IPT-EGVE Symposium (2007)
B. Fröhlich, R. Blach, and R. van Liere (Editors)

Widget Manipulation Revisited: a Case Study in Modeling Interactions Between Experimental Conditions

Jean-Bernard Martens[1], Arjan Kok[2] and Robert van Liere[2,3]

[1] Department of Industrial Design, TU. Eindhoven, The Netherlands
[2] Department of Computing Science, TU. Eindhoven, The Netherlands
[3] Center for Mathematics and Computer Science, Amsterdam, The Netherlands

Abstract

Widgets are often used to perform control tasks in three-dimensional (3D) virtual environments (VEs). Spatial interactions through widgets require precise 3D manipulations, and several design aspects of VEs contribute to the ease, accuracy, and speed with which users can perform these interactions. Throughout the years, VE researchers have studied relevant design aspects; for example, the location and size of the widgets, monoscopic versus stereoscopic viewing, the presence or absence of co-location, or the inclusion of (passive) tactile feedback, are all design aspects that have been studied. However, researchers have mostly studied design aspects in isolation and have paid little attention to possible interactions between conditions.

In this paper, we introduce a method for modeling interaction effects between experimental conditions and illustrate it using data from a specific case study, i.e., widget manipulation tasks. More specifically, we model how the effect of passive tactile feedback interacts with stereoscopic viewing for three widget manipulation tasks. We also model how these effects vary between two tasks, i.e., button and menu item selection. Models that include interaction effects between experimental conditions can be used to get a deeper understanding in the system design trade-offs of a virtual environment.

Keywords: Widgets, interaction, virtual environments.

Categories and Subject Descriptors (according to ACM CCS): B.8.2 [Hardware]: Performance and ReliabilityPerformance Analysis and Design Aids; B.4.2 [Hardware]: Input/Output and Data CommunicationsInput/Output Devices

1. Introduction

Apart from "core" 3D interactions such as object selection and manipulation in free space, and navigation through the environment, virtual environments (VEs) also need system control. System control allows the user to change operational states within the application, to set parameter values, to change interaction tools, etc. On desktop systems these tasks are most frequently performed by means of two-dimensional (2D) widget interfaces. Since many people have experience with such interfaces, and know how to operate them, it is an obvious choice to also use them for control tasks in VEs. An important design issue is hence how to integrate widget operations within 3D VEs. A large number of detailed design decisions can indeed influence the user acceptance and performance. Examples of such relevant design conditions are: the size and shape of the widgets (linear or pie menus, for example [KW04]), the location of the widgets in the environment, i.e., fixed to the surround, display or world [FMHS93], one-handed versus two-handed interaction [LZB98], monoscopic versus stereoscopic viewing [AW04], the presence or absence of co-location [KL04] or (passive) tactile feedback [LSH99], etc.

Up to now, most of these conditions have been studied individually. Little attention has been devoted to analyzing and describing interaction effects between such conditions. The risk of not understanding interaction effects is that designers might be overly optimistic about the effect that a condition might have in a new system.

In this case study, we propose a way to mathematically model both main and interaction effects using regression. We will show how the confidence intervals on the regression coefficients can be used to deduce whether or not interaction effects can be inferred. We illustrate the method using experimental data on widget manipulation tasks. We show that the method can be used to compare conditions within a single task or system, as well as between tasks and systems.

2. Widget Manipulation

The first widget systems in VEs were floating menus [JE92]. Windows float around in the VE, or pop up when needed, and are manipulated through gestures or ray casting with a virtual laser. An important limiting factor is the precision with which gestures can be made or a laser can be pointed. Another related issue is where to locate the widgets in the environment: they should be easily accessible to the user but not obscure important parts of the VE. Mine at al. [MBS97] have discussed the problems associated with floating menus and have shown that performance using hand-held widgets is better. One of the reasons for this is that the user can perform tasks more accurately with two hands than with one hand, as was demonstrated by Hinckley et al. [HPP97]. Another advantage is that users can profit from passive tactile feedback, [Bro99]. Lindeman et al. studied different combinations of system characteristics on three different 2D manipulation tasks, i.e., selection, sliding (along a line) and docking (drag-and-drop in a plane). The conditions they studied were: 1) hand-held versus world-fixed windows [LSH99], 2) the presence or absence of tactile feedback [LSH99, LST01], and 3) 2D or 3D rendering of the widgets [LST01]. The reported main effects were that their selection task was performed about 20% faster with hand-held windows than with world-fixed windows. Hand-held windows also provided higher accuracy than world-fixed windows in the docking task. There was a considerable advantage of using passive tactile feedback in all tasks, although the time gain depended on the complexity of the task, i.e., it varied from 16% for the selection task, over 25% for the slider task, to 34-47% for the docking task (the latter result varied across both studies). The differences in performance between 2D and 3D widgets were small and probably insignificant. The studies by Lindeman et al. were performed in an immersive environment and extrapolating their results towards virtual desktop environments is non-trivial. More specifically, the widgets used in their tasks were fairly large, i.e., the docking task was performed within a window of size 23 × 17 cm, while the sliding task used a slider of size 17 × 3 cm. Not only are these sizes unrealistically large for widgets in practical control applications, where many different items need to be represented in a limited window space, but these large sizes also mask potential problems due to limited tracking accuracy. We expect that in other systems with smaller widget sizes, these limited tracking accuracies will have an effect on performance.

