also available in static-target selection, are known to suffer from a decrease in positioning precision as the reaching distance increases [Bowman et al. 2004, p. 162]. The last type of pointer enhancement, pointer–target affinity techniques, has been shown to be successful in increasing selection accuracy in tasks with multiple targets by linking the pointer to targets whose angular distances [de Haan et al. 2005] and euclidean distances [Ortega 2013] to the pointer decrease more rapidly, these techniques are described in depth in Section 2.3.1.

Concerning target enhancement, Ilich describes three possible techniques: target expansion, target repositioning, and target speed decrease. Target expansion strategies, such as the *Comet* [Gunn et al. 2009; Hasan et al. 2011], and *AttachedShock* [You et al. 2012, 2014], enhance pointing by extending the selectable target area to include a movement trail left behind

each target, but suffer from clutter and overlap when the number of selectable targets is increased. Target repositioning techniques, such as target *Ghost* [Hasan et al. 2011], which creates static proxies for each target upon activation, exhibit the same trade-off between enhanced performance but increased clutter and overlap. The last target enhancement, speed decrease, has been successfully implemented in techniques that completely stop targets [Ilich 2009; Al Hajri et al. 2011], thus reducing the task to static selection. In spite of their benefits, these target enhancements may be undesirable in certain applications, especially those that strive for realism.

Apart from pointer–target affinity techniques, the enhancements proposed by Ilich provide only partial solutions to the challenges in moving-target selection, while aggravating other existing challenges, or introducing new ones. The reason pointer–target affinity stands out is that it provides a way to anticipate the intended target and enhance the pointer only with respect to that target. Such a principle is not restricted to pointer enhancement, prediction of targets and motion endpoints has been suggested to address clutter and overlap in target enhancement techniques for both static [McGuffin and Balakrishnan 2005; Lank et al. 2007; Wonner et al. 2011] and moving-target [Hasan et al. 2011] selection tasks.

1.1.3 Control strategies in moving-target selection

The pointer–target affinity techniques described above are based on the principle that the user is constantly following their intended target. According to Tresilian [2005], however, the motion required to reach a moving-target has a pre-programmed control component that may be dominant over on-line control, especially when the motion must be executed rapidly. In other words, such a motion is minimally influenced by external sensory information once the motor commands are issued.

Thus, based on the pre-programmed control theory, this research explores the possibility of predicting moving-target selection prior to action execution. More specifically, three levels of action prediction are investigated: **action performance, prospective action difficulty, and intention.** These levels of prediction are subsequently described with respect to existing action models, including the seven stages of action [Norman 1986, 2002], and the aforementioned pre-programmed control theory.

Among these predictions, performance is assessed during action execution, the lowest cognitive stage of action. It refers to quantitative measures of task execution, which can be used to evaluate users, interfaces, and interaction techniques objectively. In particular, in accordance with the main HCI body of static-target selection studies, the focus is on predicting the movement time (*MT*) required to reach the target. For static-target selection, *MT* has been shown to be positively and linearly related to the index of difficulty, $ID = \log_2(2D/W)$, where *D* and *W* are the target’s size and distance, this relation is known as Fitts’ law [Fitts 1954]. Unfortunately, as shown in Chapter 2, the existing formulae for *MT* in moving target

selection are limited to 1-D (i.e., target velocity directly toward or away from the cursor) and lack a simple expression of difficulty similar to the aforementioned *ID*. One of the aims of this research is to extend the *MT-ID* paradigm to 2-D, and 3-D moving-target selection.

Prospective difficulty (*PD*), refers to subjective assessments of difficulty evaluated prior to action execution, which typically occurs during the higher level stage of action specification. Since prospective judgments do not require action execution, they may be measured even when performance can not, due to factors such as task feasibility, low occurrence, or even practicality. Unfortunately, direct assessment of *PD* requires either task interruption, or the usage of sensors such as EEG [Kourtis et al. 2012], which may be invasive and are generally not available in all VR setups. Nonetheless, since *PD* is assessed during action preparation, its value is probably related to the task parameters that also affect performance. In particular, it has been shown that *PD*, and *ID* are related for 1-D static-target selection tasks [Slifkin and Grilli 2006; Grilli 2011], but so far this relation has not been explored in moving-target selection. Given the usefulness of *PD*, this research also aims to extend the *PD-ID* paradigm to moving-target selection.

Finally, intention, a general action that a user plans to execute to achieve a goal, precedes action specification and is therefore at the highest cognitive level among these three measurements. In this work, the scope of such intentions is limited to the target the user plans to select.¹ Similar to *PD*, direct assessment of intention requires task interruptions, or inference via proxy measures that relate user actions to targets (user-target states), such as gaze. Based on the principle of pre-programmed action, however, it is hypothesized that users form their intentions by minimizing their prospective effort.

Minimizing prospective effort to form intentions can be observed in *undirected* selection tasks, i.e., tasks in which users are free to choose an object among multiple others, as opposed to the more (experimentally) common *directed* tasks, in which users are instructed to select a specific target. In static-target selection, given the correlation between *PD* and *ID*, minimizing prospective effort is hypothetically equivalent to minimizing *ID*.

Therefore, in terms of intention, this research aims to a) test the hypothesis that users form their intentions by minimizing their prospective effort, and b) evaluate the predictive accuracy of different user-target states as proxies for intention.

¹This scope contrasts with some of the previous work in which intention refers to the usage given to a target following selection [Mandryk and Lough 2011; Song et al. 2013; Ruiz and Lank 2014].

1.2 Research Layout

The remainder of this dissertation is structured as follows,

1. Chapter 2 presents the relevant literature related to static-target and moving-target selection, prospective difficulty, and the usage of user-target states as prediction inputs.
2. Chapter 3 presents the results of the first user study. This study evaluates the hypothesis that users form their intentions by minimizing their prospective effort as described by *ID*, in undirected moving-target selection tasks in VR. Since *ID* is a task-specific feature, limited in usefulness to undirected tasks, it was complemented with generalizable user-target features that can also be used in directed tasks.
3. Chapter 4 develops *PD* models for 2-D static-target and moving-target selection. These models are validated with the results of an online user study.
4. Chapter 5 extends the *PD* models developed in Chapter 4 to 3-D, and evaluates their usefulness as predictors of *MT*, and intention.
5. Chapter 6 presents the contributions, future work and conclusions of the dissertation.

CHAPTER 2

LITERATURE REVIEW

2.1 Selection performance

Target-selection performance in HCI is usually studied using information-theoretic models, including Fitts' Law [Fitts 1954] and, to a lesser degree, the Hick-Hyman Law [Hick 1952; Hyman 1953]. Card, English, and Burr [1978] were the first to use Fitts' Law in HCI to compare the movement times (MT) of different input devices, and Card, Moran, and Newell [1983] presented both laws as part of the operating principles of the *Model Human Processor*.

2.1.1 The Hick-Hyman Law

The Hick-Hyman Law relates the reaction time (RT) required to make a choice to the number (n) and probability (p) of each of the possible choices such that,

$$RT = a + b \sum_i^n p_i \log_2 \left(\frac{1}{p_i} + 1 \right) \quad (2.1)$$

where a and b are empirically determined coefficients. If all choices have equal probability ($p = 1/n$), Equation (2.1) is reduced to

$$RT = a + b \log_2 (n + 1) \quad (2.2)$$

Notice that the RT describes only the time to make the choice, and not the MT required to execute the selection.

2.1.2 Fitts' Law

Fitts' Law relates the mean movement time (\overline{MT}) required to reach a target to the target's index of difficulty (ID), such that

$$MT = a + b ID, \quad (2.3)$$

where a and b are empirically determined coefficients. ID gives an objective measure of the difficulty involved in pointing at a target as a function of the ratio between the target's distance (D),¹ and its size (W). The two most common ID formulations are Fitts' original formulation [1954]

$$ID_F = \log_2 \left(\frac{2D}{W} \right), \quad (2.4)$$

and the so-called Shannon formulation, suggested by MacKenzie [1989],

$$ID_{Sh} = \log_2 \left(\frac{D}{W} + 1 \right) \quad (2.5)$$

Fitts' Law has been extensively used in HCI, psychology, and related fields to model human performance (for compendia see MacKenzie [1992], and Guiard and Beaudouin-Lafon [2004]), but there is an ongoing debate concerning the correctness and usefulness of these two ID formulations [Drewes 2010; Hoffmann 2013; MacKenzie 2013]. Regardless of these issues, as noted by Drewes [2010], the use of both formulae should lead to similar results and, as noted in a personal communication with Hoffmann [2014], the differences may be of statistical but not practical significance. In any case, it is important to consider that both models represent an approximation of a more complex reality [Drewes 2010]. Citing a popular phrase of the late statistician George Box, "essentially all models are wrong, but some are useful" [Box and Draper 1987, p. 424].

2.1.2.1 Two-part formulations of Fitts' Law and gain

Welford et al. [1969] suggested that pointing could be separated in two control processes: a motor, or distance-covering process, and a visual, or homing-in process. To account for these two processes, Welford and colleagues proposed the following variation of Fitts' Law:

$$MT = a + b \log_2(D) - c \log_2(W). \quad (2.6)$$

Inspired by Kopper et al. [2010], Shoemaker et al. [2012] reformulate Welford's model by introducing the ratio $k = c/b$, which encapsulates the relative impact of the two-parts of Welford's formula, yielding a model similar in form to Equation (2.3), such that

$$MT = a + b \log_2 \left(\frac{D}{W^k} \right). \quad (2.7)$$

In their analyses, Shoemaker et al. [2012] found that two-part formulations described MT better than one-part models when pointing at targets with different levels of control-display gain (G , the ratio of the cursor speed to the input speed). Additionally, based on different

¹Traditionally researchers refer to this quantity as amplitude (A), referring to the movement amplitude required to reach for the target, in this work, the term distance (D) is preferred, particularly because in moving-target selection D , but not A , is known beforehand.

regression estimates, they also observed that k increased linearly with G . Shoemaker and colleagues also noted that c “increases quite consistently with gain,” while b “stays relatively constant,” but, contrary to the k - G relation, no comment on the possible form of the b - G and c - G relations, or on the evident increase of a with G , was given.

Even though the understanding of the k - G relation may be important to characterize the relative impact of each part of Welford’s model per G level, and gauge how far is the interaction from being modeled by Fitts’ original formulation (which is essentially Equation (2.7) with $k = 1$), understanding the a - G , b - G , and c - G relations allow comparing MT between G levels, as well as generalizing the results to other experimental conditions. Regression analyses for a - G , b - G , and c - G could generalize such contributions for arbitrary levels of G .

Alternative two-part formulation. Hoffmann and Chan [2012] reformulate Welford’s two-part model (Equation (2.6)) to separate the effect of movement amplitude from that of ID , as

$$MT = a' + b' \log_2(D) + c ID, \quad (2.8)$$

where $a' = a - c$, and $b' = b - c$.

However, Hoffmann and Chan [2012] mention that the problem with Equations (2.6), and (2.8), is that taking logarithms of either D or W is invalid as these quantities are not unitless. Shoemaker et al. [2012] mention that Welford was aware of this problem and formulated normalizing constants D_0 and W_0 , but these are not explicit. To avoid this problem, and based on the near-linear relation between $\log_2(D)$ and \sqrt{D} for wide ranges of D values, Hoffmann and Chan [2012] reformulate Equation (2.8), by replacing the $\log_2(D)$ term, by \sqrt{D} , such that

$$MT = a'' + b'' \sqrt{D} + c ID,$$

where the approximate values of a'' , and b'' are given by the regression estimates of $\log_2(D)$ on \sqrt{D} for the range of D in the experiment, such that

$$\begin{aligned} a'' &\approx a' + b' \beta_1 \\ b'' &\approx b' \beta_2 \\ \log_2(D) &\sim \mathcal{N}(\beta_1 + \beta_2 \sqrt{D}, \sigma^2). \end{aligned} \quad (2.9)$$

These approximations allow prediction of the change in coefficients for regressions of the two-part Welford model [Hoffmann and Chan 2012]. The square-root of the distance had been previously shown to be linearly related to MT on purely ballistic movements [Gan and Hoffmann 1988], therefore, under this model, the distance covering phase is assumed to be mostly ballistic, whereas the homing-in phase is assumed to follow to Fitts’ ID [Hoffmann and Chan 2012].

2.1.2.2 Fitts' Law formulations for 2-D static-target pointing tasks

Fitts' Law is inherently 1-D, but in practice it is used on 2-D tasks, effectively ignoring the effects of both target angle (θ), and target shape on movement time. A 2-D target selection with a bivariate target (i.e., with different width W , and height H) is shown in Figure 2.1.




Figure 2.1. Two-dimensional task with a bivariate target. The red disc indicates the cursor; the green rectangle represents a target with width W , and height H . D is the cursor–target distance, and θ is the target angle.

Nonetheless, several works have studied the effect of θ on MT , including the early work of Card et al. [1978] and Jagacinski and Monk [1985], the later work of Boritz et al. [1991], MacKenzie and Buxton [1992], and Whisenand and Emurian [1995; 1996; 1999], as well as the more recent work of Appert et al. [2008], Grossman and Balakrishnan [2004; 2005], Hancock and Booth [2004], Murata and Iwase [2001], Phillips and Triggs [2001], and Zhang et al. [2012]. Overall, it appears that MT is longer for diagonal targets, than for horizontal and vertical ones (the most notable exceptions are the results of Murata and Iwase [2001], Whisenand and Emurian [1999], and Zhang et al. [2012], who found that vertical targets were slower than diagonal targets).

To the author's knowledge, the only modification of Fitts' ID that exclusively models the effect of θ , is that of Murata and Iwase [2001], which they derive empirically,

$$ID_3 = \log_2 \left(\frac{D}{W} + 1 \right) + c \sin \theta, \quad (2.10)$$

where c is an empirically determined constant. A desirable characteristic of this model, is that it separates the effect of θ from the \log_2 term.

In addition to θ , researchers have also studied targets of multiple dimensions and shapes. In these studies, the first category of extensions to Fitts' ID , which includes the ID_{\min} model by MacKenzie and Buxton [1992] and Hoffmann and Sheikh [1994], and the one-weight euclidean model by Accot and Zhai [2003], accounts for bivariate targets, but not explicitly for θ . A similar yet more general category, which includes the models by Sheikh and Hoffmann [1994], Murata [1999], and Grossman and Balakrishnan [2005], allows modeling for arbitrary target shapes through the use of probabilistic models, but not explicitly for θ . A third category,

which includes the ID_W model by MacKenzie and Buxton [1992], and the ID_θ model by Zhang et al. [2012] accounts for both bivariate targets *and* θ , but cannot account for θ when targets are univariate (e.g., circles). Finally, the most relevant category for the current research allows modeling for θ and multiple target dimensions separately; two models in this category are presented below.

The model by Appert et al. [2008] introduces an overall angle effect as a cosine that is added within the log term of the ID ,

$$ID_\theta = \log_2 \left(\frac{D}{W} + \frac{D}{H} + 0.6 \cos(\theta) \frac{D}{\min(W, H)} + 1 \right), \quad (2.11)$$

where H is the target's height.

The weighted-euclidean ID formulation suggested by Grossman and Balakrishnan [2004] is an extension to the one-weight euclidean model of Accot and Zhai [2003] that allows modeling the angle effect separately for each of the target's dimensions (formulated and tested in 3-D, but reduced here to 2-D),

$$ID_{WtEuc\theta} = \log_2 \left(\sqrt{f_W(\theta) \left(\frac{D}{W} \right)^2 + f_H(\theta) \left(\frac{D}{H} \right)^2 + 1} \right), \quad (2.12)$$

where $f_W(\theta)$ and $f_H(\theta)$ are empirically determined weights per angle.

Equations (2.10)–(2.12) represent the contributions of θ in three distinct ways. Since the formulae are derived empirically, or by analogy, this raises the question of the true form of the contribution of θ , which is probably why Grossman and Balakrishnan [2004] simply suggest to derive $f(\theta)$ empirically.

In terms of ease of use, Equations (2.10) and (2.11) are the simplest to solve, as a ubiquitous linear least-squares method may suffice, whereas Equation (2.12) requires a more specialized method, such as non-linear least-squares. In the case of the latter model, the fit in terms of R^2 may be artificially higher than the former due to the additional $2 \times \#(\theta)$ terms represented in $f_W(\theta)$, and $f_H(\theta)$. More importantly, the usage of R^2 may be inadequate to evaluate the goodness-of-fit of non-linear models [Spiess and Neumeyer 2010]. These last two points are important in the status quo of the Fitts literature, where model fit is customarily assessed using R^2 .

Aside from the arbitrary shape models [Sheikh and Hoffmann 1994; Murata 1999; Grossman and Balakrishnan 2005], an important limitation of the aforementioned models is that they do not account for different target orientations.

2.1.2.3 Fitts' Law formulations for 3-D static-target pointing tasks

To the authors knowledge, the only work to formulate ID models specific for 3-D is that of Grossman and Balakrishnan [2004]. These models, which are simple extensions to the 2-D

models from the previous section, fail to account for all of the spatial parameters of a 3-D task, shown in Figure 2.2.




Figure 2.2. Three-dimensional trivariate target. The red disc indicates the cursor; the green rectangular cuboid represents a target with width W , height H , and depth D . D is the cursor-target distance, α is the z - x azimuth angle, and θ is the altitude angle.

For instance, the weighted-euclidean ID formulation was presented in Equation (2.12) as a 2-D formula, but in fact it was studied with trivariate targets placed in a horizontal 2-D plane. Its complete form is

$$ID_{WtEu\alpha} = \log_2 \left(\sqrt{f_W(\alpha) \left(\frac{D}{W} \right)^2 + f_H(\alpha) \left(\frac{D}{H} \right)^2 + f_D(\alpha) \left(\frac{D}{D} \right)^2 + 1} \right), \quad (2.13)$$

where D is the target's depth, and α is the z - x azimuth angle. Notice that this formulation does not account for the full angular position (α, θ) of the target, nor does it account for different target orientations.

2.1.2.4 Fitts' Law formulations for 1-D moving-target pointing tasks

In comparison to static target selection, the research on Fitts' Law for 1-D and 2-D moving-target selection is very scarce. After the early studies of Jagacinski et al. [1980] on the applicability of Fitts' Law for moving-targets in 1-D, only two studies have evaluated and extended Fitts' ID in moving-target selection tasks [Hoffmann 1991; Al Hajri et al. 2011].

Jagacinski et al. [1980] found that Fitts' ID was a poor predictor of MT for moving targets with position control. Instead, they suggested an alternative formulation with analogous characteristics to two-part formulations, such as predicting longer MT for larger D and smaller W , but also taking in account the interaction between W and speed (V). Their proposed formulation is

$$MT = a + bD + c(V + 1) \left(\frac{1}{W} - 1 \right), \quad (2.14)$$

where the $V + 1$ term represents the $W \times V$ interaction, minimized for the widest target in their experimental design ($W = 0.92^\circ$) by subtracting 1 from $1/W$ (since $1/0.92 \approx 1$). In Jagacinski's data, this model yielded a good fit in terms of $R^2 = 0.96$.

As discussed by Jagacinski and colleagues, the “ -1 ” term can be considered as a fourth parameter, e.g., q , thus a possible reformulation of Equation (2.14) is

$$MT = a + bD + c(V + 1) \frac{1}{W} - d(V + 1), \quad (2.15)$$

where $d = cq$. With this fourth parameter, the formulation becomes more general, and extensible to other experimental designs. Compared to the other formulations presented in this work, this one stands out for not including any logarithmic, or square-root terms.

In the Jagacinski study, trials were considered successful only after the cursor remained within the target for at least 350 ms. This capture time duration (T_c) was probably appropriate for Jagacinski's input device (a joystick), but may be different for other input methods, such as a mouse click, which takes about 200 ms for the average user according to the Keystroke-Level Model (KLM) [Card et al. 1980], or pointing on a touch screen, which theoretically entails a 0 ms capture time. To take the T_c in consideration, Jagacinski et al. [1980] propose an alternative model,

$$MT = a + b \log_2 \left(\frac{2D}{W} \right) + c \log_2 \left(\frac{V}{W/T_c} + 1 \right). \quad (2.16)$$

Notice that this model simply reduces to Fitts' Law when $V = 0$, or $T_c = 0$ (e.g., in a touch screen). Unfortunately, the fit of this model on their data was low in terms of $R^2 = 0.71$, compared to the de facto standard in the Fitts literature, where $R^2 > 0.81$ [MacKenzie 1992, p. 101]. A third model was also proposed by Jagacinski and colleagues, but it contained the MT term on both sides of the equation and it did not, reportedly, result in a better fit in terms of R^2 ; therefore, it will not be taken in consideration.

Using a first-order control system, Hoffmann [1991] formally derived an ID for moving-target selection closer in form to Fitts original ID ,

$$ID_H = \log_2 \left(\frac{D \pm \frac{V}{K}}{\frac{W}{2} - \frac{V}{K}} \right), \quad (2.17)$$

where K , described as the person's gain, is determined empirically, and the \pm on the numerator is determined by the direction of V relative to the starting position: “ $+$ ” for approaching targets, and “ $-$ ” for distancing targets.² Similarly to Welford et al. [1969], Hoffmann also suggested a two-part formulation for his own model,

²Distancing targets are those moving away from the cursor.

$$MT = a + b \log_2 \left(D \pm \frac{V}{K} \right) + c \log_2 \left(\frac{W}{2} - \frac{V}{K} \right). \quad (2.18)$$

Reusing the experimental data from Jagacinski et al. [1980], Hoffmann [1991] validated these two models empirically, yielding fits with R^2 values of 0.84 for Equation (2.17), and 0.94 for Equation (2.18). Interestingly, neither model, although theoretically sound, was able to surpass Jagacinski's empirical model (Equation (2.14)). The R^2 fit of the models in Equations (2.14) and (2.18) are practically equivalent, but the latter includes an additional term (K) that also contributes to the coefficient of multiple determination.

Aside from the goodness-of-fit issues, Hoffmann's moving-target models require special considerations regarding the calculation and interpretation of the K term that are not explicitly addressed in the literature. According to Hoffmann [1991], the coefficient K serves to determine the critical speed (V_{crit}) beyond which target capture is not possible,

$$V_{crit} = \frac{WK}{2}.$$

Hoffmann [1991] calculated K in two distinct ways: using regression, and observing the (V, W) condition in which capture occurred less than $\rho = 50\%$ of the time. The regression approach requires a specialized technique such as non-linear least-squares (which involves the additional issues explained at the end of Section 2.1.2.2) and yields values of K whose domain is dependent on the experimental design.³ The alternative approach results in values of K that are dependent on the percentage of successful captures (ρ) per (D, W, V) condition, thus, its precision depends on the span and resolution on the conditions in the experimental design, as well as the device-specific time of capture. Using either method, once K is calculated, the set of experimental conditions (D, W, V) that can be mathematically modeled by Equations (2.17) and (2.18) is restricted to

$$\left\{ (D, W, V) \mid \frac{V}{W/2} < K \right\}, \quad (2.19)$$

for approaching targets, and

$$\left\{ (D, W, V) \mid \frac{V}{W/2} < K \iff \frac{V}{D} < K \right\}, \quad (2.20)$$

for distancing targets.

2.1.2.5 Fitts' Law formulations for 2-D moving-target pointing tasks

Al Hajri et al. [2011] also derived the ID_H formulation (Equation (2.17)) by applying the human processor model [Card et al. 1980] to describe moving-target selection. Subsequently, they suggested three extensions for ID_H by combining it with the $ID_{W'}$ and ID_{min} formulations

³For approaching targets, $K > \frac{\max(V)}{\min(W/2)}$; for distancing targets, $\left(\frac{\max(V)}{\min(W/2, D)} < K \right) \oplus \left(\frac{\min(V)}{\max(W/2, D)} > K \right)$. Otherwise, Equations (2.17) and (2.18) are not defined for all (D, W, V) conditions.

by MacKenzie and Buxton [1992], and the $ID_{WtEuc\theta}$ model (Equation (2.12)) by Grossman and Balakrishnan [2004], respectively. As discussed in Section 2.1.2.2, $ID_{WtEuc\theta}$ is the most relevant among these three models, therefore only the extension $ID_{HWtEuc\theta}$ is considered,

$$ID_{HWtEuc\theta} = \log_2 \left(\sqrt{f_W(\theta) \left(\frac{D \pm \frac{V_x}{K}}{\frac{W}{2} - \frac{V_x}{K}} \right)^2 + f_H(\theta) \left(\frac{D \pm \frac{V_y}{K}}{\frac{H}{2} - \frac{V_y}{K}} \right)^2 + 1} \right), \quad (2.21)$$

where V_x and V_y are the magnitudes of the x and y components of the velocity, respectively.

Even though Equation (2.21) can model pointing at bivariate targets moving in 2-D, the studies of Al Hajri et al. [2011] only included conditions in which targets were moving directly towards or directly away from the cursor, even if the initial target angle $\theta \in [0, 360)^\circ$.

2.2 Subjective difficulty

In addition to describing MT , ID has also been shown to be correlated to subjective difficulty [Delignières and Famose 1992; Slifkin and Grilli 2006; Grilli 2011; Chan and Hoffmann 2013]. There are two major types of subjective difficulty measurements in the literature:

1. *Prospective, or estimated difficulty* refers to ratings of subjective difficulty assessed without executing a task, e.g., in the works of Delignières [1990], Slifkin and Grilli [2006], and Grilli [2011]
2. *Perceived difficulty* refers to ratings of subjective difficulty assessed after executing a task, e.g., in the works of Delignières and Famose [1992], Shoemaker et al. [2012], and Chan and Hoffmann [2013].

There is some evidence that both measurements are correlated with each other [Delignières 1990; Grilli 2011], and with MT [Chan and Hoffmann 2013; Delignières 1993; Grilli 2011].

2.2.1 The importance of the rating method

The works of Delignières, described in his dissertation [1993], show that the best type of fit relating subjective difficulty to ID is heavily influenced by the rating method. In his studies, the best fit between ID and subjective difficulty was linear when the latter was measured using a 15-point scale; when using a ratio-rating technique, the best fit was exponential.

2.2.2 Subjective difficulty in moving-target selection

To the author's knowledge, there is no study directly relating subjective difficulty and ID for moving-target selection. The closest work, is that of Famose et al. [1991], in which perceived difficulty ratings in a 15-point scale were shown to be negatively and linearly related to the scores in a dart throwing task on a moving-target.

2.3 Prediction of intention

Predicting intended targets has been proposed as a solution to clutter and overlap in static-target selection techniques. Current static-target prediction techniques are based on the trajectory and velocity profiles of the pointer [Lank et al. 2007; Noy 2001; Wonner et al. 2011; McGuffin and Balakrishnan 2005]. The peak accuracy rates for prediction using these techniques require a wide window of user input—at least 80% of the pointing movement—but some of them are intended to predict endpoints [Lank et al. 2007; Wonner et al. 2011], rather than intended targets [Noy 2001; McGuffin and Balakrishnan 2005]. These techniques, however, are not adapted for moving-target prediction, in particular due to the apparent dependency of the users' velocity profiles on the targets' movement [Carnahan and McFadyen 1996], and the fact that the peak hand velocity is attained upon target selection [Tresilian 2005].

2.3.1 Scoring functions

In the context of moving-target selection, the studies from de Haan et al. [2005] and Ortega [2013] demonstrated the feasibility of predicting intended targets in complex VR scenes using scoring functions. Overall, the stages of their algorithms can be described as follows,

1. Target filtering: A subset of the targets is chosen based on a given criterion.
2. Target scoring: Scores for targets that meet the criterion are incremented, and scores of targets that do not are decremented.
3. Target highlighting: The target with the highest score gets highlighted.
4. Target selection: The user completes the selection on the highlighted target by executing a button action, or continues moving the cursor until the intended target is highlighted.

