Optimal performance in multitasking

George Farmer

Project report submitted in part fulfilment of the requirements for the degree of Master of Science (Human-Computer Interaction with Ergonomics) in the Faculty of Life Sciences, University College London, 2010.

ACKNOWLEDGMENTS

This project would not have been possible without the generous and expert support of Duncan Brumby and Chris Janssen. Thanks also to all who gave up their time to participate in the experiment.

ABSTRACT

This study investigated how people interleave their attention in a dual-task situation. Participants had to type a string of numbers for one task, and keep a randomly moving cursor inside a target area for the other task. Different strategies for interleaving the two tasks could be applied. A payoff function was used to capture performance on both tasks in a single score. Participants adapted their strategies to changes in the task environment and to changes in the payoff function. A simple mathematical model was developed that explored the possible strategies that participants could have adopted, and to identify the strategies that lead to optimal performance. In four out of the six conditions participants performed optimally according to the model. This implies that people can optimize their interleaving performance. Implications and limitations will be discussed.

LIST OF FIGURES & TABLES

CONTENTS

CHAPTER 1. INTRODUCTION

The term multitasking might conjure up images of a stressed office worker answering phones, typing e-mails and printing documents, all simultaneously. There is in fact much more variety and intrigue in multitasking than such images suggest. We live in a society where technology is mobile and ubiquitous, meaning that we can, and do, multitask anywhere, anytime. Multitasking also applies to almost all industries and areas of work, understanding how people multitask in an air traffic control centre or how a driver multitasks in a car enables us to design technology and interactions that minimise risk and maximise productivity. Our ability to, and the consequences of, multitasking will be ever more important areas for the research and interaction design communities.

Suppose you are busy multitasking and you want improve your overall performance, you might experiment with putting more effort into one of the tasks. Studies of multitasking have shown how an increase in performance on one task will typically lead to a decrease in performance on another task (Navon & Gopher, 1979). This means that people trade-off performance on one task against another.

The way in which people trade-off performance in a multitasking situation can be influenced by several factors. One factor is if a person is given instructions to prioritise one of the tasks (Brumby, Salvucci, & Howes, 2009). In this scenario people will perform better on the prioritised task and less well on the other task. This shows that people can control the trade-off balance, altering it at will. What it does not tell us is *why* they choose a particular level of trade-off against any other level.

A recent multitasking study by Janssen, Brumby, Chater & Dowell (2010) examined this question using a dual-task scenario in which participants had to type a string of numbers in one task and keep a randomly moving cursor inside a target area in the other task. The experiment was designed such that performance on both tasks generated a single score. Janssen et al. (2010) also developed a cognitive model which explored all the possible strategies available to participants. By comparing the strategies participants used with the strategies that the model explored it was possible to show that participants were using the optimal strategy in terms of maximising their payoff. Thus *why* participants adopted a strategy was shown to be in order to maximise their payoff.

The experiment described in this dissertation develops on the Janssen et al. study (2010). Several variations were made to the task environment including increasing the length of trials and providing participants with feedback each time they switched tasks. Two independent variables were manipulated. One was the cursor noise which determined how fast the cursor moved in the tracking task. The other was the payoff function which determined how points were awarded.

A mathematical model was also developed to explore the possible strategies that participants could adopt. This was compared with the participants' data to see if their performance was optimal. The key research question was firstly, would participants adapt their behaviour given variations in cursor noise and payoff function? Secondly, would participants be able to adapt their performance optimally?

The methodology used by Janssen et al. (2010) follows the cognitively bounded rational analysis (CBRA) approach (Howes, Lewis, & Vera, 2009) which advocates the use of an objective payoff function and a cognitive model to explore the strategy space available to participants. The Janssen et al. study (2010) was a slightly more complex experiment than has previously been attempted using the CBRA methodology. This experiment is further work on this recent paradigm and can serve to explore the robustness of both the Janssen et al. (2010) experiment and CBRA as a methodology.

The structure of the dissertation is as follows, chapter 2 is a literature review examining the background to cognitively bounded rational analysis and the original Janssen et al. (2010) experiment. It also examines some of the literature that motivated the changes made in this experiment. Chapter 3 details the method for the experiment and the model that was used to determine optimal performance. This is followed by the results of both the experiment and the model. Chapter 4 is the general discussion which focuses on the implications of the results as well as future work that could be carried out. See chapter 5 for the conclusion.

CHAPTER 2. LITERATURE REVIEW

This literature review is comprised of three main sections, first off it will introduce the literature around multitasking, highlighting the depth and breadth of research. The second part examines how optimal performance can be identified in multitasking scenarios. It will examine modelling approaches in particular. As this work builds on that by Janssen et al. (2010) it will also look at this study in more detail. The final section reviews the literature that speaks to the adaptations that this experiment has made to the Janssen et al. study (2010).

2.1. Multitasking

Work, or play, that requires people to switch between tasks can be seen as being on a spectrum according to how regularly people switch. (Salvucci, Taatgen, & Borst, 2009) In this way we can think about multitasking as *concurrent* at one end of the spectrum and *sequential* at the other. Concurrent multitasking involves doing things simultaneously or with very short switch times such as driving and talking, whereas sequential involves longer time intervals before switching between tasks, for example cooking and reading a book. The experiment described in this dissertation sits at the concurrent end of this spectrum since participants needed to switch between tasks every few seconds in most of the conditions.

At different positions on the multitasking continuum a variety of studies explore which facets of the tasks and their environments are important for our understanding of both human performance and the underlying cognitive processes which underpin it. Approaches include looking at, for example, how people learn when to switch tasks (Salvucci, Taatgen, & Kushleyeva, 2006), how auditory cues affect performance, (Hornof, Zhang, & Halverson, 2010), the effects of mental workload on the disruptiveness of task switching (Salvucci & Bogunovich, 2010) and people's sensitivity to rate of return (Payne, Duggan, & Neth, 2007).

Given the vast diversity of multitasking studies, and the diversity of modelling approaches Salvucci & Taatgen (2008) have tried to develop a theory that can account for findings in all these settings using a single framework called threaded cognition. This posits that multitasking behaviour is the execution of multiple task threads in which multiple processes may be carried out simultaneously as long as there is no conflict in the resources that the processes require. For example you can look at one thing and listen to something else, but you can't look at two different things simultaneously. When multitasking demands that two things are done in parallel, and those things use the same resource there is potential for a bottle-neck in the completion of the task.

The use of cognitive models in multitasking studies is widespread e.g., (Borst, Taatgen, & van Rijn, 2010; Brumby, Howes, & Salvucci, 2007; Horrey, Wickens, & Consalus, 2006; Lallement & John, 1998). Typically models are used to develop a deeper understanding of the underlying human cognitive architecture and also to infer implications for other tasks or the design of interactive systems. A recent development in modelling centres on measuring *optimal* performance i.e. the best possible way to complete a task. The following section examines the background to this approach which led to the Janssen et al. (2010) study, upon which this project is based.

2.2. Modelling optimal performance

This project was based upon the work of Janssen et al. (2010) who used a dual task setup consisting of tracking and typing tasks. Participants were rewarded monetarily for typing digits whilst keeping a randomly moving cursor within a target area. The experiment used an objective payoff function in order to determine whether participants performed optimally under different conditions. The following section examines the literature which provides the background for understanding the novelty of this approach.

An early attempt to formalise performance on multitasking was that of Navon & Gopher (1979), this consisted of establishing the bounds of performance in a Performance Operating Characteristics (POC) plot. Given the effect of putting more cognitive resources into one task over another it is possible to chart the trade-off in performance between the two. The resulting POC shows all the possible combinations of performance that can be achieved through the different possible investments of resources in each task. What a POC does not explain is which the best strategy might be. Whilst it shows the maximum output that could be achieved through a certain distribution of resources, it does not indicate which of these maximum outputs is better than any other. This is because the overall performance objective of the task is not explicit in a POC. In order to determine what the best, or optimal, strategy would be it is necessary to make explicit an objective payoff function for the task. That is to say there needs to be a way to measure the success of one strategy over another.

The trading-off of performance on one task against another is part of the variety of possible strategies open to a person performing multitasking work. With this in mind several studies have shown that people are able to prioritize tasks according to instruction e.g., (Brumby, Salvucci, & Howes, 2009; Horrey, Wickens, & Consalus, 2006; Levy & Pashler, 2008). In Brumby Salvucci & Howes' study participants were instructed to either prioritize a driving task or a dialling task, when instructed to prioritize the driving task participants adapted their strategy / trade-off in that they decreased their performance on the dialling task in order to increase their performance on the driving task. This ability to adapt has two important implications, firstly it is obvious that people are capable of adapting. Secondly, given that ability, they are making a choice about what level to adapt their performance to on each task. Given these abilities it is interesting to try and understand why people choose the particular strategy they do. One way to examine this problem is to create dual task study in which the combination of performances on each task results in a single payoff score. This way participants' adaptation can be assessed against how close they are to the optimal adaptation.

Studies on multitasking that use objective payoff functions have been carried out before. One such study (Wang, Proctor, & Pick, 2007) examined the effect of changing payoffs in order to see the impact that this would have on participants' strategies. This study did not examine however, the optimality of participants' strategy choice but instead how strategy responded to changes in payoff. Their principle finding being that participants' strategies were sensitive to changes in payoff and that there was a residual effect on strategy from previous payoff functions.

Work by Howes and colleagues (Howes, Vera, & Lewis, 2001; Howes, Vera, & Lewis, 2007; Lewis, Vera, & Howes, 2004; Smith, Lewis, Howes, Chu, Green & Vera, 2008; Vera, Howes, McCurdy, & Lewis, 2004) has been instrumental in developing an approach to studying the optimality of strategy choices. In its most recent form this approach is called Cognitively Bounded Rational Analysis (CBRA) (Howes, Lewis, & Vera, 2009). Fundamental to this work is the hypothesis that skilled human behaviour is the optimal solution to satisfying a number of constraints. These constraints come from the task environment (for instance only being able to see one task at a time), human cognitive architecture (for instance the maximum speed at which people can read a string and type it) and knowledge that people bring to a task. In order to determine whether the behaviour is optimal an objective payoff function must be used in the task. This gives a measurement by which participants' strategies can be assessed.