To our knowledge, little research has been performed on studying the interaction effects between experimental conditions. One exception can be found in Arsenault and Ware [AW04]. They have performed a 3D tapping task and compared the interaction effects between correct perspective and stereoscopic viewing. This tapping task is closely related to making button selections in a widget interface. In this study, they found a small improvement (of about 8%) on manipulation times when comparing correct perspective with incorrect perspective. This effect was substantially smaller than the effect caused by monoscopic versus stereoscopic viewing. This latter effect increased from 15 to 25 % for an index of difficulty (defined according to Fitt's law) that varied between 2 and 5 bits. Their data also seem to point at some interaction effects between stereo and co-location (and index of difficulty of the task), but these interaction effects were not explicitly mentioned, nor described.

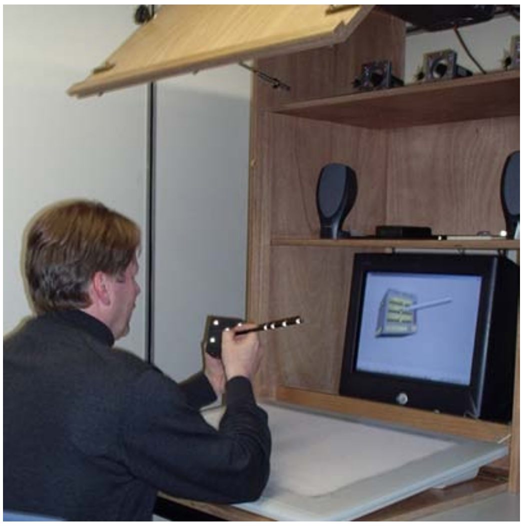

Figure 1: *The Visual Interaction Platform and the cube-widget interface.*

In summary, we intend to show how interaction effects can be measured and analyzed. We use experimental data from a widget manipulation experiment to illustrate the case. We think that both the method and the data are interesting contributions to the discussion on 3D user interfaces. In section 3 we describe the Visual Interaction Platform, the apparatus that we used in our experiment, and present the widget interface. Section 4 describes the experimental setup, while the analysis of the quantitative data is presented in section 5. The results are discussed in section 6.

3. Apparatus

The Visual Interaction Platform (VIP) [AMS*01] that was used in the experiment (see figure 1) consists of a 22 inch CRT monitor to provide a view on a virtual environment and stereo cameras with infrared light sources that track the physical objects that are located above the workspace. The display can be set to monoscopic or stereoscopic viewing modes. The VIP aims at two-handed spatial interaction. Optical tracking technology allows users to manipulate input devices, such as a pen for point selections and a cube to position and rotate models, [LM03]. The latencies caused by the optical tracker are approximately 50-60 milliseconds. The stereo cameras are located at approximately 100-120 cm from the hands of the subject, and the volume being tracked is approximately 80 × 60 × 60 cm The accuracy of the optical tracking is approximately 3 millimeters in a space up to 50 cm above the table.

The widget interface consists of a virtual cube with 2D widgets, such as buttons, menus, sliders and message boxes, and a virtual pen to perform widget selections and manipulations [KL04]. Two physical interaction devices control these virtual objects: a wooden cube and a pen. The physical cube and pen used in the experiment measure 7.5 cm and 25 cm, respectively. The widget interface provides the user with visual feedback about his actions. The virtual pen has a small sphere at its tip. The color of this sphere indicates whether the virtual pen is in contact with the virtual cube (red) or not (blue). The widgets on the surfaces of the cube change shape or color when the pen selects them (similarly as in current 2D graphical user interfaces). Most widgets, such as buttons and sliders, are displayed as 3D objects having a height and when selected, the height is reduced, just as in case of a real (pressed) button.