In *IntenSelect*, the de Haan model [2005], the filtering criterion is given by an infinite cone with aperture β_{cone} , apex located at the wand position, and orientation corresponding to wand vector \mathbf{W} . At each frame t , the wand-target angular distance α_i is calculated for each of the I targets. Targets whose $\alpha_i < \beta_{cone}$, i.e., those inside the cone, get their scores increased, whereas the rest get their scores decreased. The score for each target i is calculated following

$$angScore_i(t) = angScore_i(t-1) \cdot decay + \begin{cases} (1 - \alpha'_i / \beta_{cone}) \cdot growth & \text{if } \alpha_i \leq \beta_{cone} \\ 0 & \text{if } \alpha_i > \beta_{cone} \end{cases}$$

$$\alpha'_i = \arctan \left(\frac{D_{i,perp}}{D_{i,proj}^k} \right), \quad (2.22)$$

where $D_{i,perp}$, and $D_{i,proj}$, correspond to the perpendicular and projected distances between the wand vector \mathbf{W} , and target i . Notice that $\alpha'_i = \alpha_i$ when $k = 1$; in the de Haan experiments, however, k was heuristically selected as $4/5$, since $k = 1$ led to distant targets being easier to reach than nearby ones. Further note that when $decay = 0$, and $growth = 1$, no score accumulation occurs, resulting in the flashlight, or conical selection technique [Bowman et al. 2004, 153–154]; de Haan et al. [2005] describe balancing these parameters to obtain “a comfortable response,” but do not report the actual values.

In *Hook*, the Ortega model [2013], the filtering criterion is given by the wand–target distance D_i . At each frame t , the wand–target distance D_i is calculated for each of the I targets. Targets are then ordered ascendingly by D , their order given by $j_i = 0, \dots, I - 1$. The score for each target i is calculated following

$$dScore_i(t) = dScore_i(t - 1) + \begin{cases} (N - j_i)\Delta t & \text{if } j_i < N \\ -(decay \cdot N)\Delta t & \text{if } j_i \geq N \end{cases}$$

$$dScore_i(t) \geq 0, \tag{2.23}$$

where N is an arbitrary number of closest targets, and $decay$ is the rate with which scores decrease when a target $j_i \geq N$. In Ortega’s studies with $I = 100$ targets, $dScore_i(t)$ was used with $N = 20$, and $decay = 0.5$.

A disadvantage of de Haan’s scoring function (2.22), compared to Ortega’s scoring function, Equation (2.23), is that the former does not account for variable frame-rates. Scores are increased or decreased uniformly between frames regardless of the time difference Δt between the frames. This is a major drawback, given that variable framerates are common in VR applications

These functions are easy to implement and their performance is enhanced as the user follows each target with the pointer; however, as it happens with some of the tasks in the present work, users may not always follow the intended target with their pointer. Additionally, there is no data on the predictive accuracy (i.e., the percentage of correctly predicted targets) of such functions, or how such accuracy is affected by the target distance—it is possible that users may have made their decision before starting their pointer movement, so the prediction could be done in advance.

2.3.2 Gaze

Knowing where a person is looking is considered an indicator of what is at the “top of the stack” of a cognitive process [Just and Carpenter 1976]. With respect to object manipulation, research has shown that gaze leads hand motions [Johansson et al. 2001]. Gaze is composed of head orientation and eye orientation relative to the head [Wilson et al. 2000].

In the context of target selection, eye gaze has proven to be beneficial in assisting users during static-target selection tasks, concurrently with more traditional input devices, such as mice [Zhai et al. 1999; Blanch and Ortega 2009]. Eye-trackers, however, are expensive and may be technically challenging to integrate in CAVE-like immersive VR systems [Murray et al. 2007], like the ones considered in the current work. Furthermore, this integration may be cumbersome due to the complex calibration procedures required or the cabling limitations of certain eye trackers [Murray et al. 2007]. Some modern solutions address these problems and allow eye-tracking in VR, but their adoption is still limited and costly. Head tracking, on the other hand, is readily available in most CAVE-like systems and has been successfully integrated in large-display [Nancel et al. 2013], video-conference [Stiefelbogen 2002], mobile [Spindler et al. 2012], surface [Francone and Nigay 2011], and floor-projected [Pierard et al. 2012] interactive systems.

In a series of studies based on video recordings, Stiefelbogen and colleagues demonstrated the potential of using head orientation to detect gaze and pointing. In a meeting scenario, Stiefelbogen and Zhu [2002] showed that, on average, head orientation represents 68.9% of the overall gaze direction, and could predict the observed person, among 3 people, with 88.7% accuracy. In a subsequent set of four meeting scenarios, once again with one observer and 3 possible observed people, the usage of head pose yielded an average accuracy of 72.9%, which increased to 75.6% when combined with audio information that indicated who was speaking at each point in time Stiefelbogen [2002]. Finally, Nickel and Stiefelbogen [2003] found that head orientation was predictive of pointing direction, with a mean error angle of 22° , as well as intended target among eight targets, with 75% average accuracy.

CHAPTER 3

INTENTION IN UNDIRECTED 3-D MOVING-TARGET SELECTION

Based on the premise that users form their intentions by minimizing their prospective effort, as described by *ID*, Chapter 1 introduced the hypothesis that intention can be predicted using *ID* in undirected selection tasks.

Unfortunately, as shown in Chapter 2, the existing *ID* measures formulated specifically for moving-target selection are not compatible with this *PD-ID* framework. The Hoffmann [1991] moving-target model, presented in Equation (2.17), and its 2-D extension by Al Hajri et al. [2011], presented in Equation (2.21), require the percentage of actual captures, which is unknown in prospective action, otherwise resulting in potentially harder and unreliable calculations.

On the other hand, a review of Jagacinski's studies [1980] reveals that even if Fitts' *ID* yielded poor correlations with *MT* for moving targets, the results were marginally better among conditions with equal target speeds. Additionally, low speed conditions resulted in slightly larger *MT-ID* correlations than high speed conditions. Mindful of these limitations, this Chapter evaluates the hypothesis that *ID* may be predictive of user intention in moving-target selection tasks that have a single, relatively low *V*. Analogous to the way Jagacinski's model [1980], Equation 2.14, uses *D* and *W* separately to predict *MT*, the predictive accuracy of *ID* is compared to that of its separate *D* and *W* components.

The potential usefulness of *ID* and *W*, however, relies on the assumption that the user is free to choose their intended target, in other words, it is limited to undirected tasks. Therefore, these task-specific measures are complemented with generalizable features that can also be used in directed tasks. Specifically, given their success at predicting user intention in directed tasks as reported in the related-work chapter, the proposed generalizable features consist of user-target states.

The contents of this chapter are based on papers published by the author and some of the committee members. Specifically, the *ID*, *D*, and *W* analysis appeared on Casallas et al. [2013], and the generalizable feature extensions appeared on Casallas et al. [2014].

3.1 Hypotheses

The following hypotheses motivated the current study,

3.1.1 Task-specific features for intention prediction in undirected moving-target selection

H1. *ID* can accurately predicting intention in undirected moving-target selection tasks.

3.1.2 Generalizable features for intention prediction in undirected moving-target selection

H2. Relative user–target features can accurately predict intention in undirected moving-target selection tasks.

H3. Combining task-specific, and generalizable features yields better predictive accuracy than using either separately.

3.2 Pilot study

Prior to the execution of the main experiment, a pilot study was conducted as a class project for the Machine Learning course (cs 573) at Iowa State University during the Spring of 2012, with classmates Ashwin S. Natarajan, and Keji Hu. The goal of the pilot study was to determine the best learning algorithm between Naïve Bayes [Mitchell 1997, 177–180], Neural Networks [Bishop 2006, 225–290], and c4.5 decision trees [Quinlan 1993], as well as the optimal time to measure the *D*, and *W* features to predict the intended target.

Ten unpaid students participated in the experiment, which was deployed in METAL, a 3-surface CAVE-like VE. The experimental setup is depicted in Figure 3.1.

The results favored the usage of the c4.5 algorithm, which was second to Neural Networks in terms of accuracy but, contrary to the latter, produced easy-to-understand rules. Additionally, the results supported the usage of the *D*, and *W* features measured at the beginning of each trial over features measured at one-third, and half of the total trial time. Due to restrictions of the Iowa State University Institutional Review Board (IRB) for in-class projects, the full extent of the results is not published.

3.3 Methods

3.3.1 Participants

Twenty-six unpaid participants, from the city of Chalon-sur-Saône, France aged 23 to 47, participated in the study. There were eighteen males and eight females; only two participants



Figure 3.1. Experimental setup of the pilot study.

were left-handed.

3.3.2 Apparatus

The experiment was developed in `vr juggLua` [Pavlik and Vance 2012], a Lua wrapper for `vr juggler` and `openscenegraph (osg)`. The application was deployed in the `MOVE`, a 4-surface CAVE-like virtual environment with three walls and a floor. The $3 \times 3 \times 2.67$ m environment was projected using passive Infitec stereo [Jorke et al. 2008] at 1160×1050 pixels per face. Four infrared ART cameras tracked the pose (position, P , and orientation, Q) of the participant's head and wand, using reflective markers mounted on Infitec stereo glasses and an ART Flystick2, respectively. This allowed the participant to have an adequate 3-D perception and interact with the virtual world.

A y -up coordinate system was used, with its origin placed at ground level in the middle of the VE, z decreasing towards the front wall, and x increasing towards the right wall.

3.3.3 Procedure

After filling a short survey, the participant was asked to enter the `MOVE`, face the front wall, and stay on a circular landmark ($r = 0.25$ m) located in the middle of the VE $(0,0,0)$, while completing a series of target selection tasks. In each trial, the participant was presented with

a horizontal array of virtual spheres of different sizes, starting in front of them and flying towards them in z . All of the spheres had the same texture, scaled accordingly to the sphere's size. The participant was instructed to touch each sphere by extending their arm only to reach the spheres; as opposed to wait for the spheres with their arm already extended. If a sphere was touched, or if it got 0.5 m past the participant's head in z , it disappeared. Each trial ended when the participant had touched all of the spheres, or when the remaining spheres got past their head.

Visual and auditory feedback were used to engage with the participant and indicate their performance. A virtual counter was placed at ground level, 5 m in front of the participant at $(0, 0, -5)$, which would show the number of missed spheres during each block; the counter would be reset to zero at the beginning of each block of trials. When the participant hit a sphere, a spatialized sound, co-localized with the wand position, would be played; when the spheres got past the participant's head, a different spatialized sound, co-localized with the overall centroid of the remaining spheres, would be played. Compared to the pilot study, the recreated virtual world had enhanced depth cues, such as a grid floor, better lighting, and textures for both the terrain and the spheres, some of these differences can be identified by comparing Figures 3.1 and 3.2.



Figure 3.2. Experimental setup with an array two spheres

During each trial, at each application frame, the elapsed time (t), head pose (P_h, Q_h), wand pose (P_w, Q_w), sphere positions (P_i) and possible collisions between the wand and the spheres were recorded in a log file.

3.3.4 Design

A within-subjects, factorial design was used, with three blocks of trials, each with a different number of conditions presented in a random order. In every trial, all of the spheres appeared 0.3 m below the participant's head and 5 m in front of them ($P_{i,y} = P_{h,y} - 0.3, P_{i,z} = -5$).

The first block each trial had only one sphere, moving at a constant speed of 2.5 m/s in z . Factors were sphere radius ($r_1 \in \{0.1, 0.2\}$) and sphere position (*left*: $P_{1,x} = 0.5$, *center*: $P_{1,x} = 0$, and *right*: $P_{1,x} = -0.5$). Each of the six conditions was presented to the participant in a random order until completing five trials per condition (30 total). The first block was intended only for training, so that users could become familiar with the environment and the task.

After completing the first block, the number of spheres was increased to two and velocity was decremented to 1.5 m/s in z . The spheres were positioned 0.5 m apart in x but the pair could appear offset to the *right* ($P_{1,x} = -0.5, P_{2,x} = 0$), or *left* ($P_{1,x} = 0.5, P_{2,x} = 0$), or appear *centered* ($P_{1,x} = -0.25, P_{2,x} = 0.25$) with respect to the user (see Figure 3.3). Factors were sphere radius ($r_i \in \{0.1, 0.2\}$) and row position (*left*, *center* and *right*). Each of the 12 conditions was presented to the participant in a random order until completing five trials per condition (60 total).




Figure 3.3. Possible row positions—*left*, *center* and *right*—with respect to the user in the two-sphere block

In the last block, there were 40 trials with three spheres per row moving with the same velocity. Nevertheless, to simplify the presentation of predictive methods and results, the scope of the analyses covered in this chapter only include the 2-sphere case. Instead, the analysis of a more complex task with one, three, and six spheres, and different velocities is presented in Chapter 5.

3.3.5 Data integrity

Since the goal was to predict intended targets, and not trial performance, trials in which a participant did not touch any sphere were discarded.

3.3.6 Predictive methods

Based on the results from the pilot study, all the feature-sets were evaluated using the J48 classifier, the Weka [Hall et al. 2009] open source implementation of c4.5. The algorithm is briefly described below.

C4.5 chooses its decision nodes recursively, based on the feature (A) that yields the greatest information gain (\mathcal{I}), such that

$$\mathcal{I}(S, A) = \mathcal{H}(S) - \mathcal{H}(S|A), \quad (3.1)$$

where $\mathcal{H}(S)$ is the entropy of the training set S , given by

$$\mathcal{H}(S) = - \sum_{i=1}^I p_i \log_2 p_i, \quad (3.2)$$

where p_i is the proportion of elements in S that belong to class i , which corresponds in this experiment to the ratio of the number of trials in which sphere sph_i was chosen to the number of trials in the training set, i.e.,

$$p_i = \frac{n_i}{\#(S)}. \quad (3.3)$$

Finally, $\mathcal{H}(S|A)$ is the entropy of the training set (S) split by the values of feature A , its value corresponds to

$$\mathcal{H}(S|A) = \sum_{v \in \text{Values}(A)} \frac{\#(S_v)}{\#(S)} \mathcal{H}(S_v), \quad (3.4)$$

where S_v corresponds to the subset obtained by splitting S by the value v of feature A . In addition to using \mathcal{I} to choose its decision nodes, the c4.5 algorithm uses additional rules to simplify the resulting tree and reduce overfitting.¹

The advantage of c4.5, as stated in Section 3.2, is that it produces easy to interpret rules and a relatively high accuracy. In this study's scope, the decision trees allowed representation and analysis of the possible participant strategies to solve each task. To further avoid tree over-fitting to the experimental data, 10-fold cross validation was used on the generated tree models.

¹The sole usage of \mathcal{I} for decision-node selection is characteristic of the ID3 algorithm [Quinlan 1986], a predecessor of c4.5.

Due to the undirected nature of the studied task, participants may exhibit an overall preference for the left or right sphere. Therefore, the results obtained from c4.5 are compared to a zero-rule classifier that always predicts the most frequently chosen sphere. Such a classifier is equivalent to a 1-node decision tree generated from an empty feature-set (\emptyset).

3.3.6.1 Model comparison

Models are compared based on their accuracy (acc), or percentage of correct predictions. To give a rough sense of the uncertainty around the accuracy (acc) measurements of each model, the 95% CI of acc is approximated using

$$acc \pm z_{.95} \sqrt{\frac{acc(1-acc)}{numTrials}}, \quad (3.5)$$

where $z_{.95} \approx 1.96$ is the 97.5th percentile of the normal distribution. According to Mitchell [1997, p. 141], this 95% CI approximation using the normal distribution is adequate when $numTrials \geq 30$, which is the case of this study.

In cases where the confidence interval comparison is not sufficient to elucidate “significant” model differences, an additional test is warranted to gauge the uncertainty in the difference of accuracies. Based on Equation (5.13) from [Mitchell 1997, p. 144] the 95% CI for the difference between the accuracies of two classifiers is given by

$$(acc_a - acc_b) \pm z_{.95} \sqrt{\frac{acc_a(1-acc_a)}{numTrials} + \frac{acc_b(1-acc_b)}{numTrials}} \quad (3.6)$$

Note, however, that all feature-sets are tested on the same trials, thus, the confidence intervals given by Equation (3.6) may be too conservative [Mitchell 1997, p. 144].

3.4 Task-specific feature analysis

Based on the initial wand position (P_w), sphere diameter (W_1, W_2) and initial sphere position (P_1, P_2), different values were calculated, including wand-sphere distances,

$$D_1 = |P_w - P_1|, \quad (3.7)$$

$$D_2 = |P_w - P_2|, \quad (3.8)$$

wand-sphere indices of difficulty using Shannon’s formulation (Equation (2.5)),

$$ID_1 = \log_2 \left(\frac{D_1}{W_1} + 1 \right), \quad (3.9)$$

$$ID_2 = \log_2 \left(\frac{D_2}{W_2} + 1 \right), \quad (3.10)$$

inter-sphere distance,

$$D_{sph} = |P_2 - P_1|, \quad (3.11)$$

inter-sphere indices of difficulty,

$$ID_{1,2} = \log_2 \left(\frac{D_{sph}}{W_2} + 1 \right) \quad (3.12)$$

$$ID_{2,1} = ID_{1,2} = \log_2 \left(\frac{D_{sph}}{W_1} + 1 \right), \quad (3.13)$$

and total indices of difficulty

$$ID_{T1} = ID_1 + ID_{1,2}, \quad (3.14)$$

$$ID_{T2} = ID_2 + ID_{2,1}, \quad (3.15)$$

Using the c4.5 algorithm, feature-sets $\{ID_{T1}, ID_{T2}\}$, $\{ID_{1,2}, ID_{2,1}\}$, $\{ID_1, ID_2\}$ and $\{D_1, D_2, r_1, r_2\}$ were evaluated to predict the first selected sphere.

3.4.1 Results

Participants showed an overall preference for the right sphere ($\sim 64\%$). Thus, as shown in the last row of Table 3.1, the decision tree generated using the zero-rule approach always predicted sph_2 as the selected sphere with an accuracy of 63.81%, 95% CI [61.41%, 66.21%].

Decision trees generated with the c4.5 algorithm from feature-sets 1–4 (see Table 3.1) yielded approximately $71\% \pm 2.26\%$ accuracy on predicting the selected sphere, with a 95% confidence level, which is significantly better than both chance (50%), and the aforementioned zero-rule predictor ($64\% \pm 2.4\%$). Even though all of the feature-sets yielded similar accuracies, the generated tree for feature-set 1 was more complex than those generated for feature-sets 2–4, making it less practical and perhaps over-fitted to the data [Mitchell 1997], especially considering that its 4 non-leaf nodes were generated from only 2 attributes.

Table 3.1. Accuracy and 95% confidence intervals for the evaluated feature-sets

	Feature-set	Size	Leaves	acc	95% CI
1	ID_{T1}, ID_{T2}	9	5	70.56%	[68.28%, 72.83%]
2	$ID_{1,2}, ID_{2,1}$	5	3	71.21%	[68.95%, 73.47%]
3	ID_1, ID_2	5	3	70.95%	[68.68%, 73.21%]
4	D_1, D_2, r_1, r_2	5	3	71.21%	[68.95%, 73.47%]
5	\emptyset	1	1	63.81%	[61.41%, 66.21%]

Interestingly, the fact that feature-sets 2 and 4 had the same accuracy, 95% CI, and a similar tree configuration (3 leaves out of 5 nodes) implies that they are equivalent. This may seem surprising, given that the only relevant factors in the inter-sphere indices of difficulty ($ID_{1,2}$,

$ID_{2,1}$) which compose feature-set 2, are sphere diameters (W_1, W_2),² whereas feature-set 4 is composed not only of sphere radii (r_1, r_2), but also of wand–sphere distances (D_1, D_2). A closer look at the generated decision tree for feature-set 4 (Figure 3.4), however, reveals that the resulting tree included only sphere radii; wand–sphere distances (D_1, D_2) were probably ignored by the c4.5 algorithm on the basis of low information gain. Thus, it is reasonable to conjecture that the radii provide an equivalent information gain to feature-sets 2 and 3, since their generated trees had similar configurations and yielded an equivalent accuracy.




Figure 3.4. Decision tree for feature-set 4, suggesting that participants based their decisions only on sphere size, with a preference for the right sphere. Leaves represent prediction outcomes (sph_1 is the left sphere, and sph_2 is the right sphere), while the other nodes represent tested attributes (r_1 or r_2). The numbers in parenthesis within the leaves represent the total number of instances that fall into that leaf, over the number of incorrectly predicted instances among these instances.

The overall tendency for choosing the right sphere (sph_2) first is likely due to the majority of the participants being right-handed; unfortunately, there weren't enough left-handed participants to evaluate the effects of handedness on the generated models. According to the decision tree generated from feature-set 4 presented in Figure 3.4, participants would only choose the left sphere (sph_1) first if the right one (sph_2) was smaller; sph_2 would be selected first if its radius was greater or equal than that of sph_1 .

3.5 Generalizable feature analysis

At each frame, measurements that could relate the target positions to the participant's head pose (P_h, Q_h), and wand position (P_w) were calculated. Subsequently, these measurements

²Inter-sphere distances are equal for all of the trials ($D_{sph} = 0.5$), annulling their influence on $ID_{1,2}$ and $ID_{2,1}$ and, thus, on feature-set 2, see Equations (3.12) and (3.13).

were averaged in a time window, and different feature-sets were evaluated to predict the intended sphere.

3.5.1 Relative user-target features

First, the head-sphere vectors

$$\mathbf{P}_{ih} = P_i - P_h, \quad (3.16)$$

are calculated, where P_i corresponds to the absolute position of sphere $i \in \{1, 2\}$.

Subsequently, the dot products between the normalized head (\mathbf{H}), and \mathbf{P}_{ih} vectors

$$dot_i = \hat{\mathbf{H}} \cdot \hat{\mathbf{P}}_{ih}. \quad (3.17)$$

are calculated. The head-target dot product, dot_i , has the advantage of being an easy to interpret, normalized scalar: the closer dot_i gets to 1, the more the user's head orientation is aligned with sph_i . Since the spheres do not overlap in the user's field of view, the dot-product difference

$$\Delta dot = dot_1 - dot_2. \quad (3.18)$$

is calculated. This quantity serves to determine the relative pose of the user's head with respect to the spheres. The closer Δdot gets to 1, the more the user's head is aligned with sph_1 ; the closer the quantity gets to -1 , the more the user's head is aligned with sph_2 ; a value of 0 implies that the user's head is oriented right in the middle of both spheres.

Finally, the difference between wand-sphere distances

$$\Delta D = D_1 - D_2, \quad (3.19)$$

is calculated. Where D_1 , and D_2 are the wand-sphere distances given by Equations (3.7) and (3.8). Similar to Δdot , ΔD serves to determine the relative position of the user's wand with respect to the spheres. A positive ΔD implies that the wand is farther from sph_1 ; a negative quantity implies that the wand is farther from sph_2 ; 0 implies that the wand is equidistant from both spheres.

3.5.1.1 Distance score feature

To validate the usefulness of the proposed user-target features, their predictive accuracy is compared to the distance scoring function proposed by Ortega [2013], presented in Equation (2.23). Since the experimental environment was composed of two targets, $N = 1$ is chosen, such that only the closest target's score is increased. Compared to the $decay = 0.5$ value used by Ortega, a higher $decay = 0.9$ value is chosen to account for the fact that most of the movement happened late in each trial; a low decay rate would not have permitted to revert

the possible target-score accumulation that occurs while participants are waiting for targets to become reachable. The resulting scoring function for each target i at frame t is therefore

$$dScore_i(t) = dScore_i(t-1) + \begin{cases} (1 - j_i)\Delta t & \text{if } j_i < 1 \\ -0.9 \Delta t & \text{if } j_i \geq 1 \end{cases}$$

$$dScore_i(t) \geq 0, \quad (3.20)$$

where $j_i = 0, \dots, I-1$ is the ascending target order based on the wand-target distance D_i .

3.5.2 Time-window selection

To cope with issues of instability and inaccuracy that affect human movement [Shadmehr et al. 2010], the values of Δdot and ΔD are averaged in a time window, rather than using their instant values.

In interactive usage contexts, both the feature averaging and the scoring function start running upon user activation. The present analysis, however, is done *post hoc*, so the functions are applied to the data during a graphically determined time window. Ideally, the time window would start before the beginning of the reaching action, while the user is specifying their intentions and actions [Norman 2002], and end before the target is reached. In the scope of this study, the Δdot profile was analyzed graphically over time, to determine an appropriate window heuristically, as shown in Figure 3.5. Other possible approaches are discussed in the future work section.

Because there is no time between trials, the starting non-zero Δdot values in Figure 3.5 are likely due to participants fixating the last sphere they touched on the previous trial. The subsequent convergence towards zero, between 0 s and 1 s suggests that their gaze is shared between both spheres, probably while making their decision. After 1 s, Δdot starts diverging again, suggesting that participants' heads are oriented towards one of the two spheres. If this is the case, the increased divergence could be related to the increased separation of the spheres in the participant's field of view, as they get closer to them. Furthermore, after 1 s, the red- and blue-colored selection labels of sph_1 , and sph_2 , are more clearly clustered above and below zero, respectively.

This graphical evidence suggests that roughly 1 s, and 1.5 s are good start, and end times for the window. These times roughly correspond to 42.5% and 63.8% of the 5th percentile of the selection times (2.35 s). Within this window, both the mean dot product difference ($\overline{\Delta dot}$) and the mean wand-target distance ($\overline{\Delta D}$) are calculated.




Figure 3.5. $\Delta\dot{t}$ vs. time. Each line corresponds to a trial, colored according to the selected sphere. The plot has been trimmed to the 5th percentile of the selection times (2.35 s).

3.5.3 Evaluation

Generalizable user–target feature-sets $\{\overline{\Delta dot}\}$, $\{\overline{\Delta D}\}$, $\{\overline{\Delta dot}, \overline{\Delta D}\}$, and $\{\overline{\Delta dot}, r_1, r_2\}$, as well as their combinations with the sphere radii, $\{\overline{\Delta dot}, r_1, r_2\}$, $\{\overline{\Delta D}, r_1, r_2\}$, and $\{\overline{\Delta dot}, \overline{\Delta D}, r_1, r_2\}$ are evaluated using the c4.5 classifier to predict the first sphere selected by the user (sph_i).

Models generated from generalizable user–target features, $\overline{\Delta dot}$ and $\overline{\Delta D}$, are compared to the scoring classifier *bestDRank*, which always predicts the chosen sphere as the one with the best *dScore*, given by Equation (3.20).

In the case of feature-sets combining generalizable features, $\overline{\Delta dot}$ and $\overline{\Delta D}$, and task-specific features, r_1 and r_2 , the baseline classifier is the decision tree generated from the best feature-set from Section 3.4, i.e., $\{r_1, r_2\}$.

The performance of *bestDRank* is simply evaluated by calculating its predictive accuracy, i.e. the ratio of correct predictions to the number of trials. The performance of the remaining features is evaluated based on the 10-fold cross validation accuracy of the decision-trees generated from each feature-set.

3.5.4 Results and discussion

3.5.4.1 Generalizable user–target features

As shown in Table 3.2, all feature-sets performed better than chance and a zero-rule predictor. On average, all of the proposed feature-sets performed better than the *bestDRank* baseline classifier. Among the individual features, $\overline{\Delta D}$ yielded less average accuracy than $\overline{\Delta dot}$ with a more complex tree, making it less practical and perhaps over-fitted to the data [Mitchell 1997, p. 67]. Furthermore, since the $\overline{\Delta D}$ accuracy was only marginally higher than that of the *bestDRank* baseline, no further analyses on the former are carried out.

Table 3.2. Tree size, number of leaves, accuracy and 95% confidence intervals for the evaluated generalizable moving-target feature-sets

Feature-set	Size	Leaves	acc	95% CI
<i>bestDRank</i>	1	1	68.09%	[65.77%, 70.42%]
$\overline{\Delta dot}$	3	2	70.69%	[68.42%, 72.96%]
$\overline{\Delta D}$	5	3	68.42%	[66.1%, 70.74%]
$\overline{\Delta dot}, \overline{\Delta D}$	11	6	71.73%	[69.48%, 73.97%]

Given that their 95% CI overlap, the difference between accuracies of *bestDRank*, $\overline{\Delta dot}$, and $\overline{\Delta dot}, \overline{\Delta D}$ is calculated, as well as the 95% CI of this difference. Results are presented in Table 3.3.

The fact that feature-set $\{\overline{\Delta dot}, \overline{\Delta D}\}$ yielded the greatest average accuracy, which was significantly better than both the baseline *bestDRank* and feature $\overline{\Delta D}$ confirms the value of using head–target and wand–target relative features to predict intention in moving-target

Table 3.3. Accuracy difference and 95% confidence intervals for the evaluated generalizable moving-target feature-sets. Asterisks (*) denote a significant difference ($\alpha = 0.05$).

