

User adaptation in User-System-Interaction

Arnout R.H. Fischer

The work presented in this thesis was carried out under the auspices of the J.F. Schouten school for User-System Interaction Research.

Printing: Universiteitsdrukkerij Technische Universiteit Eindhoven

Cover realisation: Paul Verspaget

An electronic version of this thesis is available from the library of the Technische Universiteit Eindhoven at: <http://www.tue.nl/bib/ftproef.html>

© 2004 Arnout R.H. Fischer – Eindhoven – the Netherlands

CIP DATA LIBRARY TECHNISCHE UNIVERSITEIT EINDHOVEN

Fischer, Arnout R.H.

User adaptation in user-system-interaction / by Arnout Robert Hans Fischer. – Eindhoven : Technische Universiteit Eindhoven, 2004. – Proefschrift. –

ISBN 90-386-2038-1

NUR 778

Keywords: Reverse engineering / Self regulation / User system interaction / Human computer interaction / Action control / Heuristics / Affect

User adaptation in User-System-Interaction

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
Rector Magnificus, prof.dr. R.A. van Santen, voor een
commissie aangewezen door het Colleges voor
Promoties in het openbaar te verdedigen op
donderdag 8 januari 2004 om 16.00 uur

door

Arnout Robert Hans Fischer

geboren te Rotterdam

Dit proefschrift is goedgekeurd door de promotoren:

prof.dr. C.J.H. Midden

en

prof.dr. G.W.M. Rauterberg

Copromotor:

dr.ir. F.J.J. Blommaert

Acknowledgements

I would like to thank a number of people who made the time during my Ph.D. project stimulating and enjoyable.

In the first place I would like to thank Frans Blommaert for the many inspiring discussions, which often encouraged me to take my thoughts to a higher level. I would also like to thank Cees Midden, for his open-minded and fresh insights on the investigated ideas and the valuable comments while writing this thesis, and Matthias Rauterberg for initiating the project and keeping involved in and after the turbulent final years of the IPO. Furthermore, I would like to thank Gerard van Galen and Armin Kohlrausch for their careful review of the draft of this thesis and their useful comments. I am also grateful to Rachel Hoekstra for checking the English used in this text and suggesting improvements where applicable.

I would like to express my gratitude to all colleagues and students, too many to mention, who contributed to the great time I've had in Eindhoven and for their considerable help in achieving the results presented in this thesis, for example by participating, often with pleasure, in the reported experiments. Special thanks go to Janneke Staaks, whose experiment helped me in focussing my ideas; Wieneke Keller and Susanne Tak for finding participants and helping me gather data. Special thanks also go to Mark Houben for his help in compiling the music used for the experiments reported in chapters 4 and 5, and for advice on Matlab; Hilde Keuning for her help on the experiments reported in chapter 5 and for comments on early versions of that chapter; and Peter de Vries for reading and commenting a first draft version of the thesis. Finally I would like to mention my roommates over the years, Andrei Malchanau, Stéphanie van den Berg, and Fiona Montijn-Dorgelo, the many discussions with you considerably helped me shape my thoughts.

Contents

Chapter 1: Introduction	1
1.1 Reverse engineering	1
1.2 User-system-interaction	3
1.3 Self-regulation of user-system-interaction	4
1.4 Feedback control	7
1.4.1 Monitoring	8
1.4.2 Evaluation	10
1.4.3 Adjustments	12
1.5 Hypotheses and outline of the research	13
Chapter 2: Monitoring and evaluation of time-delay	17
2.1 Introduction	18
2.1.1 Requirements of monitoring and evaluation mechanisms	19
2.2 Experiments	21
2.2.1 Experiment 1	23
2.2.2 Methods and Materials	23
2.2.3 Results of experiment 1	28
2.2.4 Discussion of experiment 1	34
2.2.5 Experiment 2	37
2.2.6 Methods and Materials	38
2.2.7 Results of experiment 2	42
2.2.8 Discussion of experiment 2	45
2.3 General discussion of chapter 2	46

Chapter 3: Evaluation of visual search tasks	49
3.1 Introduction	49
3.2 Experiment	53
3.2.1 Methods and Materials	53
3.2.2 Results	58
3.2.3 Discussion of the experimental data	62
3.3 Simulation	63
3.3.1 Simulation of the evaluation mechanism	63
3.3.2 Different search strategies	65
3.3.3 Simulated strategy selector	65
3.3.4 Updating reference values for evaluation	69
3.4 General discussion of chapter 3	72
Chapter 4: Action adjustments in an ongoing gambling task	75
4.1 Introduction	75
4.2 Experiment	80
4.2.1 Methods and Materials	81
4.2.2 Results	86
4.2.3 Discussion of the experimental results	94
4.3 Simulation	97
4.3.1 Outcome and discussion of the simulation	99
4.4 General discussion of chapter 4	101
Chapter 5: Regulation of mouse-cursor-movements in different contexts	103
5.1 Introduction	103
5.1.1 Motor action regulation	106
5.2 Experiments	109
5.2.1 Methods and materials for experiment 1	109
5.2.2 Results of experiment 1	114
5.2.3 Discussion of experiment 1	124
5.2.3 Experiment 2	128
5.2.4 Methods and materials experiment 2	129