Howes, Lewis & Vera (2009) carried out a study on Psychological Refractory Period (PRP) in order to demonstrate the applicability of CBRA. The task involved in studying PRP was fairly simple (a straight forward stimulus response paradigm). The work upon which this experiment is based (Janssen et al., 2010) used the CBRA approach in assessing the optimality of participants' strategies in a task setup that was more complex since it involved discretionary interleaving of tasks.

2.3. Context for adaptations

Given the novelty of using a CBRA approach on a more complex dual task scenario there is substantial scope to introduce further adaptations to the study in order to explore the effects and demonstrate the robustness of the paradigm. The literature to support the interventions made in this experiment is reported in this final section.

The setup of the experiment requires participants to make a judgement as to how much time they can spend on one task before they must switch to the other. Taatgen, van Rijn, & Anderson, (2007) consider prospective time estimation and argue that it forms part of the human cognitive architecture. In their paper they report data from (Rakitin, Gibbon, Penney, Malapani, Hinton & Meck, 1998) in which participants estimated time intervals, this data shows how participants' estimation of time becomes progressively less accurate as the time period increases.

The experiment provides feedback during the task; studies on the effects of feedback during task performance also prove very interesting. Neth, Sims, & Gray (2006) explored the effects of providing global feedback on the human tendency to meliorate (preference for immediate reward over long-term reward). Despite providing feedback on participants' prospective global payoff they had only mixed success in countering the tendency for participants to maximise their local gain at the expense of their global gain.

Risk and probability play a large part in the experiment since the random movement of the cursor means that as time goes by there is increased probability and risk of incurring a penalty. The assessment of risk and participants' decision making as a result has been the subject of much research, a classic in this field is Prospect Theory (Kahneman & Tversky, 1979) of which a key principle is that people will seek risk in situations of sure loss and be risk averse in situations of sure gain. The theory further suggests that value is relative to the person's current position rather than an absolute. Thus a pay rise of £5,000 has more value to a person whose salary is £10,000 than to a person whose salary is £50,000, despite the absolute increase being the same. How people make decisions under risk is therefore a subject of particular interest in a multitasking context, here risk could be inherent in a payoff function or could be a consequence of poor performance on a task such as driving.

In more recent but related research on people's judgement of risk (Trommershäuser, Maloney, & Landy, 2008) it has been shown that people are better able to judge probability when using human motor action than when given simple economic tasks. In this experiment the authors took an economic decision task in which people typically misrepresent the probability of a rare event and created a mathematically equivalent task in which participants used a motor action (pointing to and touching a target area). The key findings from the study were that participants came closer to maximising their gain when using motor actions than when given a simple economic decision.

2.4. Summary

This literature review started by demonstrating the depth and breadth of research in to multitasking. The second section concentrated on how within this field there has been a concerted effort to develop cognitive models that shed light on the human cognitive architecture as well allowing us to infer implications for HCI design. This work has led to the interesting and recent paradigm of trying to determine whether people can adapt and use an optimal strategy when multitasking. Janssen et al. (2010) have shown that using the CBRA approach including objective payoff functions is an effective and novel research method.

The final section deals with areas of the multitasking literature which offer interesting adaptations that could be made to the Janssen et al. (2010) study in order to further explore the robustness of using objective payoff functions to measure optimal performance. These concentrate on small concrete changes that remain within the CBRA paradigm but could broaden our understanding in this new area.

CHAPTER 3. THE EXPERIMENT

The experiment used a dual-task setup based on that used by Janssen et al. (2010). The task setup required participants to interleave two tasks in order to maximise payoff. Participants gained points through typing a string of digits in one task and avoided a penalty by keeping a randomly moving cursor inside a target area in the other task.

The principle decision participants needed to make was how long they could spend typing (typing visit duration), and therefore earning points, before they needed to switch tasks in order to prevent the cursor from leaving the target area and thereby incur a penalty. The experiment included an objective payoff function so that performance on both tasks resulted in a single score. By creating a simple mathematical model it was possible to assess whether the score achieved by participants was near the maximum possible score and therefore whether participants were adapting optimally

The key hypotheses tested in the experiment concerned whether manipulating cursor noise and payoff function caused participants to adapt their typing visit duration strategy and whether they could do so optimally:

- Participants will adapt their mean typing visit duration strategy to changes in cursor noise

- Participants will adapt their mean typing visit duration strategy to changes in the payoff function

- Participants will adapt their mean typing visit duration strategy *optimally* to changes in cursor noise and payoff function

The question of optimal adaptation is designed to test the cognitively bound rational analysis hypothesis that skilled performance is the optimal adaptation to cognitive and environmental constraints (Howes, Lewis, & Vera, 2009).

Further to the above questions the experiment was designed to explore some other differences in design to the original Janssen et al. (2010) study. These were:

- Participants were given feedback on their score during each trial instead of just at the end of the trial.

- The string of numbers to be typed was made infinite and the duration of the trial was set at 120 seconds. The original study had a string of 20 digits.

These adaptations were not subject to formal hypotheses but were included in an attempt to further test the strength of the CBRA approach

3.1. Method

Participants

Twenty participants were recruited, 13 male and seven female. Participants were aged between 22 and 37 years of age (*M* = 27 years, *SD* = 4 years). Participants were postgraduate students at University College London recruited via an electronic forum. Participants all volunteered and were unpaid. An incentive of a £10 voucher was offered for the participant that performed best i.e. scored the most points in the experiment.

Materials

The task environment was displayed on a 17 inch flat-screen monitor at a resolution of 1024 by 1280 pixels. A Logitech Attack 3 joystick was used for the tracking task and a Logitech diNovo wireless keyboard was used for the typing task. The keyboard was positioned so that the numeric keypad was operated with the participant’s left hand; the joystick was operated with their right hand. The monitor was positioned directly in front of the participant and at eye-level.

A paper print out of the task environment was used to explain the tasks participants would need to complete (Appendix A). They were also given a questionnaire to fill out upon completion of the experiment which gathered demographic information and the participants’ perception of their own performance on the task. (Appendix B).

The task interface was written in the Python programming language and covered the entire screen area of 1024 by 1280 pixels. In every instance the same programme was run using the same computer and peripheral equipment. Figure 3.1 shows the layout of the interface, the typing task is on the left and the tracking task on the right.

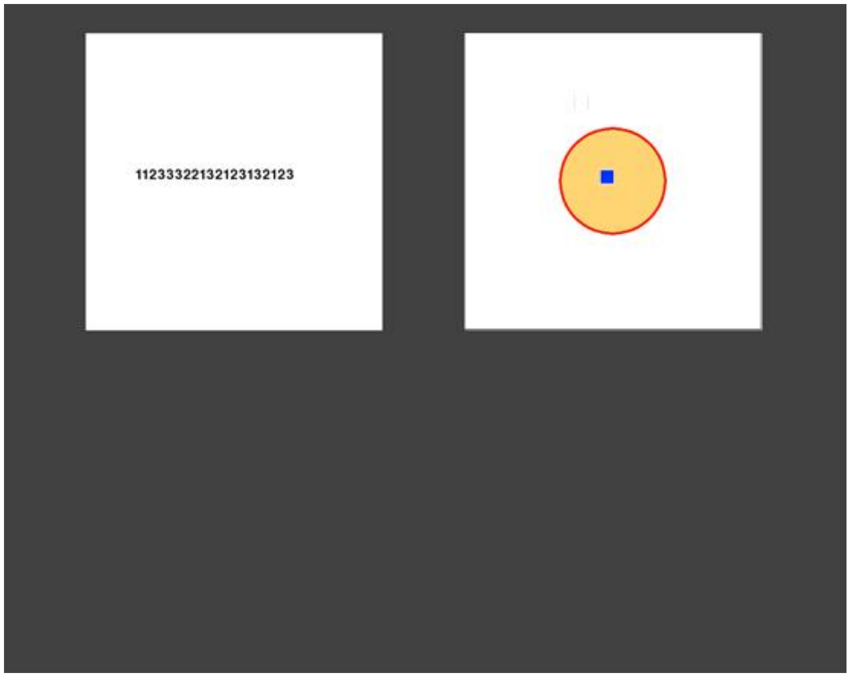

Figure 3.1 Layout of the task interface

Tracking task

During the tracking task participants could see a circular target area (the yellow area delimited by the red circle in figure 3.1) with a radius of 120 pixels on the monitor. At the beginning of the task a square cursor (10 by 10 pixels) was present in the middle of the circle. As the task progressed the cursor moved in a random fashion. The movement of the cursor was governed by a random function which had a mean of zero and a standard deviation of either three or five pixels depending on the experiment condition.

The main objective of the tracking task was to prevent the cursor from leaving the area defined by the circle (target area). When cursor noise was high the cursor took less time to exit the target area and therefore needed more careful monitoring by the participant. Participants controlled the movement of the cursor by angling the joystick in the direction they wished to move it.

Typing task

The typing task required participants to enter a string of digits shown on the screen. The string always consisted of the numbers 1, 2 and 3 in a random order. No number could appear more than three times consecutively. The display always showed 27 digits. As participants typed the left-most digit in the string, that digit would disappear and a new randomly generated digit would appear in the right-most position (see figure 3.1). The string was infinite in that it kept generating replacement numbers until the task time-limit was reached.

Dual task

In the dual task mode participants had to carry out the typing task and the tracking task simultaneously for 120 seconds. Only one of the tasks could be viewed and controlled at any one time, when one task was in view the other was covered by a grey square as shown in figure 3.2. By default the typing window was visible, when participants held down the trigger on the joystick the view switched to the tracking task. The random movement of the cursor was continuous throughout the trial whether or not it was in view. The objective of the dual task was to maximise payoff by typing in as many numbers as quickly and accurately as possible; allowing the cursor to drift outside the target area incurred a penalty.

