

Modeling Steering within Above-the-Surface Interaction Layers

Raghavendra S Kattinakere¹, Tovi Grossman², Sriram Subramanian¹

¹Department of Computer Science
University of Saskatchewan
Saskatoon, Canada

ragu.ks@usask.ca, sriram@cs.usask.ca

²Department of Computer Science
University of Toronto
Toronto, Canada
tovi@dgp.toronto.edu

ABSTRACT

Interaction techniques that utilize the space above the display surface to extend the functionalities of digitized surfaces continue to emerge. In such techniques, movements are constrained by the bounds of a layer. In addition, constraints imposed on the direction of movement within the layer may be present. Despite the presence of such techniques, there is limited understanding of human capabilities for performing the required steering task. In this paper we study and model user performance when steering through constrained and unconstrained paths in above-the-surface layers. Through a series of experiments we validate the derivation and applicability of our proposed models.

Author Keywords

3D steering, digital table, 3D, command layers.

ACM Classification Keywords

H.5.2 [User Interfaces]: Interaction styles

INTRODUCTION

Pen-based systems offer a number of advantages over the traditional mouse-based desktop metaphor, such as allowing for fluid input, and the direct manipulation of underlying data. With the rapid development of display technology, such systems now come in many forms. HCI researchers are challenged to design interfaces for a variety of pen-based systems such as PDAs [7], Tablet PCs [6, 10, 14, 15], tabletops [21], and large displays [12, 18].

Many such systems would also be termed *digitized surfaces*, meaning they are able to sense the location of the input device even when it is above the display surface, in the *tracking state*. Recent research has investigated how this tracking state can be used in the design of pen-based interaction techniques [8, 10, 11, 14, 21]. For example, Hover Widgets [11], allow users to invoke localized widgets by making gestures in the tracking state (Figure 1, left).

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

CHI 2007, April 28–May 3, 2007, San Jose, California, USA.

Copyright 2007 ACM 978-1-59593-593-9/07/0004...\$5.00.

We term an interactive area above the display surface as an *above-the-surface interaction layer*. In most systems, there will only be one such layer, the tracking state of the input device. An exception is multi-layer interaction techniques [21], which divide the space above a table into multiple discrete interaction layers (Figure 1, right).

While above-the-surface interaction layers increase the functionality of pen-based systems, they require users to steer the tip of the stylus through a constrained tunnel. For example, to move within the tracking state, the stylus must remain above the display surface without touching it, and not extend too far from the display surface so that it is out of sensing range. Currently, our understanding of the human abilities to perform such a steering task is based on the steering law proposed by Accot and Zhai [1]. However, this model applies to a 2D desktop environment, and it is not clear if and how this model can be applied to user movements in above-the-surface interaction layers.

In this paper we study and model user performance when steering through above-the-surface interaction layers, under varying levels of directional constraints. We begin with a review of related work, and a description of Accot and Zhai's original steering model. We then discuss how this model can be extended and applied to steering through above-the-surface layers, with and without the presence of a directional constraint. Through a series of four formal experiments we validate the derivation and applicability of these models. We conclude with a discussion of the implications to design and possible lines of future work.

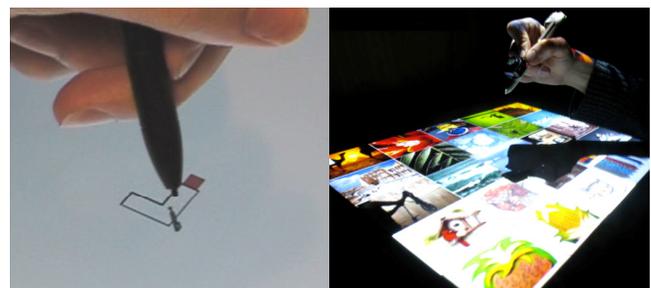


Figure 1. Above-the-surface interaction. Left: A Hover Widget [11] is used by making a gesture with the pen in the tracking state. Right: Multi-layer interaction [21] divides the space above a tabletop display into multiple interaction layers.

RELATED WORK

We divide our review of related work into two sections. We first discuss the interaction techniques which have been implemented for above-the-surface interaction. We then discuss previous efforts at modeling human performance, in both pointing and steering tasks.

Above-the-surface Interaction

Many interaction techniques have leveraged the tracking state of digitized pen-based systems to provide added functionality. These techniques all provide actual examples of steering within above-the-surface interaction layers.

Most of these techniques require the user to perform pen movements within the layer without any directional constraints. The vacuum [8] is a technique that supports reaching distant objects in a large display with a pen. By moving the pen within the tracking state, the vacuum widget can be used for interaction with multiple objects. Exiting the layer into the out-of-range zone dismisses the Vacuum. Stitching uses pen gestures that span multiple displays to seamlessly connect displays for co-located collaboration [14]. Users stitch displays by moving the pen across displays while keeping the pen in the tracking state. Tracking menus [10], are menu widgets that stay under the pen cursor and close at hand. The user can reposition the menu by dragging past its edges while in the tracking state. Multi-layer interaction techniques [21] divide an enlarged tracking state into multiple interaction layers. Users can navigate within individual layers to access different tools, and perform various commands.

Hover Widgets [11] are an example of a technique which requires users to not only steer within the tracking state, but to do so under imposed directional constraints. Gestures, defined by tunnels, are made in the tracking state to quickly access localized interface elements.

In summary, there are a number of techniques which are making use of above-the-surface interaction layers, and require users to steer within them. Despite this, there is little understanding as to human capabilities when performing such a task. We now discuss related modeling techniques which will help us obtain a better understanding.

Modeling Pointing and Steering

The derivation of Accot and Zhai's steering law is based on Fitts' law which models pointing [9]. The law states that the pointing performance is limited by the capacity of the human motor system. The commonly used form of Fitts' law [16] predicts the movement time MT to select a target of width W at a distance of A as follows:

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

where a and b are empirically determined constants. The logarithmic term is called the *index of difficulty* (ID).