5.2.6	Results of experiment 2	133
5.2.7	Comparison between explicit and implicit task regulation	137
5.2.8	Discussion of experiment 2	139
5.3	General discussion of chapter 5	141
	Chapter 6: General discussion	143
	References	151
	Appendices	159
	Appendix A: Music files	159
	Appendix B: PANAS scale	160
	Samenvatting	163
	Summary	165
	Curriculum Vitae	167

Chapter 1: Introduction

When I first started writing this thesis, I had to look at the keys of my computer keyboard to get the correct letters on the screen. Although I did not put explicit effort in learning how to type, over time my typing skills improved to the point where I no longer have to watch my fingers but only have to look at the screen. By the time I had finished writing, at least when I was alert, my typing skill had improved so much that I could even look out of the window. However, after a long day working this skill diminished and I once again needed to pay more attention to the task at hand. Apparently, in different situations and with different levels of skill, I execute the same user-system-interaction (USI) tasks in different ways. This example suggests that besides consciously carrying out tasks, the interaction process that accomplishes these tasks is controlled as well. Such control processes are often called self-regulatory (Carver and Scheier, 1998). A self-regulatory process is a process that, if necessary, guides goal-directed activities over time and through changing contexts (Karloly, 1993). Investigation of the mechanism of self-regulation of user-behaviour in the context of user-system-interaction, should lead to better understanding about the interaction properties leading to successful user-system-interaction. Practically, this allows the design of interfaces that facilitate users to achieve better interaction. Therefore, the central theme of this thesis is to try to understand and model the self-regulatory mechanism for user-system-interaction.

1.1 Reverse engineering

To study the mechanism of self-regulation, I use a form of analysis-through-synthesis called the reverse engineering approach (Dennett, 1981; Marr, 1982;

Newell, 1990), which is to synthesise a system that performs the same tasks as the real-life tasks that are studied. The rationale behind the reverse engineering approach is that synthesising a system gives often more insight into mechanisms that determine the behaviour of the system, compared with insights that follow the attempt to deduce these mechanisms based on observed behaviour alone (Braitenberg, 1984).

In the reverse engineering approach the intentional stance is adopted, which is a strategy to interpret the behaviour of an entity *'as if'* it is conducted by a rational agent (Dennett, 1999). This interpretation is possible because adequate behaviour for the completion of a certain task will be the same, independent of the origin or structure of the system that has led to that behaviour.

The synthesis of the studied system is usually done at three levels. At the highest 'goal' or 'knowledge' level the operational goals of the system are formulated. The definition of these goals is generic and can be applied to all similar problems (e.g., the hard-disk of a computer should be accessible, so a disk operating system should exist). A second level is an 'algorithmic', 'design', or 'strategic' level. At this level functions are specified that result in achieving the goals. The solutions at this level are still somewhat general, although a more specific solution emerges (e.g., interaction with a user should be facilitated so either a text based user interface such as MS DOS or a graphical user interface, such as Microsoft Windows or Mac OS is defined). At the most concrete specification level these functions are uniquely defined at a 'physical' or 'implementation' level. A distinction between two identical functional solutions of the same problem can only be made at this level (e.g., the electronics scheme of the AMD Athlon compared with that of the Intel Pentium processor). The system is sequentially synthesized and tested at each of these levels.

After synthesising a system, the researcher determines to what extent the synthesised system describes the studied behaviour, by comparing the behaviour of the synthesised system with real-life behaviour. To apply this approach it is assumed (1) that the system can be independently specified at the different levels and (2) that it is possible to specify the functions of the investigated behavioural system (Dennett, 1994).

In the following sections the reverse engineering approach is used to specify a self-regulatory mechanism for user-system-interaction. Starting with the assumed goals of this mechanism, I will specify regulation of interaction up to a point at which functions are described that can be compared with psychological functions. I will limit the specification of the system to the goal and strategic levels and postpone the detailed implementation of the psychological phenomena. In the discussion, I will address how to specify possible implementations of the control mechanism.