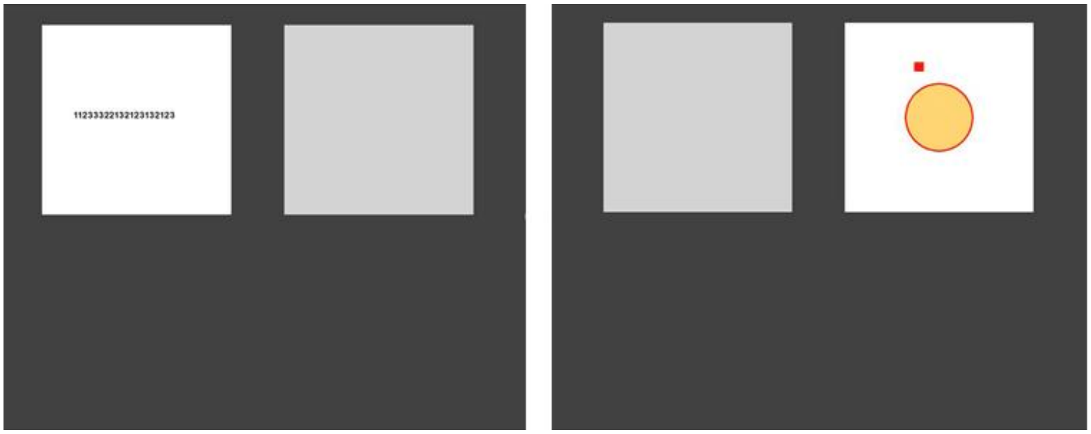

Figure 3.2 Task interface during typing visits (left) and during tracking visits (right)

An additional cue in the tracking task was the colour of the cursor, if the cursor was blue this meant that the cursor had not left the target area. If the cursor was red this indicated that the cursor had left the target area (as in figure 3.2). This colour cue was deemed necessary because the random nature of the cursor's movement meant that it was possible for it to leave the circle and then re-enter it. This might have led participants to believe that the cursor had not left the target area when in fact it had.

Design

A 2(cursor noise) x 3(payoff function) within-subjects design was used. Cursor noise had two levels and could be either 'low' or 'high'. A low cursor noise meant the cursor's random movement had a standard deviation of 3 pixels. A high cursor noise meant the cursor's random movement had a standard deviation of 5 pixels. Payoff function had three levels: lose-500, lose-all and lose-half. The lose-500 payoff function meant that the penalty for allowing the cursor to drift outside the target area was to have 500 points deducted from the points gained during the preceding typing task visit. In the lose-all condition the penalty was to lose all of the points gained in the preceding visit, and in the lose-half condition, half of the points gained in the preceding visit were deducted.

The main dependent variable was the mean of the visit durations to the typing task. This showed how long on average participants would type for, before switching to the tracking task.

Participants scored points determined by an objective payoff function that rewarded fast and accurate typing while giving a penalty when the cursor drifted outside the target area. The gain in points was 10 for every correct digit typed and -5 for every incorrect digit typed. Gain was calculated in the same way across all conditions.

Procedure

Participants were at first given verbal instructions on how to complete the tasks at the same time as being shown a print out of the task environment (Appendix A). The programme was then started allowing participants to practise the task using the joystick and keyboard. Participants practised the tracking task twice; during the practice tracking sessions participants were encouraged to explore the full range of motion of the joystick as well as to observe the speed at which the cursor moved when the joystick was not being used. There then followed two typing practice sessions. Participants were instructed to type in the number on screen as quickly and accurately as possible. The final part of the practice consisted of two dual task sessions.

Participants were informed that they could gain points through fast and accurate typing and that they could lose points through typing errors and allowing the cursor to drift outside the area defined by the circle. Participants were not informed of the amounts for either the gain or the penalties in any of the conditions but were told that the payoff function changed between conditions.

Each time a participant visited the tracking task their score from the preceding typing task visit was displayed. The score was shown above the tracking task and was simply a total amount with no breakdown of how it was calculated. At the end of each trial the total score for that trial was displayed. At the end of each condition the total score for that condition was displayed.

The experiment consisted of six conditions (two levels of cursor noise and three levels of payoff function), for each condition there were two dual task trials lasting 120 seconds each. At the beginning of each condition participants had two practice tracking sessions of 20 seconds, in order that they could estimate the speed of the cursor, and two practice typing sessions, after which the programme displayed feedback on the number of typing errors made. After completing two conditions participants got a break of approximately 120 seconds while the programme was restarted with the next set of conditions. The order of the conditions was randomized and counter-balanced across all participants.

The programme ran on a timer and could not be sped up or slowed down, therefore all participants experienced identical durations of tasks in all conditions. The programme also wrote a comma-separated-values output file for every participant recording all the dependent variables. Upon completing the experiment participants were asked to complete a brief questionnaire (Appendix B).

3.2. Model development

In order to determine whether participants were adopting an optimal strategy, i.e. one that maximised payoff, a simple mathematical model was developed in order to systematically explore the performance of every possible strategy in order to determine the best one

Consider a single trial with a length of 120 seconds. During the trial participants would type some numbers then switch to the tracking task, move the cursor to the centre of the target area and then switch back to the typing task. The longer participants spend at the typing task the more likely it is that the cursor will have drifted outside the target area in the tracking task. Therefore the mean duration of a participant's visits to the typing task represents a strategy choice on their part. A strategy which, it was hypothesised, they would adapt when changes were made to the cursor noise or the payoff function. The model was developed to calculate the score that participants would get at the end of a trial for every possible variation in the mean amount of time spent on visits to the typing task.

In order to calculate the trial payoff a number of parameters needed to be established. How these were calculated and their importance is described below.

Inter Key Press Time

The inter key press time (Kpress) refers to the amount of time between each key press during the typing task. For example if a participant has an inter key press time of 0.5 seconds they would be able to type four numbers in a two second visit (2 / 0.5 = 4). Knowing how many numbers a participant typed is crucial since this is how they earned points. Inter key press time can be thought of as a cognitive constraint, how long it takes is limited be each person's cognitive system's ability to read, process and type a number from a screen. The inter key press time was calculated by dividing the number of digits typed by the time spent on the typing task.

Error Rate

The error rate (Erate) refers simply to the number of typing errors participants made during the typing task. This amount is used to calculate the penalty for typing mistakes. The error rate can also be thought of as a cognitive constraint as it is also limited by the participant's cognitive abilities. The error rate was calculated as the number of incorrect key presses divided by the total number of key presses.

Tracking Time

Tracking time (trackingtime) is the amount of time that a participant spends on the tracking task and away from the typing task. This is important since during this time the participant cannot type and therefore cannot earn more points (although they will be preventing a penalty). Tracking time is part cognitive constraint and part task constraint. It is cognitive in the sense that speed at which it can be done is limited by the cognitive system, and it is a task constraint in the sense that it is caused by the design of the task (only one task can be controlled at a time). Tracking time was calculated as the total trial duration minus the typing time.

Probability that the cursor will cross the boundary

The probability that the cursor will cross the boundary (Pcross) is a critical part of the model. In order to model the likely gain and penalty values for a participant's visit to the typing task it was necessary to know what the probability was that the cursor would have left the target area while the participant was typing. In order to calculate this probability, the R statistical environment was used to carry out 3000 simulations of the random cursor movement over 120 seconds. This provided the probability that cursor would exit the target area for every 0.023 seconds between zero and 120 seconds. A moving average function was applied to these probabilities in order to smooth out noise in the data. 0.023 seconds was used since this was the rate at which the Python programme updated the movement of the cursor. The probability that cursor will cross the boundary is a task constraint since it is part of the task environment and unaffected by the participants' cognitive abilities.

Typing task visit time

The typing task visit time (*i*) reflects the amount of time that the participants chose to type for and therefore for the purposes of the model is the strategy which is evaluated. The following section explains how these parameters were used to model the trial payoff for each possible typing task visit time.

Model calculations

At the highest level of abstraction the model took a mean visit time (*i*) and calculated the expected value of the payoff for the trial (EVpertrial) as 120 divided by the sum of the mean visit time and the mean tracking time. Because the trial was always 120 seconds long this gave the total number of visits that the participant could make to the typing task. The number of visits to the typing task multiplied by the expected score for a single visit with a duration of *i* gives the total score for the trial as shown in the following equation.

$$EVpertrial_i = \left(\frac{120}{i + trackingtime}\right) * EVpervisit_i$$

In order to establish the expected value per visit to the typing task (EVpervisit) the model made the following calculation. For a mean visit time of *i* the expected value is the probability that the cursor has not crossed the boundary (P!cross) multiplied by the value gained if the boundary has not been crossed (V!cross). This gives the gain for an instance of mean visit time *i*. This gain is the added to the error penalty for incorrect key presses (Verror). This is then added to the probability that boundary will have been crossed (Pcross) multiplied by the penalty value for crossing the boundary (Vcross). In effect the expected value for any one visit is the gain multiplied by the probability that cursor has not crossed the boundary, which is added to the penalty multiplied by the probability that the cursor has crossed the boundary. This is shown in the following equation.

$$EVpervisit_i = ((P!cross_i * V!cross_i) + Verror_i) + (Pcross_i * Vcross_i)$$

In order to calculate the value of the gain (V!cross) the model made the following calculation. The mean visit time of *i* is divided by the amount of time it takes to press a key (Kpress). This provides the number of key-presses that participant made which is multiplied by the gain for a correct key-press (always 10 points in all conditions). The gain calculation is shown in the following equation.

$$V!cross_i = \left(\frac{i}{Kpress}\right) * 10$$

In order to calculate the typing error penalty (Verror), the mean visit time (*i*) was divided by the key-press time (Kpress) which gave the number of keys pressed. This was then multiplied by the error rate (Erate) to give the number of incorrect key-presses. This number was in turn multiplied by the typing penalty (always -5 in all conditions). The typing error penalty calculation is shown in the following equation.

$$Verror_i = \left(\left(\frac{i}{Kpress}\right) * Erate\right) * (-5)$$

The penalty value (Vcross) is straightforward to calculate since this the probability that the cursor has crossed the boundary (Pcross) multiplied by the penalty for the relevant payoff condition (PCpen). In the lose-500 condition the penalty is simply -500. In the lose-all condition the penalty is equal to -1 multiplied by the gain. In the lose-half condition the penalty is -0.5 multiplied by the gain. The penalty value calculation is shown in the following equation.

$$Vcross_i = Pcross_i * PCpen$$

3.3. Results

No data were excluded from the analysis. For each condition of the experiment participants performed two trials. Only data from each second trial has been used in the results. The purpose of using the data from each second trial was to be sure that participants had time to adapt to the payoff function before data on their performance was collected.

Typing Task – Mean Visit Time

A 2 (cursor noise) x 3 (payoff function) repeated-measures analysis of variance (ANOVA) was carried out on the dependent variable of mean visit time to the typing task, the alpha level used throughout was .05.