4. Experimental Setup

Our experiment aimed at quantifying a number of hypothesized main effects:

1. Passive tactile feedback is expected to improve the performance of widget manipulations. If the user physically feels that the pen touches the cube, he is expected to be able to perform widget tasks faster and with fewer mistakes than when he has only visual information that the pen touches the cube. This is less obvious than it looks because of the limited accuracy of the optical tracking. Instances will occur where the pen physically touches the cube, without this being (correctly) recognized by the virtual application.
2. Stereoscopic viewing is expected to improve the performance of widget manipulations. It is however unclear a priori whether or not the benefit is substantial, especially in case passive tactile feedback is provided.

These two predicted main effects are well-known from literature, but we also want to analyze their interaction, i.e., understand how they behave in isolation and in combination, and across different tasks.

4.1. Experimental Conditions

The following experimental conditions are tested:

1. Tactile feedback (T) versus no tactile (nT) feedback. The physical cube and pen object provide passive tactile feedback (T) when they touch. Since they have the same shape and size as their virtual counterparts, physical contact between pen and cube is intended to correspond with virtual contact. Assigning a small lengthwise offset (of 4cm) to the tracked position of the pen makes is possible to put the virtual pen into the virtual cube. This forces the user to operate the widget interface without tactile feedback (nT). He has to hold the real pen 4 cm away from the real cube to let the virtual pen touch the virtual cube.
2. Stereoscopic (S) versus monoscopic (nS) viewing. In the VIP, the images on the vertical screen are normally viewed without stereoscopic glasses (nS), because the use of such glasses can be experienced as cumbersome. Stereoscopic viewing (S) without head tracking can however be supported. Since no head tracking is available, the camera viewpoint is fixed to 30 cm in front and 25 cm above the center of the tablet (or equivalently, 75 cm in front of the center of the monitor).

Combining the T/nT conditions with the S/nS conditions leads to four combinations to be tested.

4.2. Tasks

Different types of 2D widgets exist, each with their own way of manipulation. In order to allow more general statements about the performance of widget tasks under the conditions mentioned in section 4.1, we designed three representative tasks that varied over a wide range of difficulty (as will be exemplified later by the results of the experiments). Two of the tasks (button and slider task) were similar to tasks performed by Lindeman et al. [LSH99, LST01], although the size of our widgets was substantially smaller.

Button task. Pressing a button is a selection task. For the button task 16 buttons are defined on one of the sides of the cube: 15 yellow, 1 blue, see figure 2. The subject has to push the blue button. When selected properly, the button turns yellow and another button turns blue. An entire task consists of 20 button selections. The subject has to perform this task in all conditions. Each condition uses another sequence of selections to avoid that the subject learns the sequence. The index of difficulty, defined as $ID = \log_2(D/W+1)$, where D is the distance between two successive buttons in the sequence and W is the button size (of 16 mm), varies between 1.5 and 2.5 bits. The total traversed distance is the same in all sequences, so that total task performance times can be compared fairly.

Slider task. Slider setting is a steering task. When a slider is selected it has to be moved along the widget until the correct value is reached. If the pen moves more than 0.5 cm away