Feature-set _a	Feature-set _b	Δacc	95% CI
<i>bestDRank</i>	$\overline{\Delta dot}$	-2.59%	$[-5.96\%, 0.67\%]$
<i>bestDRank</i>	$\overline{\Delta dot}, \overline{\Delta D}$	-3.63%	$[-7.01\%, -0.41\%]^*$
$\overline{\Delta dot}$	$\overline{\Delta dot}, \overline{\Delta D}$	-1.04%	$[-4.33\%, 2.2\%]$
$\overline{\Delta D}$	$\overline{\Delta dot}, \overline{\Delta D}$	-3.31%	$[-6.68\%, -0.08\%]^*$

selection. As previously stated, this is likely due to the inherent visuomotor nature of the moving-target selection tasks, where users need to fixate on the chosen target while moving their hands towards them. This model is presented in Figure 3.6.

Due to the task and evaluation differences with previous work on intention prediction in target selection, the results are not directly comparable to the latter, but suggest the potential of the presented approach. The time-window limits are likely to change according to the task (e.g., if the user has to search for their intended target in a cluttered environment), but it may be possible to detect patterns similar to Figure 3.5 when the intended target is fixated upon, which is more efficient than previous approaches that require large portions of the entire hand trajectory as predictive inputs [McGuffin and Balakrishnan 2005; Lank et al. 2007; Wonner et al. 2011].

Furthermore, in other tasks the generated tree nodes will likely have different split values than those presented in Figure 3.6. It is possible, however, that the split values for trees in other binary selection tasks will also be close to zero.

Finally, using a single relative head—target parameter, such as $\overline{\Delta dot}$, and a single wand—target relative parameter, such as $\overline{\Delta D}$, may not be useful or viable in tasks with more and differently positioned targets. A solution could be to create similar features for every possible pair of targets.

3.5.5 Combined task-specific and generalizable features

Tables 3.4 and 3.5 show that Feature-set $\{\overline{\Delta dot}, \overline{\Delta D}, r_1, r_2\}$ performed significantly better than all of the other feature-sets, surpassing the best task-specific feature-set, $\{r_1, r_2\}$, by almost 7%, and the best generalizable feature-set, $\overline{\Delta dot}, \overline{\Delta D}$, by about 6%. Unfortunately, the generated tree was too big (21 nodes) to fit in this paper.

Surprisingly, and contrary to the results from the previous section, combining the $\overline{\Delta dot}$ relative feature with the sphere radii (r_1, r_2) did not yield better accuracy than feature-set $\{\overline{\Delta D}, r_1, r_2\}$.

The fact that feature-set $\{\overline{\Delta D}, r_1, r_2\}$ performed marginally better than the baseline, $\{r_1, r_2\}$, suggests that a function of target size and distance, albeit different from ID , can adequately predict the selected sphere in this type of task. In the results from Section 3.4, however, the

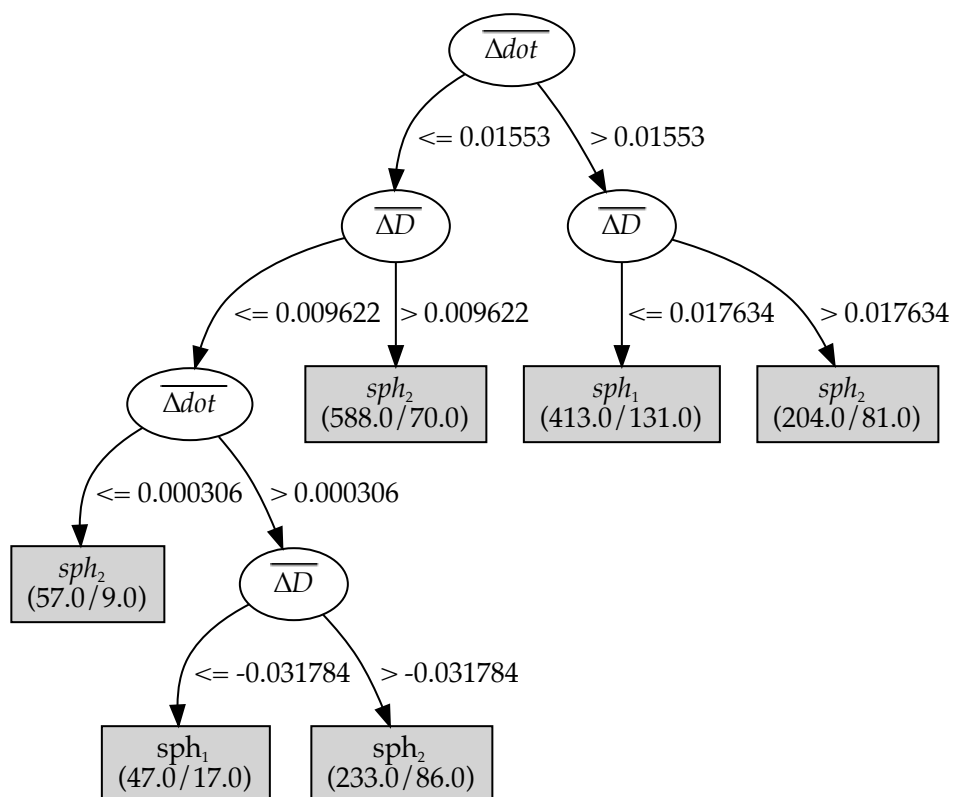


Figure 3.6. Generated decision tree for feature-set $\{\overline{\Delta dot}, \overline{\Delta D}\}$. The numbers in parenthesis within the leaves represent the total number of instances that fall into that leaf, over the number of incorrectly predicted instances among these instances.

distance D_0 —measured at the beginning of each trial—was deemed to yield less information gain than the sphere radii. The apparent increase in information gain by integrating $\overline{\Delta D}$, observed in the present work, reflects a correlation between wand and object position, as previously suggested by Ortega [2013], but only after a certain preparation time [Nieuwenhuizen et al. 2009].

Table 3.4. Tree size, number of leaves, accuracy and 95% confidence intervals for the evaluated target-based feature-sets.

Feature-set	Size	Leaves	acc	95% CI
r_1, r_2	5	3	71.21%	[68.95%, 73.47%]
$\overline{\Delta dot}, r_1, r_2$	7	4	73.35%	[71.14%, 75.55%]
$\overline{\Delta D}, r_1, r_2$	27	14	74.19%	[72.01%, 76.37%]
$\overline{\Delta dot}, \overline{\Delta D}, r_1, r_2$	21	11	78.02%	[75.95%, 80.08%]

Table 3.5. Accuracy difference and 95% confidence intervals for the target-based feature-sets. Asterisks (*) denote a significant difference ($\alpha = 0.05$), dots (.) denote a marginal difference ($\alpha = 0.1$).

Feature-set _a	Feature-set _b	Δacc	95% CI
r_1, r_2	$\overline{\Delta dot}, r_1, r_2$	-2.14%	[-5.43%, 1.04%]
r_1, r_2	$\overline{\Delta D}, r_1, r_2$	-2.98%	[-6.28%, 0.16%].
$\overline{\Delta dot}, r_1, r_2$	$\overline{\Delta D}, r_1, r_2$	-0.84%	[-4.05%, 2.32%]
$\overline{\Delta D}, r_1, r_2$	$\overline{\Delta dot}, \overline{\Delta D}, r_1, r_2$	-3.83%	[-7.04%, -0.85%]*

3.6 Discussion

Results are summarized according to the study hypotheses, each followed by a discussion.

3.6.1 Task-specific features for intention prediction in undirected moving-target selection

Considering that task-specific decision trees were built based only on the initial position of the user’s wand and the initial size and position of the spheres, predictions bore a relatively high accuracy, compared to both chance and a zero-rule predictor. It is likely, however, that the accuracy will decrease if the number of targets is increased, but it is expected that the accuracy will still be better than chance and a zero-rule predictor.

H1. *ID* can accurately predicting intention in undirected moving-target selection tasks.
Partially supported.

Even though ID served as a good task-specific feature for predicting the intended moving target, using solely sphere radii (r_1, r_2) yielded an equivalent accuracy ($\sim 72\%$). Recall that all of the evaluated ID measures, presented in Equations (3.9)–(3.15), are based on target radii. This suggests that target size yields more information gain than the initial target distance, and indicates a very basic strategy from the users in which distance does not play an important role for choosing targets.

This may be due to the fact that the spheres get closer to the user throughout each trial, eventually annulling the z component of the target's distance. Additionally, since there was some waiting time before the targets were reachable, it is possible that users prepared the starting horizontal position of their wands prior to executing the pointing task, even if instructed otherwise.

3.6.2 Generalizable features for intention prediction in undirected moving-target selection

H2. Relative user–target features can accurately predict intention in undirected moving-target selection tasks. *Supported.*

The relative head–target and wand–target features, $\overline{\Delta dot}$ and $\overline{\Delta D}$, respectively, proved successful in predicting intended targets in the studied undirected moving-target selection tasks. Combined, the features yielded a $\sim 72\%$ accuracy on predicting intended targets, which was significantly better than the isolated $\overline{\Delta D}$ feature ($\Delta acc = 3.31\%$), and the Ortega [2013] scoring function ($\Delta acc = 3.63\%$).

The relative head–target feature, Δdot , proved to be useful not only for prediction, but also for establishing the adequate time window. Currently, the window is established empirically, from the Δdot vs. t plot (Figure 3.5).

H3. Combining task-specific, and generalizable features yields better predictive accuracy than using either separately. *Supported.*

The integration of features $\overline{\Delta dot}$, and $\overline{\Delta D}$, with r_1 and r_2 , improved the predictive accuracy of intended targets in undirected moving-target selection. The combined feature-set performed significantly better than all of the other feature-sets, surpassing the best task-specific feature-set, $\{r_1, r_2\}$, by almost 7%, and the best generalizable feature-set, $\overline{\Delta dot}, \overline{\Delta D}$, by about 6%.

As opposed to the task-specific features analysis, the results revealed that combining target radii with target distances yielded better accuracies than using each feature separately, but only if the distance was measured after the trial start, within the chosen time window.

3.7 Conclusion

This chapter evaluated the usage of Fitts' ID and its separable components, D and W , as task specific features, as well as $\overline{\Delta dot}$ and $\overline{\Delta D}$, as generalizable features to predict intended targets in undirected moving-target selection tasks. All of the evaluated feature-sets performed better than chance, and a zero-rule predictor, and the combination of both types of features yielded the best accuracy ($\sim 78\%$).

The generalizable features were calculated within a time-window ending at about two-thirds of the selection time, heuristically selected based on participant gaze. Future work could explore automating this process by finding the optimal start and end window limits, by measuring different inputs. Furthermore, these times could be related to existing models, such as the Hick-Hyman Law [Hick 1952; Hyman 1953], presented in Equation (2.1). Notice, however, that in interactive contexts, the time window may not be necessary, since users may select their start and end times directly.

These results should be generalizable to different moving-target selection tasks, provided that additional factors, such as a greater number of spheres with different starting positions and different velocities, are taken into consideration. The potential of using other measures of difficulty (either objective or subjective) formulated specifically for moving-target selection to predict user intention should also be explored.

CHAPTER 4

PROSPECTIVE DIFFICULTY OF 2-D STATIC-TARGET AND MOVING-TARGET SELECTION

Chapter 3 evaluated the hypothesis presented in Chapter 1 that *ID* was predictive of user intention in undirected moving target selection, based on the premise that users form their intentions by minimizing their prospective effort, as described by *ID*. Due to the inadequacy of existing *ID* measures formulated specifically for moving-target selection, Chapter 3 resorted to using a static-target formulation of *ID*, specifically the Shannon model (Equation (2.5)) and its separate *D* and *W* components. Consistently with the results of Jagacinski et al. [1980] for the *MT* of moving-targets, the results revealed that such *ID* was limited in usefulness for predicting intention in moving-targets, thus, hinting at its inadequacy as a *PD* measure for moving-target selection as well.

In order to address these limitations, this Chapter attempts to extend the *PD-ID* paradigm to moving-target selection tasks. The proposed model is initially formulated and evaluated in 2-D tasks to avoid issues inherent to 3-D interaction, such as incorrect depth-perception, but bearing its extensibility to 3-D in mind. A general 2-D moving-target pointing task is presented in Figure 4.1, along with its relevant attributes.




Figure 4.1. Measurements relevant to the target-selection tasks considered in this chapter. The red and the green circles indicate the starting positions of the cursor and the target, respectively. *D* is the initial distance between the cursor and the target, *W* is target width, θ is the initial angle between the target and the cursor, *V* is the target speed, and ϕ is the target's movement direction relative to the initial cursor-target vector.

To achieve this goal, the current chapter introduces three models based on Fitts' *ID* for the *PD* of pointing at 2-D static targets, and 1-D and 2-D moving targets. The model formulae are presented below, in Sections 4.1.1–4.1.3.

4.1 Formulations

4.1.1 *PD* formulation for 2-D static-target pointing tasks

Two-dimensional static-target pointing tasks refer to tasks with initial target angle $\theta \in [0, 360)^\circ$. For this type of task, the model proposed in this chapter separates the target-angle (θ) effect from the $ID = \log_2(2D/W)$ term, such that

$$PD = a_\theta + b \log_2 \left(\frac{2D}{W} \right),$$

where a_θ and b are empirically derived coefficients; a_θ is the per-angle θ intercept. This formulation is based on Fitts original formulation for *MT*,¹ and the 2-D *ID* extensions by Murata and Iwase [2001], Appert et al. [2008], and Grossman and Balakrishnan [2004]. This equation is developed in Section 4.2.2.

4.1.2 *PD* formulation for 1-D moving-target pointing tasks

One-dimensional moving-target pointing tasks refer to tasks with initial target angle $\theta \in \{0, 180\}^\circ$, and target moving-direction $\phi \in \{0, 180\}^\circ$. For this type of task, the proposed model is based on the two-part model by Hoffmann and Chan [2012] with per-speed (V) regression coefficients, such that

$$PD = a_V + b_V \sqrt{D} + c_V \log_2 \left(\frac{2D}{W} \right),$$

where a_V , b_V , and c_V are empirically derived coefficients. Based on an analogy between control–display gain (G) and V , this formulation is inspired by the work of Shoemaker et al. [2012], who found that two-part models can be used to describe static-target pointing with different levels of G .

This formula is developed in Section 4.2.3. Additionally, as shown in Section 4.4.2.2, coefficients a_V , b_V , and c_V appear to be linearly related to the speed V of the target.

4.1.3 *PD* formulation for 2-D moving-target pointing tasks

Two-dimensional moving-target pointing tasks refer to tasks where the target's moving direction $\phi \in [0, 360)^\circ$, regardless of the initial target angle θ . For this type of task, the

¹The choice of Fitts formulation for *ID* is explained in Section 4.2.1.

proposed model is an extension of the 1-D moving-target model, separating the initial target distance D into two components, such that

$$PD = a_V + b \sqrt{D_s} + c_V \sqrt{D_m} + d \log_2 \left(\frac{2D_m}{W} \right)$$

$$D_s = |D \sin \phi|$$

$$D_m = |D \cos \phi|,$$

where a_V , b , c_V , and d are empirically derived coefficients; a_V and c_V are per-speed V coefficients. The development of this formula is both analytical, as explored in Section 4.2.4, and empirical, as explored in Section 4.4.2.2. Additionally, as shown in Section 4.4.2.2, a_V and c_V appear to be linearly related to the speed V of the target.

4.2 Modeling

Based on the literature review presented in Chapter 2, and motivated by the goal of extending the PD - ID paradigm to 2-D moving-target pointing tasks, this section presents an incremental development of such a model. First, the choice of the ID formula for the PD of 1-D static-target tasks is presented, followed by the extension to 2-D static-target tasks, 1-D moving-target tasks, and, finally, 2-D moving-target tasks.

4.2.1 Choice of ID formulation

In contrast with the majority of the post-1990's Fitts-related literature in HCI, Fitts' ID formulation is chosen over the Shannon formulation. The most important reason for this choice is the simplicity in deriving the formulae proposed to model prospective difficulty, notably the 2-D model presented in Section 4.2.2. Another important reason for this choice is the compatibility of Fitts' ID with prospective pointing tasks: if there were a prospective task in which the pointer started at the boundary of a static target (i.e., $D = W/2$), the expected difficulty for an imagined movement should be 0, allowing one to hypothesize the intercept $a = 0$ in Equation (4.1). The latter is not the case for the Shannon ID , since the intercept represents a situation in which the pointer starts at the center of the target. Tasks in which $D < W/2$, however, are out of the scope of this study as they may be hard to interpret prospectively, besides not being adequately modeled in the Fitts paradigm [Soukoreff and MacKenzie 2004, p. 768]. Lastly, choosing the Fitts ID is consistent with the subjective difficulty literature, allowing comparison of the current results to previous work in the area.

Therefore, following the previous work in subjective difficulty and Fitts' Law [Delignières and Famose 1992; Slifkin and Grilli 2006; Grilli 2011; Chan and Hoffmann 2013], the model for the PD of pointing at a 1-D static target is

$$\begin{aligned}
PD &= a + b ID \\
ID &= \log_2 \left(\frac{2D}{W} \right).
\end{aligned} \tag{4.1}$$

4.2.2 Formulations of prospective difficulty in 2-D static-target pointing tasks

To the author's knowledge, there is no extension to Fitts' original ID that explicitly accounts for the target angle θ . Indeed, the 2-D extensions to Fitts' ID presented in Equations (2.10)–(2.12) are based on the Shannon formulation. However, analogous models based on Fitts original ID can be formulated as well.² For example, Equation (2.10), by Murata and Iwase [2001], is analogous to

$$ID_{MI} = \log_2 \left(\frac{2D}{W} \right) + c \sin \theta, \tag{4.2}$$

Equation (2.11), by Appert et al. [2008], is analogous to

$$ID_{ACB} = \log_2 \left(\frac{2D}{W} + \frac{2D}{H} + 0.6 \cos(\theta) \frac{2D}{\min(W, H)} \right), \tag{4.3}$$

and Equation (2.12), by Grossman and Balakrishnan [2004], is analogous to

$$ID_{GB} = \log_2 \left(\sqrt{f_W(\theta) \left(\frac{2D}{W} \right)^2 + f_H(\theta) \left(\frac{2D}{H} \right)^2} \right). \tag{4.4}$$

These three formulations are easier to manipulate than their Shannon counterparts. For example, by isolating ID from the angle and shape effects in Equations (4.3) and (4.4), the three equations can be expressed in a similar form,

$$ID_{2-D} = \log_2 \left(\frac{2D}{W} \right) + f \left(\theta, \frac{W}{H} \right). \tag{4.5}$$

For Equation (4.2) there is no shape effect, only an angle effect expressed as

$$f \left(\theta, \frac{W}{H} \right) = f(\theta) = c \sin \theta. \tag{4.6}$$

For Equation (4.3), the angle and shape effects are expressed as

$$f \left(\theta, \frac{W}{H} \right) = \log_2 \left(\frac{W}{H} + 0.6 \cos(\theta) \frac{W}{\min(W, H)} + 1 \right). \tag{4.7}$$

²This is done by multiplying D by 2, and removing the “+1” terms from the Shannon-derived formulae. A similar approach was taken by Shoemaker et al. [2012] to reformulate the distant pointing model by Kopper et al. [2010].

Finally, the angle and shape effects for Equation (4.4) are expressed as

$$f\left(\theta, \frac{W}{H}\right) = \log_2 \left(\sqrt{f_W(\theta) + f_H(\theta) \left(\frac{W}{H}\right)^2} \right). \quad (4.8)$$

With these reformulations, Fitts' ID becomes nested within ID_{2-D} (i.e., $ID_{2-D} = ID$, given $f(\theta, W/H) = 0$). Based on this equivalency, the linear relation between PD and ID found in the literature also extends to ID_{2-D} , thus

$$PD = a + b ID_{2-D}, \quad (4.9)$$

which can also be formulated as a model with one intercept per $(\theta, W/H)$ condition, such that

$$PD = a_{\theta \times W/H} + b ID, \quad (4.10)$$

where $a_{\theta \times W/H} = a + b f(\theta, W/H)$. For univariate targets, where $W = H$, the formula is further simplified to

$$PD = a_{\theta} + b ID, \quad (4.11)$$

where $a_{\theta} = a + b f(\theta)$. If no specific form of $f(\theta, W/H)$, or $f(\theta)$ is assumed, Equations (4.10) and (4.11) can be solved using the ubiquitous linear least-squares method, which does not suffer from the shortcomings of its non-linear counterpart described at the end of Section 2.1.2.2. For example, by using indicators, or dummy variables (x_i) for the i levels of θ , Equation (4.11) can be represented as

$$PD = \alpha_1 + \alpha_2 x_2 + \alpha_3 x_3 + \dots + \alpha_{i-1} x_{i-1} + \alpha_i x_i + b ID, \quad (4.12)$$

such that

$$(x_2, x_3, \dots, x_{i-1}, x_i) = \begin{cases} (0, 0, \dots, 0, 0) & \text{if } \theta = \theta_1 \\ (1, 0, \dots, 0, 0) & \text{if } \theta = \theta_2 \\ (0, 1, \dots, 0, 0) & \text{if } \theta = \theta_3 \\ \dots & \dots \\ (0, 0, \dots, 1, 0) & \text{if } \theta = \theta_{i-1} \\ (0, 0, \dots, 0, 1) & \text{if } \theta = \theta_i, \end{cases} \quad (4.13)$$

where $\theta_1, \dots, \theta_i$, are the different levels of θ .

4.2.3 Formulations of prospective difficulty in 1-D moving-target pointing tasks

As discussed in the Related Work section, the existing formulations for predicting MT in moving-target selection are not compatible with the current $PD-ID$ framework. Equation (2.15), by Jagacinski et al. [1980], is simple but cannot be reduced to Fitts' Law, so its use would imply modeling static and moving targets differently; Equations (2.17) and (2.18), by Hoffmann [1991], are derived from Fitts' Law, but require the percentage of actual captures (unknown in prospective action), otherwise resulting in potentially harder and unreliable calculations. Therefore, it seems necessary to introduce another formulation to model the PD of moving-target selection, that is both simple and compatible with Fitts' ID .

To this end, an analogy between static-target selection under the effect of control-display gain, and moving target selection without gain is given. Under control-display gain,

$$G = \frac{V_{display}}{V_{control}}, \quad (4.14)$$

the movement time (MT) required to reach for a static target located at position x is

$$\begin{aligned} MT &= \frac{x}{V_{display}} \\ &= \frac{x}{G \times V_{control}}. \end{aligned} \quad (4.15)$$

On the other hand, with $G = 1$ (i.e., $V_{control} = V_{display}$) the MT required to reach for a target located at position x , and moving with velocity V_{target} is

$$MT = \frac{x}{V_{control} - V_{target}}. \quad (4.16)$$

Combining (4.15) and (4.16) gives the equivalency

$$G = 1 - \frac{V_{target}}{V_{control}}. \quad (4.17)$$

Thus, reaching for a static target with $G > 0$ is equivalent to reaching for a moving target with $|V_{control}| > |V_{target}|$. Reaching for a static target with $G = 0$ (i.e., when the display does not react to the control input) is equivalent to reaching for a moving target with $V_{control} = V_{target}$ —in both cases the target always remains at the same distance and pointing at it is impossible. Likewise, reaching for a static target with $G < 0$ (i.e., when the display's velocity is opposite to the control input) is equivalent to reaching for a moving target with $|V_{control}| < |V_{target}|$ —in both cases the target gets farther from the cursor as time elapses. The major shortcoming of this analogy is that it assumes that the user's control velocity is constant.

This is the basis for hypothesizing that two-part formulations model 1-D moving-target selection, analogous to the way Shoemaker et al. [2012] apply a two-part model to 1-D static-target selection with gain. Furthermore, based on the approximate equivalency between the

two-part formulations of Welford and Hoffmann, described in Section 2.1.2.1, the latter should also model moving-target selection with the added benefit of mathematical correctness and, more importantly, the apparent linear relation between the regression coefficients and V .³

For the prospective difficulty of 1-D moving-target selection tasks, the proposed relationship is

$$PD = a_V + b_V \sqrt{D} + c_V \log_2 \left(\frac{2D}{W} \right), \quad (4.18)$$

where V is the target velocity. The form of this model is similar to that of Equation (2.15), by Jagacinski et al. [1980], with $a_V = a - d(V + 1)$, and $c_V = c(V + 1)$. Aside from the log and square-root terms, the biggest difference in both models is that Jagacinski's model does not allow the effect of D to be different for each V . However, it seems logical that higher target speeds will require faster reaction times, hence more ballistic-type motions, which are represented by the \sqrt{D} term in Hoffmann's two-part model.

4.2.4 Formulations of prospective difficulty in 2-D moving-target pointing tasks

The motion required to point at a target moving in 2-D can be divided in two components, one perpendicular to the target's velocity, modeled by Equation (4.11), and one parallel to the target's velocity, modeled by Equation (4.18). Under these assumptions, the corresponding model is hypothesized as

$$PD = a_{V \times \phi} + b_V \log_2 \left(\frac{2D_s}{W} \right) + c_V \sqrt{D_m} + d_V \log_2 \left(\frac{2D_m}{W} \right), \quad (4.19)$$

where D_s is the perpendicular, or "static distance," and D_m is the parallel, or "moving distance." Since the true D_s and D_m are unknown in prospective movements, it is assumed that D_s , D_m , and D form a right triangle as shown in Figure 4.2, so that $D_s = |D \sin \phi|$, and $D_m = |D \cos \phi|$, where ϕ is the target's movement direction relative to the initial cursor-target vector.⁴




Figure 4.2. Two-dimensional moving-target model. The red circle is the starting point, the dark green circle is the moving target.

³It can be shown that the "Welford" coefficients in [Shoemaker et al. 2012, Table XI], with the corresponding transformations described in Equations (2.8)–(2.1.2.1), vary linearly with G .

⁴In the more general case, when θ and ϕ are known, $D_s = D \frac{\sin \phi}{\sin(\phi + \theta)}$ and $D_m = D \frac{\sin \theta}{\sin(\phi + \theta)}$.

This simplification is compensated by having different b_V , c_V , and d_V parameters per speed, and a different intercept $a_{V \times \phi}$ per velocity (V, ϕ), allowing D_s and D_m to be rescaled, which approximates the more general case in which D_s , D_m , and D do not form a right triangle. The per-velocity intercept also allows for asymmetry on the effect of ϕ .

Dropping the denominator within the $\log_2(2D_s/W)$ term, and using the approximate log-square-root equivalence suggested by Hoffmann and Chan [2012] results in,

$$PD = a_{V \times \phi} + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V \log_2 \left(\frac{2D_m}{W} \right). \quad (4.20)$$

Under this model the complete motion would be described in three parts—two ballistic motions, one to align the cursor with the target’s movement axis and one in-line with the target’s movement axis, and a homing in motion in-line with the target’s movement axis.

4.3 Methods

4.3.1 Apparatus

To evaluate the proposed models, an experiment was implemented using Scalable Vector Graphics (svg) [W3C 2011] and deployed in an anonymous, online Qualtrics survey.

4.3.2 Participants

Participants were recruited by distributing the survey link using social media, and yielded 49 respondents who answered the entire survey, out of 83 partial attempts—only the responses from the full attempts were retained. Twenty-one participants reported being students, and it was determined via timestamps that 25 participants were in the Americas, one in India, and the rest in either Europe or Africa.

Participants were aged 20 to 48 years old ($M = 29.65$, $Mdn = 29$); there were 14 females and 35 males. Most respondents reported being right handed (45), no respondents reported being ambidextrous.

4.3.3 Procedure

Upon opening the survey website, each participant was asked to accept a study consent form. If agreed, an animation of a moving object was presented. After playing the animation, the participant was asked if a moving object was visible on the screen—this question served as a browser check. After successfully completing the browser check, the participant was asked to fill in a short questionnaire concerning background and demographic information.

Subsequently, the participant was presented with two blocks of questions. In each question, the participant was asked to assess the prospective difficulty of touching a circular target from a starting point as fast as possible, without actually touching the screen or making any finger


movements. The participant was advised to maximize the survey window prior to the start of each block, to avoid distractions and reduce scrolling. In the first block each target was static; in the second block each target was moving.

Each question was presented on a separate page and rated using radio buttons, as suggested by Van Schaik and Ling [2007]. In each question the participant judged pointing difficulty using a 9-point “Likert-type” scale, where 1 was labeled as “Very easy” and 9 was labeled as “Very difficult.”