The first result of interest was whether or not participants had adapted the mean duration of their visits to the typing task in response to variations in the cursor noise. As shown in figure 3.3, in the low cursor noise conditions the mean visit time to the typing task was longer (4.40 seconds, SD = 0.42) than when the cursor noise was high (2.45 seconds, SD = 0.23), $F(1,19) = 61.12$, $p<.001$.

Equally of interest was whether participants had adapted the mean duration of their visits to the typing task in response to changes in the payoff functions. Changes in payoff function also had a significant effect on mean visit time. Figure 3.4 shows that the mean visit time decreased from the lose-half conditions (4.03 seconds, SD = 0.52) to the lose-all conditions (3.21 seconds, SD = 0.29) and again to the lose-500 conditions (2.85 seconds, *SD* = 0.26), $F(2,38) = 5.1$, $p<0.05$. There was no significant interaction, $F(2,38) = 0.25$. A pairwise comparison with Bonferroni correction showed that there was only a significant effect from the payoff function of lose-500 to lose-half, $p=.02$

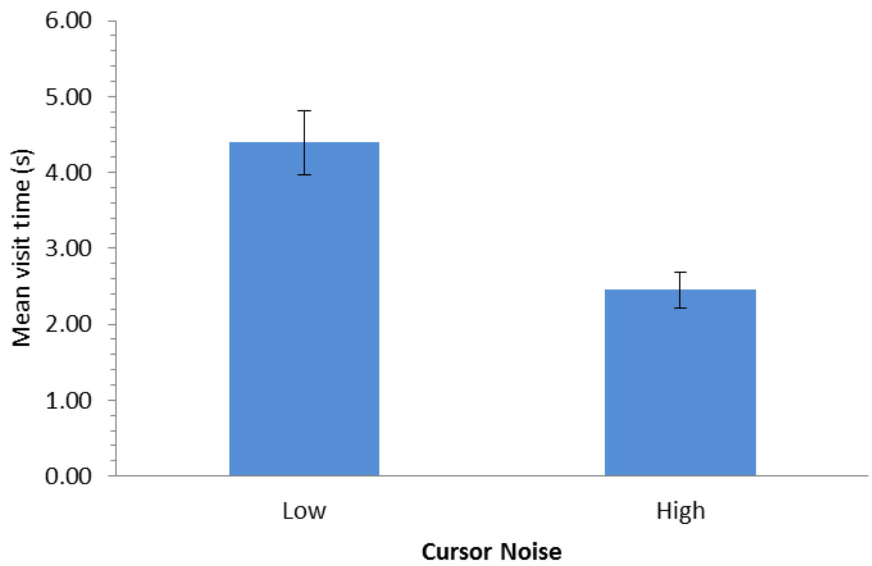

Figure 3.3 Mean duration of visits to the typing task for the different cursor noise values. Error bars depict standard error

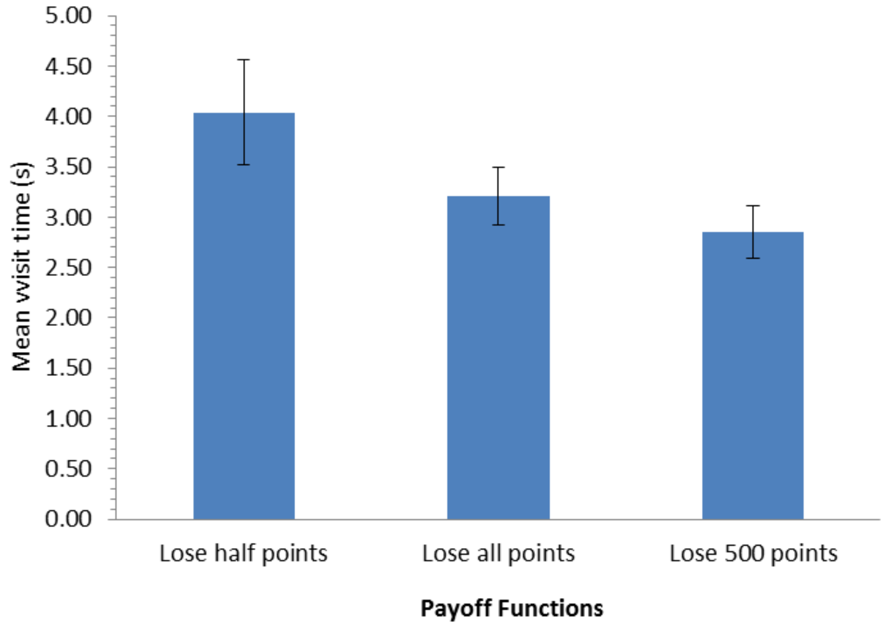

Figure 3.4 Mean duration of visits to the typing task by payoff function. Error bars depict standard error

Overall these results show that participants were adapting their dual task strategy in response to changes in both cursor noise and payoff function. In order to determine whether participants were adapting optimally the following section describes the results of the mathematical model.

3.4. Model Results

Given that participants clearly adapted their behaviour to changes in the cursor noise and the payoff penalty, the next step was to determine whether the strategies they adopted were optimal. The first step in achieving this was determining the fit of the mathematical model to the data gathered during the experiment.

Model fit

In order to determine the fit of the model data to the participant data, each participant's mean values for typing visit time, tracking visit time, inter key-press and error rate were passed to the model as parameters. The total points predicted by the mathematical model and the total points participants actually achieved were then plotted together in order to determine how good a fit the model was. In order to calculate the fit of the model to the human data a linear line of best fit with the intercept set to zero was used. The intercept was set to zero since, in theory, the model score and the participant score should have been identical.

The R^2 values of the lines of best fit are shown in table. 3.1. The nearer the R^2 values are to 1 the better the fit. The model performed best in four out of the six conditions where the payoff functions were either lose-all or lose-half. In these conditions the R^2 values were all above 0.84. The model performed less well in the two conditions in which the payoff penalty was to lose-500 points. One reason for this poor performance is that if the standard deviation of the mean typing visit time passed to the model was relatively high then impact on the payoff accuracy would be severe given the loss of 500 points for every instance of a tracking penalty being incurred.

Table 3.1 Each condition's R^2 values for the line of best fit (intercept set to zero)

Payoff Function	**Cursor Noise**	**R^2 value**	**y**
Lose-500	High	$R^2 = 0.3468$	y = 0.531x
	Low	$R^2 = 0.4812$	y = 0.8907x
Lose-all	High	$R^2 = 0.8351$	y = 0.8883x
	Low	$R^2 = 0.925$	y = 0.9266x
Lose-half	High	$R^2 = 0.8616$	y = 0.9069x
	Low	$R^2 = 0.9142$	y = 0.9421x

Figure 3.5 shows the lines of best-fit for the best and worst fits of the model data to the human data. The worst fit is for the condition where the payoff function is lose-500 and the cursor noise is high. The best fit is for the condition where the payoff function is lose-all and the cursor noise is low.

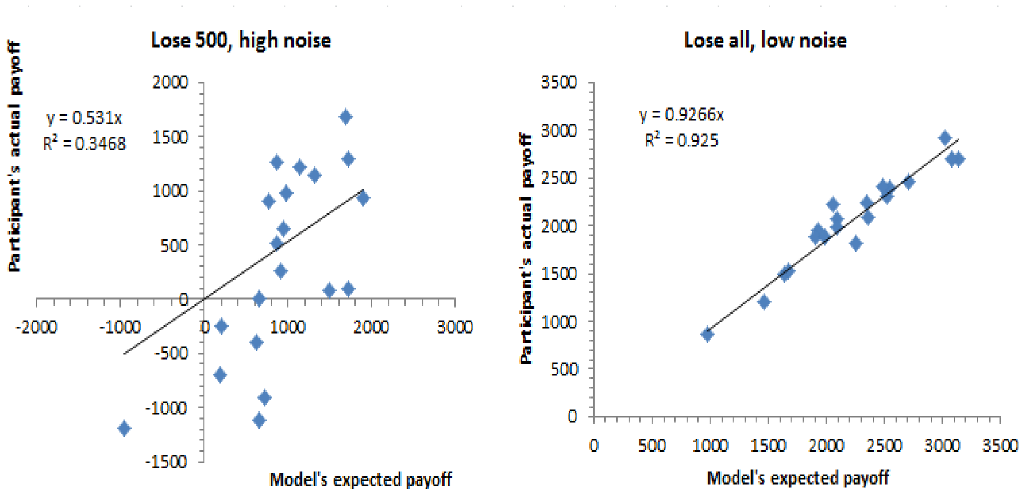

Figure 3.5 Plot of the worst (left) fit of the model to the human data and the best fit (right)

Were participant strategies optimal?

In order to determine whether participants performed optimally mean human visit time was compared to the predicted mean visit time at which the mathematical model predicted optimal payoff. Figure 3.6 shows the mean human visit time plotted with the model's predicted trial payoff by mean visit time for each condition.

In the lose-all and lose-half conditions of the experiment participants' mean visit time was near to optimum mean visit time predicted by the model. In both of these conditions participants' mean visit time afforded them a trial payoff which was within 2% of the maximum predicted by the model.

In the lose-500 points conditions of the experiment participants did not perform optimally according to the model. In these conditions the way to optimize the score would have been to ignore the tracking task and spend the whole 120 seconds on the typing task. Although the penalty would almost certainly have been incurred this would have been outweighed by the accumulation of points in the typing task over 120 seconds. Instead participants typically adopted a mean visit duration which attempted to prevent the imposition of the penalty at all.

In five of the six conditions participants adopted a mean visit time which was shorter than the optimal mean visit time predicted by the model. In the Lose-500 and high cursor noise condition however, participants' mean typing visit duration was longer at 2.18 seconds (*SD* = 0.24) than the optimal time predicted by the model of 1.45 seconds for the first peak in the plot (see figure 3.6).

These results show that participants did adopt an optimal strategy, achieving within 2% of the maximum possible score. However in the lose-500 conditions there are two interesting observations to make. Firstly the true optimal mean visit duration to the typing task would be the entire 120 seconds. No participants did this however, they did adopt a strategy that is close to optimal with regard to the first peak in the payoff plot (see figure 3.6). Also interesting is that the lose-500 with high cursor noise condition is the only one in which participants had a longer mean duration than the model's predicted optimal visit duration. The optimal modelled duration for this condition is also the shortest of all six conditions. The implications of these results are discussed in more detail in the general discussion that follows.