Mackenzie and Buxton [17] extended Fitts' original model for a 2D target acquisition task. They examined several formulas for the index of difficulty for a rectangular target of width W and height H , and found that the *min* model, which only considers the smaller of the two dimensions, to have the highest correlation with their experimental data. Addressing some of the difficulties associated with the *min* model, Accot and Zhai later refine this model into a weighted *Euclidian* model, expressed by:

$$MT = a + b \log_2 \left(\sqrt{\left(\frac{A}{W}\right)^2 + \eta \left(\frac{A}{H}\right)^2} + 1 \right) \quad (2)$$

where η is empirically determined [4]. The addition of the parameter η allows the model to weight the effect of the height differently from the effect of the width.

Another task prominent in graphical user interfaces is steering, or tunneling, which can be described as the task of moving through a constrained path, such as when a user navigates through a hierarchical cascading menu [5]. To derive a model for this task, Accot and Zhai [1] first consider a goal crossing task, where a user must travel a distance A , and then cross a goal with length W . They found this task can be accurately modeled with Fitts' Law (Equation 1).

Based on this, they derived a model, which predicts the time necessary to steer through a path. In the most basic case, the path has a constant width, W , and length, A . The law is derived by considering the task to be an infinite series of goal crossing tasks. The resulting model reduces to:

$$MT = a + b \left(\frac{A}{W} \right). \quad (3)$$

In the more complex task, the path width varies along its length. In this case, the task can be modeled by the equation:

$$MT = a + b \int_C \frac{ds}{W(s)} \quad (4)$$

where C is the path and $W(s)$ is the width of the path at the point s [1]. Follow-up studies have investigated the effects of scale [2] and sharp corners [19] on the steering task.

A notable contribution of Accot and Zhai's original steering law work, is that it has been used in follow up work to model actual user interface tasks [5, 13]. The original work has also inspired new interaction techniques [3, 6]. It is our goal to extend these discussed models, such that they can be applied to the task of steering within an above-the-surface interaction layer. In the next section, we present our newly proposed models, which we then validate in a series of four formal experiments.

3D STEERING MODEL

The first task which we wish to model is steering within an above-the-surface layer, where the only constraint is the size of the layer. Examples of this scenario are seen in interaction techniques where the input device must stay within an above-the-surface layer while traveling from one point to another [8, 14]. We define the size of the layer, or *thickness*, T , as the distance between the bottom and top planes which define the layer. Figure 2a illustrates the task.

To derive a model for this task, we first consider the analogous goal crossing task, in which the goal is defined as the plane perpendicular to the display surface, extending from the bottom to the top boundaries of the layer (Figure 2b). We hypothesize that this task will be modeled by Accot and Zhai's original goal crossing formulation, where we simply replace the width variable with our thickness variable T . This gives the equation:

$$MT = a + b \log_2 \left(\frac{A}{T} + 1 \right). \quad (5)$$

Assuming this model does accurately model the goal crossing task depicted in Figure 2b, then it is reasonable to assume that the tunneling task depicted in Figure 2a, can be derived using Accot and Zhai's method, which would give the following model:

$$MT = a + b \left(\frac{A}{T} \right). \quad (6)$$

A more complex scenario is steering through a tunnel within a layer, where the movement is constrained not only by the layer thickness, T , but also by a path which is imposing a directional constraint, W . An example of this scenario is seen when the user activates a Hover Widget [11], as the input device must make a specific gesture, defined by a tunnel, in the display's tracking state. In this case, the size of the tracking state defines the layer thickness, and the width of the tunnel defines the directional constraint. Such a task is illustrated in Figure 3a.

To derive a model for this scenario we again consider the analogous goal crossing task, in which the goal is defined by a rectangle, perpendicular to the display surface, with a height T , and width W (Figure 3b). It is interesting that we now have a bivariate, or two-dimensional, goal crossing task. Such a task has never been studied before, but we hypothesize that it can be modeled as a bivariate pointing task. We use the Euclidian model for bivariate pointing (Equation 2) to derive the following model:

$$MT = a + b \log_2 \left(\sqrt{\left(\frac{A}{W} \right)^2 + \eta \left(\frac{A}{T} \right)^2} + 1 \right). \quad (7)$$

Under the assumption that this model can accurately predict movement times for the 2D goal crossing task, depicted in Figure 3b, we can use it to derive a model for the 2D

tunneling task depicted in Figure 3a. We again use the same methodology presented by Accot and Zhai [1]. We first break the tunneling task into a series of N goal crossing tasks (Figure 4). By summing the associated ID values for each of these N tasks, which have distances of A/N , we get:

$$ID_N = N \log_2 \left(\sqrt{\left(\frac{A}{NW} \right)^2 + \eta \left(\frac{A}{NT} \right)^2} + 1 \right) \quad (8)$$

As N approaches infinity, the task becomes the desired 2D tunneling task. To obtain the index of difficulty for this tunneling task, we take the limit of ID_N as N approaches infinity. Using a first order Taylor series expansion of $\log_2(1+x)$, we obtain:

$$ID = \lim(N \rightarrow \infty) ID_N = \frac{1}{\ln(2)} \sqrt{\left(\frac{A}{W} \right)^2 + \eta \left(\frac{A}{T} \right)^2} \quad (9)$$

These leads to our model for the movement time as:

$$MT = a + b \sqrt{\left(\frac{A}{W} \right)^2 + \eta \left(\frac{A}{T} \right)^2} \quad (10)$$

In the following sections we describe a series of experiments, with the goal of understanding human capabilities when steering within above-the-surface interaction layers. In doing so we also hope to validate the proposed models presented here.

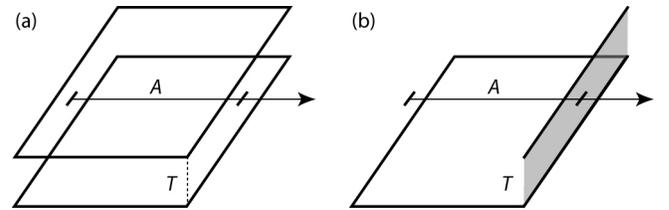


Figure 2. (a) Steering through a layer, constrained only by the layer thickness, T . (b) The analogous goal crossing task.