To give a scoring reference, a sample image (the anchor) corresponding to the middle difficulty (5), was given both before each block and in the header of each question. The same static-target anchor was given for both static and moving targets.

To prevent the participant from trying to execute the pointing task, the mouse pointer disappeared when hovering over each image. There was no time limit for each answer and, in the case of moving targets, the animations could be replayed at the participant’s discretion. A screenshot of one of the moving-target questions is presented in Figure 4.3.

Given the following image, consider that the difficulty of touching the blue target starting from the red point is 5.



What is the difficulty of touching the green target **as fast as possible** starting from the red point?



play

Very easy 1 2 3 4 5 6 7 8 Very hard 9

>>

Survey Powered By [Qualtrics](#)

Figure 4.3. Screenshot of a moving-target question with $\phi = 45^\circ$. The anchor is represented with a blue target.

To speed up survey completion, once the participant selected a *PD* rating, the survey would automatically advance to the next question, as opposed to selecting an answer and clicking on the “next” button; the “next” button served only to skip questions, and there was

no “back” button to revisit previous answers. The trade-off was that the participant could not correct misclicked responses, but this was mitigated by implementing “auto-advance” in the consent and browser-check questions, so that the participant would expect the same behavior when answering horizontally-arranged multiple-choice questions. Nevertheless, two participants commented on a few misclicks they made during the *PD* ratings.

At the end of each block, the participant was asked to rate how seriously they answered the preceding block, using a 1–5 “Likert-type” scale, where 1 was labeled as “Not seriously at all” and 5 was labeled as “Very seriously.”

4.3.4 Design

A within-subjects, factorial design was used, with two blocks of trials.

The first block consisted of static targets. Factors were target distance, $D \in \{200, 400\}$ pixels (px); target radius, $r \in \{20, 40, 80\}$ px; and target angle, $\theta \in \{-45, 0, 45\}^\circ$. Notice that θ was measured clockwise, consistently with the *y-down* screen coordinate system. Each of the 18 conditions was presented once to the participant in a random order. A sample of this task, with a depiction of its parameters, is shown in Figure 4.4.




Figure 4.4. Static-target task. The red circle is the starting point, the green circle is the target.

The second block consisted of moving targets. Factors were initial target distance, $D \in \{200, 400\}$ px; target radius, $r \in \{20, 40, 80\}$ px; target speed, $V \in \{200, 400, 800\}$ px/s; and initial target-movement-direction, $\phi \in \{-45, 0, 45\}^\circ$; the initial target angle θ was always 0° . Each of the 54 conditions was presented once to the participant in a random order. A sample of this task, with a depiction of its parameters, is shown in Figure 4.5.




Figure 4.5. Moving-target task. The red circle is the starting point, the green circle is the moving target.

The anchor for both static- and moving-target tasks was set as ($D = 200\text{ px}$, $r = 20\text{ px}$, $\theta = 0^\circ$, $V = 0\text{ px/s}$). The choice of the same static-target anchor for both types of task was based on the notion that moving-target selection is an extension to static-target selection.

4.3.5 Data integrity

Within the 49 complete survey attempts, two respondents rated more than 75% of their static-target responses with the minimum value (1), which suggests that they did not relate to the given anchor to rate their answers, therefore, all of their responses were discarded. Out of the 47 remaining attempts, six respondents rated more than 75% of their 1-D moving-target responses with the minimum value (1), which suggests that they may have been executing a “crossing” task, in which they waited for the moving target to go over the cursor rather than touching the target as fast as possible, as instructed,⁵ therefore, all of their responses were removed. Within the 41 retained attempts, there were 2 missing answers in the static block, and none in the moving block. In total, there were 736 complete answers for the static block, and 2214 for the moving block.

4.3.6 Statistical methods

In general, the analyses in this chapter are conducted using Bayesian methods using Markov Chain Monte Carlo (MCMC). Bayesian methods have several advantages compared to traditional frequentist methods, including the use of prior knowledge, the generation of full parameter distributions, the ability to test complex models, and robustness against unbalanced samples and multiple comparisons [Kruschke et al. 2012].

By generating full distributions for the parameters of each model, Bayesian methods allow the formulation of probabilistic statements on the parameters and the models, as opposed to relying on approximate confidence intervals, or the seemingly unreliable p -values [Kruschke et al. 2012; Kaptein and Robertson 2012; Dragicevic et al. 2014]. Within this framework, testing for the “statistical significance” of null, or point hypothesis are avoided; instead, the emphasis is placed on the estimation of the different parameters, and the assessment of their credible values using the 95% highest density intervals (HDI) of their distributions. Following the recommendations from Gelman et al. [2013], the point estimates of each parameter are given using the median value of the posterior distribution.

4.3.6.1 Models for inferential statistics

Since participant responses are recorded in an ordinal scale, to evaluate the main effects of the experimental conditions and their interactions, on PD , the ordered probit regression model with per-subject intercept of the form

⁵This behavior was identified in the comments section of the survey, wherein a respondent reported being unsure if touching the target meant “crossing” it or “landing” in it, ultimately opting for the former.

$$\begin{aligned}
PD_i &= \begin{cases} 1 & \text{if } PD_i^* < \gamma_1 \\ 2 & \text{if } \gamma_1 < PD_i^* \leq \gamma_2 \\ \dots & \\ 8 & \text{if } \gamma_7 < PD_i^* \leq \gamma_8 \\ 9 & \text{if } \gamma_8 < PD_i^* \end{cases} \\
PD_i^* &= \alpha_i + \mathbf{X}_i \boldsymbol{\beta} + \epsilon_i \\
\alpha_i &\sim \mathcal{N}(0, \sigma_u^2) \\
\epsilon_i &\sim \mathcal{N}(0, 1)
\end{aligned} \tag{4.21}$$

is used, where α_i is the per-subject intercept, and $\boldsymbol{\beta}$ is the coefficient vector for \mathbf{X}_i that includes different predictors according to the experimental block.

PD^* and γ_k are both defined in the probit scale, however, for simplicity, all the estimates in the inferential statistics sections are presented in the original 1–9 measurement scale.⁶ Once again, the emphasis is on estimating the effect sizes, corresponding to the coefficients of the different predictors within $\boldsymbol{\beta}$, rather than on testing the statistical significance of those predictors.

The probit models are fit using MCMC sampling, using Stan [2014] via R [2014]. For each model, four MCMC chains of 5,000 iterations each, including 2,500 warm-up iterations, are drawn, for a total of 10,000 saved simulations. The convergence of the chains, and the goodness-of-fit of the model is verified to ensure the validity of the inference, but these metrics are not presented in the main body of the chapter.

4.3.6.2 Models for Regression Analysis

Unless otherwise indicated, a modified version of the Bayesian linear regression scripts by Kruschke et al. [2012] are used to regress the proposed models for PD . Each model is fit in Stan [2014], with four MCMC chains of 30,000 iterations, including 5,000 warm-up iterations, for a total of 100,000 saved iterations.

Model fit is assessed using the posterior distribution of the model's standard deviation σ , and R^2 , calculated as the percentage of explained variance [Gelman and Hill 2007, p. 41],

$$R^2 = 1 - \frac{\sigma^2}{s_y^2}, \tag{4.22}$$

where s_y is the standard deviation of y . The central value of this distribution is close to the so-called adjusted r-squared (R_{adj}^2).⁷ As opposed to the traditional R^2 , used in most of the

⁶The scaling procedure consists in dividing each coefficient in $\boldsymbol{\beta}$ by $M(\gamma_2 - \gamma_1, \dots, \gamma_8 - \gamma_7)$.

⁷The equivalency between the adjusted and unadjusted R^2 is given by, $R_{adj}^2 = 1 - (1 - R^2) \frac{n-1}{n-p-1}$, where n is the number of data points, and p is the number of fitted parameters, not including the intercept.

Fitts' literature, this estimate does not necessarily increase with the addition of new predictors, which facilitates direct model comparison between nested models of different complexities.

The point estimates of R^2 may also be lower than those obtained using least-squares regression, since the R^2 value is optimized in the latter, but not in Bayesian linear regression. The difference in point estimates, however, is likely small, and less interesting than the 95% credible intervals of the R^2 distribution.

4.3.7 Hypotheses

All of the following hypotheses assume univariate targets in a 2-D space, i.e., circles. Thus, the difference between task dimensions (1-D, and 2-D) is given by the angles θ and ϕ of the target, with respect to the starting position.

4.3.7.1 Static-target selection

H1. Fitts' *ID* model with per-angle intercept, described in Equation (4.11), models the prospective difficulty of two-dimensional static-target selection tasks better than the classic Fitts' *ID* model with single intercept, described in Equation (4.1).

4.3.7.2 Moving-target selection

H2. Hoffmann's two-part model with per-velocity coefficients, described in Equation (4.18), accurately models the prospective difficulty of one-dimensional moving-target selection tasks.

H2.1. The coefficients a_V , b_V , and c_V of Equation (4.18) vary linearly with V .

H3. The three-part model with per-speed coefficients, and per-velocity intercept, described in Equation (4.20), accurately models the prospective difficulty of two-dimensional moving-target selection tasks.

H3.1. In Equation (4.20) the coefficients b_V , c_V , and d_V vary linearly with V , as well as the coefficients $a_{V \times \phi}$ within the same angle ϕ .

4.4 Analysis

4.4.1 Static-target block

4.4.1.1 Inferential statistics

To evaluate the effects on *PD* of D , r , θ , and their interactions, the ordered probit regression of Equation (4.21) is fit with the following 13 predictors:

- An indicator for $D = 400$

- Two indicators for $r \in \{40, 80\}$
- Two indicators for $\theta \in \{-45, 45\}$
- Two indicators for $D \times r$
- Two indicators for $D \times \theta$
- Four indicators for $r \times \theta$,

which implies that the base values, i.e., $\mathbf{X}_i = \mathbf{0}$, are set to $D = 200$, $r = 20$, and $\theta = 0$.

Main Effects. The estimates in Figure 4.6 show that there are strong main effects of D , and r , but relatively small effects of θ .




Figure 4.6. Main-effect coefficient plot for the fitted ordered probit model on static-target PD ratings. Points, thick lines, and thin lines represent the medians, 50%, and 95% HDI for each of the coefficients in β .

As expected, PD increases with D , and decreases with larger radii r . Similarly, downward targets ($\theta = 45$) were perceived as more difficult than horizontal ones, by about 0.56, 95% HDI [0.13, 0.98]. There is also some evidence that the angle effect on PD is asymmetrical, with downward angles being perceived as more difficult than upward angles, $p(\beta_{\theta=45} > \beta_{\theta=-45}) = 0.89$; however, the average difference in PD between upward and downward angles is very small, $Mdn(\beta_{\theta=45} - \beta_{\theta=-45}) = 0.26$, 95% HDI [-0.15, 0.69]. This is surprising, and contradicts the findings of Whisenand and Emurian [1995], who found that upward angles took more time to be reached; this discrepancy could indicate that PD is not completely consistent with MT .

Two-way Interactions. As shown in Figure 4.7, the interaction effects have small magnitudes compared to the main effects, indicating that most of the changes in PD are explained by the latter.

The $D \times r$ interaction shows that, at $D = 400$, the main, negative effect of $r = 80$, $Mdn(\beta_{r=80}) = -3.66$, 95% HDI [-4.17, -3.1], is slightly attenuated by about 0.62, 95% HDI [0.15, 1.07].




Figure 4.7. Two-way interaction coefficient plot for the fitted ordered probit model on static-target *PD* ratings. Points, thick lines, and thin lines represent the medians, 50%, and 95% HDI for each of the coefficients in β .

Similarly, the $D \times \theta$ interaction shows that, at $D = 400$, both main angle effects are considerably reduced and made almost symmetrical,

$$\begin{aligned} \text{Mdn}(\beta_{\theta=-45} + \beta_{D=400 \times \theta=-45}) &= 0.19, \text{ 95\% HDI } [-0.22, 0.61], \text{ vs.} \\ \text{Mdn}(\beta_{\theta=45} + \beta_{D=400 \times \theta=45}) &= 0.15, \text{ 95\% HDI } [-0.27, 0.56]; \end{aligned}$$

suggesting that at high distances, angle differences of $\theta = \pm 45$ contribute less to *PD*.

On the contrary, on targets with radii $r > 20$ the asymmetrical effect of θ observed in the main effects is exacerbated,

$$\begin{aligned} \text{Mdn}[(\beta_{\theta=45} + \beta_{r=40 \times \theta=45}) - (\beta_{\theta=-45} + \beta_{r=40 \times \theta=-45})] &= 0.45, \text{ 95\% HDI } [0.03, 0.9], \text{ and} \\ \text{Mdn}[(\beta_{\theta=45} + \beta_{r=80 \times \theta=45}) - (\beta_{\theta=-45} + \beta_{r=80 \times \theta=-45})] &= 0.56, \text{ 95\% HDI } [0.11, 1.04]. \end{aligned}$$

4.4.1.2 Regression analysis

Prospective difficulty ratings are grouped by (D, r, θ) condition, and summarized using means.

First, following Equation (4.1), a simple linear regression of the form

$$\overline{PD} \sim \mathcal{N}(a + b \text{ID}, \sigma^2) \quad (4.23)$$

is fit. The results, presented on Table 4.1, show that the fit is quite good in terms of R^2 . However, Figure 4.8 shows that the regression line consistently underestimates downward

targets ($\theta = 45^\circ$, represented by the green dots), and overestimates some of the horizontal targets angles ($\theta = 0^\circ$, represented by the red dots).




Figure 4.8. Regression for $\overline{PD} = a + b ID$. The thick black line represents the posterior median for the intercept and slope, the gray lines represent 100 posterior simulations from intercept and slope.

Table 4.1. Posterior medians (first row) and 95% HDI (second row) of the of the regression estimates for $\overline{PD} = a + b ID$.

a	b	σ	R^2
-0.08	1.39	0.24	0.97
$[-0.44, 0.29]$	$[1.27, 1.51]$	$[0.16, 0.34]$	$[0.94, 0.99]$

These results suggest that a model accounting for target direction might model the PD of pointing at 2-D static-targets more accurately. Thus, following Equation (4.11), a linear regression with per-angle intercept of the form

$$\overline{PD} \sim \mathcal{N}(a_\theta + b ID, \sigma^2) \quad (4.24)$$

is fit.

The regression, displayed in Figure 4.9, shows a great fit, with a very good $R^2 = 0.99$, and no visible non-random patterns around the regression lines. The results, presented in Table 4.2, show estimates of σ and R^2 that are better (lower for σ , and higher for R^2), and

more precise than those of the simple linear model, shown in Table 4.1. This is reflected in Figure 4.9, which shows that all of the points in the dataset fall within the lines formed by the posterior draws of a_θ and b , with no visible non-random patterns.




Figure 4.9. Regression for $\overline{PD} = a_\theta + b ID$. The thick black lines represent the posterior median for each intercept and slope, the gray lines represent 100 posterior simulations from intercept and slope.

Table 4.2. Posterior medians (first row) and 95% HDI (second row) of the of the regression estimates for $\overline{PD} = a_\theta + b ID$.

$a_{\theta=0}$	$a_{\theta=-45}$	$a_{\theta=45}$	b	σ	R^2
-0.27	-0.11	0.13	1.39	0.16	0.99
$[-0.55, 0.02]$	$[-0.4, 0.17]$	$[-0.16, 0.4]$	$[1.3, 1.48]$	$[0.11, 0.24]$	$[0.97, 0.99]$

4.4.1.3 Participant performance and self assessment

The median time participants spent in each question was 5 s, 95% HDI $[1, 20]$, and only in five trials the survey window went out of focus. The median seriousness rating with which participants assessed their performance was 4, 95% HDI $[3, 5]$, with 1 being “Not seriously at all,” and 5 “Very seriously.”

4.4.2 Moving-target block

4.4.2.1 Inferential statistics

To evaluate the effects on PD of D , r , V , and ϕ , and their interactions, the ordered probit regression of Equation (4.21) is fit with the following 25 predictors:

- One indicator for $D = 400$

- Two indicators for $r \in \{40, 80\}$
- Two indicators for $V \in \{-400, -800\}$
- Two indicators for $\phi \in \{-45, 45\}$
- Two indicators for $D \times r$
- Two indicators for $D \times V$
- Two indicators for $D \times \phi$
- Four indicators for $r \times V$
- Four indicators for $r \times \phi$
- Four indicators for $V \times \phi$,

implying that the base values, i.e., $\mathbf{X}_i = \mathbf{0}$, are set to $D = 200$, $r = 20$, $V = -200$, and $\phi = 0$.

Main Effects. The estimates in Figure 4.10 show that there are strong main effects of r , V , and ϕ . As in the static block, PD decreased with larger r , however, as opposed to the static block, PD was virtually unaffected by D , $Mdn = -0.02$, 95% HDI $[-0.35, 0.33]$. Additionally, PD increased with $|V|$, and almost symmetrically with $\phi = \pm 45$, $Mdn(\beta_{\phi=45} - \beta_{\phi=-45}) = 0.07$, 95% HDI $[-0.32, 0.44]$. Given the G - V analogy presented in Section 4.2.3, this is consistent with the subjective assessments of Shoemaker et al. [2012], who found that higher G levels were perceived as more difficult.




Figure 4.10. Main-effect coefficient plot for the fitted ordered probit model on moving-target PD ratings. Points, thick lines, and thin lines represent the medians, 50%, and 95% HDI for each of the coefficients in β .

Two-way Interactions. Figure 4.11 shows that interaction effects have small magnitudes compared to the main effects, indicating that most of the changes in *PD* are explained by the latter.




Figure 4.11. Two-way interaction coefficient plot for the fitted ordered probit model on moving-target *PD* ratings. Points, thick lines, and thin lines represent the medians, 50%, and 95% HDI for each of the coefficients in β .

Even though the main effect of *D* was very weak, the interactions between *D* and the other main effects show some interesting, non-additive relations. For levels of $r < 80$, *D* had a close-to-zero effect on *PD*; however, when $r = 80$, *D* added about 0.55, 95% HDI [0.22, 0.86], to the prospective difficulty of the task, which indicates that the effect of *r* is lower on larger distances (similar to the static block). Likewise, for levels of $|V| < 800$, *D* had a close-to-zero effect on *PD*; however, when $V = -800$, *D* actually made the task easier on average by 0.55,

95% HDI [0.22, 0.86], indicating perhaps that D gives more reaction time to participants on high speed scenarios, and this is perceived as easier. Also, the $D \times \phi$ interaction had a positive effect on PD , which indicates that non-zero angles are perceived as more difficult with higher distances.

Concerning the interactions between r and the other main effects, $r > 20$ mostly decrease the effects of V and ϕ , $Mdn(\beta_{r>20 \times (V \neq -200, \phi \neq 0)}) = -0.24$, 95% HDI $[-0.78, -0.01]$. The interaction effects of $r = 80 \times V = -800$, and $r = 40 \times \phi = -45$, however, are close to 0, $Mdn(\beta_{r=80 \times V=-400}, \beta_{r=40 \times \phi=-400}) = 0.05$, 95% HDI $[-0.35, 0.43]$.

Finally, concerning the $V \times \phi$ interaction, when $|V| > 200$, the angle effects are somehow exacerbated. At $V = -400$, the $\phi = \pm 45$ effects increase by about 0.22, 95% HDI $[-0.26, 0.52]$, and at $V = -800$, the $\phi = \pm 45$ effects increase by about 0.49, 95% HDI $[0.08, 0.86]$.

Symmetry of target movement-direction effects. In general, the main and interaction effects of ϕ seem to be symmetrical, which suggests that the three-part model with per $V \times \phi$ intercept, described in Equation (4.20), is unnecessarily complex, and a simpler, nested model with intercepts varying only by V , such that

$$PD = a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m, \quad (4.25)$$

may be sufficient. Indeed, this model already includes the interactions between D , r , V , and ϕ in its non-constant terms.

4.4.2.2 Regression analysis

Prospective difficulty ratings are grouped by (D, r, V, ϕ) condition, and summarized using means.

One-dimensional moving-target model. A subset of the moving-target tasks, where $\phi = 0$, is fit using Hoffmann's two-part model with per-velocity coefficients, described in Equation (4.18). Therefore, a linear regression of the form

$$\overline{PD} \sim \mathcal{N}(a_V + b_V \sqrt{D} + c_V ID, \sigma^2) \quad (4.26)$$

is fit. The resulting estimates, presented in Table 4.3, show that the fit is very good in terms of R^2 , with credible values above 0.92, and distributions of a , b , and c that increase in magnitude with V . To interpret the results graphically, Formula (4.26) is reparametrized to yield one simple regression line per (D, V) condition, each with intercept $\alpha_{D,V} = a_V + b_V \sqrt{D}$, and slope c_V .

The six regression lines, presented in Figure 4.12, show that the model fits the data very well, without any noticeable non-random patterns around the regression lines. Based on the intercepts of the reparametrized regression lines, the graph suggests that the $PD-ID$ baseline increases with $|V|$, and decreases with D . Based on the slopes, the graph suggests that the rate of PD change per ID change increases slightly with $|V|$.




Figure 4.12. Regression for $\overline{PD} = a_V + b_V \sqrt{D} + c_V ID$ reparametrized to yield six simple regression lines per (D, V) condition, each with intercept $\alpha_{D,V} = a_V + b_V \sqrt{D}$, and slope c_V . The thick black lines represent the posterior median for each intercept and slope, the gray lines represent 100 posterior simulations from intercept and slope.

Table 4.3. Regression estimates for $\overline{PD} = a_V + b_V \sqrt{D} + c_V ID$.

Estimand	<i>Mdn</i>	95% HDI
$a_{V=-200}$	2.49	[1.32, 3.64]
$a_{V=-400}$	3.5	[2.33, 4.63]
$a_{V=-800}$	5.26	[4.1, 6.42]
$b_{V=-200}$	-0.16	[-0.24, -0.08]
$b_{V=-400}$	-0.19	[-0.27, -0.11]
$b_{V=-800}$	-0.26	[-0.34, -0.18]
$c_{V=-200}$	1.12	[0.88, 1.36]
$c_{V=-400}$	1.19	[0.95, 1.43]
$c_{V=-800}$	1.38	[1.14, 1.62]
σ	0.22	[0.13, 0.36]
R^2	0.97	[0.92, 0.99]

Change of coefficients per V. Consistently with the k -G analyses from Shoemaker et al. [2012], linear regressions of the point estimates from Table 4.3— $\hat{a}_V = \text{Mdn}(a_V)$, $\hat{b}_V = \text{Mdn}(b_V)$, and $\hat{c}_V = \text{Mdn}(c_V)$ —on target velocity V are calculated using the default least-squares method in R [2014]. The results, given in Table 4.4, show that there is some evidence of the linear relation between V and \hat{a} , V and \hat{b} , as well as V and \hat{c} .

Table 4.4. Least-squares estimates for the regressions of the point estimates from Table 4.3— $\hat{a}_V = \text{Mdn}(a_V)$, $\hat{b}_V = \text{Mdn}(b_V)$, and $\hat{c}_V = \text{Mdn}(c_V)$ —on target velocity V .

Model	<i>Intercept</i>	<i>Slope</i>	$\hat{\sigma}$	\hat{R}^2
$\hat{a}_V \sim \mathcal{N}(\alpha_1 + \alpha_2 V \times 10^{-3}, \sigma_a^2)$	1.61	-4.58	0.07	1
$\hat{b}_V \sim \mathcal{N}(\beta_1 + \beta_2 V \times 10^{-3}, \sigma_b^2)$	-0.12	0.17	0	1
$\hat{c}_V \sim \mathcal{N}(\zeta_1 + \zeta_2 V \times 10^{-3}, \sigma_c^2)$	1.03	-0.43	0.01	0.99

Similar to the analyses of Shoemaker et al. [2012], however, the least-squares standard deviation estimates, $\hat{\sigma}_a$, $\hat{\sigma}_b$, and $\hat{\sigma}_c$, are too optimistic, as they ignore the uncertainty on a , b , and c . This prevents the calculation of accurate confidence/credible intervals for α_1 , α_2 , β_1 , β_2 , ζ_1 , and ζ_2 .

To give a better sense of uncertainty, a hierarchical model, which simultaneously estimates the parameters in Tables 4.3 and 4.4, is fit. Following the recommendations from Gelman [2006] for standard deviation parameters on hierarchical models, especially those with a low number of groups (in the current study, for example, $\#(V) = 3$), half-Cauchy priors⁸ are chosen for σ_a , σ_b , σ_c , and σ , with scale parameter $A = 5$, which is wide with respect to the range of the data (1–9). Since the center of the α_1 , α_2 , β_1 , β_2 , ζ_1 , and ζ_2 distributions are

⁸The half-Cauchy distribution is equivalent to the right half of a 1-degree-of-freedom t distribution.

not expected to have a considerable change with respect to the least-squares estimates from Table 4.4, the latter are used as mean parameters for their priors; still, a wide prior variance parameter (10^2) on their distribution is used to allow for some flexibility both in the mean and the variance of the marginal posteriors. These weakly-informative priors give reasonable constraints to the posterior distributions, improving the MCMC sampler convergence, while minimally affecting the inferences from the data. In any case, the goal of this hierarchical model is not to get accurate point estimates, especially of a , b , c , α , β , and ζ , for which rough estimates are already known; rather, the goal is to get more realistic uncertainty around α , β , and ζ , and better estimates of σ_a , σ_b , and σ_c .

The model is fit using Stan [2014], with four MCMC chains of 5,000 iterations, including 2,500 warm-up iterations, for a total of 10,000 saved iterations.⁹ The levels and estimates of the model are shown in Table 4.5.

Table 4.5. Levels and estimates of the hierarchical regression $\overline{PD}_i \sim \mathcal{N}(a_{V[i]} + b_{V[i]} \sqrt{D_i} + c_{V[i]} ID_i, \sigma^2)$. $\hat{\alpha}_1$, $\hat{\alpha}_2$, $\hat{\beta}_1$, $\hat{\beta}_2$, $\hat{\zeta}_1$, and $\hat{\zeta}_2$ correspond to the least-squares estimates from Table 4.4.

Model	Estimand	<i>Mdn</i>	95% HDI
$a_V \sim \mathcal{N}(\alpha_1 + \alpha_2 V \times 10^{-3}, \sigma_a^2)$	$a_{V=-200}$	2.49	[1.46, 3.49]
	$a_{V=-400}$	3.47	[2.55, 4.38]
	$a_{V=-800}$	5.25	[4.15, 6.32]
$b_V \sim \mathcal{N}(\beta_1 + \beta_2 V \times 10^{-3}, \sigma_b^2)$	$b_{V=-200}$	-0.16	[-0.23, -0.09]
	$b_{V=-400}$	-0.19	[-0.26, -0.13]
	$b_{V=-800}$	-0.26	[-0.34, -0.19]
$c_V \sim \mathcal{N}(\zeta_1 + \zeta_2 V \times 10^{-3}, \sigma_c^2)$	$c_{V=-200}$	1.12	[0.89, 1.34]
	$c_{V=-400}$	1.2	[0.98, 1.39]
	$c_{V=-800}$	1.38	[1.15, 1.6]
$\sigma \sim \text{half-Cauchy}(5)$	σ	0.21	[0.13, 0.34]
$\alpha_j \sim \mathcal{N}(\hat{\alpha}_j, 10^2)$	α_1	1.62	[-2.42, 6.31]
	α_2	-4.53	[-11.88, 3.67]
$\sigma_a \sim \text{half-Cauchy}(5)$	σ_a	0.86	[0, 4.72]
$\beta_j \sim \mathcal{N}(\hat{\beta}_j, 10^2)$	β_1	-0.13	[-1.11, 0.74]
	β_2	0.16	[-1.48, 2.05]
$\sigma_b \sim \text{half-Cauchy}(5)$	σ_b	0.11	[0, 0.98]
$\zeta_j \sim \mathcal{N}(\hat{\zeta}_j, 10^2)$	ζ_1	1.03	[-2.45, 3.75]
	ζ_2	-0.41	[-6.03, 5.34]
$\sigma_c \sim \text{half-Cauchy}(5)$	σ_c	0.42	[0, 3.69]

The hierarchical-model results show that the standard deviations σ_a , σ_b , and σ_c are much higher than those from Table 4.4. Correspondingly, the intervals around α_1 , α_2 , β_1 , β_2 ,

⁹Due to the time it took the MCMC sampler to run the iterations compared to the linear regression models, the number of iterations was reduced with respect to the latter. Nevertheless, the 10,000 iterations were sufficient for the sampler to converge.