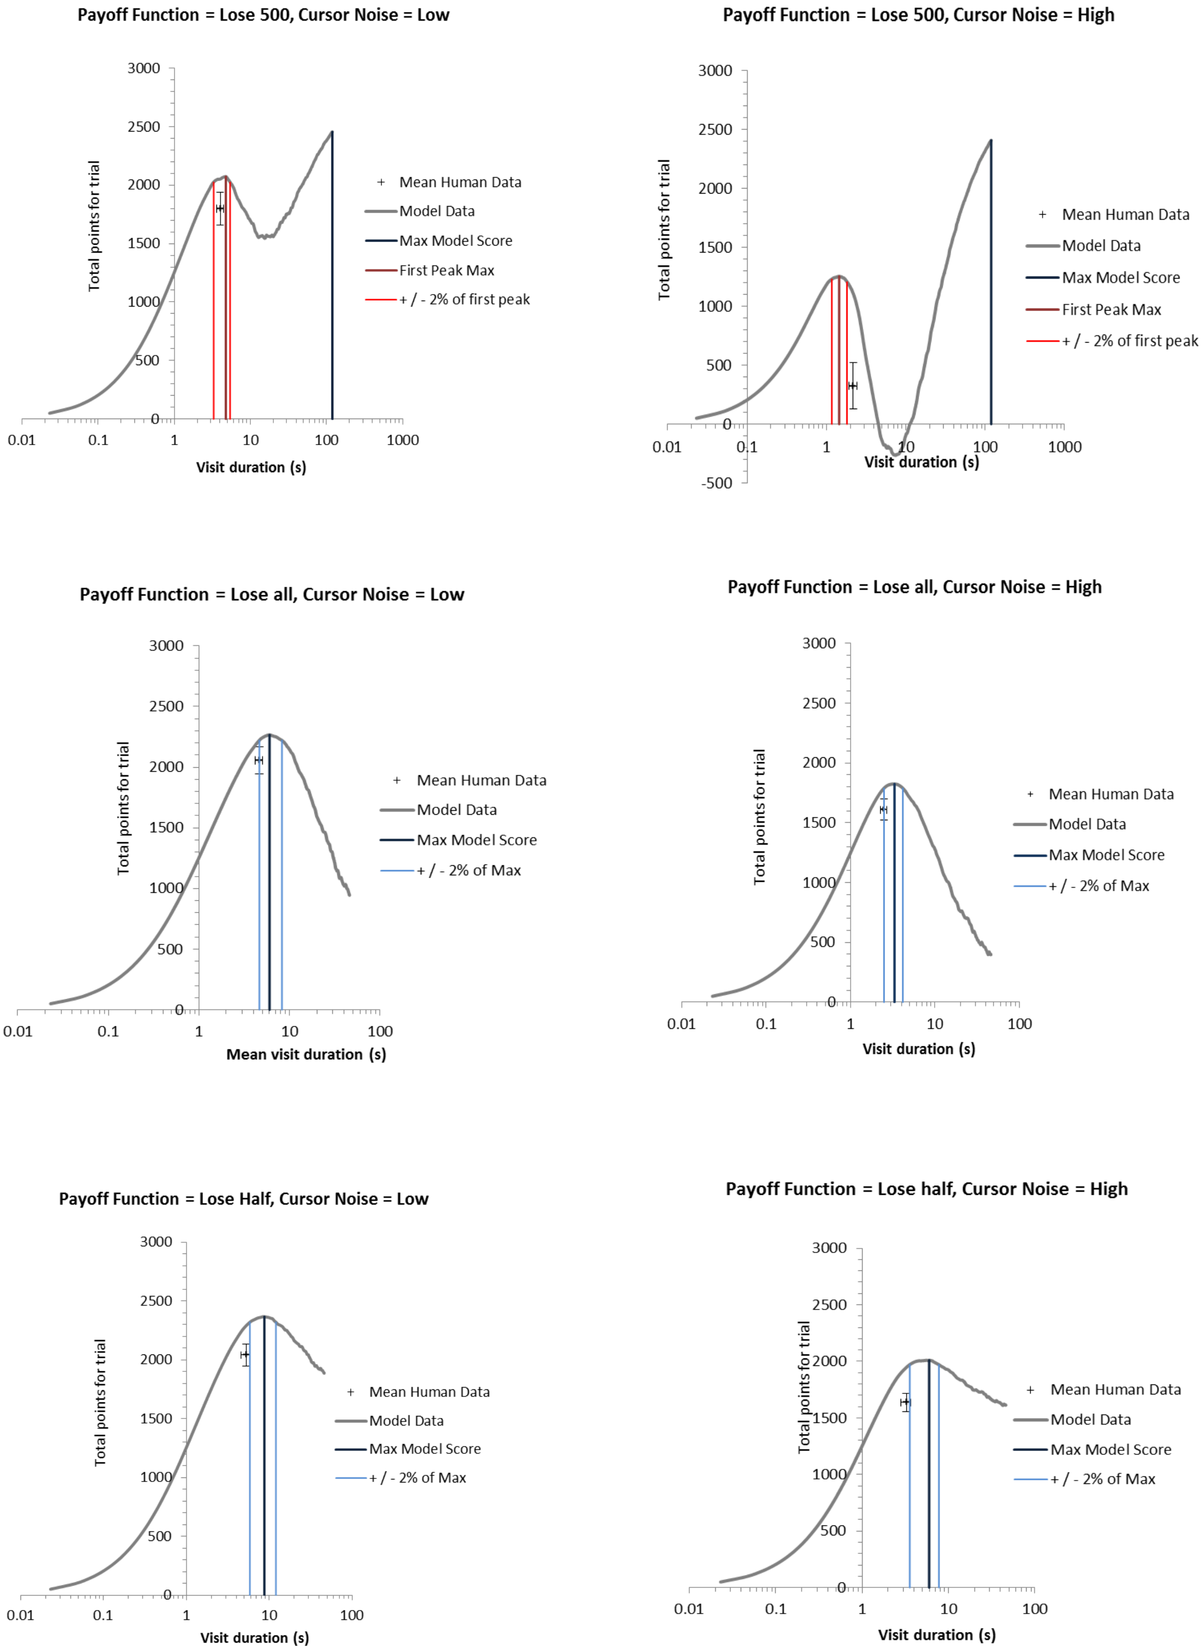

Figure 3.6 Plots showing participants' mean visit duration and trial payoff, with the modelled mean visit durations and payoffs.

CHAPTER 4. GENERAL DISCUSSION

4.1. Summary

The experiment data showed clearly that participants adapted their performance in response to manipulations of cursor noise and payoff function. An increase in cursor noise resulted in shorter visit durations to the typing task whilst changing the payoff function also resulted in differing visit durations to the typing task. These results establish the fact that participants were able to adapt, and did adapt their performance to changing task conditions.

A mathematical model was developed in order to systematically explore the strategy space that participants had available to them. This model showed in each condition what the optimal mean visit duration to the typing task should be in order to maximise the payoff for the trial. Data from the experiment revealed a good fit to the model in four out of the six conditions.

In both the high and low cursor noise for the lose-all and lose-half conditions the model showed that participants had adapted their strategies to achieve within two per cent of the maximum possible payoff. In these conditions participants not only adapted to the changing conditions of the experiment but did so optimally lending support to Howes' CBRA approach (Howes, Lewis, & Vera, 2009).

The two remaining conditions of high and low cursor noise for the lose-500 payoff condition show some interesting variations in that participants did not achieve the maximum payoff predicted by the model. Drawing conclusions from this result is however complicated by the worse fit of this data to the model. These results are discussed more fully in the following section.

4.2. Cursor noise, payoff function and optimal strategies

The decrease in mean visit time to the typing task that resulted from increasing the noise of the cursor was expected. Increasing the noise meant that the cursor was likely to leave the target area sooner and so needed to be checked more often. This result is consistent with the study by Janssen et al. (2010) who showed that increasing cursor noise led to fewer digits being typed per visit. There is an added robustness to this finding in that in the Janssen et al. (2010) study participants had just 20 digits to type whereas in this experiment the only limit on the number of digits that could be typed was the 120 second time-limit of the trial. Despite this difference in the task setup, the effect of increased cursor noise was the same.

A significant effect on the mean duration of the visits to the typing task was seen by changing the payoff functions. The use of the different payoff functions was a key difference between this study and the Janssen et al. (2010) study. A further adaptation was that this experiment provided participants with feedback after each visit to the typing task and not just at the end of the trial i.e., this experiment provided both local and global feedback. The implications of the effect that these variations had are interesting and can be better understood with reference to the model.

The mathematical model developed was very simple and used few parameters. Despite this there was a good fit to the experimental data in four out of the six conditions. In these conditions participants adapted optimally to achieve within 2% of the maximum possible score. This puts into context the adaptation of strategy that participants made in response to changes in the payoff function. Without a mathematical model of optimum performance we might have deduced that participants spent less time typing when the payoff functions had a more severe tracking penalty because they were risk-averse as shown in work such as prospect theory (Kahneman & Tversky, 1979). What the model allows us to see is that in the majority of conditions participants adapted optimally to the change in payoff. This gives us a deeper understanding of behaviour since the model shows that adapting mean typing time in the face of different payoff conditions was in fact the optimal strategy.

This does not mean that participants were not risk-averse since in five out of the six conditions participants adopted a mean visit time that was shorter than the best possible visit time (but always within 2% of optimal). This raises the interesting prospect that participants' risk aversion was a constant across the payoff conditions rather than increasing or decreasing with changes in payoff. Such a conclusion would need further research but it is the use of a model that allowed the question to be raised because it made it possible to compare participants' performance with the optimal performance thereby answering *why*, not just showing *how* their strategies were chosen.

In general, adapting to payoff functions has been shown before (Wang, Proctor, & Pick, 2007). However, in this case, as in the Janssen et al. (2010) study, there was evidence of optimal adaption. As well as using different payoff functions the experiment also differed from Janssen et al. (2010) in providing local as well as global feedback, i.e., participants got feedback on their score after each visit to the typing task not just at the end of the trial. Given the optimal adaptation in the majority of conditions there is again further evidence of the robustness of the original Janssen et al. (2010) work and the principle of using objective payoff functions to determine optimal performance (Howes, Lewis, & Vera, 2009).