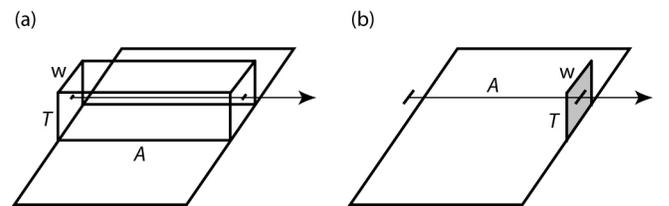


Figure 3. (a) Steering through a layer, constrained by the layer thickness, T , and the directional constraint of the path, W . (b) The analogous goal crossing task.

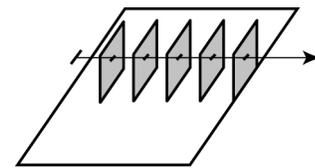


Figure 4. A series of goal crossing tasks.

APPARATUS

All experiments which are described below were performed on the same apparatus. The experiments ran on an Intel Pentium(R) 4 CPU 3.20 Ghz PC with 1 GB RAM. Participants sat at a 124.5cm x 158cm tabletop surface. A 1024x768 pixel image was projected onto the surface using a ceiling mounted projector. A stylus was used for input, and was tracked using a Polhemus Liberty Motion Tracker. The motion tracking system provided positional data at a rate of 240 Hz. The pen was calibrated to report X and Y values in pixels on the tabletop surface, and Z values in centimeters, corresponded to the height above the tabletop surface. All users were seated comfortably and controlled the pen with their dominant hand. Users were allowed to rest their hand while completing the task, much like the hand rests while writing with a pen. The pen controlled the cursor position using a direct 1 to 1 mapping. Figure 5 illustrates the apparatus. In all experiments the bottom bound of the layer was 0.2cm above the display surface.

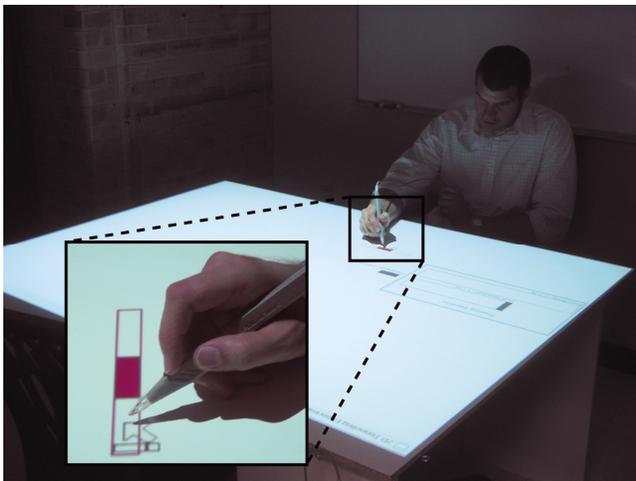


Figure 5. Apparatus used for the experiments.

LAYER VISUALIZATION

For users to be able to efficiently steer through an above-the-surface interaction layer, they must be aware of the position of their stylus within the layer, as well as the thickness of the layer. This is especially important for our experimental procedure, as the layer thickness will be an independent variable, and thus changing from one trial to the next. We use a cursor visualization to provide this information to the user, similar to how pressure widgets can be used to convey pressure information [20]. Through informal usage observations and iterative design we converged on the following visualization.

The visualization consists of a 5x20 pixel rectangle, with the top and bottom of this rectangle representing the top and bottom boundaries of the current layer respectively. Within this rectangle, a small 5 pixel horizontal line originates from the base of the cursor extending to the right (Figure 6a). This line represents the current height of the input device within the layer. The position of the line never

changes in relation to the position of the cursor. It is the position of the rectangle that changes with the height of the stylus, such that the position of the line relative to the rectangle always indicates the position of the stylus in relation to the top and bottom boundaries of the layer. Therefore, moving the pen down will move the rectangle upwards (Figure 6b), and moving the pen up will move the rectangle downwards (Figure 6d). If the cursor leaves the lower or upper bounds of the layer, then the rectangle will jump to be above or below the line accordingly (Figure 6c, e), and the rectangle will be filled red. This was done because in our experiment, an error state is entered if the stylus leaves the layer. Outside of our experimental paradigm, other design approaches could be investigated, especially if multiple layers exist that the stylus could move between, such as in the work by Subramanian et al. [21].

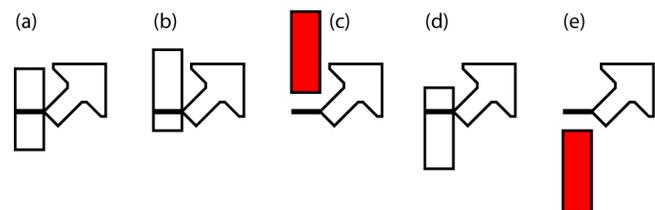


Figure 6. Cursor visualization.

EXPERIMENT 1: A Pilot Study of 1D Goal Passing

In this experiment we investigate a goal passing task, where the goal to be passed is a vertical region extending above the display surface. The experimental parameters will be the thickness of this region, or layer, and the distance between the goals which are to be passed.

We do not see this as being a task which will normally be carried out in pen-based interfaces. The reason for the experiment is for theoretical purposes, as the derivation of our model of the more applicable tunneling task is based on the model for this task. As such, we only ran two participants through this experiment, enough to ensure that the movement times will follow our proposed model (Equation 5).

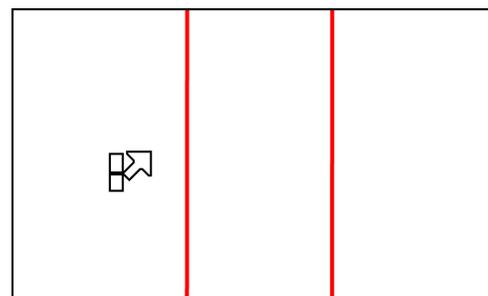


Figure 7. The 1D goal crossing task used in Experiment 1.

Participants

Two volunteers (both male), aged 18 and 19 participated in the experiment. Both participants were right handed and controlled the stylus with their right hand. None of the subjects had previous experience with using large digital tables. Both subjects were tested individually.

Procedure

The goal passing task was accomplished by passing a start and end goal from left to right. Each goal was depicted as a vertical red line spanning the display, separated by a distance of A (Figure 7).