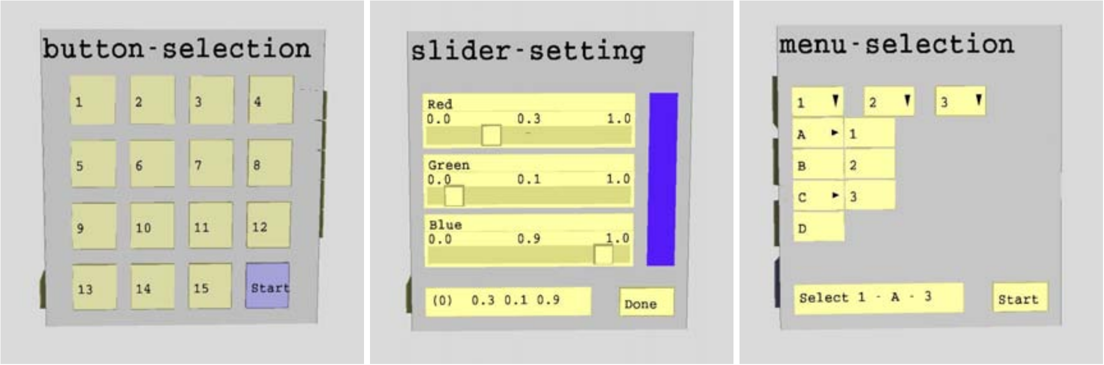

Figure 2: *Illustrations of the tasks in the 2D widget experiment.*

from the axis of the slider, then contact between the pen and slider is broken and the slider needs to be re-selected before it can be moved again. For the slider task one of the sides of the cube contains three sliders, a color indicator (right), a message box (bottom left) and a button (bottom right), see figure 2. The values of the sliders define the color that is shown in the color indicator. The sliders can have values from 0.0 to 1.0 in steps of 0.1. The task starts by pressing the button. The color (red, green and blue value) to be defined by the subject is shown in the message box. When the subject has moved all three sliders to the correct values, he has to push the button to indicate that the color definition is finished. If the color is defined correctly, the message box shows a new selection. Otherwise, the subjects needs to adjust the settings until they are correct. An entire task consists of 5 color definitions. During a single color definition, each slider has to be changed by at least 0.1, and the sum of the changes of all three sliders is exactly 1. This assures that times for color definitions can be compared fairly.

Menu task Navigation through a pop-up menu can be performed as either a selection or a steering task. It can consist of discrete sub-selections (of menu or submenu items) and/or continuous movements (steering) between sub-selections. For the menu task one of the sides of the cube contains three drop-down menus, a message field (bottom left) and a start button (bottom right), see figure 2. Menus "1" and "2" contain both menu and submenu items (hence, contain 2 layers of menu items), while menu "3" contains only menu items (only a single layer of menu items). When the start button is pushed, the message box shows the selection to be made. After finishing the proper menu selection, the message box displays the next selection. During movement between the sub-selections, it is possible to make an incorrect selection (e.g. by accidentally hitting or releasing the menu at a non-requested entry), in which case the menu selection has to be restarted. An entire task consists of 12 correct menu selections, 5 of which are sub-menu selections (such as "1-A-2"), instead of primary menu selections (such as "1-B"). In each task the same 12 selections have to be made. To avoid learning the sequence, each time the subject performs this task the selections are presented in a different order.

The menu and the submenu selections will be analyzed separately since the complexity of the interactions involved is quite distinct. Menu selections such as "1-B" can be performed in two ways. The first way is by making two distinct point selections, i.e., select drop-down menu "1", lift the pen and make a second selection on menu item "1-B". The second way is by combining a point selection with a steering operation, i.e., first select drop-down menu "1" and drag the pen without lifting to the required menu item "1-B". The submenu selections can be performed by three point selections or by a combination of point selections and steering operations. For example "1-A-2" can be reached by a point selection of the drop-down menu "1", followed by a linear drag downwards to item "1-A" which will open the submenu, followed by a linear drag to the right to submenu item "1-A-1" and a linear drag downwards to the requested submenu item "1-A-2". Such a succession of three steering operations in different directions is obviously much more difficult than a single linear drag operation required for (main) menu selections.

4.3. Subjects

Twenty-two different subjects participated in the experiment: 20 right-handed subjects and 2 left-handed subjects. Seven of the subjects were female, and all were experienced computer users (human-computer interaction scientists). Three subjects had previous experience using the VIP, but none of them had used the cube-widget interface before.

The subjects performed all 3 tasks in a given condition, before moving on to the next condition. The order in which the subjects went through the conditions was varied, in order to average out practice and fatigue effects.

4.4. Procedure

Before the actual experiment started, the principles and operation of the VE were explained to the subject, and the subject was told what he was expected to do during the experiment. A training session allowed the subject to get acquainted with the interface.

Each task (button, menu and slider) contained a separate "start" button, so that the subjects could decide themselves when to start a task. The time for each separate widget manipulation within a task (for the slider task: the time to define a color with three sliders) and the time for completion of the entire task were logged automatically. The number and type of incorrect selections were also recorded in case of the button and the menu task, although they will not be analyzed in this paper.

5. Data Analysis

We collected 1760 button selections, 440 color definitions (one color definition consists of three slider settings and one button push) and 1056 menu selections.