The two conditions which did not show optimal performance merit a detailed discussion. Both of these conditions were the lose-500 payoff function. This payoff function showed the worst fit to the mathematical model so drawing conclusions from the discrepancy between participant performance and the modelled optimal performance needs to be done with caution. The poor fit of the model is most likely caused by two factors. Primarily, the severity of the penalty (-500) meant that if during a trial a participant twice experienced a random cursor movement that deviated from the likely probability, they could end up with 1000 points less than the model would predict. Alternatively if the participant's mean typing visit time had a high standard deviation i.e., they weren't consistent in the amount of time spent on each typing visit, it is again likely, because of the penalty size, that their score would be different from that predicted by the model.

Given the model's fit in the other four conditions it seems likely that the optimal score it predicts for each of the possible mean typing visit durations is accurate but that the model is sensitive to the consistency of participants' visit duration to the typing task. A problem exacerbated by the heavy potential loss in the lose-500 condition.

A further problem in the lose-500 condition is that the model shows that the optimal strategy is to type continuously and not switch to the tracking task at all, i.e., the optimal typing task visit duration would be 120 seconds (the entire length of the trial). None of the participants did this. Instead participants switched between tasks in a manner that suggests they were trying to avoid incurring the penalty at all. As shown in figure 4.1 participants adopted a strategy which appears to be targeting the first peak in the model data line, this first peak shows the optimum strategy if a switching between tasks approach is adopted. The model data line shows that if participants did not switch at all they would achieve the highest possible payoff at the end of the trial.

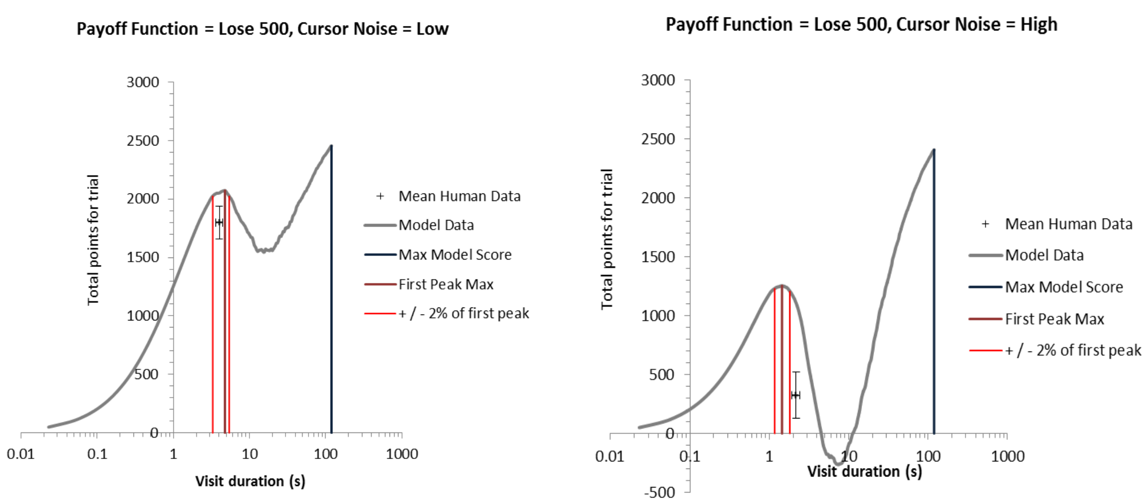

Figure 4.1 Plot showing participants' mean typing visit durations and payoff with the modelled typing visit durations and their payoffs.

In a sense there are two optimal strategies in the lose-500 conditions one which involves switching between the tasks and one that doesn't. From the position of the mean human data point on these graphs it seems clear that participants were trying to optimise payoff based on a switching tasks strategy. In the lose-500 with low cursor noise condition participants were close to the optimum (within 2 % as was the case in the other four conditions). However, interestingly in the lose-500 with high cursor noise condition participants appear to have adopted a switching tasks strategy but have failed to get within 2% of the optimum (as they did in all other conditions). Furthermore it is only in this condition where the mean human data shows a longer mean duration of the typing task than is optimal.

Since the lose-500 with high cursor noise condition required the shortest visit duration (1.45 seconds) to achieve optimal performance and this was the only condition in which participants spent longer than the optimal duration. This raises the question of whether this condition was too difficult for participants in that they couldn't achieve what they were aiming for assuming they were aiming (as in all other conditions) for a slightly shorter than optimal duration. If this is the case then participants' inability to achieve their objective could be deemed a cognitive constraint and future models would benefit from including this constraint. The next section details other improvements and future research that could be carried out to strengthen and further our understanding of the results obtained by this experiment.

4.3. Limitations and Future Research

Experiment limitations and future research

A possible limitation in the design of the study was the absence of any consideration of possible learning effects. This problem applies at several levels, firstly participants may simply have gotten better at the task toward the end of the study and therefore performed better in the last condition. In a study with a similar task setup (Hornof, Zhang, & Halverson, 2010) the experiment was carried out over three days and participants' performance improved noticeably from one day to the next.

Another issue to do with learning is that the CBRA paradigm (Howes, Lewis, & Vera, 2009) hypothesis concerns skilled behaviour, that is behaviour that has had time to learn the skills involved and become practised. Given the improvement in performance recorded by Hornof, Zhang & Halverson (2010) it is reasonable to assume that had the participants in this experiment had more opportunity to learn and practice, the experiment might have provided a fairer test of the CBRA hypothesis. Furthermore, Erev & Gopher (1999) advocate that the actual learning process itself is important to understanding multitasking behaviour and performance. Future work could look at modelling how participants learn during the course of the experiment.

A further possible design flaw in the experiment was the possibility of ordering effects. It has been shown that whilst people do adapt strategy to payoff there is also a residual effect of the previous payoff condition on the strategy that is adopted following the introduction of a new condition (Wang, Proctor, & Pick, 2007). Wang, Proctor & Pick recommend that people are given an opportunity to practice under new payoff conditions in order that they can develop a new strategy.

This experiment was counter-balanced with respect to payoff conditions and participants' data was only analysed from the second of their two 120 second trials in each condition however, a single trial may not have been sufficient time for participants to adapt fully to the new condition. Although not analysed, in running the experiment it was clear that some participants carried forward conservative strategies following the most severe payoff condition. This suggests that in future work it may be worth adopting a between subjects design in evaluating the effect of different payoff conditions.

In the study by Janssen et al. (2010) it was noted that in the large radius, low noise condition people switched to the tracking task unnecessarily – they had time simply to complete the typing task without tracking. The authors posited that this may be because participants had become biased to switching between the two tasks. In this experiment extending the length of the trial to two minutes was meant to overcome this issue, since over the course of two minutes it would always be necessary to switch at some point during the trial. However, in the lose 500 points condition of this experiment the optimal strategy would have been not to switch at all as well. No participant adopted this strategy and it is difficult to know to what extent they may have been biased into switching, not just from other conditions, but also because the experiment is explained of consisting of two tasks. Thus although participants were instructed to try and maximise their payoff they may also have felt that they were in some sense duty-bound to continue with the tracking task.

The more likely reason for participants switching unnecessarily between tasks is that they were unable to deduce from the local feedback what the optimal strategy should be. It was mathematically and logically possible to deduce from the local feedback (feedback shown at every switch between tasks) that the loss of 500 points for incurring the tracking penalty could have been outweighed by continuing to accrue points through typing. Neth, Sims & Gray (2006) raised the question in their study of whether a different combination of local and global feedback might have lowered demands on participants' working memory and enabled them to adopt a maximising strategy. This question is also interesting for any future research done on this experiment, for instance it might be possible to model the working memory demands of interpreting feedback.

It should be noted at this point that three of the 20 participants did adopt strategies that were unusual and suggest that it was at least possible to deduce the optimum strategy. Foremost amongst these was the participant who after completing the experiment asked if she should not have switched, which would indeed have been the optimal strategy. She elaborated that she had not done so because the stress of the situation had prevented her from being sure and therefore had simply carried on.

The highest scoring participant in the study deliberately incurred the penalty so that she could make a judgement as to what strategy she would adopt. She nevertheless failed to adopt the optimal strategy in the lose-500 payoff condition. The final interesting strategy undertaken by a participant during the experiment was to 'bank' points by switching to the tracking task for a split second without actually moving the cursor. This was an interesting strategy since it brought the time cost of switching between the tasks down to a few milliseconds and had the knock-on effect of providing more regular feedback on the position of the cursor in the tracking task. These different approaches highlight the flexible and adaptive nature of human multitasking even within a simple task environment. This makes it all the more interesting that given the variety of strategies available to people, in most conditions they adopted the optimal one, something which can be shown because of the CBRA methodology used.

The design of the payoff function also creates some difficulties in the analysis of the data. Because the payoff functions were categorically different rather than just quantitatively different it is difficult to infer with confidence the nature of the effect they had on participants' performance. The flip-side of this argument is that the categorical difference of the payoff functions makes participants' near optimal adaptation to them all the more interesting. Nevertheless a possible adaptation for future work would be to implement a payoff function which altered along a consistent scale and thereby allow the experimenter to build a clearer picture of the relationship between the payoff function and the strategy of the participant. This principle is clearly demonstrated with the other independent variable in this study where it is clear that an increase in cursor noise will result in a decrease in visit time.

Model limitations and future research

A clear problem with the model is that it uses time elapsed as a factor in calculating the gain that participants will get. In reality participants scored points for every correct key press. Whilst the number of key presses correlates with time elapsed, this can be problematic, especially for shorter time periods. For example if a participant visited the typing task for one second and their key press interval was 0.38 seconds they would only have the opportunity to type two digits earning them 20 points. The model, on the other hand, would have calculated (1/0.38)*10 which would return 26.32 points and therefore would have overestimated the likely score. If a discrepancy of 6.32 points is multiplied out over the course of a 120 second trial, in which for argument's sake, there might have been 50 visits then the total discrepancy becomes (50*6.32) 316 points. An unequivocal enhancement to the model for any future work would be to ensure that it calculated the maximum possible key presses within an elapsed time period rather than just the elapsed time period itself.

Perhaps the main problem with the model was that because it used the *mean* visit duration to the typing task there was no way of accounting for a participant's variation in strategy within a trial. If a participant experimented with an extremely long visit to the typing task and then for the remainder of the trial used many short visits the mean visit time may not have accurately reflected what happened during the task and therefore the model would produce an inaccurate forecast score. For future work this problem is somewhat harder to deal with however one possible adaptation would be to pass the participant's standard deviation as well as the mean visit time as parameters to the model and use this to calculate a range within which their score could be expected to lay.