ζ_1 , and ζ_2 are wider than those that would be obtained from the least-squares regression. Nevertheless, the similarity in the distributions of parameters a , b , c , and σ between Tables 4.3 and 4.5 suggests that, in spite of the wide uncertainty, linear relations between V and a , V and b , as well as V and c , are plausible and congruent with the data.

Two-dimensional moving-target model. As suggested in the inferential statistics section, the three-part model with per-speed intercepts described in Equation (4.25) is evaluated for all the moving-target tasks. Therefore, a linear regression of the form

$$\overline{PD} \sim \mathcal{N}(a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m, \sigma^2) \quad (4.27)$$

is fit. The resulting estimates are presented in Table 4.6.

Table 4.6. Regression estimates for $\overline{PD} = a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m$.

Estimand	<i>Mdn</i>	95% HDI
$a_{V=-200}$	2.98	[2.23, 3.7]
$a_{V=-400}$	3.59	[2.86, 4.33]
$a_{V=-800}$	5.96	[5.25, 6.72]
$b_{V=-200}$	0.08	[0.06, 0.1]
$b_{V=-400}$	0.1	[0.09, 0.12]
$b_{V=-800}$	0.1	[0.09, 0.12]
$c_{V=-200}$	-0.22	[-0.27, -0.17]
$c_{V=-400}$	-0.24	[-0.29, -0.18]
$c_{V=-800}$	-0.31	[-0.36, -0.26]
$d_{V=-200}$	1.31	[1.16, 1.47]
$d_{V=-400}$	1.41	[1.26, 1.56]
$d_{V=-800}$	1.41	[1.26, 1.56]
σ	0.26	[0.21, 0.32]
R^2	0.97	[0.96, 0.98]

Table 4.6 shows that all of the coefficients are “statistically significant” in the sense that their 95% HDI do not include 0. However, the low values of b_V indicate that $\sqrt{D_s}$ contributes very little to PD , and its coefficients show very small variations among levels of V , indicating that the contribution of the “static distance” D_s to PD is minimally affected by V . There is no credible difference between $b_{V=-400}$ and $b_{V=-800}$, and the difference between $b_{V=-400}$ and $b_{V=-200}$ is so small, that a target starting at $D = 400$, moving with $(V = -800, \phi = \pm 45)$ would be perceived as harder than a target with the same starting position, but moving with $(V = -200, \phi = \pm 45)$, only by 0.37 units, 95% HDI $[-0.06, 0.8]$.

It is possible that the differences between values of b_V , and d_V will be larger for targets with different configurations, however, with the present experimental-design and observed data this is not the case. This suggests that a model with fixed b , and d , such that

$$PD = a_V + b \sqrt{D_s} + c_V \sqrt{D_m} + d ID_m, \quad (4.28)$$

may be sufficient to adequately fit the data. Therefore, its corresponding linear regression,

$$\overline{PD} \sim \mathcal{N}(a_V + b \sqrt{D_s} + c_V \sqrt{D_m} + d ID_m, \sigma^2), \quad (4.29)$$

is fit. The resulting estimates are presented in Table 4.7. The estimates are very similar to those in Table 4.6; the previous differences between $b_{V=-200}$ and $b_{V \in \{-400, -800\}}$, and between $d_{V=-200}$ and $d_{V \in \{-400, -800\}}$ are now represented in a_V , whose values present wider changes between levels of V than in the previous model. In terms of goodness of fit both models are very similar, there is no credible R^2 difference, and the difference of 0.01 in σ , in both the median and the upper bound of the 95% HDI, is not of practical importance in the 1–9 measurement scale.

To interpret the results graphically, Formula (4.29) is reparametrized to yield one simple regression line per $(D, V, |\phi|)$ condition, each with intercept $\alpha_{D,V,|\phi|} = a_V + b \sqrt{D} |\sin \phi| + c_V \sqrt{D} |\cos \phi|$, and slope d . The 12 regression lines presented in Figure 4.13 show that the model fits the data very well, with no noticeable non-random patterns around the regression lines.

Table 4.7. Regression estimates for $\overline{PD} = a_V + b \sqrt{D_s} + c_V \sqrt{D_m} + d ID_m$.

Estimand	Mdn	95% HDI
$a_{V=-200}$	2.74	[2.04, 3.43]
$a_{V=-400}$	3.72	[3.03, 4.42]
$a_{V=-800}$	6.08	[5.39, 6.76]
b	0.1	[0.09, 0.11]
$c_{V=-200}$	-0.22	[-0.27, -0.17]
$c_{V=-400}$	-0.23	[-0.28, -0.19]
$c_{V=-800}$	-0.3	[-0.35, -0.26]
d	1.38	[1.29, 1.47]
σ	0.27	[0.21, 0.33]
R^2	0.97	[0.96, 0.98]

Based on the intercepts of the reparametrized regression lines, the graph suggests that the PD – ID baseline increases with ϕ and $|V|$, and decreases with D . The slopes suggest that the rate of PD change per ID stays roughly similar for all conditions.

Change of coefficients per V . In a manner analogous to the analyses presented in the 1-D moving-target section, linear regressions on the per- V parameter estimates from Table 4.7— $\hat{a}_V = \text{Mdn}(a_V)$, and $\hat{c}_V = \text{Mdn}(c_V)$ —are first calculated using the default least-squares method in R [2014]. The results, shown in Table 4.8, give some evidence in support of the linear relation between V and \hat{a} , as well as V and \hat{c} .




Figure 4.13. Regression for $\overline{PD} = a_V + b\sqrt{D_s} + c_V\sqrt{D_m} + d ID_m$ reparametrized to yield twelve simple regression lines per $(D, V, |\phi|)$ condition, each with intercept $\alpha_{D,V,|\phi|} = a_V + b\sqrt{D}|\sin \phi| + c_V\sqrt{D}|\cos \phi|$, and slope d . The thick black lines represent the posterior median for each intercept and slope, the gray lines represent 100 posterior simulations from intercept and slope.

Table 4.8. Least-squares estimates for the regressions of the point estimates from Table 4.7— $\hat{a}_V = \text{Mdn}(a_V)$, and $\hat{c}_V = \text{Mdn}(c_V)$ —on target speed V .

Model	Intercept	Slope	$\hat{\sigma}$	\hat{R}^2
$\hat{a}_V \sim \mathcal{N}(\alpha_1 + \alpha_2 V \times 10^{-3}, \sigma_a^2)$	1.56	-5.61	0.1	1
$\hat{c}_V \sim \mathcal{N}(\zeta_1 + \zeta_2 V \times 10^{-3}, \sigma_c^2)$	-0.19	0.14	0.01	0.92

Once again, the least-squares estimates $\hat{\sigma}_a$, and $\hat{\sigma}_c$ are too optimistic, as they ignore the uncertainty on a , and c . Thus, to have a better sense of uncertainty around the estimates, and assess the plausibility of a linear relation between V and a , as well as V and c , a hierarchical model which simultaneously estimates the parameters in Tables 4.7 and 4.8, is fit. The hierarchical model reuses the least-squares estimates $\hat{\alpha}_1$, $\hat{\alpha}_2$, $\hat{\zeta}_1$, and $\hat{\zeta}_2$ from Table 4.8 as priors for the mean parameters of α_1 , α_2 , ζ_1 , and ζ_2 , respectively, as well as the distribution medians $\text{Mdn}(b)$, and $\text{Mdn}(d)$ from Table 4.7, as priors for the mean parameters of b , and d , respectively; the wide prior variance (10^2) allows for some flexibility both in the mean and the variance of the marginal posteriors. Similar to the 1-D moving-target hierarchical model, half-Cauchy priors are used for the variance parameters σ_a , σ_c , and σ .

The model is fit using Stan [2014], with four MCMC chains of 5,000 iterations, including 2,500 warm-up iterations, for a total of 10,000 saved iterations. The levels and estimates of the model shown in Table 4.9, again, show higher estimates for σ_a , and σ_b than those obtained in Table 4.8, and correspondingly, wider intervals around $\hat{\alpha}_1$, $\hat{\alpha}_2$, $\hat{\zeta}_1$, and $\hat{\zeta}_2$ than those that would be obtained from the least-squares regression. Nevertheless, the minimal change in the distributions of a , b , c , d , and σ between Tables 4.7 and 4.8 suggests that, in spite of the wide uncertainty, the linear relations between V and a , and V and c are plausible and congruent with the data.

4.4.2.3 Participant performance and self assessment

The median time participants spent in each question was 3 s, 95% HDI [1, 11], and only in 11 trials the survey window went out of focus. Participants replayed the animation in 15.31% of the trials, but within those trials the median number of replays was 1, 95% HDI [1, 3]. The median seriousness rating with which participants assessed their performance was 4, 95% HDI [3, 5], with 1 being “Not seriously at all,” and 5 “Very seriously.”

4.5 Discussion

Results are summarized according to the study hypotheses, each followed by a discussion. After the hypotheses discussions, the limitations of the current study and possible future work are presented.

Table 4.9. Levels and estimates of the hierarchical regression $\overline{PD}_i \sim \mathcal{N}(a_{V[i]} + b \sqrt{D_{si}} + c_{V[i]} \sqrt{D_{mi}} + d ID_{mi}, \sigma^2)$. $\hat{\alpha}_1$, $\hat{\alpha}_2$, $\hat{\zeta}_1$, and $\hat{\zeta}_2$ correspond to the least-squares estimates from Table 4.8. \hat{b} , and \hat{d} correspond to the distribution medians for b , and d from Table 4.7.

Model	Estimand	<i>Mdn</i>	95% HDI
$a_V \sim \mathcal{N}(\alpha_1 + \alpha_2 V \times 10^{-3}, \sigma_a^2)$	$a_{V=-200}$	2.7	[2.06, 3.33]
	$a_{V=-400}$	3.78	[3.19, 4.36]
	$a_{V=-800}$	6.06	[5.4, 6.74]
$b \sim \mathcal{N}(\hat{b}, 10^2)$	b	0.1	[0.09, 0.11]
$c_V \sim \mathcal{N}(\zeta_1 + \zeta_2 V \times 10^{-3}, \sigma_c^2)$	$c_{V=-200}$	-0.22	[-0.26, -0.18]
	$c_{V=-400}$	-0.24	[-0.28, -0.2]
	$c_{V=-800}$	-0.3	[-0.35, -0.26]
$d \sim \mathcal{N}(\hat{d}, 10^2)$	d	1.38	[1.28, 1.46]
$\sigma \sim \text{half-Cauchy}(5)$	σ	0.26	[0.22, 0.33]
$\alpha_j \sim \mathcal{N}(\hat{\alpha}_j, 10^2)$	α_1	1.58	[-2.41, 5.3]
	α_2	-5.61	[-12.19, 1.22]
$\sigma_a \sim \text{half-Cauchy}(5)$	σ_a	0.64	[0, 4.18]
$\zeta_j \sim \mathcal{N}(\hat{\zeta}_j, 10^2)$	ζ_1	-0.18	[-0.7, 0.43]
	ζ_2	0.15	[-0.94, 1.28]
$\sigma_c \sim \text{half-Cauchy}(5)$	σ_c	0.08	[0, 0.6]

4.5.1 Prospective difficulty of static-target selection

H1. Fitts' *ID* model with per-angle intercept, described in Equation (4.11), models the prospective difficulty of two-dimensional static-target selection tasks better than the classic Fitts' *ID* model with single intercept, described in Equation (4.1). *Partially supported.*

Both the fixed intercept model of Equation (4.1),

$$PD = a + b ID,$$

and the per-angle intercept model of Equation (4.11),

$$PD = a_\theta + b ID,$$

yielded very good fits with credible values of $R^2 = [0.94, 0.99]$, and $R^2 = [0.97, 0.99]$, respectively. However, compared to the results of the fixed intercept model, presented in Table 4.1, the results of the per-angle intercept model presented in Table 4.2 showed slightly better, and more precise estimates of σ and R^2 . The difference in precision, $Mdn(\sigma_1 - \sigma_2) = 0.07$, 95% HDI [-0.04, 0.2], however represents less than 1% of the total 1–9 scale range.

Even though the 2-D model is “significantly better” in statistical terms,¹⁰ the difference in precision between both models is not of practical significance. The importance of the model, however, lies in its flexibility and extensibility to other angles, therefore, *H1* is partially supported.

In any case, the fact that *ID* was linearly related to *PD*, supports the usage of the chosen 9-point rating scale, with labels at the poles, and a middle anchor, to measure prospective difficulty. This point is crucial to validate the *PD* ratings for moving-target selection using this scale.

4.5.2 Prospective difficulty of moving-target selection

H2. Hoffmann’s two-part model with per-velocity coefficients, described in Equation (4.18), accurately models the prospective difficulty of one-dimensional moving-target selection tasks. *Supported*.

The 1-D moving-target selection model of Equation (4.18),

$$PD = a_V + b_V \sqrt{D} + c_V ID,$$

yielded a very good fit, with credible values of $R^2 = [0.92, 0.99]$, presented in Table 4.3, and regression lines that closely follow the data with no visible non-random patterns in the residuals, as shown in Figure 4.12. This was expected given the structural similarity between Equation (4.18) and Jagacinski’s “general” model in Equation (2.15), as well as the $G-V_{target}$ analogy, presented in Section 4.2.3.

H2.1. The coefficients a_V , b_V , and c_V of Equation (4.18) vary linearly with V . *Supported*.

The results suggest that coefficients a_V , b_V , and c_V of Equation (4.18) vary linearly with V . The analogy with gain, for which these approximated coefficients also seem to increase linearly, also supports this hypothesis. Using the credible intervals presented in Table 4.5 allows *PD* to be predicted using different V , while still propagating the uncertainty in these estimates.

H3. The three-part model with per-speed coefficients, and per-velocity intercept, described in Equation (4.20), accurately models the prospective difficulty of two-dimensional moving-target selection tasks. *Supported*.

It was determined that the three-part model with per-speed coefficients, and per-velocity intercept of Equation (4.20),

$$PD = a_{V \times \phi} + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m,$$

¹⁰Which can also be verified using a classical F -test, not shown.

was unnecessarily complex for the evaluated dataset; however, the simpler model with per-speed intercept and coefficient c described in Equation (4.28),

$$PD = a_V + b \sqrt{D_s} + c_V \sqrt{D_m} + d ID_m,$$

yielded a very good fit, with credible values of $R^2 = [0.96, 0.98]$, presented in Table 4.7, and regression lines that closely follow the data, with no visible non-random patterns in the residuals, as shown in Figure 4.13.

Since Equation (4.28) is nested within Equation (4.20), this implies that the latter also fits the data well, therefore $H3$ is supported. Even though the simpler model was sufficient for the purposes of this chapter, Equation (4.20) should not be discarded for future analyses, especially those having a wider variety of angles, in which angle asymmetry might have a larger effect.

H3.1. In Equation (4.20) the coefficients b_V , c_V , and d_V vary linearly with V , as well as the coefficients $a_{V \times \phi}$ within the same angle ϕ . *Partially Supported.*

The results suggest that coefficients a_V , and c_V of Equation (4.28) increase linearly with V . Contrary to what was hypothesized, b , and d were almost constant across V . For b , this suggests that the prospective ballistic motion parallel to the target movement is minimally affected by V . For d , this suggests that the effect of a same target width does not change by V , as shown in the interaction plot in Figure 4.11. Using the credible intervals presented in Table 4.9, this allows PD to be predicted using different V , while still propagating the uncertainty in these estimates.

4.6 Conclusion

Contrary to most Fitts-related studies in HCI, which have focused mostly on modeling performance in static-target pointing tasks, the main goal of this study was to explore the PD of moving-target pointing tasks in 2-D. This goal was achieved by developing and empirically evaluating models for 2-D static-target pointing, and 1-D and 2-D moving-target pointing. To the author's knowledge, this is the first study to evaluate pointing tasks in which the target is not moving directly toward, or directly away from the cursor, making the models closer to generalization than some of the work in the literature.

The proposed models have, in the author's opinion, relatively simple formulations. These models are inspired on previous work on the MT of 2-D static-target tasks [Murata and Iwase 2001; Appert et al. 2008; Grossman and Balakrishnan 2004], as well as two-part models [Welford et al. 1969; Hoffmann and Chan 2012], which have proven effective to model static-target selection with different levels of G [Shoemaker et al. 2012]. Incidentally, the empirical data analysis was made using Bayesian methods, less prevalent in HCI yet very powerful—in this

chapter, in particular, hierarchical Bayesian regressions allowed to propagate the uncertainty on the main model parameters to secondary regressions made on those parameters themselves.

As will be discussed in Chapter 6, there are many opportunities to extend the current work. Hopefully, this will incentivize and facilitate further HCI research on moving-target selection and subjective difficulty, which remain overshadowed by static-target selection and performance. Additionally, it is hoped that the analytical methods presented in this chapter will be considered as a viable alternative to the prevailing frequentist methods in HCI, and motivate more analyses to go beyond “statistical significance.”

CHAPTER 5

PERFORMANCE AND INTENTION IN DIRECTED 3-D MOVING-TARGET SELECTION

Chapters 1–3 hypothesized that ID was predictive of user intention in moving target selection, based on the premise that users form their intentions by minimizing their prospective effort, as described by ID . The results from Chapter 3, however, revealed that separating the W and D components was more effective for predicting intention, in a similar way as the models proposed in Chapter 4 separate D from ID to predict PD .

This similarity suggests that the proposed models may also be predictive of user intention, but this relation needs to be further explored. In particular, in order to predict user intention in VR, these models need to be extended to 3-D, and evaluated during action execution. To achieve these goals, this chapter explores the applicability of the PD models from Chapter 4 to predict selection performance. The performance predictions are, in turn, used to predict user intention by integrating them into the framework described in Chapter 3.

In order to generalize the results from Chapters 3 and 4, this chapter includes a wider span of experimental factors, including more levels of velocity, target angle, and number of spheres. Additionally, this chapter includes tasks in which users do not have to wait for the targets to be within their reach, and tasks in which users must choose a specific target among multiple others.

5.1 Extending the 2-D formulae to 3-D

One advantage of the 2-D model proposed in Chapter 4 is that it can theoretically be used to model 3-D target-selection tasks. Consider the general 3-D moving-target pointing task presented in Figure 5.1. Even though the target is moving in 3-D, the target velocity (\mathbf{V}) and initial target position (\mathbf{D}) define a right triangle with sides D , D_m , and D_s , which, together with W , are the inputs in the model in Equation (4.20).

In addition to the target altitude θ , present in 2-D tasks, 3-D tasks require an additional angle to define the target position, in this case a z - x azimuth angle α is chosen. Due to depth-perception problems in stereoscopic displays [Grossman and Balakrishnan 2004; Bowman et al.




Figure 5.1. Three-dimensional moving-target model. The red and green spheres indicate the starting positions of the cursor and the target, respectively. The initial target point (**D**) is determined by the initial cursor–target distance (D), z – x azimuth angle (α), and altitude angle (θ).

2004], targets with identical (V, D) , but displayed at different screen depths, given by (θ, α) , may be perceived with different PD . This effect has previously been observed for the MT required to acquire 3-D static targets [Teather and Stuerzlinger 2011].

Analogous to 2-D tasks, the initial angular coordinates (α, θ) could be represented with different $a_{V \times \alpha \times \theta}$ intercepts. Yet again, it is hypothesized that, as long as the target is defined within the user’s field of view, the initial α and θ will have a minimal effect, since these angles change as the target moves.

Likewise, the target orientation can no longer be solely defined with ϕ . In fact, as shown in Figure 5.2, any rotation γ of **V** around **D** defines right triangles with equal dimensions D_s , and D_m . Since their parameters are identical, all of these triangles are considered equal in the 2-D models of the previous chapter; if the effect of γ is large, however, the 2-D moving-target model needs to be extended in order to be used effectively in 3-D.




Figure 5.2. Identical model parameters V , D_m , and D_s for three different γ rotations of **V** around **D**.

Finally, to generalize the model to 3-D, the calculations of D_m , and D_s can be simplified by using vector algebra. At any given instant, the displacement vector between the target's position P , and the user's wand position P_w is given by

$$\mathbf{D}' = P_w - P. \quad (5.1)$$

Given this, D_m , and D_s correspond to the lengths of the projection and rejection of \mathbf{D}' on the target's unit velocity $\hat{\mathbf{V}}$, respectively, such that

$$D_m = |\mathbf{D}' \cdot \hat{\mathbf{V}}| \quad (5.2)$$

$$D_s = \|\mathbf{D}' - (\mathbf{D}' \cdot \hat{\mathbf{V}}) \hat{\mathbf{V}}\|. \quad (5.3)$$

Alternatively, the value of D_m can be reused to calculate D_s , which can be less expensive computationally if the calculations are executed serially, such that

$$D_s = \sqrt{\mathbf{D}' \cdot \mathbf{D}' - D_m^2}. \quad (5.4)$$

5.2 Subjective difficulty and performance

In Chapter 1 it was argued that prospective and performance assessments are related, but the former can complement or even supersede the latter because they are not dependent on action execution. Nonetheless, measures of performance, such as MT or percentage of errors are still important to evaluate users and interfaces objectively. Additionally, measuring the PD of a given task requires either direct participant assessment, which interrupts the task, or the usage of sensors such as EEG [Kourtis et al. 2012], which may be invasive and are not available in all VR setups.

Based on the correlation between subjective difficulty (prospective, and perceived) and MT explored in previous work, it is hypothesized that the proposed PD formulae should also predict MT , such that

$$MT = a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V \log_2 \left(\frac{2D_m}{W} \right). \quad (5.5)$$

The relation between PD and MT , however, may be different for moving targets than for static ones. In particular, in Chapter 4 it was observed that PD increases with V even for targets that are moving toward the cursor, but in this case the MT should *decrease* with V as long as the target has not been missed. This may imply that the signs in some of the regression coefficients in the proposed formulae for PD may be contrary in MT formulae. Therefore, this chapter investigates the relation between the PD models of 2-D moving-target pointing introduced in Chapter 4, and the performance of 3-D moving-target selection tasks.

5.3 Predicting user intention in directed tasks

Chapter 3 evaluated the possibility of predicting intention in undirected tasks using *ID*, based on the premise that users form their intentions by minimizing their prospective effort, as described by *ID*. In directed tasks, such as the ones studied in this chapter, the user must select the given target regardless of the prospective effort required to reach it. The principle of minimizing prospective effort, however, may still apply when choosing the path required to reach the specified target, similar to the way people choose routes in everyday life [Christenfeld 1995; Bailenson et al. 2000].

According to the initial segment strategy (iss) [Christenfeld 1995; Bailenson et al. 2000], people tend to choose routes whose initial segments are straight, and where turns occur last—both characteristics are found in the model for 2-D moving-target selection proposed in Section 4.2.4. Therefore, the hypothesis that the target with the minimal \widehat{MT} as predicted by Equation 5.5 is the intended target is evaluated in this chapter.

Since prediction is based only on one feature, a scoring function based on the one proposed by Ortega [2013], presented in Equation (2.23), is used. At each frame t , \widehat{MT} is calculated for each of the I targets using Equation (5.5). Targets are then ordered ascendingly by \widehat{MT} , their order given by $i = 0, \dots, I - 1$. The score for each target i is calculated following

$$\begin{aligned} \widehat{T}Score_i(t) &= \widehat{T}Score_i(t-1) + \begin{cases} (N-i)\Delta t & \text{if } i < N \\ -(decay \cdot N)\Delta t & \text{if } i \geq N \end{cases} \\ \widehat{T}Score_i(t) &\geq 0 \end{aligned} \quad (5.6)$$

where N is an arbitrary number of faster-to-reach targets, and *decay* is the rate with which scores decrease when a target $i \geq N$. In Ortega's studies with 100 spheres, a distance-based scoring function, $dScore_i(t)$, was used with $N = 20$, and $decay = 0.5$, in Chapter 3, $dScore_i(t)$ was used with $N = 1$, and $decay = 0.9$. To validate the usefulness of $\widehat{T}Score_i(t)$, its predictive accuracy is compared to $dScore_i(t)$

5.4 Hypotheses

The following hypotheses guided the design and analysis of the current study,

5.4.1 Moving-target selection performance in 3-D

- H1. The three-part model with per-speed coefficients, and per-velocity intercept, described in Equation (5.5), accurately models the movement time of three-dimensional moving-target selection tasks.

H1.1. The azimuth α , altitude θ , and γ angles, not modeled in Equation (5.5), have a small effect on MT .

H1.2. In Equation (5.5) the coefficients a_V , b_V , c_V , and d_V vary linearly with V .

5.4.2 Predicting user intention in 3-D moving-target directed-selection tasks

H2. At any given instant, the minimal movement time estimated by Equation (5.5) can be used to predict the intended target in a 3-D moving-target directed-selection task.

H2.1. The accuracy of the predictions based on estimated movement time will be greater than those based on wand-target distance.

These hypotheses assume univariate targets in a 3-D space, i.e., spheres.

5.5 Methods

5.5.1 Participants

Participants from Iowa State University were recruited through the Psychology Department research participant pool (SONA), through word of mouth, and through the HCI mailing list.

There were 33 participants, aged 18 to 39 years old ($M = 23.48$, $Mdn = 21$); ten participants were females, and only one participant was left handed. Most participants (31) were students, and nine participants reported having past experience in VR.

5.5.2 Apparatus

The user study took place in the Virtual Reality Applications Center, at Iowa State University. The code-base from Study 1 was reused and modified to implement this experiment. The VR JuggLua application was deployed in the C6, a six-surface CAVE-like virtual environment. All of the faces of the 3.05 m^3 VE, except the back wall, were projected using active stereo at 4096×4096 px. The participant's head and wand pose (P, Q) were tracked using a 6-DOF Intersense IS-900 inertial-ultrasonic hybrid tracker, allowing them to adequately perceive the 3-D world, and interact with it.

5.5.3 Procedure

Upon arriving to the study site, the participant was asked to read and accept a study consent form. If consent was granted, the participant was asked to answer a Qualtrics survey regarding their video game, VR, sports, and demographic background. The participant was also asked to complete the "Perspective Taking/Spatial Orientation Test" [Hegarty and Waller 2004], and the Simulator Sickness Questionnaire (SSQ) [Kennedy et al. 1993]. Subsequently, the participant's

visual acuity was evaluated using the Interactive Visual Acuity Chart (IVAC) [Olitsky et al. 2006].

After completing the pre-survey questionnaire, the participant was asked to enter the C6 and step on a circular landmark located in the middle of the VE $(0,0,0)$, facing the front wall. After inquiring if the participant's perception of the virtual world was adequate, the experimenter asked the participant to choose a comfortable starting position for their wand ($P_{w,0}$), specifying that they should return to this position at the beginning of each trial.

In each trial, the participant was presented with an array of equally-textured virtual spheres of different sizes, initially located in front of them, and moving with different velocities. The spheres would start moving when the participant's wand was placed within a green translucent sphere, sph_0 , of radius 0.05 m centered at $P_{w,0}$.

The participant was asked to keep the wand within sph_0 until the goal sphere became highlighted. If they did not follow this instruction, sph_0 would turn red. Once the goal sphere became highlighted, sph_0 would disappear and the participant was asked to extend their arm to touch the goal sphere. If any sphere was touched during the selection process, whether highlighted or not, it disappeared. The trial, however, would end only when the highlighted sphere was touched, or if its distance to the user's head was greater than 2 m, in which case it was assumed to be unreachable.

Visual and auditory feedback were used to indicate participant performance. A virtual counter was placed in front of the participant at $(-1, -1, -10)$, which showed the number of highlighted spheres missed by the participant during each experimental block. The counter would be reset to zero at the beginning of each block of trials. Two different sounds were played: one when the participant hit any sphere, and another when the trial ended due to the highlighted sphere being too far away, respectively.

At each frame, the elapsed time, the pose of the participant's head (P_h, Q_h) and wand (P_w, Q_w), each of the sphere positions (P_i), and the possible wand-sphere collisions were recorded in a log file. The experimental setup is depicted in Figure 5.3.

After all the trials were complete, the participant was asked to exit the C6 and fill a post-experiment survey, which included the ssQ. The total length of the study was less than 60 minutes.

5.5.4 Design

A within-subjects, factorial design was used, with three blocks of trials, each with a different number of conditions presented in a random order without replacement.