Many experimental papers still analyze strictly positive measurements such as time duration D for a task using techniques such as ANOVA (Analysis of Variance). Such analyses are based on two assumptions, i.e., that the data are normally distributed and have constant variance. In practice, such positively-valued data are mostly heavily skewed and the variance increases with amplitude. Taking the logarithm of the data ($\log D$) usually results in a much closer correspondence with the above-mentioned assumptions, as is demonstrated in Figure 3 for the duration measurements in the button selection task. Similar results were obtained for the other tasks.

From the estimated standard deviations on $\log D$ confidence intervals can be constructed for the average $\log D$, using the t-distribution, [DS98]. The average values and the confidence intervals on $\log D$ can subsequently be mapped to duration D to allow for an easy interpretation of the results. Following [ML03], we prefer to use confidence intervals $[L_D, U_D]$ that correspond to $1/\sqrt{2}$ times the 95 % confidence intervals of the mean. This choice allows to easily perform hypothesis testing (with the most commonly used threshold probability of $\alpha = 0.05$), i.e., two averages are significantly different if the confidence intervals do not overlap. The results are graphically illustrated in Figures 4 and 5. For the menu task, it turns out that there is a significant difference in the times needed for primary menu selections and the times needed for submenu selections.

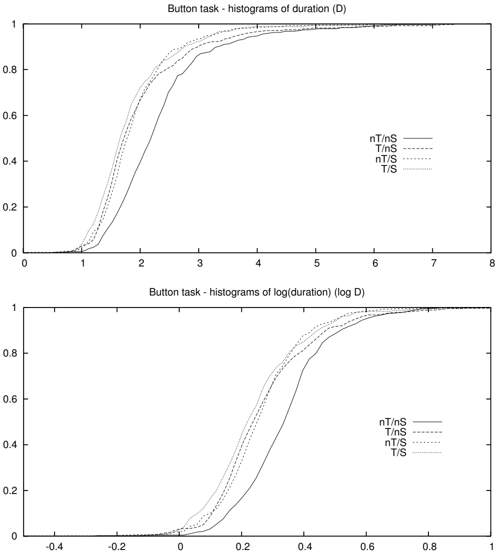

Figure 3: *Cumulative histograms of duration D and* $\log(D)$ *for the button selection task in four experimental conditions.*

A linear regression model can be used to model both the main effects and the interaction effects of the different conditions. Such a model summarizes the data and makes it easier to interpret them. More specifically, we model the logarithm of the duration times (for each task separately) by the following linear model

$$\log D = \log D_o + \log f_T \cdot T + \log f_S \cdot S + \log f_{TS} \cdot T \cdot S,$$

which corresponds to a nonlinear model for the duration times themselves, i.e.,

$$D = D_o \cdot (f_T)^T \cdot (f_S)^S \cdot (f_{TS})^{T \cdot S},$$

where

- D_o is the average duration in the nT/nS condition
- f_T is the reduction factor caused by tactile feedback
- f_S is the reduction factor caused by stereo viewing
- f_{TS} describes the interaction between tactile feedback and stereo
- T is 1 if tactile feedback is on (T), 0 otherwise (nT)
- S is 1 if stereo is enabled (S), 0 otherwise (nS)

Note that an additional advantage of modelling $\log D$ instead of D itself now becomes evident: the regression variables f_T, f_S and f_{TS} express *relative* rather than *absolute* gains. The relative gain of using tactile feedback alone, stereo alone or combined are f_T, f_S and $r_{TS} = f_{TS} \cdot f_T \cdot f_S$, respectively.

When estimating these linear regression models from the data, we do not only estimate the regression variables themselves but also their 95 % confidence intervals [DS98]. In case the 95 % confidence interval for the interaction term f_{TS} includes 1, we estimate a simpler regression model without an interaction term ($f_{TS} = 1$). The reason to do this is that we think it is interesting to distinguish cases were systems conditions interact from cases where they don't. Non-interacting system conditions are obviously easier to inter-