To begin a trial, participants had to position the stylus within 10 pixels to the left of the start goal, and above the surface, such that it was within the bounds of the current trial layer. Once this was done, participants had to dwell for 0.6s and then click a button on the stylus. At this point the color of the goals would turn green indicating that the participant could proceed to crossing the goals. These starting constraints were added to control the initial velocity of the pen when the trial began, and to prevent users from going from one trial to the next without regard to accuracy.

When the first goal was crossed it would change color and when the second goal was crossed the trial ended. Both goals had to be successfully crossed from left to right in the correct order for the trial to be completed. An audio cue was given each time a goal was successfully crossed.

Because this was a goal crossing task, the stylus only had to be within the layer bounds when the goals were crossed. If the stylus was not in the layer when a goal was crossed, the participant would have to back-track and repeat the crossing for that goal. If this happened with the end goal, the trial would be counted as an error. The total error rate was displayed during the experiment, and participants were told to balance speed and accuracy such that their error rate remained at approximately 4%.

Design

A repeated measures within-participant design was used. The independent variables were layer thickness, T (1, 1.5, 2, and 2.5 centimeters), and the distance between the goals, or amplitude, A (5, 15, 25, and 35 centimeters). This design resulted in ID values ranging from 1.58 to 5.17 as determined by Equation 5. A fully crossed design resulted in 16 combinations of T and A .

The experiment was divided into 4 blocks. The blocks were ordered by thickness, with all trials for one thickness being completed before moving on to the next. This was done to prevent confusion of constantly changing layer thicknesses.

Within each block and for each thickness, trials for each of the 4 lengths were presented 9 times in random order, resulting in a total of 576 trials. The ordering of layer thickness was balanced by being reversed for the second participant. Before the first block, a practice session was given, consisting of each of the 16 conditions presented in random order.

Results

Movement time, MT , was the main measure for the experiment, defined as the time between crossing the start and end goals. In order to get a precise measurement of the time when the pen crossed each goal, we linearly

interpolated between the last event when the pen was reported to be to the left of the goal, and the first event when the pen was reported to be to the right of the goal. In our analysis of movement time we removed trials in which errors occurred. We also removed outliers more than 3 standard deviations from the group mean movement time, 2.5% of the data.

Repeated measures analysis of variance showed a main effect for T ($F_{3,3} = 29$, $p < .0001$) and A ($F_{3,3} = 173$, $p < .0001$). Movement times for each thickness were 1.2s for $T = 1$ cm, 0.82s for $T = 1.5$ cm, 0.81s for $T = 2$ cm, and 0.67s for $T = 2.5$. This confirms our belief that movement times will be constrained by the layer thickness.

Figure 8 plots the movement times by the index of difficulty, defined by Equation 5. Linear regression analysis showed that the data fit to the model with an R^2 value of 0.83. The equation for MT is given by:

$$MT = -852 + 503 \log\left(\frac{A}{T} + 1\right) \quad (11)$$

The R^2 value is somewhat lower than desired, but it is reasonable to expect that with more participants, the movement times would continue to conform to our model, with a higher fit.

The overall error rate for the experiment was 2.4%, which is slightly lower than the desired 4% error rate. The condition that seemed to have the most effect on error rate was with $A = 35$ cm, where the error rate was 7.3%. This gives more explanation as to why our model did not have a higher fit to the data. Indeed if we remove this condition from the data, the R^2 value increases to 0.92.

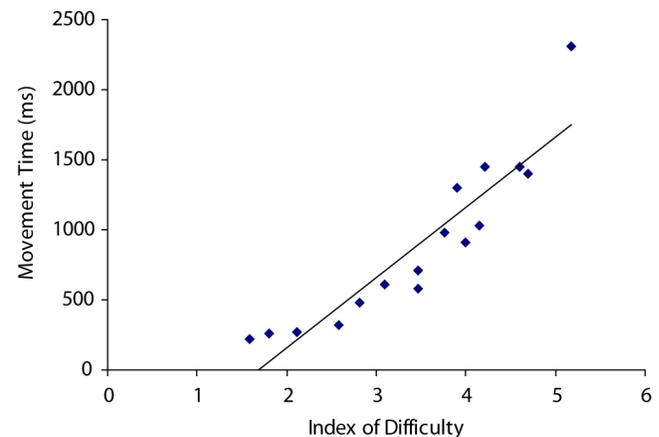


Figure 8. Movement times by ID for the 1D goal passing task.

Overall, the data provides the necessary confirmation that a goal crossing task for which the goal is constrained by a layer thickness can be modeled using Fitts' Law. This validates the derivation of our model for a tunneling task that is constrained by a layer thickness (Equation 6). We validate this model in the following experiment.

EXPERIMENT 2: 1D Tunneling

In the previous experiment we found that the layer thickness will affect movement times in a goal crossing task as we would predict from Fitts' Law. While we may not expect to see the experimental task in actual applications, it was necessary to validate our derivations of the other models which we will be testing in this work.

In this experiment, we focus on the task of steering within an above-the-surface interaction layer, without the presence of directional constraints. An example of such a task is seen when a user must move from one point to another while staying in the tracking state to maintain a mode [8, 14]. Along with investigating human capabilities when performing this task, we will also test the ability of our proposed model (Equation 6) to predict movement times.

Participants

Twelve volunteers (8 male, 4 female), aged 21 to 35 participated in the experiment. Participants were right handed and controlled the stylus with their right hand. Four subjects had previous experience with using large digital tables. All participants were tested individually.

Procedure

The general procedure was the same as in the previous experiment. In this case the task was to steer through a tunnel of thickness T , over a distance A . The task was again accomplished by passing a start and end goal from left to right, however in this case the stylus had to remain within the layer during the entire trial. The tunnel area was depicted as a solid red rectangle, spanning the height of the display. The left edge of the rectangle was the start goal, and the right edge of the rectangle was the end goal. The goals were again centered with the participants seating position. Figure 9 illustrates the task.

The procedure to begin a trial was the same as in Experiment 1. When the trial could begin, the rectangle would turn green, and when the first goal was crossed the tunnel turned orange. As in experiment 1, both goals had to be successfully crossed from left to right in the correct order for the trial to be completed.