The simplicity of the model also meant that there were relevant parameters which were not used. For instance participants would typically not return the cursor to the exact middle of the target area. Since the probability of the cursor exiting the area was based on the assumption of the cursor starting from the exact centre the probability calculations would always be slightly out. Since the accuracy with which participants returned the cursor to the centre of the circle could be seen as a strategic choice on their part this could be a useful parameter to include in future models.

Perhaps one of the most interesting areas to look at for future work and models would be a more thorough examination of participants' strategies for the typing task. The current model only examined the visit duration, the accuracy and the inter key press rate (amount of time between each key press). These three parameters perhaps only scratch the surface of how participants approached the typing task. Many participants reported after the experiment that before switching to the tracking task, they would encode (memorise) the next chunk of numbers in the string. This meant that when they returned to the typing task they were already prepared to type the next set of numbers from memory. Once this cycle was started it was self-perpetuating as participants could type from memory and while doing so encode the next chunk to be typed at the next visit. It is plausible therefore that in some situations participants were only ever typing from memory and not directly from the string that was on screen.

A similar strategy was noted by Smith et al. (2008) in their cognitively bound rational analysis of a typing task. In this study participants (and the model) chose to encode a chunk of a string of numbers while typing another chunk. By doing this participants were able to parallelise aspects of the task and thus save time over all. Feedback from participants in this experiment suggests that a similar process may have been occurring. Possible future research could therefore involve gathering eye-tracking data to get a more accurate picture of how participants chunked the typing task.

Parallelisation and chunking show just some of the ways in which the typing task could be further investigated. There are plenty of other aspects which could also be taken into account. If participants were indeed encoding chunks then this suggests that working memory capacity will form a cognitive constraint on the participant's strategy choice. Combined with the task constraint of the amount of time participants had available to visit the typing task there is an opportunity to carry out an interesting cognitively bound rational analysis of this experiment to see to what extent participants adopt the optimal strategy when these factors are included. These adaptations would however perhaps require a cognitive model to be built using an actual cognitive architecture rather than just a simple mathematical model. To explore this experiment further at a simpler level, manipulations could include showing only one digit at a time in the typing task or showing an entirely new string each time participants returned to the typing task.

4.4. Applications & Implications

Perhaps the main implication of this study is that having further tested the principle of using an objective payoff function to determine optimal performance, this method could be rolled out to other different and / or more complex tasks. Part of the value of this method is that by limiting the possible range of strategies a participant could adopt it helps assess whether empirical data that matches model data is creditable to the cognitive architecture of the model or down to the strategy that the participant chose.

When designing multitasking interfaces or work practices, there are several interesting aspects to this study which are worth considering. The way in which feedback on performance is given needs to be carefully thought out. Can users understand the meaning of the feedback for overall performance? Do they have the time and ability to make sense of it?

Another interesting possibility is whether multitasking work can be designed such that there is a single performance measure. There is the possibility then that people's ability to perform optimally could be exploited for better productivity. These applications are somewhat abstract; in reality further research in to the questions raised by this experiment might be the most productive way of yielding more concrete design implications.

CHAPTER 5. CONCLUSION

When multitasking, people have to make choices about how to interleave attention between tasks. Understanding why people make these choices can be effectively explored by using an objective payoff function and a model of possible strategies. In this way strategies can be shown to be optimal as they lead to the maximum payoff. This experiment varied a dual task scenario by manipulating the cursor noise and payoff function. Participants adapted their strategies to these manipulations. Furthermore, by using an objective payoff function and through the development of a simple mathematical model, this study has shown that in four out of six conditions participants adopted the optimal strategy.

These results are in line with previous work using a similar task environment (Janssen et al., 2010). This experiment significantly altered the task setup from that used by Janssen et al., nonetheless broadly the same results were found in the adoption of optimal strategies. This provides evidence that supports the Janssen et al. study and the CBRA paradigm which now offers exciting opportunities to explore further applications of this method. The two conditions where participants did not make the optimal strategy choice need further, and potentially valuable, research. Two key research questions emerged, one around the effects of feedback and how this could be delivered in a way such that participants use the optimal strategy. Alternatively a more complex cognitive model might be developed including new constraints that could show alternative optimal strategies in these conditions, given the bounds of participants' cognitive abilities.

REFERENCE S

Borst, J. P., Taatgen, N., & van Rijn, H. (2010). The problem state: a cognitive bottleneck in multitasking. *Journal of experimental psychology. Learning, memory, and cognition*, *36*, 363-382.

Brumby, D., Howes, A., & Salvucci, D. D. (2007). A cognitive constraint model of dual-task trade-offs in a highly dynamic driving task. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 233-242). New York, NY: ACM Press.

Brumby, D., Salvucci, D. D., & Howes, A. (2009). Focus on driving: How cognitive constraints shape the adaptation of strategy when dialing while driving. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1629-1638.

Erev, I., & Gopher, D. (1999). A cognitive game-theoretic analysis of attention strategies, ability, and incentives. In D. Gopher & A. Koriat, *Attention and performance XVII: Cognitive regulation of performance: Interaction of theory and application* (pp. 343-371). Cambridge, MA: MIT Press.

Hornof, A. J., Zhang, Y., & Halverson, T. (2010). Knowing where and when to look in a time-critical multimodal dual task. In S. E. Hudson, G. Fitzpatrick, W. K. Edwards, T. Rodden, & E. Mynatt, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York, New York, USA: ACM Press.

Horrey, W. J., Wickens, C. D., & Consalus, K. P. (2006). Modeling drivers' visual attention allocation while interacting with in-vehicle technologies. *Journal of experimental psychology. Applied*, *12*(2), 67-78.

Howes, A., Lewis, R., & Vera, A. (2009). Rational adaptation under task and processing constraints: implications for testing theories of cognition and action. *Psychological review*, *116*(4), 717-51.

Howes, A., Vera, A., & Lewis, R. (2007). Bounding rational analysis: Constraints on asymptotic performance. In W. Gray, *Integrated models of cognitive systems*. Oxford University Press.

Howes, A., Vera, A., Lewis, R., & McCurdy, M. (2004). Cognitive Constraint Modeling : A Formal Approach to Supporting Reasoning About Behavior. In K. Forbus, D. Gentner, & T. Regier, *Proceedings of the Proceedings of the 26th annual meeting of the Cognitive Science Society* (pp. 595-600). Mahwah, NJ: Lawrence Erlbaum Associates.

Janssen, C., Brumby, D., Dowell, J., & Chater, N. (2010). A Cognitively Bounded Rational Analysis Model of Dual-Task Performance Trade-Offs. In D. D. Salvucci & G. Gunzelmann, *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 103-108). Philadelphia, PA: Drexel University.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263-291.

Lallement, Y., & John, B. E. (1998). Cognitive Architecture and Modeling Idiom : An Examination of Three Models of the Wickens Task. In *Proceedings of the twentieth annual conference of the Cognitive Science Society* (pp. 597-602).

Levy, J., & Pashler, H. (2008). Task Prioritisation in Multitasking during Driving : Opportunity to Abort a Concurrent Task Does Not Insulate Braking Responses from Dual-Task Slowing. *Applied Cognitive Psychology*, *22*, 507-525.

Lewis, R., Vera, A., & Howes, A. (2004). A constraint-based approach to understanding the composition of skill. In M. Lovett, C. Schunn, C. Lebiere, & P. Munro, *Proceedings of the Sixth International Conference on Cognitive Modeling* (pp. 148-153). Mahwah, NJ: Lawrence Erlbaum.

Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review*, *86*(3), 214-255.

Neth, H., Sims, C., & Gray, W. (2006). Melioration dominates maximization: Stable suboptimal performance despite global feedback. In R. Sun & N. Miyake, *Proceedings of the 28th annual meeting of the cognitive science society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Payne, S. J., Duggan, G. B., & Neth, H. (2007). Discretionary task interleaving: heuristics for time allocation in cognitive foraging. *Journal of experimental psychology. General*, *136*(3), 370-88.

Rakitin, B., Gibbon, J., Penney, T., Malapani, C., Hinton, S., Meck, W., et al. (1998). Scalar expectancy theory and peak-interval timing in humans. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*(1), 15–33.

Salvucci, D. D., & Taatgen, N. (2008). Threaded cognition: an integrated theory of concurrent multitasking. *Psychological review*, *115*(1), 101-30.

Salvucci, D. D., Taatgen, N., & Borst, J. P. (2009). Toward a unified theory of the multitasking continuum. In *Proceedings of the 27th international conference on Human factors in computing systems - CHI '09* (pp. 1819-1828). New York, New York, USA: ACM Press.

Salvucci, D. D., Taatgen, N., & Kushleyeva, Y. (2006). Learning when to switch tasks in a dynamic multitasking environment. In *Proceedings of the seventh international conference on cognitive modeling* (p. 268–273). Oxford, UK: Taylor & Francis/Psychology Press.

Salvucci, D., & Bogunovich, P. (2010). Multitasking and monotasking: the effects of mental workload on deferred task interruptions. In S. E. Hudson, G. Fitzpatrick, W. K. Edwards, T. Rodden, & E. Mynatt, *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 85-88). New York: ACM Press.

Smith, M. R., Lewis, R. L., Howes, A., Chu, A., Green, C., Vera, A., et al. (2008). More than 8,192 Ways to Skin a Cat : Modeling Behavior in Multidimensional Strategy Spaces. In B. C. Love, K. McRae, & V. M. Sloutsky, *Proceedings of the 30th annual conference of the Cognitive Science Society* (pp. 1441-1446). Austin, TX: Cognitive Science Society.

Taatgen, N., van Rijn, H., & Anderson, J. (2007). An integrated theory of prospective time interval estimation: the role of cognition, attention, and learning. *Psychological review*, *114*(3), 577-98.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in cognitive sciences*, *12*(8), 291-7.

Vera, A., Howes, A., McCurdy, M., & Lewis, R. (2004). A constraint satisfaction approach to predicting skilled interactive cognition. In E. Dykstra-Erickson & M. Tscheligi, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 121-128). New York, New York, USA: ACM Press.