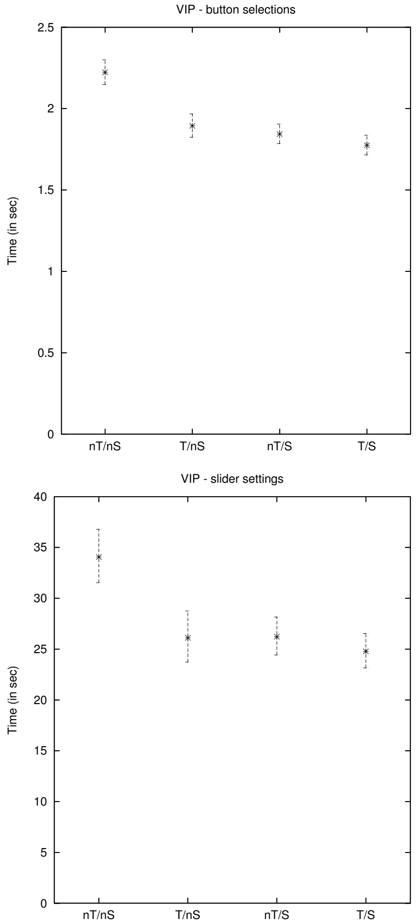

Figure 4: *Average performance times with confidence intervals for the button and slider tasks.*

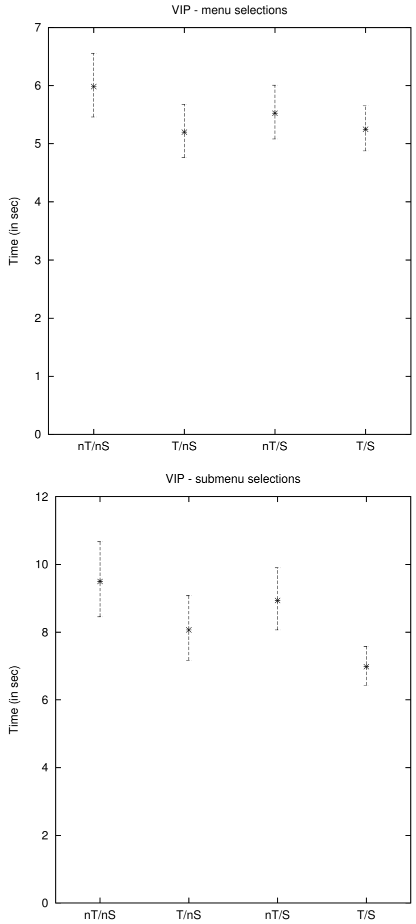

Figure 5: *Average performance times with confidence intervals for the menu and submenu tasks.*

pret. The results of the regression analysis are summarized in table 1. Cases in which 1 is contained in the confidence interval $[L_E, U_E]$ of an estimate (except for D_o) imply that the condition has no statistically significant effect on task duration. The relative advantages r_{TS} of combined techniques is also added to the tables for easy reference.

6. Discussion

We discuss various insights that are obtained from the regression model:

- Conditions in isolation (main effects)
  The conclusions are most straightforward for the performance times on the pure selection (i.e., button) task and for the steering (i.e., slider) task. For both tasks, there is a significant effect of passive tactile feedback on performance time. The performance gain (modeled as f_T) is approximately 8% for the button task and 23% for the slider task. The gain of tactile feedback in case of the button task is substantially smaller than the gain for the comparable selection task (16%) reported in [LSH99], while the gain for the slider task is about the same in both studies. The effect of stereo without tactile feedback (f_S) is 17% for the button task, which is very close to the result found by Arsenault and Ware [AW04] for their 3D tapping task in a VR environment. The gain for the slider task (23%) is even higher.
  The conclusions from the menu task are less straightforward, and also depend on whether primary menu items or submenu items were selected. The effect of tactile feedback dominates over the effect of stereo ($f_T < f_S$ for both menus and submenus), and both effects are independent. Unlike in the button and slider tasks, stereo offers a gain of 10 % in performance time on top of passive tactile feedback in case of submenu selections. The strategy that is adopted by the subjects, i.e., to use primarily selection

	D_o	f_T	f_S	f_{TS}	r_{TS}
button					
E	2.22	0.852	0.830	1.129	0.798
L_E	2.15	0.811	0.790	1.054	
U_E	2.30	0.896	0.871	1.210	
slider					
E	34.07	0.767	0.770	1.233	0.728
L_E	31.47	0.685	0.690	1.056	
U_E	36.88	0.858	0.859	1.439	
menu					
E	5.85	0.911	0.966		0.880
L_E	5.43	0.838	0.888		
U_E	6.29	0.990	1.050		
submenu					
E	9.71	0.813	0.902		0.733
L_E	8.86	0.733	0.813		
U_E	10.64	0.901	1.001		

Table 1: *Estimates E of regression parameters together with their 95 % confidence intervals* $[L_E, U_E]$ *for the model that describes the effects of passive tactile feedback (T) and stereo viewing (S). The four rows correspond to the different tasks being performed.*

or steering, seems to have a significant effect on performance, however. Steering is often not very successful, especially when the 3D scene is viewed in mono. After a number of mistrials, subjects therefore tend to switch from steering to subsequent selections. This switching is most pronounced in the case of submenus.