1.2 User-system-interaction

The user-system-interaction process is specified as a sequence of information processing cycles to describe the control of user-system-interaction, (Blommaert and Janssen, 1999). For the research in this thesis, I specify this cycle with four processors. As the focus of the research is on the user side, I define only a single processor for the application. On the user side I use the three human processors of the total behavioural system that were defined by Newell (1990): a perception processor, a cognition processor, and an action processor (figure 1.1). Ongoing interaction is modelled as a sequence of these interaction cycles. This sequence is repeated until a target is met.

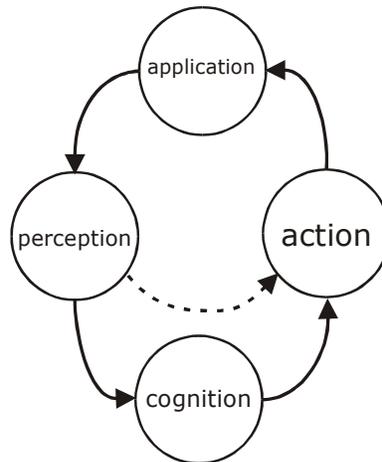


Figure 1.1: User-system-interaction as an information-processing cycle with a perception, cognition, action, and an application processor.

In human-computer-interaction research, interaction is often described as a multi-level architecture (e.g., Brinkman, 2003; Newell, 1990; Norman, 1988; Rasmussen, 1983) that allow the specification of targets and goals, as well as the strategies that are needed to achieve them. In the description of the interaction process the three levels of interpretation of an intentional system are distinguished (Newell, 1990). At the highest level the goals of the interaction process are specified. Strategies and action plans aimed at achieving these goals are defined in the middle levels. These strategies are implemented as basic interaction processes or action sequences that are executed at the lowest interaction level, i.e., that of action executions (Norman, 1984). When reading a book, for example, the goal of the perception process is to gather the relevant information. The strategy to achieve this is to start at the top left corner of the page and follow the lines. The implementation is to run a motor program for the eyes.

The interpretation of user-system-interaction in three levels implies that user-system-interaction can be controlled at these three levels. An independent control system has to be specified to independently regulate the goals, strategies, and action-sequences (Dennett, 1994).

The research in this thesis focuses on understanding the control of the action level of interaction. Once an understanding of the action level of interactions is achieved, it will be possible to extend the understanding of control of interaction to the higher levels in future research. However, the settings at the action execution level are determined by goals and strategies. Therefore, to understand the control of the action execution level in the context of the whole of the interaction, knowledge of the relationships between the levels is needed. Occasional sidesteps to the other levels will be made to acquire this knowledge.

1.3 Self-regulation of user-system-interaction

Different control mechanisms are defined for the separate levels to regulate interaction, based on the assumption that the levels are independent (Dennett, 1994). In the adopted approach of describing three computational levels, it is assumed that the control mechanisms can be defined as rational agents at the intentional level. However, a strict definition of rationality would demand that the control mechanism 'knows' all the consequences of its behaviour, which would require logical omniscience (Cherniak, 1999). To enable the study of real

phenomena, previous researchers have proposed less stringent definitions of rationality, such as bounded rationality (Simon, 1990). When using the intentional stance, a pragmatic approach to rationalism is usually adopted, namely that rationality means something like the best possible or optimal weighing of interaction process parameters (Dennett, 1981). In this thesis rationality or optimality refers to a ‘realistic’ type of rationality rather than to ‘perfect’ rationality.

Using a realistic definition of rationality, the goals of the synthesised system for self-regulation of user-system-interaction can be defined. The main goals are to find an optimal, or most rational, sequence of interaction cycles and if necessary to change properties of the interaction process so that they lead towards that optimal interaction. To do this observed behaviour is treated as the consequence of a self-regulatory mechanism (Carver and Scheier, 1998).

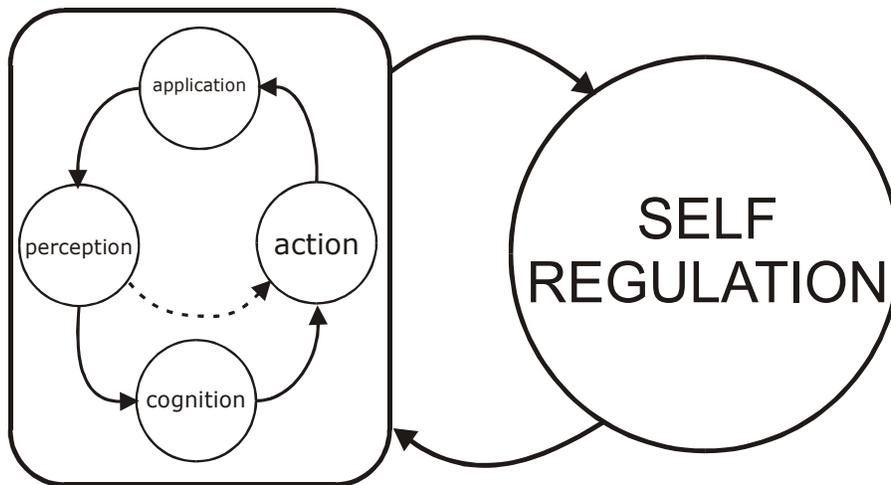


Figure 1.2: A self-regulatory control system. The self-regulatory mechanism continuously improves the user-system-interaction process.