Because this was now a tunneling task, the stylus had to be within the layer bounds through the entire trial. If the stylus left the layer at any time once the trial began then the trial would be counted as an error.

Design

The design was the same as in Experiment 1. A repeated measures within-participant design was used. The independent variables were layer thickness, T (1, 1.5, 2, and 2.5 centimeters), and the distance between the goals, or amplitude, A (5, 15, 25, and 35 centimeters). This design resulted in ID values ranging from 2 to 35 as determined by Equation 6. A fully crossed design resulted in 16 combinations of T and A .

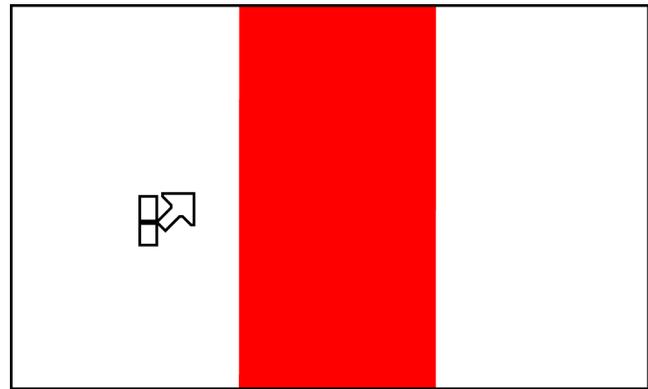


Figure 9. The 1D tunneling task used for Experiment 2.

Four blocks were again ordered by thickness, with all trials for one thickness being completed before moving on to the next. The ordering of layer thickness was counterbalanced between participants using a 4x4 balanced Latin Square design. Before the first block, a practice session was given, consisting of each of the 16 conditions.

Results

Movement time, MT , was again the main dependent measure and had the same definition as in Experiment 1. We again removed trials in which errors occurred, as well as outliers which were more than 3 standard deviations from the group mean movement time, 1.6% of the data.

Repeated measures analysis of variance showed a main effect for T ($F_{3,33} = 243$, $p < .0001$) and A ($F_{3,33} = 2477$, $p < .0001$), and a significant $T \times A$ interaction ($F_{9,99} = 39.7$, $p < .0001$). Movement times for each thickness were 0.61 for $T = 1$ cm, 0.48s for $T = 1.5$ cm, 0.42s for $T = 2$ cm, and 0.41s for $T = 2.5$. Post hoc analysis shows that all pairs are significantly different except for $T = 2$ cm and $T = 2.5$ cm. Figure 10 illustrates the interaction between T and A . It can be seen that the effect of T becomes stronger for higher values of A . Even so, there is no difference between $T = 2$ and $T = 2.5$ even for the largest distance. This indicates that when travel distances are expected to be in the range tested in our experiment, a layer thickness within the 2 to 2.5cm range would be appropriate.

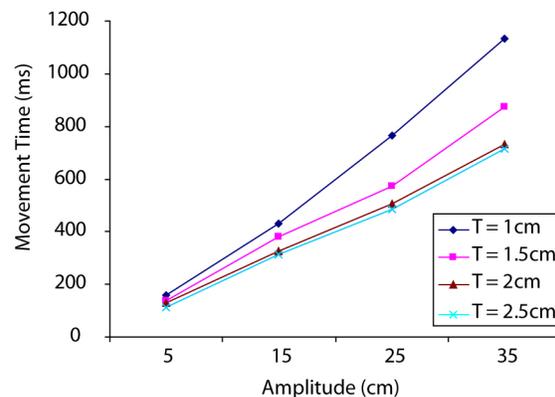


Figure 10. Movement times by amplitude and layer thickness.

Figure 11 plots the movement times by the index of difficulty, defined by Equation 6. Linear regression analysis showed that the data fit to the model with an R^2 value of 0.92. The movement time MT , is given by the equation:

$$MT = 84.7 + 31.25 \left(\frac{A}{T} \right) \quad (12)$$

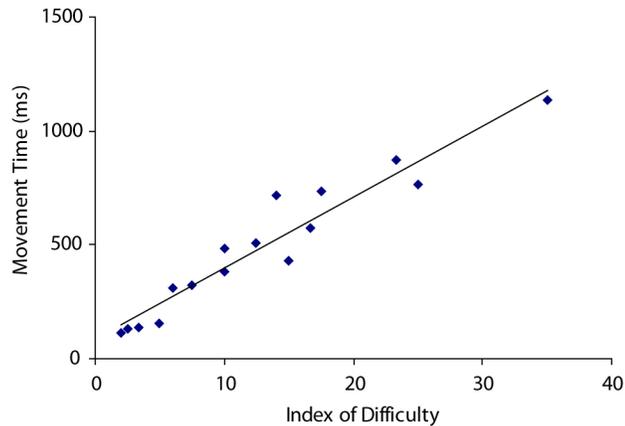


Figure 11. Movement times by ID for the 1D steering task.

The overall error rate for the experiment was only 1%. While this is lower than the desired 4% level, it does seem to indicate that our cursor visualization allows participants to steer within the layer boundaries. Error rates were slightly higher for larger values of A and smaller values of T , but remained under 3% across all conditions except for $A = 35\text{cm}$, $T = 1\text{cm}$, for which the error rate was 8.1%.

The results of Experiments 1 and 2 have provided useful information about steering within above-the-surface interaction layers when the only constraint is the thickness of the layer. In the following experiments we investigate what happens when there are also directional constraints imposed on the movements.

EXPERIMENT 3: A Pilot Study of 2D Goal Passing

In the previous sections we have investigated a tunneling task when movements are constrained by the layer thickness. We validated that our model, which is based on Accot and Zhai's steering law [1], can be used to predict movement times. We now turn our focus to the task of steering within a layer along a path which imposes a directional constraint. Such a task is seen in previously developed interaction techniques, such as Hover Widgets, where users make specific gestures defined by tunnel boundaries in the tracking state of a pen-based system [11].