In the first block there was one sphere per trial.¹ The target was placed in one of the front vertices of an imaginary tetrahedron whose sides were 1 m, with one vertex placed at $P_{w,0}$, and the three others in front of the participant. Factors were sphere position, $P \in$

¹Notice that participants were still instructed to stay at $P_{w,0}$ before this sole sphere became highlighted.

$\{up, bottom-left, bottom-right\}$, radius, $r \in \{0.1, 0.2\}$ m, speed, $V \in \{0.5, 1, 1.5\}$ m/s, and initial movement direction, $\phi \in \{15, 30, 45\}^\circ$. Since there were already 54 conditions, the value of the initial γ angle was set randomly between 180° , and 360° ; there was one trial per condition, for a total of 1,782 trials.

In the second block there were three spheres per trial. The targets were placed in one of the front vertices of an imaginary tetrahedron whose sides were 1 m, with one vertex placed at $P_{w,0}$, and the three others in front of the participant. Factors were sphere position, $P \in \{up, bottom-left, bottom-right\}$, radius, $r \in \{0.1, 0.2\}$ m, speed, $V \in \{0.5, 1, 1.5\}$ m/s, and initial movement direction, $\phi \in \{15, 30, 45\}^\circ$. The value of the initial γ angle was set randomly between 180° , and 360° , and the (P, r, V, γ) conditions of the two remaining spheres were also set randomly within the possible experimental values. There was one trial for each of the 54 conditions in this block, for a total of 1,782 trials.

In the last block there were six spheres per trial. The targets were placed in one of the front vertices of two imaginary tetrahedrons, each with sides of 1 m, with one vertex placed at $P_{w,0}$, and the three others in front of the participant; the front vertices of the second tetrahedron correspond to the vertically-mirrored positions of the first tetrahedron's front vertices (this arrangement is shown in Figure 5.3, middle). Factors were sphere position, $P \in \{up, up-left, bottom-left, bottom, bottom-right, up-right\}$, radius, $r \in \{0.1, 0.2\}$ m, speed, $V \in \{0.5, 1, 1.5\}$ m/s, and initial movement direction, $\phi \in \{15, 30, 45\}^\circ$. The value of the initial γ angle was set randomly between 180° , and 360° , and the (P, r, V, γ) conditions of the five remaining spheres were also set randomly within the possible experimental values. There was one trial for each of the 108 conditions in this block, for a total of 3,564 trials.

Regardless of the block, in all trials the initial sphere positions were set at $P_i - \mathbf{V}_i * 1$ s. The goal sphere became highlighted after getting past its original tetrahedral position P , i.e. after approximately 1 s, when approximately at $D = 1$ m.² The criterion for highlighting the sphere was $D < 1$ m. Notice that, at the moment of highlight all spheres were approximately at their tetrahedral position P_i , at about the same distance from $P_{w,0}$. This was done to give participants some time to perceive the motion of the different spheres, and to make the sphere arrangements seemingly aleatory.

5.6 General results

5.6.1 Experimental issues

There were several technical issues, as well as some participant wellness problems during the execution of the experiment.

²These values are approximate due to the variable frame rate.



Figure 5.3. Experimental setup with six spheres. Left, the sphere starting positions, middle the spheres approximately at their controlled positions, right the spheres after the goal sphere gets highlighted. The green sphere represents sph_0 .

5.6.1.1 Technical issues

Incomplete logging. Because of logging problems, the data from the first two participants was incomplete, so they were excluded from the performance and intention analyses.

Tracking interruptions. In some trials, the wand tracker lost track of the participant's wand causing a delay between the real and the virtual wand or, in the worst case scenario, a complete cessation of the tracking. In the latter cases, the experimenter asked the participants to stop the motion of the wand to allow the tracking to resume. This was most noticeable on participants with brusque movements. Fortunately, by virtue of the experimental design, in which trials would not start before the participant placed the wand at sph_0 , this issue did not affect consecutive trials. Finally, in several trials the head tracker stopped functioning, an issue that was most often detected and fixed at the end of each block; nonetheless, the participants did not seem to notice this issue. No trials were removed from the data due to tracking interruptions.

Projector problems. During five experimental sessions, one or two of the C6 projectors could not be started. In three instances the affected projector was in the right screen, in one instance the affected projector belonged to the front screen, and in the last instance both the right and front faces were affected. These sessions were, nonetheless, included in the analyses.

5.6.1.2 Participant-wellness issues

Two participants had noticeable wellness problems. The first of these participants had a cast on their right arm upon arrival, but decided to participate using with their left hand. The second participant decided to take a break during the three-sphere block, and declared having some discomfort on their right shoulder; even though the experimenter suggested to terminate the session without any consequences for the participant, they decided to continue

using their left hand. This last participant commented having some nausea by the end of the experiment. These sessions were retained in the analyses.

5.6.1.3 Simulator sickness questionnaire

The ssq ratings were scored using the procedure described by Kennedy et al. [1993, p. 212]. The results, shown in Figure 5.4, show that the score distributions assessed before and after the experiment are very similar. Notice that each of the categories, nausea, oculomotor, disorientation, as well as the total score, have different scales.




Figure 5.4. Distributions of the pre- and post-experiment assessment scores of the ssq by scale. Lower and upper lines of the boxes represent the first and third quantiles, their distance called the inter-quantile range (IQR), thick box lines represent the median values, upper and lower whiskers represent values that are within 1.5 IQR of the box hinges. Points represent values that are greater than the third quartile plus 1.5 IQR.

Even though the main goal of this study is not to explore simulator sickness, these results give some evidence against a possible effect of simulator sickness in trial performance and participant action. This evidence, however, contradicts the conjecture that simulator sickness is worsened in moving-target selection presented in Chapter 1.

5.7 Performance analysis

5.7.1 Exploratory data analysis

In accordance with some of the ideas of Tukey [1980], a brief exploratory analysis is carried out before attempting to apply statistical-modeling techniques to summarize the data. This analysis serves the purpose of presenting the general trends in the data, informally assessing the aforementioned hypotheses, and possibly unveiling (or letting the reader discover) unexpected findings. Following Tukey’s [1980, p. 24] statement that “the picture-examining eye is the best finder we have of the wholly unanticipated,” this analysis is mostly carried out using data visualization techniques; numerical methods to quantify the magnitude and uncertainty of the different effects are instead presented in the subsequent section.

5.7.1.1 Between-block trial performance

As shown on the left Figure 5.5, the percentage of successful captures ρ per participant decreased with the increase of Spheres per block. An effect of fatigue is discarded, given the similarity between the pre- and post-experiment scores of the ssQ presented above.

It is likely, however, that this difference is due to a delayed reaction time RT as the number of spheres increased, which is congruent with the Hick-Hyman law [Hick 1952; Hyman 1953]. This is consistent with the increase in T that follows the increment of spheres per block displayed on the right of Figure 5.5.

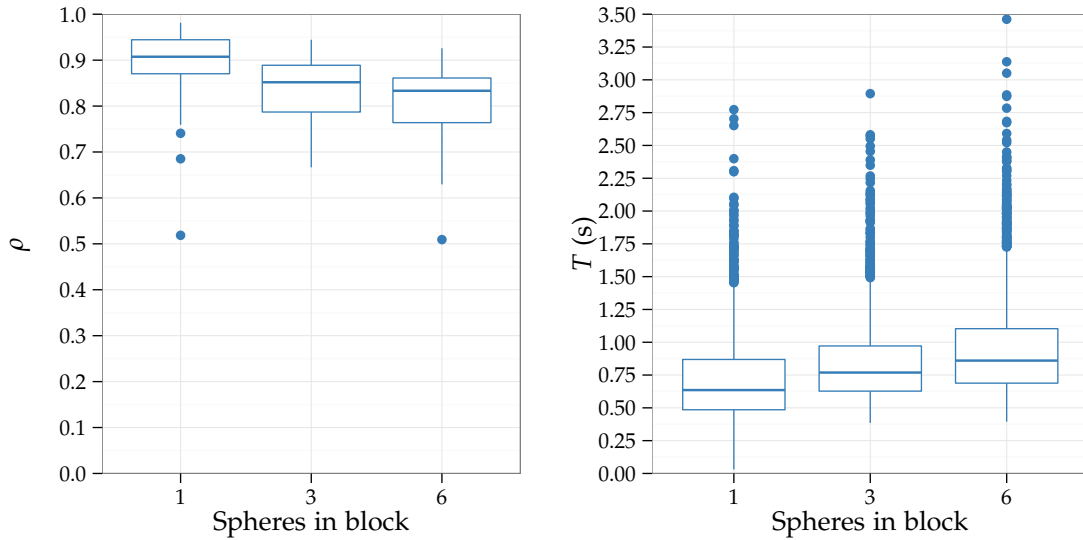


Figure 5.5. Distribution of participant success rates (left), and successful trial completion times (right) per experimental block.

T is used as a proxy for RT , which cannot be measured directly from the data. This

substitution is not perfect since participants may have selected other spheres prior to attaining the target sphere in trials where the number spheres was greater than one. In the 1-sphere trials, additionally, some participants may have “cheated” by approaching the target sphere before it was highlighted, which corresponds to the values of T close to zero at the tip of the lower whisker of the 1-sphere T boxplot.

5.7.1.2 Effects on MT for the successful trials in the 1-sphere block

It is clear that the number of spheres per trial affected participant performance. However, the main goal of the current section, in terms of performance prediction, is to identify the factors that affect the MT required to reach a single moving target. Therefore, the scope of this analysis covers only the successful trials of the 1-sphere block; the extension to multiple targets is left for future work.

In this section, a distinction is made between *controlled factors*, i.e., P , r , V , and ϕ that were systematically assigned to each trial, and *experimental covariates*, i.e., factors that were either assigned randomly, or not controlled for during each trial.

Controlled factors. Among the main factors, V had the strongest effect on the MT of the successful trials, as shown in Figure 5.6. As expected, MT decreased with the increase of V . In low velocity trials, additionally, the span of MT was also larger.

Target radius r also had a visible effect on MT . Over all, bigger targets had shorter MT , although this decrease is visibly smaller for targets with high V , suggesting an interaction between both. Additionally, in low V trials, the span of MT was smaller for targets with $r = 0.2$, than for targets with $r = 0.1$.

Targets whose position P was up with respect to sph_0 had larger MT than targets located below sph_0 . This increase in MT was smaller for targets with high V , suggesting an interaction between P , and V .

Increases in ϕ resulted only in small increases of MT . In general, this increase was more visible for small targets located above sph_0 .

Experimental covariates. Due to the experimental design, there are certain factors that were not controlled for, which could have an effect on MT , including γ , and the initial wand position at the highlight frame in each trial.

The effect of γ on MT . To avoid generating too many experimental conditions, the angle γ was randomly assigned between 180° and 360° . Yet, in order to generalize the 2-D model to 3-D without further modification, its effect on MT must remain small.

As seen at the top of Figure 5.7, the distribution of angles is close to uniform in the $[180, 360]^\circ$ range. As highlighted by the superimposed smoothing curve, on average, MT remained very similar across the range of assigned γ angles. This suggests a very small effect of γ on MT , which supports the usage of the 2-D model for 3-D tasks.




Figure 5.6. Distributions of movement times for the successful trials per experimental factor in the 1-sphere block.




Figure 5.7. Movement times per γ angle for the successful trials in the 1-sphere block. The top bars show the distribution of γ binned every 10° .

Wand positions at the highlight frame. To avoid hindering the flow of the study, only a soft constraint was placed on the wand position before the highlight frame, i.e., sph_0 which would change from green to red whenever it did not encircle the tip of the wand. Even though this allowed participants to approach the target sphere before it became highlighted, the three position density plots of Figure 5.8 show that the highest density of wand positions was within sph_0 , $p(\|P_w - P_{w,0}\| < 0.05) = 0.5$. In a few trials, however, participants seem to have “cheated” by placing their wand beyond sph_0 , thus approaching the target, $p(P_{w,z} - z_0 < -0.05) = 0.17$. Similarly, the negative skew of the wand positions along the y axis reflects the fact that in some trials participants got closer to the bottom targets by lowering their wands below sph_0 ; this phenomenon was less visible for trials where the target was initially above the user. Finally, the distribution of wand positions along the x axis, shown on the left part of Figure 5.8, was symmetrical with respect to x_0 . The fact that participants were more compliant with their left-right and up positions relative to $P_{w,0}$ may reflect a problem of depth perception and occlusion of the wand tip by sph_0 .

The effect of the different P_w at the highlight frame on each trial’s MT can be seen in Figure 5.9. Overall, an increase in the wand-sphere distance $\|P_w - P\|$ resulted in an increase of about 0.88 s in MT , as highlighted by the trend line. This implies that even within the bounds of sph_0 , depicted by the vertical orange lines, the average change in MT is about 0.09 s. The trend line was purposely chosen as a least-squares linear model to represent the approximate




Figure 5.8. Front and right views of the distribution of wand positions at the frame where the sphere was highlighted in the 1-sphere trials. Red, blue, and green points represent the wand positions in trials where the sphere was at the bottom-right, bottom-left, and up, respectively. Correspondingly colored stars represent the sphere positions in those frames. The orange circle represents sph_0 . The top, and side black ribbons represent the densities of the distributions of x_0 and z_0 , and y_0 , overlaid on the length of sph_0 along that coordinate shown in orange. Coordinates are relative to $P_{w,0}$.

increase in MT per $\|P_w - P\|$, but it is clear that such a simple model would yield poor MT predictions.³




Figure 5.9. Movement times for each of the distances between the target sphere and the wand at the highlight frame $\|P_w - P\|$, for the successful trials in the 1-sphere block. The top bars show the distribution of $\|P_w - P\|$ binned every 0.025 m. The orange lines represent the distance between the target sphere and $P_{w,0} \pm 0.05$ m.

5.7.2 Regression analysis

Successful trials in the 1-sphere block are grouped by (P, r, V, ϕ) condition. In accordance with the typical Fitts literature, each condition's MT is summarized using means. Due to their variation with respect to the controlled conditions, the wand and sphere positions per condition are set as their mean at the highlight frame, i.e., $\bar{P}_w = (\bar{P}_{w,x}, \bar{P}_{w,y}, \bar{P}_{w,z})$, and $\bar{P} = (\bar{P}_x, \bar{P}_y, \bar{P}_z)$. Using \bar{P}_w , and \bar{P} , D_m , and D_s are calculated using Equations (5.2), and (5.4).

Following Equation (5.5), a linear regression of the form

$$\overline{MT} \sim \mathcal{N}(a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m, \sigma^2) \quad (5.7)$$

is fit. The resulting estimates, presented in Table 5.1, show a very good fit in terms of R^2 , with credible values in $[0.89, 0.96]$. The distributions of a , b , c , and d decrease in magnitude

³Such a linear model would suffer from several deficiencies, which includes explaining only a small percentage of the variance of MT ($R^2 = 0.15$, portrayed by the narrow confidence interval compared to the span of the data [Gelman and Hill 2007, p. 62]), and the clearly heterogeneous variance in the residual distribution.

with V , and the non-intercept parameter with highest magnitude is c , potentially indicating the importance of the ballistic component of the movement along the target trajectory on MT . The regression is represented graphically in Figure 5.10.




Figure 5.10. Residuals $(\overline{MT} - \widehat{MT})$ vs. fitted values (\widehat{MT}) for the $\overline{MT} \sim \mathcal{N}(a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m, \sigma^2)$ regression. The dotted lines represent $\pm \hat{\sigma}$.

The residual plot of the regression, presented in Figure 5.10, suggests that the residual variance increases with the decrease in V . This implies that \overline{MT} predictions for low speed targets, $V = 0.5$, will be less accurate than predictions for high speed targets, $V = 1.5$.

Such a result is consistent with existing theories on moving target interception, which suggest that tasks with less temporal constraints (e.g., when targets have low V) are more prone to sensory-motor noise, thus resulting in more variable MT [Tresilian 2005, 134–135]. Therefore, given that the lack of precision in \overline{MT} predictions for lower V may be due to irreducible noise, and that there are only two visible outliers, the model is not further modified.

5.7.2.1 Change of coefficients per V

Consistently with the PD analyses from Chapter 4, and the gain analyses of Shoemaker et al. [2012], linear regressions on the per- V parameter estimates from Table 5.1— $\hat{a}_V = Mdn(a_V)$, $\hat{b}_V = Mdn(b_V)$, $\hat{c}_V = Mdn(c_V)$, and $\hat{d}_V = Mdn(d_V)$ —are first calculated using the default least-squares method in R [2015]. The results, shown in Table 5.2, give some evidence in support of the linear relation between V and \hat{a} , V and \hat{b} , V and \hat{c} , and, to a lesser degree V and \hat{d} .

Table 5.1. Regression estimates for $\overline{PD} = a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V ID_m$.

Estimand	Mdn	95% HDI
$a_{V=0.5}$	-3	[-4.75, -1.21]
$a_{V=1.0}$	-2.48	[-3.93, -1.04]
$a_{V=1.5}$	-1.82	[-3.14, -0.51]
$b_{V=0.5}$	1.25	[0.83, 1.67]
$b_{V=1.0}$	1.03	[0.61, 1.46]
$b_{V=1.5}$	0.83	[0.4, 1.24]
$c_{V=0.5}$	3.09	[1.36, 4.76]
$c_{V=1.0}$	2.47	[1.07, 3.88]
$c_{V=1.5}$	1.74	[0.54, 2.94]
$d_{V=0.5}$	0.16	[0.1, 0.23]
$d_{V=1.0}$	0.11	[0.04, 0.17]
$d_{V=1.5}$	0.09	[0.03, 0.16]
σ	0.07	[0.05, 0.08]
R^2	0.93	[0.89, 0.96]

Table 5.2. Least-squares estimates for the regressions of the point estimates from Table 5.1— $\hat{a}_V = \text{Mdn}(a_V)$, $\hat{b}_V = \text{Mdn}(b_V)$, $\hat{c}_V = \text{Mdn}(c_V)$, and $\hat{d}_V = \text{Mdn}(d_V)$ —on target speed V .

Model	Intercept	Slope	$\hat{\sigma}$	\hat{R}^2
$\hat{a}_V \sim \mathcal{N}(\alpha_1 + \alpha_2 V, \sigma_a^2)$	-3.62	1.19	0.06	0.99
$\hat{b}_V \sim \mathcal{N}(\beta_1 + \beta_2 V, \sigma_b^2)$	1.46	-0.42	0.01	1
$\hat{c}_V \sim \mathcal{N}(\zeta_1 + \zeta_2 V, \sigma_c^2)$	3.78	-1.35	0.04	1
$\hat{d}_V \sim \mathcal{N}(\delta_1 + \delta_2 V, \sigma_d^2)$	0.19	-0.07	0.02	0.79

As discussed in Section 4.4.2.2, the least-squares estimates on regression coefficients tend to be too optimistic. Thus, to better assess the uncertainty around the estimates of Table 5.2, a hierarchical model which simultaneously estimates the latter parameters and those in Table 5.1, is fit. The hierarchical model reuses the least-squares estimates $\hat{\alpha}_1, \hat{\alpha}_2, \hat{\beta}_1, \hat{\beta}_2, \hat{\zeta}_1, \hat{\zeta}_2, \hat{\delta}_1$, and $\hat{\delta}_2$ from Table 5.2 as priors for the mean parameters of $\alpha_1, \alpha_2, \beta_1, \beta_2, \zeta_1, \zeta_2, \delta_1$, and δ_2 , respectively; the wide prior variance (2^2), with respect to the range of \overline{MT} (0.38–1.36), allows for some flexibility both in the mean and the variance of the marginal posteriors. Given the low number of levels of $\#(V) = 3$, and following the recommendations of Gelman [2006] for standard deviation parameters on hierarchical models, half-Cauchy priors are chosen for $\sigma_a, \sigma_b, \sigma_c, \sigma_d$, and σ , with scale parameter $A = 1$, which, again, is wide with respect to the range of \overline{MT} .

The model is fit using Stan [2015], with four MCMC chains of 10,000 iterations, including 5,000 warm-up iterations, for a total of 20,000 saved iterations. The levels and estimates of the model shown in Table 5.3, show higher estimates for $\sigma_a, \sigma_b, \sigma_c$, and σ_d than those obtained in

Table 5.2, and correspondingly, wider intervals around $\hat{\alpha}_1, \hat{\alpha}_2, \hat{\zeta}_1, \hat{\zeta}_2, \hat{\delta}_1$, and $\hat{\delta}_2$ than those that would be obtained from the least-squares regression. Nevertheless, the minimal change in the distributions of a, b, c, d , and σ between Tables 5.1 and 5.3 suggests that, in spite of the wide uncertainty, the linear relations between V and a, b, c , and d are plausible and congruent with the data.

Table 5.3. Levels and estimates of the hierarchical regression $\overline{MT}_i \sim \mathcal{N}(a_{V[i]} + b_{V[i]} \sqrt{D_{Si}} + c_{V[i]} \sqrt{D_{mi}} + d_{V[i]} ID_{mi}, \sigma^2)$. $\hat{\alpha}_1, \hat{\alpha}_2, \hat{\beta}_1, \hat{\beta}_2, \hat{\zeta}_1, \hat{\zeta}_2, \hat{\delta}_1$, and $\hat{\delta}_2$ correspond to the least-squares estimates from Table 5.2.

Model	Estimand	Mdn	95% HDI
$a_V \sim \mathcal{N}(\alpha_1 + \alpha_2 V, \sigma_a^2)$	$a_{V=0.5}$	−3.08	[−4.26, −1.93]
	$a_{V=1.0}$	−2.48	[−3.39, −1.64]
	$a_{V=1.5}$	−1.84	[−2.94, −0.75]
$b_V \sim \mathcal{N}(\beta_1 + \beta_2 V, \sigma_b^2)$	$b_{V=0.5}$	1.27	[0.96, 1.59]
	$b_{V=1.0}$	1.04	[0.75, 1.31]
	$b_{V=1.5}$	0.83	[0.48, 1.2]
$c_V \sim \mathcal{N}(\zeta_1 + \zeta_2 V, \sigma_c^2)$	$c_{V=0.5}$	0.16	[0.1, 0.22]
	$c_{V=1.0}$	0.11	[0.05, 0.17]
	$c_{V=1.5}$	0.09	[0.03, 0.16]
$d_V \sim \mathcal{N}(\delta_1 + \delta_2 V, \sigma_d^2)$	$d_{V=0.5}$	0.16	[0.1, 0.22]
	$d_{V=1.0}$	0.11	[0.05, 0.17]
	$d_{V=1.5}$	0.09	[0.03, 0.16]
$\sigma \sim \text{half-Cauchy}(1)$	σ	0.07	[0.05, 0.08]
$\alpha_j \sim \mathcal{N}(\hat{\alpha}_j, 2^2)$	α_1	−3.7	[−5.69, −1.76]
	α_2	1.22	[−0.47, 2.98]
$\sigma_a \sim \text{half-Cauchy}(1)$	σ_a	0.28	[0, 1.34]
$\beta_j \sim \mathcal{N}(\hat{\beta}_j, 2^2)$	β_1	1.48	[0.21, 2.71]
	β_2	−0.44	[−1.57, 0.77]
$\sigma_b \sim \text{half-Cauchy}(1)$	σ_b	0.2	[0, 1.14]
$\zeta_j \sim \mathcal{N}(\hat{\zeta}_j, 2^2)$	ζ_1	3.86	[1.99, 5.8]
	ζ_2	−1.39	[−3.14, 0.27]
$\sigma_c \sim \text{half-Cauchy}(1)$	σ_c	0.28	[0, 1.36]
$\delta_j \sim \mathcal{N}(\hat{\delta}_j, 2^2)$	δ_1	0.19	[−0.74, 1.19]
	δ_2	−0.07	[−0.96, 0.82]
$\sigma_d \sim \text{half-Cauchy}(1)$	σ_d	0.11	[0, 0.86]

5.8 Predictive analysis

The $\widehat{T\text{Score}}_i(t)$ scoring function used in this section is calculated using the \widehat{MT} predictions given by Equation 5.5 with the parameter estimates from Table 5.1.

5.8.1 Three-sphere trials

5.8.1.1 Wand positions at the highlight frame

Before attempting to predict the intended sphere in each trial, the wand positions at the highlight frame are visualized to identify potential issues, and differences with the 1-sphere trials, that could affect the predictions.

As shown in Figure 5.11, compared to the 1-sphere block of trials, there was a larger percentage of wand positions within sph_0 at the highlight frame, $p(||P_w - P_{w,0}|| < .05) = 0.64$. “Cheating” was not possible in this block since the goal sphere was unknown before it became highlighted, as opposed to the 1-sphere block. Thus, as shown in the top right, and right-most distributions of Figure 5.11, instead of placing their wand beyond sph_0 , participants who did not have their wand within sph_0 tended to place their wand closer to their bodies, $p(P_{w,z} - z_0 > .05) = 0.22$, and below sph_0 , $p(P_{w,y} - y_0 < -.05) = 0.13$.




Figure 5.11. Front and right views of the distribution of wand positions at the frame where the goal sphere was highlighted in the 3-sphere trials. Red, blue, and green points represent the wand positions in trials where the sphere was at the bottom-right, bottom-left, and up, respectively. Correspondingly colored stars represent the sphere positions in those frames. The orange circle represents sph_0 . The top, and side black ribbons represent the densities of the distributions of x_0 and z_0 , and y_0 , overlayed on the length of sph_0 along that coordinate shown in orange. Coordinates are relative to $P_{w,0}$.

5.8.1.2 Predictive accuracy

At each frame (t) after the goal sphere is highlighted, $\widehat{T}Score$, and $dScore$ are calculated, with $N \in \{1, 2\} \times decay \in \{0.5, 0.9\}$. For each of these scores, the predicted sphere corresponds to the sphere with the maximum score. Subsequently, nine frames per trial are extracted, corresponding to the last frame before $0.1T, 0.2T, \dots, 0.9T, T$, where T is the time required to complete each trial. Finally, the accuracy across trial percentages pT is calculated, the results are presented in Figure 5.12.

The results indicate that the $\widehat{T}Score$ predictor is better than the $dScore$ predictor before 50% of the total trial time T has elapsed. Before $0.5T$, the $dScore$ predictor is less reliable than chance; after $0.5T$, the $dScore$ predictor's accuracy gets gradually better than that of $\widehat{T}Score$. This change in accuracy may reflect an increasing error in the \widehat{MT} predictions yielded by Equation 5.5 with the parameter estimates from Table 5.1. The increased error, however, is expected given that the D_s , and D_m predictors at $T > 0$ are outside of the range of the D_s , and D_m used to train the model, i.e., those at $T = 0$. Alternatively, this change could indicate a shift to a ballistic phase in which $D \simeq D_m$ is most influential on T . Finally, notice that at $1.0T$, none of the predictors have 100% accuracy, suggesting that the decay rates may be too low for the explored tasks.

5.8.2 Six-sphere trials

5.8.2.1 Wand positions at the highlight frame

As shown in Figure 5.13, the percentage of wand positions within sph_0 at the highlight frame was very similar to the 3-sphere block, $p(\|P_w - P_{w,0}\| < 0.05) = 0.6$. As in the 3-sphere block, participants who did not have their wand within sph_0 tended to place their wand closer to their bodies, $p(P_{w,z} - z_0 > 0.05) = 0.25$, but also to the right of sph_0 , $p(P_{w,x} - x_0 > 0.05) = 0.13$. Given that most participants were right handed, this suggests a search for a more comfortable waiting position.

5.8.2.2 Predictive accuracy

The accuracy of $\widehat{T}Score$, and $dScore$ across trial percentages $0.1T, 0.2T, \dots, 0.9T, T$ is calculated following the same procedure as in the 3-sphere trials. The results are presented in Figure 5.12.

The results indicate that the $\widehat{T}Score$ predictor is better than the $dScore$ predictor before 40% of the total trial time T has elapsed. Before $0.4T$, the $dScore$ predictor is less reliable than chance; instead, between $[0.5T, 0.7T]$, the $\widehat{T}Score$ predictor is less reliable than chance; and between $[0.8T, 1.0T]$, both predictors have fairly similar accuracy except when $N = 1$, and $decay = 0.9$, in which case the $dScore$ predictor is about 5% better. Notice, once again, that at $1.0T$, none of the predictors get 100% accuracy, suggesting that the decay rates may be too low for the explored tasks.