Self-regulatory mechanisms have been studied as generic control mechanisms in the regulation of different types of processes such as cybernetics (Powers, 1973), behavioural science (Carver and Scheier, 1998; Karoly, 1993), management and manufacturing control (Kotler, 1991; Van der Aalst and Van Hee, 1997), and artificial intelligence (Sloman, 1999). A self-regulatory mechanism for user-system-

interaction receives information about the status of the interaction process and adjusts the interaction process towards an optimum (figure 1.2). In practice this means that the self-regulatory mechanism stabilises optimal interactions and initiates a change towards a better interaction in sub-optimal interactions.

In adopting the reverse engineering approach it was assumed that the synthesised system is optimising the interaction. To apply this to human behaviour it is necessary to assume that humans do indeed optimise behaviour. Self-regulatory optimisation is inherent to the way living beings have evolved (Carver and Scheier, 1998). Evolutionary theorists state that more optimal behaviour should be interpreted as behaviour that leads to the creature's higher evolutionary fitness (Dennett, 1995), at least as long as a higher fitness of the creature will benefit its genes (Dawkins, 1976). Creatures with any kind of guidance mechanism that regulates behaviour so as to improve its chances of survival, have higher fitness than creatures that do not have such a guidance mechanism. Therefore, creatures with guidance mechanisms will occupy a larger proportion of the population over the generations than similar creatures without a guidance mechanism. Mechanisms will evolve that regulate behaviour towards better achievement of survival goals. For example, in foraging behaviour a guidance system will evolve to aid predators in catching their prey. Such a system will ensure that foxes effectively hunt rabbits instead of ineffectively trying to take down elephants. Interpreted in process terms, mechanisms will evolve to guide behaviour to be effective for survival. Once effectiveness is achieved, those creatures that hunt the most nutritious prey will get fed the best, and therefore have a higher fitness. In process terms this means that mechanisms will evolve resulting in both effective and efficient behaviour. The fox that chases ladybirds will be worse off than its sibling who catches the same number of rabbits. The same argument also explains the evolution of multi-functionality in behaviour. Rabbit-hunting foxes that possess a mechanism that allows them to notice lions and change ongoing behaviour based on this information, have an obvious survival advantage. In order to evolve, the guidance mechanism of multifunctional behaviour must guarantee that multi-functional behaviour is at least as good as a single response. Once a guidance mechanism has evolved, the creature with the guidance mechanism that improves the behaviour of the creature the most has the evolutionary advantage. Multifunctional creatures will therefore evolve self-

regulatory mechanisms that optimise complex multi-functional behaviour. Where a single-purpose creature can suffice with inborn knowledge of optimal behaviour, truly multifunctional creatures, of which humans are the most explicit example (Dennett, 1995), cannot. The self-regulatory mechanism for multifunctional behaviour can only be successful in optimising behaviour, if it is able to figure out the optimal interaction for a wide range of situations.

However, the definition of optimality used above is still hopelessly fuzzy. To be able to test whether behaviour is optimised, a concept of optimal interaction is needed. The interpretation of the evolved control mechanisms in process terms allows the provisional interpretation of optimality in terms of effectiveness and efficiency. These are the same measures that determine good interaction in quality evaluation norms for software applications (ISO 9421 - 11, 1997). The process interpretation also introduces a biological mechanism that regulates behaviour towards objective survival optima, which can be interpreted as the best possible cost-benefit trade-off. This objective success can be interpreted as a cost-benefit relationship in which the achieving of survival goals are the benefits and the resources spent to do this are the costs.

1.4 Feedback control

To specify a self-regulatory system for user-system-interaction, the user-system-interaction process is described as a sequence of interaction cycles. To optimise a sequence of these user-system-interaction cycles, the self-regulatory mechanism needs to receive information about the interaction, and needs a mechanism to change the interaction process for the better. The relationship between the information and useful changes of action is made by a mechanism that interprets the information, an evaluator. An obvious self-regulatory mechanism for improving ongoing processes is some kind of feedback mechanism (e.g., Miller, Pribram and Galagher, 1960). Feedback control (figure 1.3) consists of three components: (1) monitoring, (2) evaluation, and (3) process adjustment (see Carver and Scheier, 1998, among others).