Wang, D., Proctor, R. W., & Pick, D. F. (2007). Acquisition and Transfer of Attention Allocation Strategies in a Multiple-Task Work Environment. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *49*(6), 995-1004.

APPENDIX A

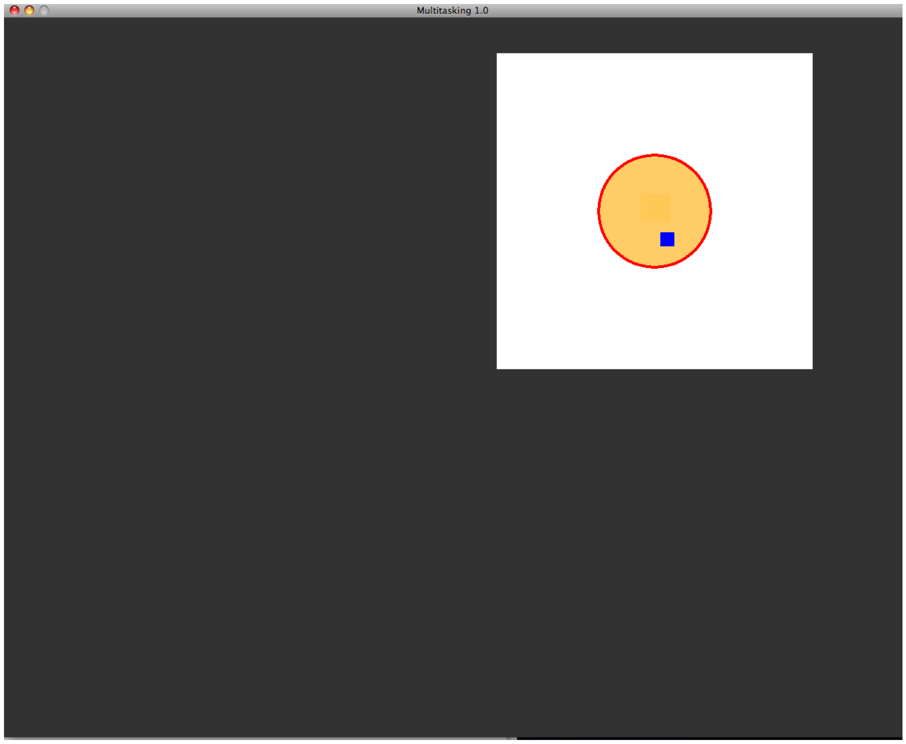

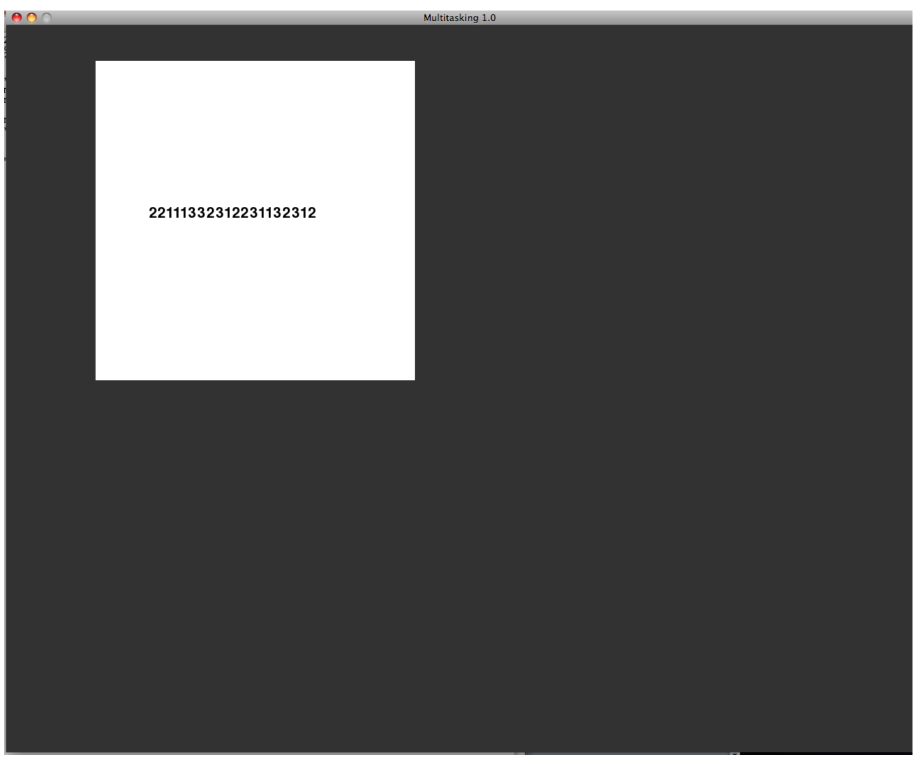

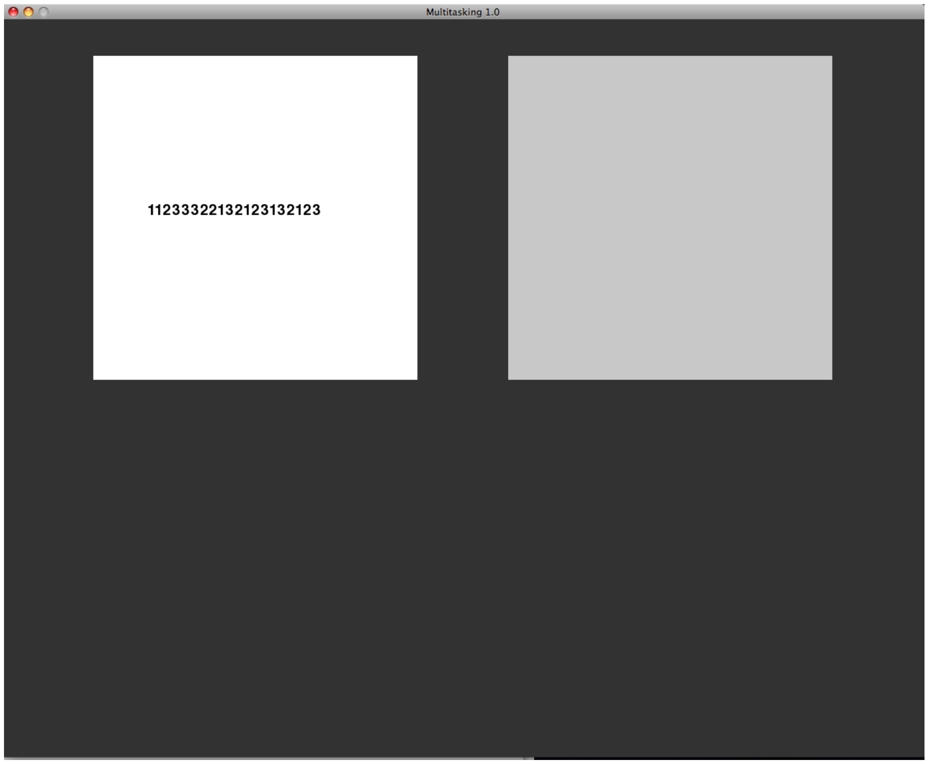
Multitasking 1.0
11233322132123132123

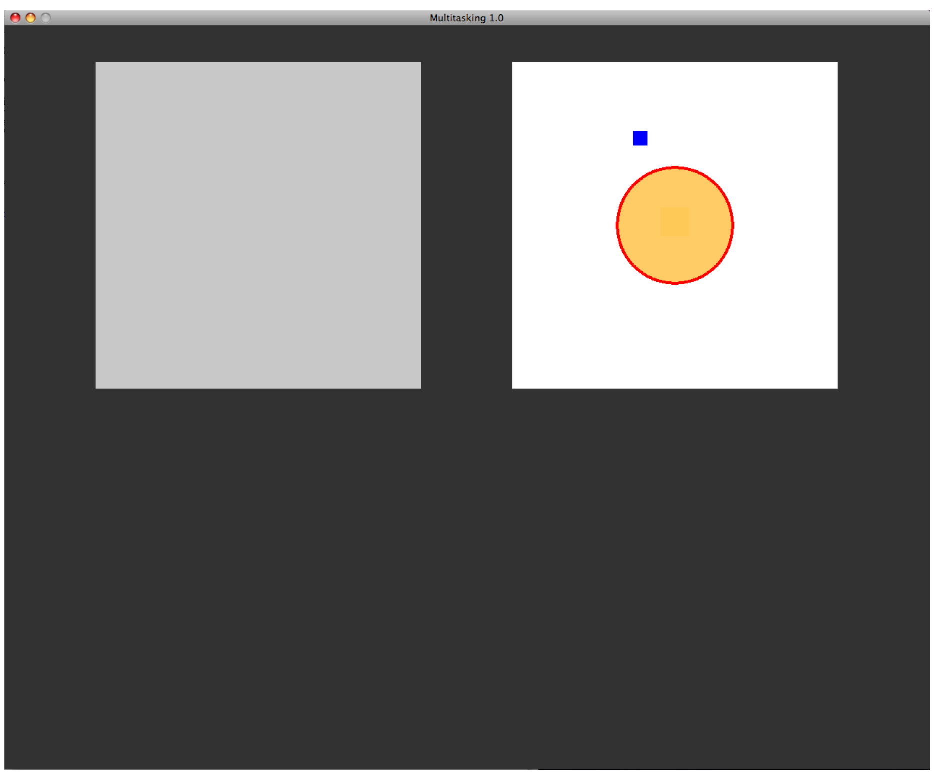
Multitasking 1.0

APPENDIX B

PPID: ______
July 2010

Please fill in the answer that most suits your situation

Sex: female / male	
Age: _____ years	
I am: left-handed / right-handed	
Nationality: ________________________	
I have a qualification in touch-typing: yes / no	
How many hours a day do you spend typing text on a computer on average? Take into account both formal (e.g., project reports) and informal (e.g., chatting, e-mail) texts. ____ hours a day	

How often do you play video games that involve a joystick?

very often 1 2 3 4 5 never

How well did you think you performed, and why?
……………………………………………………………………………………………………
……
……………………………………………………………………………
……
……………………………………………………………………………

Have you used a particular strategy/strategies during the experiment? If so, could you explain your strategy/strategies?
……
……
……
……
……
……
……………………………………………………………………………

Do you have any additional comments about the experiment?
……
……
……
……
……
……
……