- Interactions between conditions
  The interaction condition $f_{TS} > 1$ indicates that the effects of tactile feedback and stereoscopic viewing are not independent, i.e., the effect of tactile feedback is f_T in case of monoscopic viewing, and $r_T = f_T \cdot f_{TS}$ in case of stereoscopic viewing. For the button task, the gain is 15% ($f_T = 0.85$) and 4% ($r_T = 0.96$), respectively, while for the slider task, the gain is 23% ($f_T = 0.77$) and 5% ($r_T = 0.95$), respectively. Hence, we can conclude that the benefit of passive tactile feedback virtually disappears when stereo viewing is present.
  The above conclusion does not imply that passive tactile feedback is not important, quite the contrary. Indeed, the interaction effect $f_{TS} > 1$ also implies that the influence of stereo when passive tactile feedback is present, i.e., $r_S = f_S \cdot f_{TS}$, is also limited for both the button task ($r_S = 0.94$) and the slider task ($r_S = 0.95$). This indicates that stereo viewing is not strictly necessary for these tasks when passive tactile feedback is provided.
  We can also model interactions between task conditions by extending the model to

$$D = D_o \cdot (f_T)^T \cdot (f_S)^S \cdot (f_{TS})^{T \cdot S} \cdot (f_M)^M \cdot (f_{TM})^{T \cdot M} \cdot (f_{SM})^{S \cdot M}$$

  where $M = 0, 1$ is used to signal two different task conditions. For instance, we can expect menu selections to be approximately equivalent to two button selections. If we jointly model the data from the button task (M=0) and the menu task (M=1), we however find that the main effect is $f_M = 2.72$, with a 95% confidence interval equal to $[2.54, 2.90]$. This loss in efficiency (since f_M is significantly larger than 2) is most likely due to the fact that not all subjects adopt a two-button selection strategy, but a slider strategy that is less efficient. The interaction effects are modeled by $f_{TM} = 1.09$, with confidence interval $[1.01, 1.18]$, and $f_{SM} = 1.00$, with confidence interval $[0.93, 1.08]$. Since f_{SM} is not significantly different from 1, we can conclude that the relative advantage of stereo is the same for both tasks. Since f_{TM} is significantly larger than 1, we can conclude that the relative advantage of tactile feedback is smaller in the menu task than in the button task.
- Conditions across Virtual Environment Platforms.
  We can also use the mathematical model to make comparisons between systems. Since the same widget manipulation experiment was performed on the Personal Space Station (PSS) [KL04], we can compare the effect of tactile feedback between closely related conditions on the PSS (i.e., excluding co-location) and on the VIP (i.e., including stereo). The result is shown in table 2, where S=0 refers to the VIP system and S=1 to the PSS system.

	D_o	f_T	f_S	f_{TS}	r_{TS}
button					
E	1.85	0.954	0.653		0.623
L_E	1.80	0.922	0.632		
U_E	1.91	0.986	0.676		
slider					
E	26.23	0.945	0.744	0.804	0.565
L_E	24.51	0.859	0.672	0.696	
U_E	28.07	1.040	0.824	0.928	
menu					
E	5.64	0.912	0.582		0.531
L_E	5.28	0.843	0.537		
U_E	6.09	0.987	0.630		
submenu					
E	8.94	0.781	0.512	1.225	0.490
L_E	8.18	0.690	0.448	1.014	
U_E	9.76	0.885	0.586	1.481	

Table 2: *Estimates E of regression parameters together with their 95 % confidence intervals* $[L_E, U_E]$ *for the model that describes the effects of passive tactile feedback (T) across platforms (i.e., S=0 is the VIP system with stereo viewing, while S=1 is the PSS system without co-location).*

Note that there are no interaction effects for the button and menu task, and opposite interaction effects for the more difficult submenu and slider tasks. More specifically, since $f_{TS} > 1$ in case of the submenu task, we conclude that the effect of tactile feedback in reducing task duration is smaller in case of the PSS than in case of the VIP for this task. Since $f_{TS} < 1$ in the case of the slider task, the situation is reversed for this task.