Up to this point the feedback mechanism is defined as a control loop that can be applied to a wide range of processes. In the following sections the functional aspects of these mechanisms are specified for user-system-interaction. A separate feedback mechanism has to be defined in order to optimise that specific level for

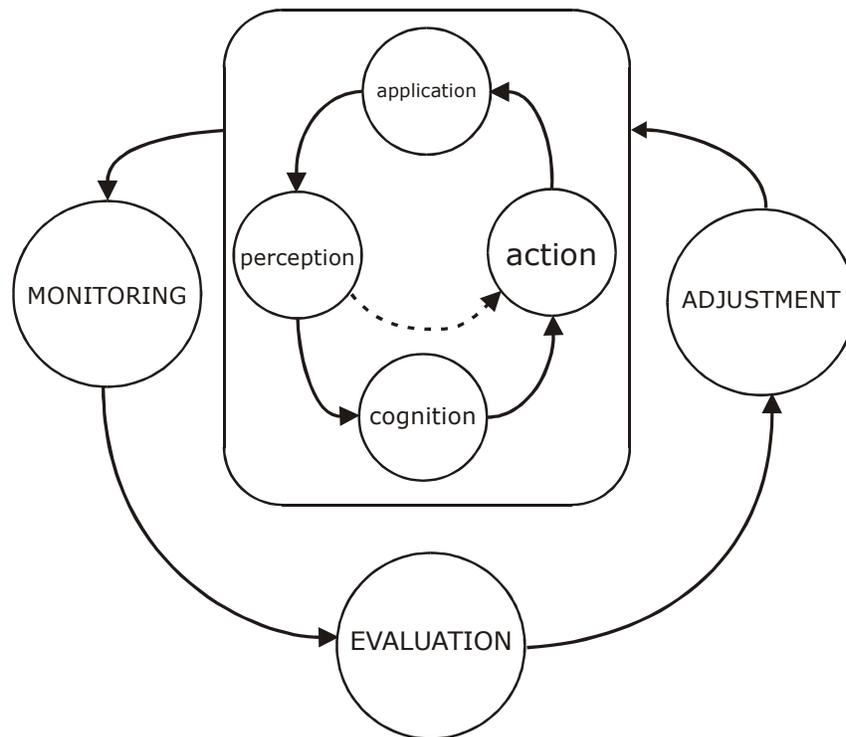


Figure 1.3: A feedback loop controlling user-system-interaction.

each of the different levels of the user-system-interaction, (goal, strategy and action execution level of interaction). In this thesis, I will focus on the self-regulatory system for the action execution level.

1.4.1 Monitoring

In user-system-interaction the goal of the monitoring mechanism is to record the status of the interaction process and pass this information on to the evaluator. Monitoring is executed at each of the three description levels of the interaction process. Information about the achievement of the goals, strategy and implementation is independently recorded and passed on to independent evaluators (Dennett, 1994).

The first strategy of the monitoring mechanism should be to record only those elements of the interaction that are relevant for the optimisation of the interaction

process. In the interpretation of optimal interaction as the best possible combination of effectiveness and efficiency, these relevant elements might be: spent time, used force, and distance to the goal. To be able to record the relevant interaction properties the monitoring function first has to learn what the relevant properties for interaction control are. To do this, a control process for the monitor must exist that receives and uses information about the success of the selection of interaction elements. It is likely that the relevant elements of interaction will not be determined every time they are needed, but that information from memory will be used. The specified mechanism that selects the relevant properties of interaction can start monitoring interaction from three initial states of knowledge about the interaction process: (1) there is no knowledge at all; a completely new interaction process is encountered and the monitor has to learn from scratch what properties to record, (2) there is previous knowledge about similar situations; the interaction is somewhat like a known interaction and the monitor starts out with recording the properties for that interaction. After this initial phase, fine-tuning for that particular interaction follows, and (3) the monitor knows exactly what to record; the interaction is well practised.

The second goal of monitoring is to pass the recorded information to the evaluator. To do this the monitored interaction-process properties have to be transformed into a format that the evaluator can use. To achieve this, a transformation function is required. For example, when working on a very tight schedule it is useful to focus on differences of only a few seconds but when considering actions for next week, hours suffice (Katar, Britton, and Nehaniv, 2000). In this example the monitor should therefore pass on detailed information about time spans in the very near future and should compress larger time spans in the more distant future.

The obvious psychological process that is available as monitor is the perception process. There is evidence that perception processes execute the first goal of the monitor, i.e., the selection of relevant information (Nijenhuis and Blommaert, 1997). Perception processes also execute the second goal of the monitor, which is the transformation of physical values into information that can be evaluated (Janssen and Blommaert, 2000; Johnston, 1999).