Before investigating this specific task, we will first look at the constrained version of the goal crossing task which was used in Experiment 1. It is again imperative to do this as the model for this task is used to derive the model for the tunneling task. In validating our proposed model for the constrained goal crossing task, we will also be validating our derivation of the model proposed for the more practical

constrained tunneling task. As with Experiment 1, this task is theoretical in nature, so we again only run two participants through the experiment.

Participants

Two volunteers (both males) aged 19 participated in the experiment. Both participants were right handed and controlled the stylus with their right hand. Neither of the subjects had previous experience with using large digital tables. However, both of them had used a Tablet PC before. Both subjects were tested individually.

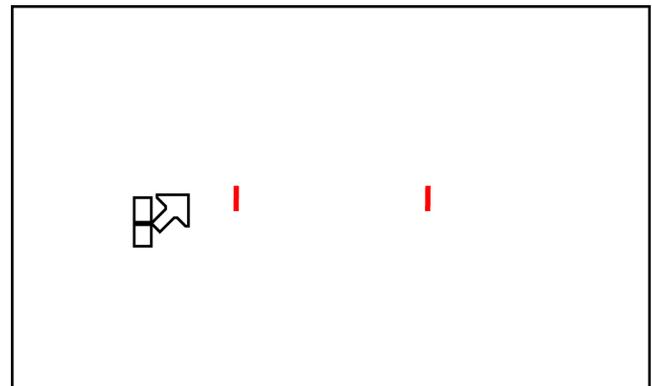


Figure 12. The 2D goal passing task used for Experiment 3.

Procedure

The general procedure for this experiment was the same as in Experiment 1. However, in this case the goals had a finite width, W , as they were used to impose the directional constraint. Figure 12 illustrates the task. As with Experiment 1, users had to successfully cross both goals from left to right to complete the task. However, in this experiment, a successful cross required the pen to be within the layer bounds, and within the extents of the goal. As with the previous experiments, if the cross was unsuccessful, then the participant would have to back-track and repeat the crossing for that goal.

Design

A repeated measures within-participant design was used. The independent variables were layer thickness, T (1, 1.5, 2, and 2.5 centimeters), goal width, W (1, 1.5, 2, and 2.5 centimeters), and the distance between the goals, or amplitude, A (5, 15, 25, and 35 centimeters). The resulting range of ID values, as calculated by Equation 7, will depend on the value of η , which will be determined by the results obtained in this experiment. A fully crossed design resulted in 64 combinations of T , W , and A .

The experiment was divided into 3 blocks, each ordered by thickness. Within each block and for each thickness, trials for each of the 16 W and A combinations were presented 5 times in random order, resulting in a total of 960 trials. The ordering of layer thickness was balanced by being reversed for the second participant. Before the first block, a short practice session was given.

Results

Movement time, MT , was again the main dependent measure, and we removed trials in which errors occurred, as well as outliers which were more than 3 standard deviations from the group mean movement time, 0.86% of the data. The overall error rate for the experiment was 3.3%, with higher error rates for smaller values of T and W and larger values of A .

Repeated measures analysis of variance showed a main effect for T ($F_{3,3} = 2.7$, $p < .05$), W ($F_{3,3} = 294$, $p < .0001$) and A ($F_{3,3} = 1380$, $p < .0001$). The weaker significance for T indicates that the layer thickness may not have as much impact when a task is also constrained by a tunnel width. Movement times for the four values of T were between 0.61s and 0.65s, while they ranged from 0.19s to 1.1s for the values of A and 0.48s to 0.87s for the values of W .

By a least-squares fit method, we estimated the value of η for our model in Equation 7 to be 0.002. This indicates that the impact of T is almost negligible in comparison to the effect of W . Using this value of η , linear regression analysis gives an R^2 value of 0.88. Because of the low value of η , we also tested a naïve model, which only considers A and W :

$$ID = \log\left(\frac{A}{W} + 1\right) \quad (13)$$

Interestingly enough, this model provides a higher fit to the data, with an R^2 value of 0.90, illustrated in Figure 13. The movement time MT , is given by the equation:

$$MT = -604.4 + 361 \log\left(\frac{A}{W} + 1\right) \quad (14)$$

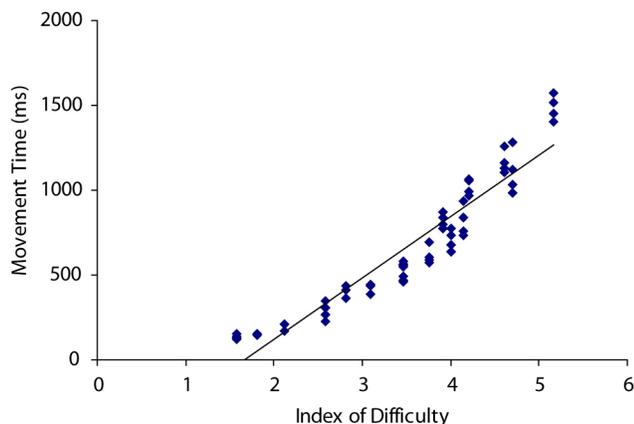


Figure 13. Movement times by ID for the 2D goal passing task.

This result indicates that during the goal crossing task, it is much easier to control the height of the pen, in comparison to staying within the boundaries defined by the directional constraints. We should take into account that this was only a 2-participant experiment, and in the next experiment we will revisit the issue. However, the result does tell us that for the 2D tunneling task, which will be presented in the next experiment, in addition to testing the model presented

in Equation 10, we should also consider the naïve model which ignores the layer thickness:

$$ID = \frac{A}{W} \quad (15)$$

EXPERIMENT 4: 2D Tunneling

In this experiment, we investigate the 2D tunneling task, where the user must steer through an above-the-surface layer, while following a specific path that imposes a directional constraint. As with the task used in Experiment 2, this is an important task to understand, as it is an element of an existing interaction technique [11]. We hope to gain an understanding of how the layer thickness, T , and tunnel width, W , affect movement time, and how these effects compare to one another. The results of Experiment 3 indicate that W will be the dominantly constraining variable. We will also test the validity of our originally proposed model for this task (Equation 10), along with the naïve form of this model (Equation 15), proposed based on our results of Experiment 3.