7. Conclusion

We have introduced a method for modelling main and interaction effects between experimental conditions for interaction tasks in a VE. The merits of these models are that the effects are expressed in terms of relative rather than absolute gains and that significant effects can be distinguished from insignificant ones. A system designer can use the model to make informed trade-offs when designing a virtual environment. Also, the model can be used to compare main and interaction effects across interaction tasks and systems.

In this case study, we have applied such models for analyzing three different widget manipulation tasks. In the case of the Visual Interaction Platform, we have shown how the effects of passive tactile feedback interacts with stereo viewing. When analyzed in isolation, the effect of tactile feedback dominates over the effect of stereo. However, when analyzing the interactions between these effects, we can conclude that the benefit of tactile feedback virtually disappears when stereo viewing is present (or vice versa). In addition, we have compared these conditions between the VIP and PSS systems. We have shown that the effect of tactile feedback is smaller for the PSS in the case of the submenu task. However, for the slider task, we show that the effect of tactile feedback is greater for the PSS.

This case study shows that a systems designer should be very cautious when reasoning about experimental conditions for interaction tasks in a VE. Not only is it important to understand the effects of experimental conditions in isolation, but also to understand the interactions between these conditions. In addition, the systems designer should also be cautious when reasoning about experimental conditions across platforms.

For future work, we plan to use these models to analyze different conditions, tasks and performance metrics. More specifically, we plan to analyze the effects that tracker accuracy and widget size have on the efficiency of interaction tasks such as steering. Also, we plan to also use error frequencies as a performance metric next to task duration time.

References

[AMS*01] ALIAKSEYEU D., MARTENS J., SUBRAMANIAN S., VROUBEL M., WESSELINK W.: Visual interaction platform. In *Proceedings of Human-Computer Interaction - INTERACT 01* (2001), pp. 232–239.

[AW04] ARSENAULT R., WARE C.: The importance of stereo and eye coupled perspective for eye-hand coordination in fish tank vr. *Presence 13* (2004), 549–559.

[Bro99] BROOKS F.: What's real about virtual reality? *IEEE Computer Graphics and Applications 19*, 6 (Nov./Dec. 1999), 16–27.

[DS98] DRAPER N. R., SMITH H.: *Applied Regression Analysis*. Wiley-Interscience Publication, 1998.

[FMHS93] FEINER S., MACINTYRE B., HAUPT M., SOLOMON E.: Windows on the world: 2d windows for 3d augmented reality. In *Proceedings of UIST '93* (1993), pp. 145–155.

[HPP97] HINCKLEY K., PAUSCH R., PROFFITT D.: Attention and visual feedback: The bimanual frame of reference. In *Proceedings of 1997 Symposium on Interactive 3D Graphics* (1997), pp. 121–126.

[JE92] JACOBY R. H., ELLIS S. R.: Using virtual menus in a virtual environment. In *Proceedings of Visual Data Interpretation* (1992), SPIE, pp. 39–48.

[KL04] KOK A., LIERE R.: Co-location and tactile feedback for 2D widget manipulation. In *Proceedings of the IEEE Virtual Reality Conference 2004* (2004), pp. 233–234.

[KW04] KOMERSKA R., WARE C.: A study of haptic linear and pie menus in a 3d fish tank vr environment. *haptics 00* (2004), 224–231.

[LM03] LIERE R., MULDER J.: Optical tracking using projective invariant marker pattern properties. In *Proceedings of the IEEE Virtual Reality Conference 2003* (2003), pp. 191–198.

[LSH99] LINDEMAN R. W., SIBERT J. L., HAHN J. K.: Towards usable vr: An empirical study of user interfaces for immersive virtual environments. In *Proceedings of CHI '99* (1999), pp. 64–71.

[LST01] LINDEMAN R. W., SIBERT J. L., TEMPLEMAN J. N.: The effect of 3d widget representation and simulated surface constraints on interaction in virtual environments. In *Proceedings of IEEE VR 2001* (2001), pp. 141–148.

[LZB98] LEGANCHUK A., ZHAI S., BUXTON W.: Manual and cognitive benefits of two-handed input: An experimental study. In *ACM Transactions on Computer-Human Interaction (TOCHI)* (1998), pp. 326–359.

[MBS97] MINE M. R., BROOKS JR. F. P., SEQUIN C. H.: Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Computer Graphics (Proceedings of SIGGRAPH '97)* (1997), vol. 31, pp. 19–26.

[ML03] MASSON M. E. J., LOFTUS G. R.: Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology 57* (September 2003), 203–220.