1.4.2 Evaluation

The evaluator receives information from the monitor and compares it with an optimal interaction state. If this comparison shows it is necessary, the evaluation mechanism will then pass on commands for the adjustment of the interaction process to the adjustment mechanism. Evaluation is also executed at the three levels of description of the interaction process.

To compare the interaction with an optimal interaction state, the evaluator first has to know or learn the optimal state of interaction. In a similar way to the monitoring mechanism, the evaluator either starts from scratch, starts from a similar situation or has precise information about the specific interaction process. In order to learn the optimal state, the specification of the optimal state of interaction also has to be controlled. This can be achieved by continuously updating the goal state, based on experience (Mellers, Schwarz and Ritov, 1999; Mellers, 2000).

The commands that are passed to the adjustment mechanism should consist of information how much the interaction is besides the goal state, which indicates how much the interaction should be adjusted. The distance to the optimal interaction, as determined by the evaluator, can be expressed in an overall value for the quality of the interaction, expressed as satisfaction or pleasure (Cabanac, 1992). Information about the specific interaction element at which the difference between optimal and actual interaction occurs, should also be sent to the action adjustment function. This information will enable the interaction to be functionally adjusted.

To study these processes, a possible implementation of the feedback loop will be investigated in more detail. In this implementation, the monitoring function passes on separate subjective representations for each of the monitored elements to the evaluator. The evaluator's knowledge of the optimal interaction state consists of a set of one-dimensional elements that together describe anticipated behaviour. The distance between the monitored interaction elements and the goal state is arrived at by simply calculating the difference between each monitored element and its anticipated state. To do this it is important that the monitor has already transformed the signal to account for the relative importance of each signal. The outcomes of these comparisons can be specified as affective values that are referred

to as hedonic tones (Johnston, 1999). Hedonic tones as defined in this thesis have the same units, and can be either positive or negative. If the ongoing interaction property is better than anticipated, the hedonic tone is positive; if it is worse than anticipated, the hedonic tone is negative. The different hedonic tones are aggregated into a single evaluation value, referred to as pleasure (Cabanac, 1992). This evaluation value is either negative if the interaction is less than optimal, or positive if the interaction is at least as good as the anticipated optimum. Negative hedonic tones indicate which element of the interaction could be improved. If necessary, a signal can then be given to the action mechanism exactly where to apply more effort. Only then will a feeling about the adequacy of the situation become aware. With this process view of emergent emotions, the proposed self-regulatory system for interaction control resembles the functional account of emotions introduced by Frijda (1986).

The proposed implementation can be executed without conscious control, thus following contemporary insights that in addition to cognitive elements, affective elements are also important in behaviour control (e.g., Chen and Chaiken, 1999; Petty and Wegener, 1999; Sloman, 1996). This is specifically the case when immediate action is required (LeDoux, 1996) or when people lack sufficient information (Damasio, 1994). The accounts of Simon (1967) and later cognitive emotion research (Oatley and Johnson-Laird, 1987) specify emotions that can be interpreted as qualitative signals for the regulation of behaviour. Four of these can be related straightforward to the evaluation of interaction processes. (1) Good interaction leads to the experience of happiness or satisfaction which signals continuation of the ongoing process. (2) Interaction that is not, or not yet, good enough leads to frustration and anger, which in turn leads to increased effort and aggression to overcome obstacles. (3) Interaction for which optimisation is needed but cannot be achieved leads to sadness, which signals abandoning the task at hand to free resources for new goals. (4) Occurrence of emergencies evokes a fear response that (temporarily) frees all resources to remove the threat (Oatley and Jenkins, 1996). The self-regulatory mechanism specified so far deals with the first two occasions, which will be studied. The third and fourth case involve changing the operational goals of the interaction process and relate to higher levels of interaction control; these are outside the scope of this thesis.

1.4.3 Adjustments

All three levels of the interaction process are adjusted, based on the signals from the evaluation. The evaluation gives a signal and a direction of adjustment (e.g., more effort or continue as before). The action adjustment mechanism knows or learns what the relevant actions are and a useful change is initiated (Wolpert, 1997). The information from the evaluator can be used to initiate any form of action adjustment, such as motor action, cognitive effort or perceptual attention.

All the mechanisms of the feedback control (monitoring, evaluation, adjusting actions) have been specified as intentional systems that know or learn what elements of interaction are of importance. To apply the relevant elements in the feedback control, the specified mechanisms are assumed to gather the relevant values for these element from a memory that is modelled one step removed from the actual feedback mechanisms (figure 1.4).