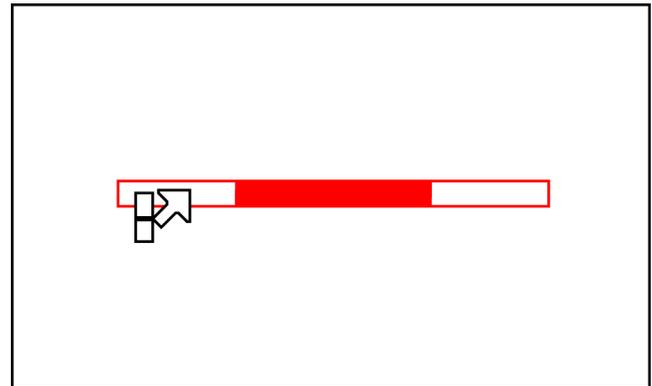


Figure 14. The 2D tunneling task used for Experiment 4.

Participants

Twelve volunteers (10 male, 2 female), aged 21 to 35 participated in the experiment. Participants were right handed and controlled the stylus with their right hand. Four subjects had previous experience with using large digital tables, and 3 of them had used a Tablet PC before. All participants were tested individually.

Procedure

The general procedure for this experiment was the same as in Experiment 2. However, in this case the tunnel had a finite width, W , imposing a directional constraint on the users movement. For extra visual feedback, the starting area was rendered on both sides of the tunnel as a rectangle. Figure 14 illustrates the task. Users had to successfully cross both goals from left to right, while staying within the bounds of the layer and tunnel. If a goal cross was unsuccessful, the participant would have to back-track and repeat the crossing for that goal. The trial was counted as an error if at any time during the trial the pen exited the bounds of the layer, or the width of the tunnel.

Design

The design was the same as in Experiment 3. A repeated measures within-participant design was used. The independent variables were layer thickness, T (1, 1.5, 2, and 2.5 centimeters), goal width, W (1, 1.5, 2, and 2.5 centimeters), and the distance between the goals, or amplitude, A (5, 15, 25, and 35 centimeters). The resulting range of ID values, as calculated by our originally proposed model (Equation 10), will depend on the value of η , which we will be determined in this experiment.

Three blocks were again ordered by thickness. Within each block and for each thickness, trials for each of the 16 W and A combinations were presented 5 times in random order, resulting in a total of 960 trials. The ordering of layer thickness was counterbalanced between participants using a 4x4 balanced Latin Square design. Before the first block, a practice session was given, consisting of 16 random trials.

Results

Movement time, MT , was again the main dependent measure, and we removed trials in which errors occurred, as well as outliers which were more than 3 standard deviations from the group mean movement time, 1.29% of the data. The overall error rate for the experiment was 2.6%, and as with the previous experiment higher error rates resulted from smaller values of T and W and larger values of A .

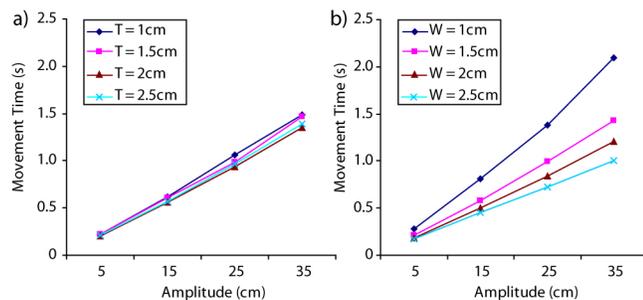


Figure 15. Interaction effects observed in Experiment 4. a) $T \times A$ interaction. b) $W \times A$ interaction.

Repeated measures analysis of variance showed a main effect for T ($F_{3,33} = 35.2$, $p < .0001$), W ($F_{3,33} = 1366$, $p < .0001$) and A ($F_{3,33} = 6544$, $p < .0001$). We also found significant $A \times T$ ($F_{9,99} = 4.48$, $p < .0001$) and $A \times W$ ($F_{9,99} = 199$, $p < .0001$) interaction effects on MT . These effects are illustrated in Figure 15. It can be seen that the effect of both T and W become stronger when A is increased. However, by comparing the two figures we again see that the effect of W on MT is much stronger than T , especially for when A is greater than 5cm. Overall movement times for the four values of T were all between 0.75s and 0.84s, while they ranged from 0.21s to 1.4s for the four values of A and 0.59s to 1.1s for the four values of W . It is also interesting to compare Figure 15a to Figure 10 from Experiment 2. We see that the presence of the directional constraint in this task drastically reduces the effects of T .

By a least-squares fit method, we estimated η for our model in Equation 10 to be 0.1638. This is larger than its value for the goal crossing task discussed in the previous experiment, indicating that our Euclidian model may be more appropriate for this task. Using this value, linear regression analysis gives a high R^2 value of 0.989 (Figure 16). The movement time, MT is given by the equation:

$$MT = 20.59 + 55.85 \sqrt{0.16 \left(\frac{A}{T} \right)^2 + \left(\frac{A}{W} \right)^2}$$

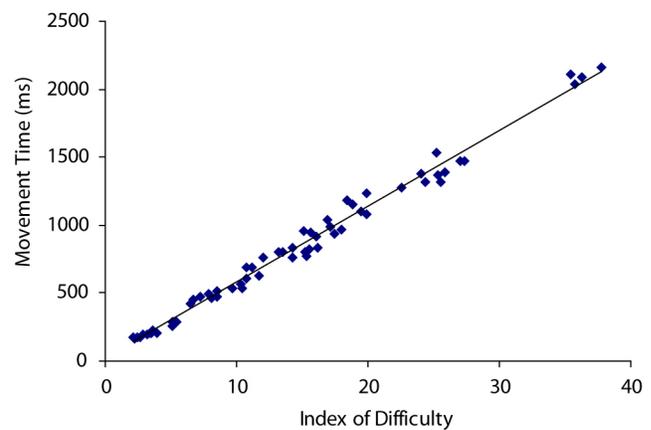


Figure 16. Movement time by ID for the 2D steering task.