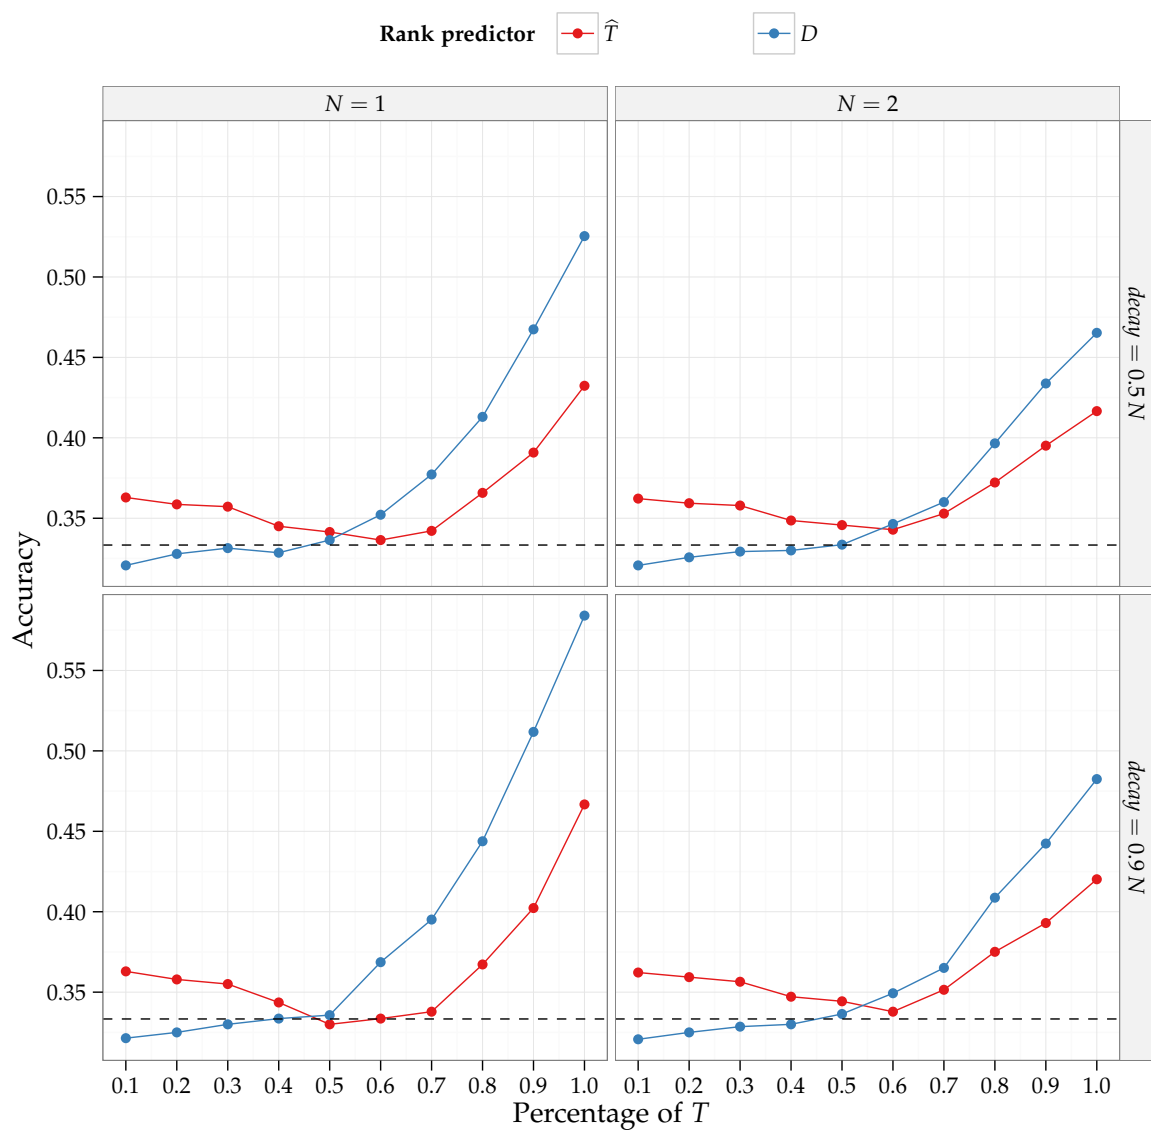


Figure 5.12. Three-sphere accuracy comparison for the $\hat{T}Score$, and $dScore$ predictors for trial percentages $0.1 T, 0.2 T, \dots, 0.9 T, T$, using different N , and $decay$ parameters. The dotted lines represent the accuracy given by chance, i.e. $1/3$.




Figure 5.13. Front and right views of the distribution of wand positions at the frame where the goal sphere was highlighted in the 6-sphere trials. Red, pink, blue, brown, green, and purple points represent the wand positions in trials where the sphere was at the bottom-right, bottom, bottom-left, up-left, up, and up-right, respectively. Correspondingly colored stars represent the sphere positions in those frames. The orange circle represents sph_0 . The top, and side black ribbons represent the densities of the distributions of x_0 and z_0 , and y_0 , overlayed on the length of sph_0 along that coordinate shown in orange. Coordinates are relative to $P_{w,0}$. Only a sample of 1,000 trials out of the 2,705 successful trials is shown.

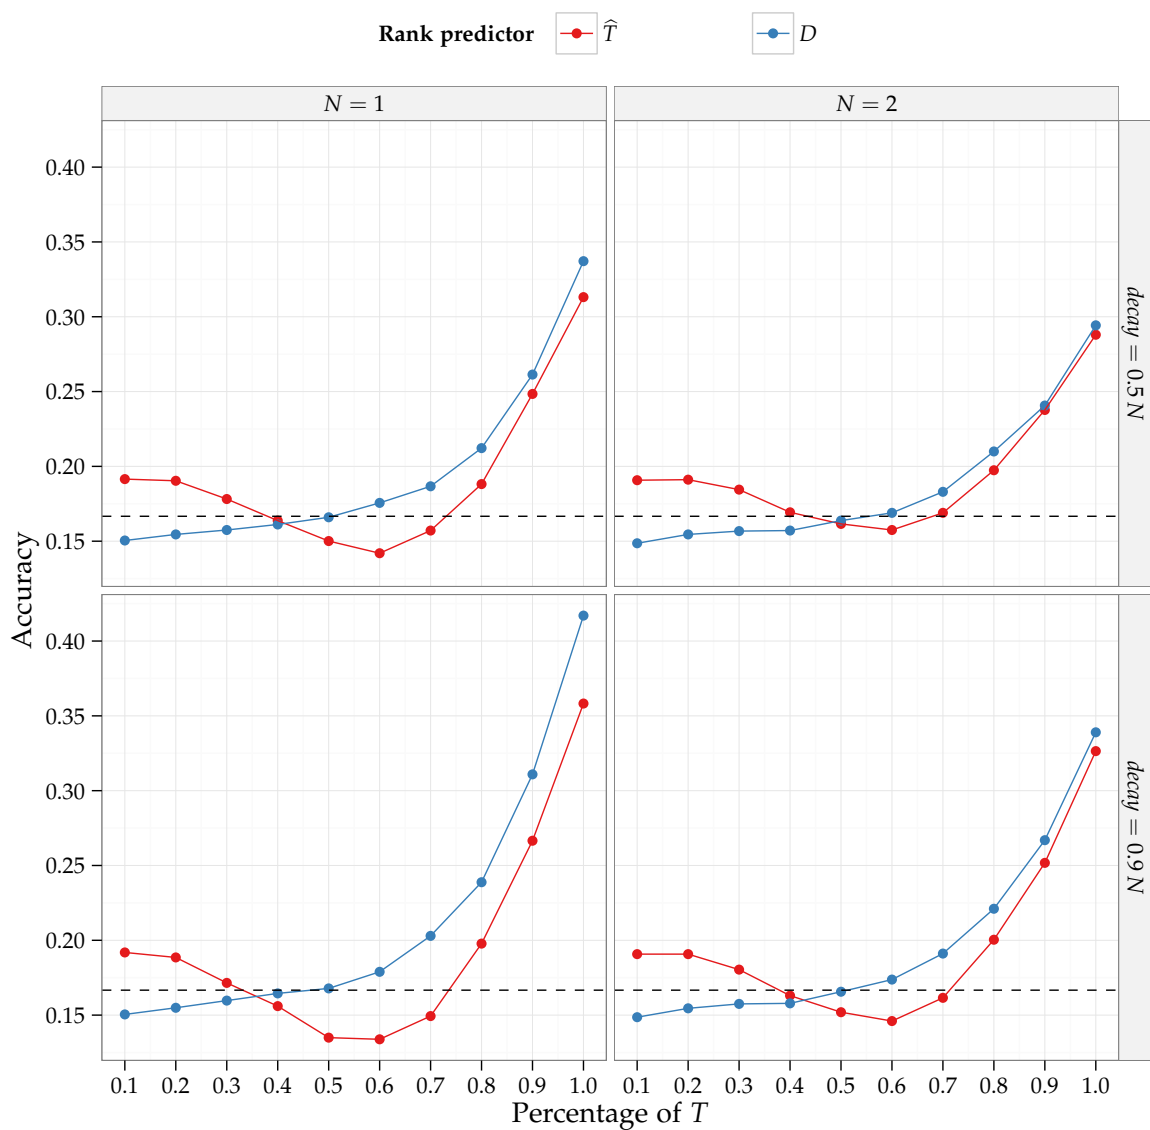


Figure 5.14. Three-sphere accuracy comparison for the $\hat{T}Score$, and $dScore$ predictors for trial percentages $0.1T, 0.2T, \dots, 0.9T, T$, using different N , and $decay$ parameters. The dotted lines represent the accuracy given by chance, i.e. $1/3$.

5.9 Discussion

Results are summarized according to the study hypotheses, each followed by a discussion.

5.9.1 Moving-target selection performance in 3-D

H1. The three-part model with per-speed coefficients, and per-velocity intercept, described in Equation (5.5), accurately models the movement time of three-dimensional moving-target selection tasks. *Supported.*

The three-part moving-target selection model of Equation (5.5),

$$MT = a_V + b_V \sqrt{D_s} + c_V \sqrt{D_m} + d_V \log_2 \left(\frac{2D_m}{W} \right),$$

yielded a good fit, with credible values of $R^2 \in [0.89, 0.96]$, presented in Table 5.1. The residuals, displayed in Figure 5.10, indicated a wider variance for $V = 0.5$, compared to $V \in 1.0, 1.5$, but given that there were only two visible outliers, no further modifications to the model were considered. Given that a very similar model was used to predict PD , these results give further evidence to the relation between PD , and MT .

H1.1. The azimuth α , altitude θ , and γ angles, not modeled in Equation (5.5), have a small effect on MT . *Partially supported.*

The α , and θ angles were varied simultaneously as sphere positions up, bottom-right, and bottom-left, which impeded the evaluation of their effects separately. Figure 5.6 suggested that targets located upward with respect to sph_0 had a larger MT than targets located below sph_0 , but this effect was not visible in the residual plot (Figure 5.10).

Finally, as seen in Figure 5.7, MT remained very similar across the range of assigned γ angles. This suggests a small effect of γ on MT .

H1.1. In Equation (5.5) the coefficients a_V , b_V , c_V , and d_V vary linearly with V . *Supported.*

The results suggest that coefficients a_V , b_V , c_V , and d_V of Equation (5.5) decrease linearly with the increase in V . Using the credible intervals presented in Table 5.3 allows MT to be predicted using different V , while still propagating the uncertainty in these estimates.

5.9.2 Predicting user intention in 3-D moving-target directed-selection tasks

H2. At any given instant, the minimal movement time estimated by Equation (5.5) can be used to predict the intended target in a 3-D moving-target directed-selection task. *Partially supported.*

Even though it was possible to predict the intended sphere using the estimated MT via the $\widehat{TScore}_i(t)$ scoring function with mostly above random accuracy, the resulting accuracy was very low. In certain cases in the 6-sphere block, the accuracy was actually worse than chance, as shown in Figure 5.14.

Notice, however, that the parameter estimates used to calculate \widehat{MT} were obtained from the 1-sphere block, which included a different range of D_s , and D_m inputs, had $RT = 0$, and did not have issues of target clutter or occlusion, as opposed to the 3- and 6-sphere blocks. Even though the Hick-Hyman law allows modeling $RT = \log_2(I + 1)$, where I is the number of targets in the scene, as the RT elapses, the D_m is reduced by $RT V$, and the initial wand position may change, possibly resulting in a different D_s . Therefore, integrating the Hick-Hyman law in the current model is not straightforward.

A more viable way to enhance \widehat{MT} predictions in blocks with multiple targets, and presumptively the accuracy of $\widehat{TScore}_i(t)$, would be to include subject variations in the parameters of Equation (5.5). Indeed, the parameter estimates of Table 5.1 are given for the average MT per-condition, and not for all trials.

H2.1. The accuracy of the predictions based on estimated movement time will be greater than those based on wand-target distance. *Partially supported.*

The $\widehat{TScore}_i(t)$ scoring function yielded better accuracies than $dScore_i(t)$ only for trials before $0.5 T$ and $0.4 T$ in the 3- and 6-sphere blocks, respectively. In all subsequent times, $dScore_i(t)$ was consistently more accurate. As previously suggested, this inversion in prediction accuracy may indicate a shift in movement to a ballistic phase as the trial advances, in which $D \simeq D_m$ is most influential on T .

5.10 Conclusion

This chapter explored the possibility of using the 2-D moving-target models for PD presented in Chapter 4 to predict performance and intention in 3-D moving-target selection. The model yielded a good fit on the \overline{MT} values of the 1-sphere block, but it exhibited poor, yet above random, accuracy predicting intended targets in the 3- and 6-sphere blocks. Future work should explore possible extensions to this model to accommodate for user variations. Additionally, as shown in Chapter 3, combining this predictor with other features may result in considerable increases in accuracy.

To the author's knowledge, this is the first study to model MT in 3-D moving-target selection tasks, as well as the first study to model MT for moving-target selection tasks of any dimensionality in which the target is not moving directly toward, or directly away from the cursor.

CHAPTER 6

CONCLUSION

The current work provides an in-depth, empirical investigation on moving-target selection, a task that is increasingly common, yet largely understudied in human–computer interaction. In particular, predictive models were developed for three of the stages involved in the action of selecting a moving target: the intended selection, the prospective difficulty selection, and the selection performance.

Intention prediction models were developed using decision trees and scoring functions with features specific to undirected selection tasks (i.e., tasks in which the users are free to select an object among multiple others), and generalizable features that can also be used in directed selection tasks (i.e., the more common experimental tasks in which users are instructed to select a specific object). As shown in Chapter 3, target size (W) was deemed as the most predictive feature for intention in undirected selection, with an accuracy of ($\sim 71\%$) for two targets with equal velocities, but different initial positions—this was contrary to the expectation that ID , measured at the start of each trial, would be the most predictive feature in this type of task. In terms of generalizable features in undirected selection, results presented in Chapter 3 indicate that a combination of head-target relative gaze (Δdot), and cursor-target relative distance (ΔD), averaged in a 1–1.5 s time window, were predictive of selection intention with an accuracy of ($\sim 72\%$). The combination of these task-specific and generalizable features in undirected selection resulted in an accuracy of ($\sim 78\%$). Finally, in Chapter 5, scoring functions based on the predicted movement time (\widehat{MT}), and cursor–target distance, yielded poor, yet above random, accuracy in predicting the intended target in a directed selection task with 3 and 6 spheres.

Prospective difficulty (PD) models were developed and evaluated in 1-D, and 2-D moving-target selection tasks in Chapter 4. These models describe the motion required to attain a target in three parts: two ballistic motions, one to align the cursor with the target’s movement axis and one in-line with the target’s movement axis, and a homing in motion in-line with the target’s movement axis. The ballistic motions were represented with square-root distances, whereas the homing-in motion was represented with Fitts’ ID . Overall, PD was shown to increase with target speed (V), ID , and the cursor–target-axis distance, and to decrease with

the cursor–target distance along the target’s movement axis. Contrary to the expectations, the ballistic motion required to align the cursor with the target’s movement axis, and the homing-in motion along the target’s movement axis were minimally affected by target speed.

Performance prediction was explored in Chapter 5 by evaluating the usage of the 2-D model of *PD* presented in Chapter 4 as a mean movement time (\overline{MT}) predictor in 3-D moving-target selection. The three-part model yielded a good fit with the observed data, but, contrary to the *PD* model, the coefficients related to ballistic and homing-in motions decreased as the target speed increased. Such a difference was expected given that faster targets require shorter movement times to be successfully selected.

6.1 Limitations and future work

6.1.1 Intention prediction models

Overall, the main limitation of the studied intention prediction models is that they were generated based on a post hoc analysis of moving-target selection. Nevertheless, their usefulness to address the challenges of moving-target selection needs to be assessed in an interactive context, where predictions can be integrated to enhance the moving-target task (e.g., using predictions as inputs of the enhancements described in Section 1.1.2).

Decision trees can integrate different predictive features and can be interpreted as simple if-else rules, allowing them to be implemented in real-time. If the predictions were to be adapted during execution, however, the major difficulty would be to recalculate the trees in real-time without impacting performance. Future work should explore the possible optimizations that would allow adaptation of decision trees during interaction.

Scoring functions, on the other hand, use only one input feature, but can adapt in real-time to changes in user and target states with a small computational overhead. Unfortunately, the accuracy of such functions was very poor in directed tasks, which could indicate the limitations of using a single feature for scoring. Future work should explore accuracy enhancements to the proposed scoring functions, which could include integration of different features as composite scores, as well as parameter tuning.

6.1.2 *PD* prediction models

In order to develop the models for *PD* in 2-D moving-target selection mentioned earlier in this chapter, the existing *PD* model for 1-D static-target selection was first extended to 2-D. Thus, the limitations and future work related to the formulation of this 2-D static-target model are presented prior to those related to the 1-D, and 2-D moving-target models.

6.1.2.1 Two-dimensional static-target tasks

Target angle θ . The main limitation of the 2-D static target tasks evaluated in Chapter 4 was the small range of explored target angles θ . Although the literature shows that both diagonal and vertical targets are the hardest to attain, this study explored only diagonal, right-handed angles, for which the angle effect was very small. It is possible, however, that vertical targets could have a larger impact on PD , due to the tendency to perceive vertical distances as longer than horizontal distances [Higashiyama 1996]. To have a better understanding of the effect of θ , future work on 2-D static-targets should include a wider range of target angles.

Underlying function for a_θ . Perhaps related to the small range of explored angles, was the inability to find an underlying function for a_θ . Neither of the functional forms of $f(\theta)$ suggested by Murata and Iwase [2001] or Appert et al. [2008] gave a reasonable approximation to a_θ . This prevents generalization of the 2-D model, and making PD predictions for other angles based on it. Yet, assuming that the target-angle effect is small for most angles $\theta \in [0, 360)^\circ$, it may be possible to get reasonable predictions using the 1-D static-target formula.

Future work should continue to explore the underlying function for the angle effect represented in a_θ . Once again, this exploration may be facilitated by including a wider variety of θ angles in the experimental conditions.

6.1.2.2 Moving-target tasks

Target velocity (V, ϕ) . Similar to static-target tasks, the main limitation in the evaluated moving-target tasks was the small range of explored target speeds V , and movement directions ϕ .¹ With the current range of evaluated velocities, it is difficult to generalize the results to additional experimental conditions; even though the wide credible intervals obtained from the hierarchical models allow propagating this uncertainty to predictions on new data. Future work should include a wider variety of velocities and compare the current models' predictions to the actual observations.

Initial target angle θ . In all of the evaluated moving-target pointing tasks, the initial target angle $\theta = 0$. Based on the 2-D static-target model, different target angles could be represented with different intercepts, which would result in a $a_{V \times \theta}$ term. It is hypothesized, however, that the θ effect will have a minimal impact on PD , especially for tasks where $\phi \notin \{0, 180\}^\circ$, since in those cases the θ angle is changing at every frame. Future work should explore these claims with experiments including initial target angles $\theta \in [0, 360)^\circ$.

Inter-formula compatibility. The 1-D moving-target pointing model is not nested within the 2-D model, i.e., there is no constraint on the 2-D model that would lead to the 1-D one. Specifically, the ID term varies per speed V for the 1-D moving-target formula, but it stays

¹It is worth noting that previous experiments on moving-target selection included only two [Jagacinski et al. 1980] or three [Al Hajri et al. 2011] levels of non-zero speeds, and none of them explored target movement directions $\phi \notin \{0, 180\}^\circ$.

almost constant in the 2-D formula. This may imply that there is a shortcoming in the derivation of the 2-D formula, or that 1-D moving-target tasks are inherently different from 2-D ones (e.g., because users may be more inclined to wait for the target in the former), and this difference is reflected in an interaction between ID and V that does not happen in 2-D.

Nevertheless, the six bottom panes of Figure 4.13 indicate that the 1-D moving-target data can be adequately described by the 2-D moving-target model. Once again, future work should evaluate these formulae with an extended range of velocities to evaluate this apparent incongruence and explore the theoretical and practical limitations of the proposed models.

6.1.2.3 *Methods for assessing PD*

A possible shortcoming in the evaluation of the proposed PD models concerns the reliability of the prospective assessments given by the user on each task. By using a discrete scale to evaluate a continuous quantity, there is an inherent loss of precision in each assessment. Regardless of the scale, it may be difficult for participants to quantify their prospective judgments. Therefore, future work should also attempt to assess the proposed models of PD in moving-target selection using measures that do not rely on self assessment, such an approach has been explored in static-target selection using EEG [Kourtis et al. 2012].

6.1.3 Performance prediction models

6.1.3.1 *Target velocity*

Similar to PD , the study of MT for tasks with a wider range of V , and ϕ values would allow further generalization of the proposed models. In particular, more precise estimates for the linear relation between V and the coefficients of Equation (5.5) can be obtained by studying selection tasks with additional levels of V .

6.1.3.2 *Target distance*

In the design of the experiment described in Chapter 5, D was kept constant at 1 m, and the change in D_m , and D_s , was ensured by changing the ϕ angle. This constraint allowed the study of other experimental conditions without making the main study too long, but led to heavily correlated D_m , and D_s values. Even though the variability in initial wand positions mitigated this effect, such a correlation impedes a clear distinction of the effects of D_m , and D_s on MT . Therefore, future work should explore different D values.

6.1.3.3 *Additional measures of performance*

In accordance with the main body of target selection in HCI, the analysis of performance was mostly limited to MT . However, considering that there is a critical velocity beyond which selection is impossible, it seems important to quantify the effect of V , and other experimental factors, on target success rate (ρ). Predicting ρ would allow a reevaluation of

the Hoffmann [1991] moving-target model,² as well as the refinement of intention prediction models, which are currently limited to successful trials. Thus, future work on moving-target selection should also evaluate and attempt to model ρ , and other measures of performance.

6.1.4 General issues

6.1.4.1 Target shape

Target shape was not considered in the development of any of the models. Nevertheless, a formulation for 2-D bivariate targets (e.g., rectangles, where $W \neq H$) was suggested for *PD*, in which each $\theta \times (W/H)$ combination should have its own intercept, as specified in Equation (4.10). The major drawback of this solution is that the degrees of freedom of the model are reduced with each $\theta \times (W/H)$ combination, requiring a large number of data points to be evaluated. This approach can also be generalized to 3-D.

For more complex shapes a different *ID* term, such as those proposed in [Sheikh and Hoffmann 1994; Murata 1999; Grossman and Balakrishnan 2005], may be necessary. Since these *ID* terms implicitly include the target's angular position effect, the per-angle intercept may add unnecessary redundancy to the model. Therefore, future work should explore the effect of target shape on *PD* and *MT*.

6.1.4.2 Target acceleration

In the explored moving-target selection tasks, the target velocity was always constant. Such simplification is acceptable in certain tasks where targets have nearly constant velocities (e.g., in air traffic control systems, or in tasks where objects are moving on a conveyor belt) but may not be reasonable for most motions that follow real-world dynamics. Thus, it is likely that the proposed models for *PD*, and *MT* will yield less accurate responses for targets whose velocities are markedly non-constant. On the other hand, intention predictions based on inputs measured throughout a time window are less likely to suffer from a loss in accuracy due to changes in velocity.³ To evaluate these claims, future work should investigate the effect of target acceleration on selection performance, prospective difficulty, and intention detection.

²As described in Chapter 2, the usage of the Hoffmann moving-target model without ρ results in potentially harder and unreliable calculations.

³Indeed, in Ortega's experiment [Ortega 2013], target movement directions were constantly changing.

BIBLIOGRAPHY

- Johnny Accot and Shumin Zhai. 2003. Refining Fitts' Law Models for Bivariate Pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, Ft. Lauderdale, Florida, USA, 193–200. DOI:<http://dx.doi.org/10.1145/642611.642646>
- Abir Al Hajri, Sidney Fels, Gregor Miller, and Michael Ilich. 2011. Moving Target Selection in 2D Graphical User Interfaces. In *Proceedings of the 13th IFIP TC 13 International Conference on Human-Computer Interaction - Volume Part II (INTERACT '11)*, Pedro Campos, Nicholas Graham, Joaquim Jorge, Nuno Nunes, Philippe Palanque, and Marco Winckler (Eds.). Springer-Verlag, Lisbon, Portugal, 141–161. DOI:http://dx.doi.org/10.1007/978-3-642-23771-3_12
- Caroline Appert, Olivier Chapuis, and Michel Beaudouin-Lafon. 2008. Evaluation of Pointing Performance on Screen Edges. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '08)*. ACM, Napoli, Italy, 119–126. DOI:<http://dx.doi.org/10.1145/1385569.1385590>
- Ferran Argelaguet and Carlos Andujar. 2013. A survey of 3D object selection techniques for virtual environments. *Computers & Graphics* 37, 3 (2013), 121–136. DOI:<http://dx.doi.org/10.1016/j.cag.2012.12.003>
- Jeremy N. Bailenson, Michael S. Shum, and David H. Uttal. 2000. The initial segment strategy: a heuristic for route selection. *Memory & Cognition* 28, 2 (March 2000), 306–18. DOI:<http://dx.doi.org/10.3758/BF03213808>
- Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning* (9 ed.). Springer, New York, New York, USA. 738 pages.
- Renaud Blanch and Michaël Ortega. 2009. Rake cursor: improving pointing performance with concurrent input channels. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, Boston, Massachusetts, USA, 1415–1418. DOI:<http://dx.doi.org/10.1145/1518701.1518914>
- James Boritz, Kellogg S. Booth, and William B. Cowan. 1991. Fitts's Law Studies of Directional Mouse Movement. In *Graphics Interface '91 (GI '91)*. Canadian Information Processing Society, Calgary, Alberta, Canada, 216–223.

- Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, and Ivan Poupyrev. 2004. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, Boston, Massachusetts, USA. 512 pages.
- George E. P. Box and Norman R. Draper. 1987. *Empirical Model-Building and Response Surfaces*. John Wiley & Sons, Inc., New York, New York, USA. 669 pages.
- Stuart K. Card, William K. English, and Betty J. Burr. 1978. Evaluation of Mouse, Rate-Controlled Isometric Joystick, Step Keys, and Text Keys for Text Selection on a CRT. *Ergonomics* 21, 8 (1978), 601–613. doi:<http://dx.doi.org/10.1080/00140137808931762>
- Stuart K. Card, Thomas P. Moran, and Allen Newell. 1980. The Keystroke-Level Model for User Performance Time with Interactive Systems. *Commun. ACM* 23, 7 (July 1980), 396–410. doi:<http://dx.doi.org/10.1145/358886.358895>
- Stuart K. Card, Thomas P. Moran, and Allen Newell. 1983. *The Psychology of Human-Computer Interaction*. Vol. 97. Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey. 625 pages. doi:<http://dx.doi.org/10.2307/1422176>
- Heather Carnahan and Bradford J. McFadyen. 1996. Visuomotor control when reaching toward and grasping moving targets. *Acta Psychologica* 92, 1 (1996), 17–32. doi:[http://dx.doi.org/10.1016/0001-6918\(95\)00006-2](http://dx.doi.org/10.1016/0001-6918(95)00006-2)
- Juan Sebastián Casallas, James H. Oliver, Jonathan W. Kelly, Frédéric Merienne, and Samir Garbaya. 2013. Towards a Model for Predicting Intention in 3D Moving-target Selection Tasks. In *Proceedings of the 10th International Conference on Engineering Psychology and Cognitive Ergonomics: Understanding Human Cognition - Volume Part I (EPCE '13)*, Don Harris (Ed.). Springer-Verlag, Las Vegas, Nevada, USA, 13–22. doi:http://dx.doi.org/10.1007/978-3-642-39360-0_2
- Juan Sebastián Casallas, James H. Oliver, Jonathan W. Kelly, Frédéric Merienne, and Samir Garbaya. 2014. Using Relative Head and Hand-Target Features to Predict Intention in 3D Moving-Target Selection. In *Proceedings of the 2014 IEEE Virtual Reality Conference (VR 2014)*. IEEE, Minneapolis, Minnesota, USA, 51–56. doi:<http://dx.doi.org/10.1109/VR.2014.6802050>
- Alan H. S. Chan and Errol R. Hoffmann. 2013. Subjective Difficulty of Movements With Ongoing Visual Control. *Journal of Motor Behavior* 45, 6 (2013), 507–517. doi:<http://dx.doi.org/10.1080/00222895.2013.833081>
- Nicholas Christenfeld. 1995. Choices from Identical Options. *Psychological Science* 6, 1 (1995), 50–55. doi:<http://dx.doi.org/10.1111/j.1467-9280.1995.tb00304.x>

- Sue V. G. Cobb, Sarah Nichols, Amanda Ramsey, and John R. Wilson. 1999. Virtual Reality-Induced Symptoms and Effects (VRISE). *Presence: Teleoperators & Virtual Environments* 8, 2 (April 1999), 169–186. doi:<http://dx.doi.org/10.1162/105474699566152>
- Gerwin de Haan, Michal Koutek, and Frits H. Post. 2005. IntenSelect: using dynamic object rating for assisting 3D object selection. In *Proceedings of the 11th Eurographics conference on Virtual Environments (EGVE'05)*. Eurographics Association, Aalborg, Denmark, 201–209. doi:http://dx.doi.org/10.2312/EGVE/IPT_EGVE2005/201-209
- Didier Delignières. 1990. *La difficulté en escalade : exigences objectives et perception des exigences des tâches motrices*. Diplôme de l'I.N.S.E.P. Institut national du sport et de l'éducation physique.
- Didier Delignières. 1993. *Approche psychophysique de la perception de la difficulté dans les tâches perceptivo-motrices*. Ph.D. Dissertation. Université Paris V.
- Didier Delignières and Jean-Pierre Famose. 1992. Perception de la difficulté, entropie et performance. *Science & Sports* 7, 4 (1992), 245–252. doi:[http://dx.doi.org/10.1016/S0765-1597\(05\)80097-8](http://dx.doi.org/10.1016/S0765-1597(05)80097-8)
- Pierre Dragicevic, Fanny Chevalier, and Stephane Huot. 2014. Running an HCI Experiment in Multiple Parallel Universes. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)*. ACM, Paris, France, 607–618. doi:<http://dx.doi.org/10.1145/2559206.2578881>
- Heiko Drewes. 2010. Only One Fitts' Law Formula Please!. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems (CHI EA '10)*. ACM, Atlanta, Georgia, USA, 2813–2822. doi:<http://dx.doi.org/10.1145/1753846.1753867>
- J.P. Famose, D. Delignières, and J. Genty. 1991. Perception de la difficulté dans une tâche d'anticipation-coïncidence. In *IVèmes Journées Internationales d'Automne de L'A.C.A.P.S. Lille*, France, 1–4.
- Paul M. Fitts. 1954. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. *Journal of Experimental Psychology* 47, 6 (1954), 381–391. doi:<http://dx.doi.org/10.1037/h0055392>
- Jérémie Francone and Laurence Nigay. 2011. Using the user's point of view for interaction on mobile devices. In *23rd French Speaking Conference on Human-Computer Interaction (IHM '11)*. ACM, Sophia Antipolis, France, 4:1–4:8. doi:<http://dx.doi.org/10.1145/2044354.2044360>
- Khai-Chung Gan and Errol R. Hoffmann. 1988. Geometrical conditions for ballistic and visually controlled movements. *Ergonomics* 31, 5 (1988), 829–839. doi:<http://dx.doi.org/10.1080/00140138808966724>

- Andrew Gelman. 2006. Prior distributions for variance parameters in hierarchical models (Comment on Article by Browne and Draper). *Bayesian Analysis* 1, 3 (2006), 515–534. doi:<http://dx.doi.org/10.1214/06-BA117A>
- Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2013. *Bayesian Data Analysis* (3 ed.). CRC press, Boca Raton, Florida, USA. 675 pages.
- Andrew Gelman and Jennifer Hill. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models* (1 ed.). Cambridge University Press, New York, New York, USA. 648 pages.
- Suzanne M. Grilli. 2011. *Perceived Difficulty in a Fitts Task*. Master’s Thesis. Cleveland State University. http://rave.ohiolink.edu/etdc/view?acc_num=csu1322544972
- Tovi Grossman and Ravin Balakrishnan. 2004. Pointing at Trivariate Targets in 3D Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’04)*. ACM, Vienna, Austria, 447–454. doi:<http://dx.doi.org/10.1145/985692.985749>
- Tovi Grossman and Ravin Balakrishnan. 2005. A Probabilistic Approach to Modeling Two-dimensional Pointing. *ACM Transactions on Computer-Human Interaction* 12, 3 (Sept. 2005), 435–459. doi:<http://dx.doi.org/10.1145/1096737.1096741>
- Yves Guiard and Michel Beaudouin-Lafon. 2004. Fitts’ law 50 years later: applications and contributions from human-computer interaction. *International Journal of Human-Computer Studies* 61, 6 (Dec. 2004), 747–750. doi:<http://dx.doi.org/10.1016/j.ijhcs.2004.09.003>
- Tyler J Gunn, Pourang Irani, and John Anderson. 2009. An evaluation of techniques for selecting moving targets. In *CHI ’09 Extended Abstracts on Human Factors in Computing Systems (CHI EA ’09)*. ACM, New York, NY, USA, 3329–3334. doi:<http://dx.doi.org/10.1145/1520340.1520481>
- Mark Hall, Hazeltine National, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA Data Mining Software : An Update. *SIGKDD Explorations Newsletter* 11, 1 (2009), 10–18. doi:<http://dx.doi.org/10.1145/1656274.1656278>
- Mark S. Hancock and Kellogg S. Booth. 2004. Improving Menu Placement Strategies for Pen Input. In *Proceedings of Graphics Interface 2004 (GI ’04)*. Canadian Human-Computer Communications Society, London, Ontario, Canada, 221–230. <http://dl.acm.org/citation.cfm?id=1006058.1006085>

- Khalad Hasan, Tovi Grossman, and Pourang Irani. 2011. Comet and Target Ghost: Techniques for Selecting Moving Targets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, Vancouver, British Columbia, Canada, 839–848. doi:<http://dx.doi.org/10.1145/1978942.1979065>
- Mary Hegarty and David Waller. 2004. A dissociation between mental rotation and perspective-taking spatial abilities. *Intelligence* 32, 2 (2004), 175–191. doi:<http://dx.doi.org/10.1016/j.intell.2003.12.001>
- William Edmund Hick. 1952. On the rate of gain of information. *Quarterly Journal of Experimental Psychology* 4, 1 (1952), 11–26. doi:<http://dx.doi.org/10.1080/17470215208416600>
- Atsuki Higashiyama. 1996. Horizontal and vertical distance perception: The discorded-orientation theory. *Perception & Psychophysics* 58, 2 (1996), 259–270. doi:<http://dx.doi.org/10.3758/BF03211879>
- Errol R. Hoffmann. 1991. Capture of moving targets: a modification of Fitts' Law. *Ergonomics* 34, 2 (1991), 211–220. doi:<http://dx.doi.org/10.1080/00140139108967307>
- Errol R. Hoffmann. 2013. Which Version/Variation of Fitts' Law? A Critique of Information-Theory Models. *Journal of Motor Behavior* 45, 3 (2013), 205–215. doi:<http://dx.doi.org/10.1080/00222895.2013.778815>
- Errol R. Hoffmann. 2014. *Validity of the Shannon form of Fitts' law: a response to MacKenzie (2013)*. Technical Report. 22 pages. errol@tpg.com.au
- Errol R. Hoffmann and Alan H. S. Chan. 2012. Underwater movement times with ongoing visual control. *Ergonomics* 55, 12 (2012), 1513–1523. doi:<http://dx.doi.org/10.1080/00140139.2012.719038>
- Errol R. Hoffmann and Ilyas H. Sheikh. 1994. Effect of varying target height in a Fitts' movement task. *Ergonomics* 37, 6 (1994), 1071–1088. doi:<http://dx.doi.org/10.1080/00140139408963719>
- Christian Holz and Patrick Baudisch. 2010. The Generalized Perceived Input Point Model and How to Double Touch Accuracy by Extracting Fingerprints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, Atlanta, Georgia, USA, 581–590. doi:<http://dx.doi.org/10.1145/1753326.1753413>
- Ray Hyman. 1953. Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology* 45, 3 (March 1953), 188–196. doi:<http://dx.doi.org/10.1037/h0056940>
- Michael Victor Ilich. 2009. *Moving Target Selection in Interactive Video*. Master's Thesis. The University of British Columbia. <http://hdl.handle.net/2429/17444>

- Richard J. Jagacinski and Donald L. Monk. 1985. Fitts' Law in Two Dimensions with Hand and Head Movements. *Journal of Motor Behavior* 17, 1 (1985), 77–95. doi:<http://dx.doi.org/10.1080/00222895.1985.10735338>
- Richard J. Jagacinski, Daniel W. Repperger, Sharon L. Ward, and Martin S. Moran. 1980. A Test of Fitts' Law with Moving Targets. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 22, 2 (1980), 225–233. doi:<http://dx.doi.org/10.1177/001872088002200211>
- Roland S. Johansson, Göran Westling, Anders Bäckström, and J. Randall Flanagan. 2001. Eye-hand Coordination in Object Manipulation. *The Journal of Neuroscience* 21, 17 (Sept. 2001), 6917–32. <http://www.ncbi.nlm.nih.gov/pubmed/11517279>
- Helmut Jorke, Arnold Simon, and Markus Fritz. 2008. Advanced Stereo Projection Using Interference Filters. In *3DTV Conference: The True Vision - Capture, Transmission and Display of 3D Video, 2008 (3DTV 2008)*. IEEE, Istanbul, Turkey, 177–180. doi:<http://dx.doi.org/10.1109/3DTV.2008.4547837>
- Marcel Adam Just and Patricia A. Carpenter. 1976. Eye fixations and cognitive processes. *Cognitive Psychology* 8, 4 (Oct. 1976), 441–480. doi:[http://dx.doi.org/10.1016/0010-0285\(76\)90015-3](http://dx.doi.org/10.1016/0010-0285(76)90015-3)
- Maurits Kaptein and Judy Robertson. 2012. Rethinking Statistical Analysis Methods for CHI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, Austin, Texas, USA, 1105–1114. doi:<http://dx.doi.org/10.1145/2207676.2208557>
- Robert S. Kennedy, Norman E. Lane, Kevin S. Berbaum, and Michael G. Lilienthal. 1993. Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness. *The International Journal of Aviation Psychology* 3, 3 (1993), 203–220. doi:http://dx.doi.org/10.1207/s15327108ijap0303_3
- Regis Kopper, Doug A. Bowman, Mara G. Silva, and Ryan P. McMahan. 2010. A human motor behavior model for distal pointing tasks. *International Journal of Human-Computer Studies* 68, 10 (Oct. 2010), 603–615. doi:<http://dx.doi.org/10.1016/j.ijhcs.2010.05.001>
- D. Kourtis, N. Sebanz, and G. Knoblich. 2012. EEG correlates of Fitts's law during preparation for action. *Psychological Research* 76, 4 (July 2012), 514–24. doi:<http://dx.doi.org/10.1007/s00426-012-0418-z>
- John K. Kruschke, Herman Aguinis, and Harry Joo. 2012. The Time Has Come: Bayesian Methods for Data Analysis in the Organizational Sciences. *Organizational Research Methods* 15, 4 (2012), 722–752. doi:<http://dx.doi.org/10.1177/1094428112457829>

- Edward Lank, Yi-Chun Nikko Cheng, and Jaime Ruiz. 2007. Endpoint Prediction Using Motion Kinematics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, San Jose, California, USA, 637–646. doi:<http://dx.doi.org/10.1145/1240624.1240724>
- I. Scott MacKenzie. 1989. A Note on the Information-Theoretic Basis for Fitts' Law. *Journal of Motor Behavior* 21, 3 (1989), 323–330. doi:<http://dx.doi.org/10.1080/00222895.1989.10735486>
- I. Scott MacKenzie. 1992. Fitts' Law as a Research and Design Tool in Human-Computer Interaction. *Human-Computer Interaction* 7, 1 (March 1992), 91–139. doi:http://dx.doi.org/10.1207/s15327051hci0701_3
- Ian Scott MacKenzie. 2013. A Note on the Validity of the Shannon Formulation for Fitts' Index of Difficulty. *Open Journal of Applied Sciences* 3, 6 (2013), 360–368. doi:<http://dx.doi.org/10.4236/ojapps.2013.36046>
- I. Scott MacKenzie and William Buxton. 1992. Extending Fitts' Law to Two-dimensional Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '92)*. ACM, Monterey, California, USA, 219–226. doi:<http://dx.doi.org/10.1145/142750.142794>
- Regan L Mandryk and Calvin Lough. 2011. The Effects of Intended Use on Target Acquisition. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, Vancouver, British Columbia, Canada, 1649–1652. doi:<http://dx.doi.org/10.1145/1978942.1979182>
- Michael J. McGuffin and Ravin Balakrishnan. 2005. Fitts' law and expanding targets: Experimental studies and designs for user interfaces. *ACM Transactions on Computer-Human Interaction* 12, 4 (Dec. 2005), 388–422. doi:<http://dx.doi.org/10.1145/1121112.1121115>
- Tom M. Mitchell. 1997. *Machine learning* (1 ed.). McGraw-Hill, Boston, MA. 432 pages.
- David Mould and Carl Gutwin. 2004. The effects of feedback on targeting with multiple moving targets. In *Proceedings of Graphics Interface 2004 (GI '04)*. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 25–32. <http://dl.acm.org/citation.cfm?id=1006058.1006062>
- Atsuo Murata. 1999. Extending Effective Target Width in Fitts' Law to a Two-Dimensional Pointing Task. *International Journal of Human-Computer Interaction* 11, 2 (1999), 137–152. doi:http://dx.doi.org/10.1207/S153275901102_4
- Atsuo Murata and Hirokazu Iwase. 2001. Extending fitts' law to a three-dimensional pointing task. *Human Movement Science* 20, 6 (2001), 791–805. doi:[http://dx.doi.org/10.1016/S0167-9457\(01\)00058-6](http://dx.doi.org/10.1016/S0167-9457(01)00058-6)

- Norman Murray, Dave Roberts, Anthony Steed, Paul Sharkey, Paul Dickerson, and John Rae. 2007. An assessment of eye-gaze potential within immersive virtual environments. *ACM Transactions on Multimedia Computing, Communications and Applications* 3, 4 (Dec. 2007), 8:1–8:17. doi:<http://dx.doi.org/10.1145/1314303.1314311>
- Mathieu Nancel, Olivier Chapuis, Emmanuel Pietriga, Xing-Dong Yang, Pourang P Irani, and Michel Beaudouin-Lafon. 2013. High-precision pointing on large wall displays using small handheld devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, Paris, France, 831–840. doi:<http://dx.doi.org/10.1145/2470654.2470773>
- Kai Nickel and Rainer Stiefelhagen. 2003. Pointing gesture recognition based on 3D-tracking of face, hands and head orientation. In *Proceedings of the 5th International Conference on Multimodal Interfaces (ICMI '03)*. ACM, Vancouver, British Columbia, Canada, 140–146. doi:<http://dx.doi.org/10.1145/958432.958460>
- Karin Nieuwenhuizen, Jean-Bernard Martens, Lei Liu, and Robert Van Liere. 2009. Insights from Dividing 3D Goal-Directed Movements into Meaningful Phases. *IEEE Computer Graphics and Applications* 29, 6 (Nov. 2009), 44–53. doi:<http://dx.doi.org/10.1109/MCG.2009.121>
- Donald A. Norman. 1986. Cognitive Engineering. In *User Centered System Design: New Perspectives on Human-Computer Interaction*, Donald A. Norman and Stephen W. Draper (Eds.). Lawrence Erlbaum Associates, Hillsdale, New Jersey, USA, Chapter 3, 31–61.
- Donald A. Norman. 2002. *The Design of Everyday Things*. Basic Books, New York, New York, USA. 272 pages.
- David Noy. 2001. Predicting User Intentions in Graphical User Interfaces Using Implicit Disambiguation. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems (CHI EA '01)*. ACM, Seattle, Washington, USA, 455–456. doi:<http://dx.doi.org/10.1145/634067.634330>
- Scott Olitsky, Henry Lee, and Edward Young. 2006. IVAC—Interactive Visual Acuity Chart. (2006). <http://www.smb.su.buffalo.edu/oph/ped/IVAC/IVAC.html>
- Michaël Ortega. 2013. Hook: Heuristics for Selecting 3D Moving Objects in Dense Target Environments. In *Proceedings of the IEEE 8th Symposium on 3D User Interfaces (3DUI 2013)*. IEEE, Orlando, Florida, USA, 119–122. doi:<http://dx.doi.org/10.1109/3DUI.2013.6550208>
- Ryan A. Pavlik and Judy M. Vance. 2012. VR JuggLua: A framework for VR applications combining Lua, OpenSceneGraph, and VR Juggler. In *2012 5th Workshop on Software Engineering and Architectures for Realtime Interactive Systems (SEARIS)*. IEEE, Singapore, 29–35. doi:<http://dx.doi.org/10.1109/SEARIS.2012.6231166>

- Andriy Pavlovych and Carl Gutwin. 2012. Assessing Target Acquisition and Tracking Performance for Complex Moving Targets in the Presence of Latency and Jitter. In *Proceedings of Graphics Interface 2012 (GI' 12)*. Canadian Information Processing Society, Toronto, Ontario, Canada, 109–116. <http://dl.acm.org/citation.cfm?id=2305276.2305295>
- James G. Phillips and Thomas J. Triggs. 2001. Characteristics of cursor trajectories controlled by the computer mouse. *Ergonomics* 44, 5 (2001), 527–536. doi:<http://dx.doi.org/10.1080/00140130121560>
- Sébastien Pierard, Vincent Pierlot, Antoine Lejeune, and Marc Van Droogenbroeck. 2012. I-see-3D! An Interactive and Immersive System that dynamically adapts 2D projections to the location of a user's eyes. In *International Conference on 3D Imaging (IC3D)*. Liège, Belgium. <http://orbi.ulg.ac.be/handle/2268/134431>
- Ivan Poupyrev, Suzanne Weghorst, Mark Billinghurst, and Tadao Ichikawa. 1997. A framework and testbed for studying manipulation techniques for immersive VR. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '97)*. ACM, Lausanne, Switzerland, 21–28. doi:<http://dx.doi.org/10.1145/261135.261141>
- J. R. Quinlan. 1986. Induction of Decision Trees. *Machine Learning* 1, 1 (1986), 81–106. doi:<http://dx.doi.org/10.1023/A:1022643204877>
- John Ross Quinlan. 1993. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, California, USA. 302 pages. <http://books.google.com/books?hl=en&lr=&id=HExncpbYroC&pgis=1>
- R Core Team. 2014. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.r-project.org/>
- R Core Team. 2015. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.r-project.org/>
- Martin Reddy. 1994. Reducing lags in virtual reality systems using motion-sensitive level of detail. In *Proceedings of the 2nd UK VR-SIG Conference*, Robin Hollands (Ed.). Theale, Berkshire, UK, 25–31. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.50.15&rep=rep1&type=pdf#page=27>
- Jaime Ruiz and Edward Lank. 2014. Analyzing Intended Use Effects in Target Acquisition. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces (AVI '14)*. ACM, Como, Italy, 145–152. doi:<http://dx.doi.org/10.1145/2598153.2598158>
- Reza Shadmehr, Maurice A. Smith, and John W. Krakauer. 2010. Error correction, sensory prediction, and adaptation in motor control. *Annual Reviews of Neuroscience* 33, 1 (2010), 89–108. doi:<http://dx.doi.org/10.1146/annurev-neuro-060909-153135>

- Ilyas H. Sheikh and Errol R. Hoffmann. 1994. Effect of target shape on movement time in a Fitts task. *Ergonomics* 37, 9 (1994), 1533–1547. doi:<http://dx.doi.org/10.1080/00140139408964932>
- Garth Shoemaker, Takayuki Tsukitani, Yoshifumi Kitamura, and Kellogg S. Booth. 2012. Two-Part Models Capture the Impact of Gain on Pointing Performance. *ACM Transactions on Computer-Human Interaction* 19, 4 (Dec. 2012), 28:1–28:34. doi:<http://dx.doi.org/10.1145/2395131.2395135>
- João Silva, Diogo Cabral, Carla Fernandes, and Nuno Correia. 2012. Real-time Annotation of Video Objects on Tablet Computers. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia (MUM '12)*. ACM, Ulm, Germany, 19:1–19:9. doi:<http://dx.doi.org/10.1145/2406367.2406391>
- Andrew B. Slifkin and Suzanne M. Grilli. 2006. Aiming for the future: prospective action difficulty, prescribed difficulty, and Fitts' law. *Experimental Brain Research* 174, 4 (Oct. 2006), 746–753. doi:<http://dx.doi.org/10.1007/s00221-006-0518-3>
- Dan Song, Nikolaos Kyriazis, Iason Oikonomidis, Chavdar Papazov, Antonis Argyros, Darius Burschka, and Danica Kragic. 2013. Predicting human intention in visual observations of hand/object interactions. In *Proceedings of the 2013 IEEE International Conference on Robotics and Automation (ICRA 2013)*. IEEE, Karlsruhe, Germany, 1608–1615. doi:<http://dx.doi.org/10.1109/ICRA.2013.6630785>
- R. William Soukoreff and I. Scott MacKenzie. 2004. Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI. *International Journal of Human-Computer Studies* 61, 6 (Dec. 2004), 751–789. doi:<http://dx.doi.org/10.1016/j.ijhcs.2004.09.001>
- Andrej-Nikolai Spiess and Natalie Neumeyer. 2010. An evaluation of R2 as an inadequate measure for nonlinear models in pharmacological and biochemical research: a Monte Carlo approach. *BMC Pharmacology* 10, 1 (2010), 6. doi:<http://dx.doi.org/10.1186/1471-2210-10-6>
- Martin Spindler, Wolfgang Büschel, and Raimund Dachsel. 2012. Use Your Head: Tangible Windows for 3D Information Spaces in a Tabletop Environment. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces (ITS '12)*. ACM, Cambridge, Massachusetts, USA, 245–254. doi:<http://dx.doi.org/10.1145/2396636.2396674>
- Stan Development Team. 2014. Stan: A C++ Library for Probability and Sampling, Version 2.4.0. (2014). <http://mc-stan.org/>
- Stan Development Team. 2015. Stan: A C++ Library for Probability and Sampling, Version 2.6.0. (2015). <http://mc-stan.org/>

- Rainer Stiefelhagen. 2002. Tracking focus of attention in meetings. In *Fourth IEEE International Conference on Multimodal Interfaces (ICMI '02)*. IEEE Comput. Soc, Pittsburgh, Pennsylvania, USA, 273–280. doi:<http://dx.doi.org/10.1109/ICMI.2002.1167006>
- Rainer Stiefelhagen and Jie Zhu. 2002. Head Orientation and Gaze Direction in Meetings. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems (CHI EA '02)*. ACM, Minneapolis, Minnesota, USA, 858–859. doi:<http://dx.doi.org/10.1145/506443.506634>
- Robert J. Teather and Wolfgang Stuerzlinger. 2011. Pointing at 3D Targets in a Stereo Head-Tracked Virtual Environment. In *Proceedings of the 2011 IEEE Symposium on 3D User Interfaces (3DUI 2011)*. IEEE, Singapore, 87–94. doi:<http://dx.doi.org/10.1109/3DUI.2011.5759222>
- James R. Tresilian. 2005. Hitting a moving target: Perception and action in the timing of rapid interceptions. *Perception & Psychophysics* 67, 1 (2005), 129–149. doi:<http://dx.doi.org/10.3758/BF03195017>
- John W. Tukey. 1980. We Need Both Exploratory and Confirmatory. *The American Statistician* 34, 1 (1980), pp. 23–25. doi:<http://dx.doi.org/10.1080/00031305.1980.10482706>
- Paul Van Schaik and Jonathan Ling. 2007. Design Parameters of Rating Scales for Web Sites. *ACM Transactions on Computer-Human Interaction* 14, 1 (May 2007), 35 pages. doi:<http://dx.doi.org/10.1145/1229855.1229859>
- W3C. 2011. Scalable Vector Graphics (SVG). (2011). <http://www.w3.org/TR/2011/REC-SVG11-20110816/>
- A. T. Welford, A. H. Norris, and N. W. Shock. 1969. Speed and accuracy of movement and their changes with age. *Acta Psychologica* 30, 0 (1969), 3–15. doi:[http://dx.doi.org/10.1016/0001-6918\(69\)90034-1](http://dx.doi.org/10.1016/0001-6918(69)90034-1)
- Thomas G. Whisenand and Henry H. Emurian. 1995. Some Effects of Angle of Approach on Icon Selection. In *Conference Companion on Human Factors in Computing Systems (CHI '95)*. ACM, Denver, Colorado, USA, 298–299. doi:<http://dx.doi.org/10.1145/223355.223683>
- Thomas G. Whisenand and Henry H. Emurian. 1996. Effects of angle of approach on cursor movement with a mouse: Consideration of Fitt's law. *Computers in Human Behavior* 12, 3 (1996), 481–495. doi:[http://dx.doi.org/10.1016/0747-5632\(96\)00020-9](http://dx.doi.org/10.1016/0747-5632(96)00020-9)
- T. G. Whisenand and H. H. Emurian. 1999. Analysis of cursor movements with a mouse. *Computers in Human Behavior* 15, 1 (1999), 85–103. doi:[http://dx.doi.org/10.1016/S0747-5632\(98\)00036-3](http://dx.doi.org/10.1016/S0747-5632(98)00036-3)
- Hugh R. Wilson, Frances Wilkinson, Li-Ming Lin, and Maja Castillo. 2000. Perception of head orientation. *Vision Research* 40, 5 (Jan. 2000), 459–72. doi:[http://dx.doi.org/10.1016/S0042-6989\(99\)00195-9](http://dx.doi.org/10.1016/S0042-6989(99)00195-9)

- Jonathan Wonner, Jérôme Grosjean, Antonio Capobianco, and Dominique Bechmann. 2011. SPEED : Prédiction de cibles. In *23rd French Speaking Conference on Human-Computer Interaction (IHM '11)*. ACM, Sophia Antipolis, France, 19:1–19:4. doi:<http://dx.doi.org/10.1145/2044354.2044378>
- Chuang-Wen You, Yung-Huan Hsieh, and Wen-Huang Cheng. 2012. AttachedShock: Facilitating Moving Targets Acquisition on Augmented Reality Devices Using Goal-crossing Actions. In *Proceedings of the 20th ACM International Conference on Multimedia (MM '12)*. ACM, Nara, Japan, 1141–1144. doi:<http://dx.doi.org/10.1145/2393347.2396403>
- Chuang-Wen You, Yung-Huan Hsieh, Wen-Huang Cheng, and Yi-Hsuan Hsieh. 2014. AttachedShock: Design of a crossing-based target selection technique on augmented reality devices and its implications. *International Journal of Human-Computer Studies* 72, 7 (2014), 606–626. doi:<http://dx.doi.org/10.1016/j.ijhcs.2014.03.001>
- Shumin Zhai, Carlos Morimoto, and Steven Ihde. 1999. Manual and Gaze Input Cascaded (MAGIC) Pointing. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI '99)*. ACM, Pittsburgh, Pennsylvania, USA, 246–253. doi:<http://dx.doi.org/10.1145/302979.303053>
- Xinyong Zhang, Hongbin Zha, and Wenxin Feng. 2012. Extending Fitts' Law to Account for the Effects of Movement Direction on 2D Pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 3185–3194. doi:<http://dx.doi.org/10.1145/2207676.2208737>