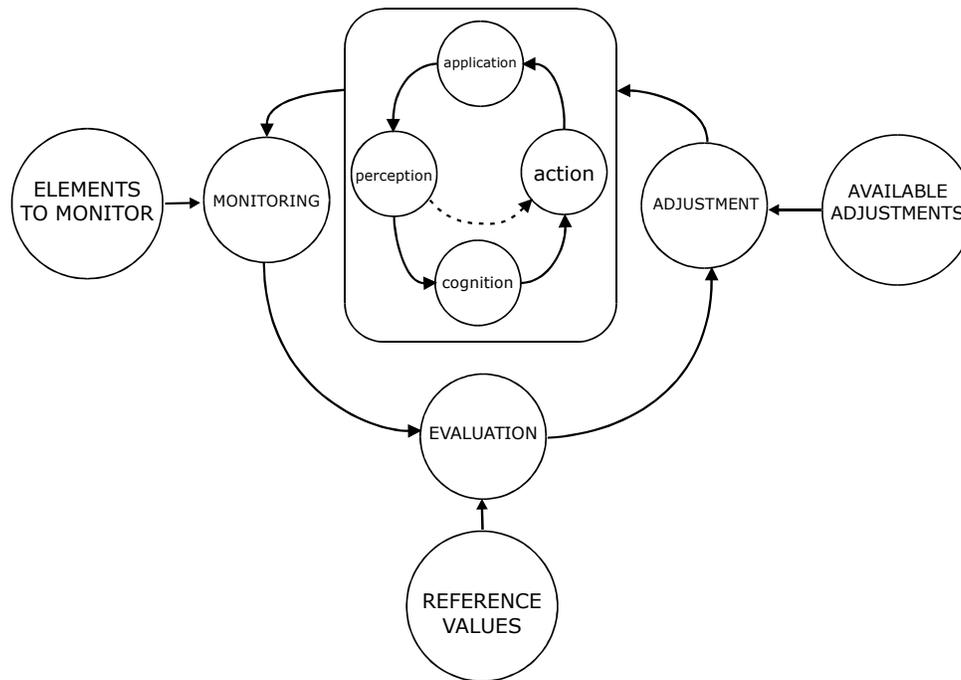


Figure 1.4: The information that the mechanisms of the feedback loop need to 'know'.

1.5 Hypotheses and outline of the research

Given this broad specification of a self-regulatory guidance mechanism for the user-system-interaction process in terms of psychological functions, specific hypotheses can be formulated based on the behaviour of the synthesised control system. Behaviour that is observed in empirical research should coincide with the hypothesised behaviour of the synthesised system. To achieve this, every time the synthesised regulatory system makes an unambiguous statement about observable behaviour, a match between the behaviour of the self-regulatory mechanism and real-life behaviour is hypothesised.

In section 1.3 the goal of the self-regulatory mechanism has been specified. In doing so it was assumed that there is indeed a functional self-regulatory mechanism that regulates interaction in a way that changes it towards a stable pattern by approaching a more optimal interaction state. To achieve this, the self-regulatory mechanism should be able to determine an optimal interaction state, and to adjust actions towards this optimum.

***Hypothesis 1:** User actions are adjusted towards an optimum.*

In the specification of the self-regulatory mechanism in section 1.4 three main functions of feedback control were introduced: monitoring, evaluation and action adjustment. In human behaviour these functions are provided by perception, cognitive or affective evaluation and changes in cognitive or motor behaviour.

***Hypothesis 2:** Perception processes provide the necessary monitoring of interaction.*

***Hypothesis 3:** Evaluation determines the difference between the optimal and the actual interaction.*

***Hypothesis 4:** Adjustments of the applied human resources result in a smaller difference between the optimal and the current interaction.*

Each of these hypotheses follows from the synthesised system and allows more detailed experimental hypotheses to be proposed, which will be formulated and tested in four empirical studies. The experimental studies in this thesis focus on the control of the physical action level of the interaction, although sidesteps to look at the control of the strategy and the goal level are made in order to find evidence for interaction control in its context.

To test the separate components of the feedback control (monitoring, evaluation, adjustment), predictions are made about the effects of changes in the efficiency and effectiveness of the interaction process on the feedback mechanisms. These predictions are used to test hypotheses 2, 3, and 4 in a number of experiments, which will be reported sequentially. This series of experiments will focus on each of the sub-processes of the interaction cycle (system, perception, cognition, action), to support the general applicability of the feedback control mechanism over the whole interaction cycle (section 1.4). All experiments are interpreted in relation to hypothesis 1, which concerns the ability of the interaction control system to increase the adequacy of interaction in a more general way. More specific evidence for hypothesis 1 can be found in the later experiments that focus on observable behaviour.