We also tested the naïve model (Equation 15), proposed based on the results of Experiment 3. It also performed well, with slightly lower R^2 value of 0.97. Even though the Euclidian model provides a higher R^2 value, it is again interesting that the naïve model provides such a high fit. This may in part be due to the fact that participants were able to rest their hand on the display surface, which helps them physically constrain the stylus height. However, we must recall that in Experiment 2, T had a much stronger effect, with the same values being tested, so we have not just chosen values of T which were too “easy”. We have demonstrated that in the presence of a directional constraint, layer thickness has much less of an effect on movement time.

DISCUSSION AND FUTURE WORK

The experimental results show that our model can be used to effectively predict movement times when steering through constrained paths in above-the-surface interaction layers. The experimental model also shows that when users are able to rest their hands on a physical surface the thickness of the layer has little effect on performance when a directional constraint is also present. We also observed that there was no significant difference in user performance between layers of thickness 2cm and 2.5cm. This suggests that for interaction techniques that leverage multiple layers, like in Subramanian et al. [21], the layers could be as small as 2cm when the layer is close to the display surface.

Our results show that users were more error prone when steering through tunnels of larger path lengths, particularly with a tunnel length of 35cm. This indicates that if large

movements within a layer are required, then designers should increase the thickness of the layer. On the most part however, overall error rates were quite low, indicating that users were able to comprehend the provided cursor visualization. The informal subjective feedback which we received also suggested that the cursor visualization provided an effective indication of the stylus location.

In the future, it would be interesting to investigate other possible visualizations for displaying the stylus position within a layer. It would be particularly interesting to look at the scenario where multiple layers exist.

As this was an initial study on human performance for above-the-surface interactions, we chose to limit our focus to the variables which allowed us to form and validate our theoretical models. In the future, it would be useful to investigate how some of the unexplored factors would affect the results which we obtained.

For example, in our study we allowed users to rest their hand on the display during the steering task. When using a Tablet PC or table top system, users are generally able to do this. However, with a vertical display such as an electronic whiteboard, or a small PDA, users may not be able to rest their hand. This would likely reduce the user's ability to control the position of the input device.

Other factors which should also be explored are the required arm reach of the user, the orientation and direction of the tunnel, and the height of the tunnel above the surface. Furthermore, it would be interesting to study how the shape of the required path affects performance. In the case when the path is not straight, unlike our experimental task, the gesture may superimpose finger movements on the hand movement, resulting in increased steering difficulty. Understanding this could even result in a generalized model of steering through paths, or tubes, in free 3D space [22].

In summary, we have taken a first step towards understanding human capabilities when steering through above-the-surface interaction layers. We have proposed several models, and validated them through a series of experiments. Our work will be a significant contribution to the HCI field, as interaction techniques which use above-the-surface layers continue to emerge.

ACKNOWLEDGEMENTS

This research was carried out as part of the NECTAR research network, and is supported by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- Accot, J. and S. Zhai. (1997). Beyond Fitts' Law: Models for trajectory-based HCI tasks. *ACM CHI*. p 295-302.
- Accot, J. and S. Zhai. (2001). Scale effects in steering law tasks *ACM CHI*. p 1-8.
- Accot, J. and S. Zhai. (2002). More than dotting the i's - foundations for crossing-based interfaces. *ACM CHI*. p 73-80.
- Accot, J. and Zhai, S. (2003). Refining Fitts' law models for bivariate pointing. *ACM CHI*. p. 193-200.
- Ahlstrom, D. (2005). Modeling and improving selection in cascading pull-down menus using Fitts' law, the steering law and force fields. *ACM CHI*. p. 61-70.
- Apitz, G. and F. Guimbretière. (2004). CrossY: a crossing-based drawing application. *ACM UIST*. p. 3-12.
- Baudisch, P., X. Xie, C. Wang and W.-Y. Ma (2004). Collapse-to-zoom: viewing web pages on small screen devices by interactively removing irrelevant content *ACM UIST*. p. 91-94.
- Bezerianos, A. and R. Balakrishnan (2005). The vacuum: facilitating the manipulation of distant objects. *ACM CHI*. p. 361-370.
- Fitts, P. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology* 47: 381-391.
- Fitzmaurice, G., A. Khan, R. Pieke, B. Buxton and G. Kurtenbach (2003). Tracking menus. *ACM UIST*. p 71-79.
- Grossman, T., K. Hinckley, P. Baudisch, M. Agrawala and R. Balakrishnan (2006). Hover widgets: using the tracking state to extend the capabilities of pen-operated devices. *ACM CHI*. p. 861-870.
- Guimbretière, F., M. Stone and T. Winograd (2001). Fluid interaction with high-resolution wall-size displays. *ACM UIST*. p. 21-30.
- Gutwin, C. and A. Skopik (2003). Fisheyes are good for large steering tasks. *ACM CHI*. p. 201-208.
- Hinckley, K., G. Ramos, F. Guimbretiere, P. Baudisch and M. Smith (2004). Stitching: Pen gestures that span multiple displays. *Conference on Advanced Visual Interfaces*. p. 23-31.
- Hinckley, K., P. Baudisch, G. Ramos and F. Guimbretiere (2005). Design and analysis of delimiters for selection-action pen gesture phrases in scriboli. *ACM CHI*. p. 451-460
- MacKenzie, S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction* 7: 91-139.
- MacKenzie, S. and W. Buxton (1992). Extending Fitts' law to two-dimensional tasks. *ACM CHI*. p. 219-226.
- Moran, T. P., P. Chiu, W. v. Melle and G. Kurtenbach (1997). Pen-based interaction techniques for organizing material on an electronic whiteboard. *ACM UIST*. p. 127-136.
- Pastel, R. (2006). Measuring the difficulty of steering through corners. *ACM CHI*. p. 1087-1096.
- Ramos, G., M. Boulos and R. Balakrishnan (2004). Pressure widgets. *ACM CHI*. p. 487-494.
- Subramanian, S., D. Aliakseyeu and A. Lucero (2006). Multi-Layer Interaction on Digital Tables. *ACM UIST*. 269-272.
- Zhai, S., Accot, J., and Woltjer, R. (2004). Human action laws in electronic virtual worlds: an empirical study of path steering performance in VR. *Presence: Teleoper. Virtual Environ.* 13(2): 113-127.