To explore a possible implementation of the control system, the function of affect as a heuristic in the synthesised self-regulatory system is investigated. In order to gain more detailed understanding of the observed behaviour, I also investigate simulations of possible underlying mechanisms for some of the experimental tasks.

To conclude this introduction, a more detailed description of the contents of the following chapters is given.

In chapter 2, two experiments will be reported. The aims of these experiments are to find evidence for the idea that perception processes take care of monitoring (hypothesis 2). The experiments will be executed by asking participants to report their perception of time-delay, an interaction element that is related to interaction adequacy. It is hypothesised that the monitor can record the elements of the interaction cycle that are relevant to notice a decrease in interaction efficiency. In this study, the manipulation of time delay will cause the decrease in interaction efficiency. A second aim of the experiments is to find evidence that decrease in interaction efficiency is evaluated as negative (hypothesis 3), which is measured by asking participants to report satisfaction.

In chapter 3, an experiment will be reported that was conducted to find evidence that the evaluation mechanism determines the adequacy of ongoing interaction compared to anticipated optimal interaction (hypothesis 3), and that it can adapt the anticipation of optimal interaction when the original estimate proves to be

incorrect. To do this, both noise and the amount of redundant information in a visual search task are varied, in which redundant information introduces a better than expected adequacy for large search fields. A simulation for the evaluation mechanism will be developed based on the results.

In chapter 4, an experiment will be presented that was used to investigate whether participants adjust interaction when information about the interaction is accumulated (hypotheses 1 and 4). This experiment focused on the strategy of interaction, by asking participants to find a profitable strategy for a card game. This sidestep to look at the strategy level gives a better understanding of the emergence of behaviour. The role of affect as a heuristic in associative interaction control is studied by investigating changes in the interaction-adjustment of participants following the manipulation of the participant's mood. The card-game experiment will be simulated, using the mechanisms developed in chapter 3, to investigate the mechanisms behind the increasing knowledge of a new interaction process.

In chapter 5, two experiments will be presented in which adjustments at the action execution level are studied. To do so, interaction targets of mouse cursor movements are varied. The relationship between observable adjustments of actions and optimal interaction will be studied (hypotheses 1 and 4). If the self-regulatory system optimises interaction subconsciously, users should act efficiently even when efficiency is not the explicit task. To test this idea, in the last experiment, participants were not told that cursor movements were recorded. To gather broader evidence for the influence of affect on interaction control, the effect of mood manipulation on mouse cursor movement will also be studied.

In the final chapter I will discuss the results of the studies in relation to the research hypotheses posed in this introduction. Practical and theoretical implications of self-regulation to predict user adaptation in user-system-interaction will be given. The general discussion will conclude with some limitations of the research approach and open issues for the future.

Chapter 2: Monitoring and evaluation of time-delay

Abstract¹

The assumption that users act as if their behaviour is controlled by a self-regulatory system (Carver and Scheier, 1998) is applied to user-system interaction. The self-regulatory system is specified as a feedback control mechanism consisting of monitoring, evaluation and action adjustment mechanisms. Two experiments were carried out to test whether perception processes take care of monitoring, and whether evaluation, based on a comparison of the results of monitoring with reference values, gives an indicator of the adequacy of the interaction. In these experiments a typical system property, time-delay, was manipulated. In a first experiment, the relationship between the estimated task duration and the actual time-to-task-completion were studied. These findings were replicated in a second experiment. The consistent findings support the idea that perception processes, interpreted as a monitoring mechanism, can keep track of physical task-time. A monotonic, positive relationship was found between the time-to-task-completion and the user satisfaction, which supports the assumption that objective efficiency influences the evaluation of interaction. A linear relationship between the estimated duration and user satisfaction indicates that the monitoring mechanism transforms physical signals into a format that can be used in the evaluation process. An increased variability in time-delay negatively influenced user satisfaction. In the second experiment the variability of time-delay is investigated further by offering specific patterns of variations of time-delay to the participants. With the same variability, some patterns received higher satisfaction scores, indicating that patterns influence satisfaction by changing the internal reference frame of the evaluation mechanism.

¹ Part of this research has been published as Fischer and Blommaert, 2001.

2.1 Introduction

Most regular users of the Internet will notice that the response time of servers can differ considerably. Long delays can be annoying (Dellaert and Kahn, 1999). The user may avoid very slow sites when next using the Internet. To make such decisions, a user has to notice when interaction is not efficient so that he or she can act to avoid inefficiency in the future. If this occurs, the interaction process is optimised in some way. Such an optimisation process can be described as the result of a self-regulatory control system that controls interaction through a feedback control mechanism (Carver and Scheier, 1998).

To better understand this feedback control mechanism, I adopt the top-down research approach of reverse engineering (Dennett, 1994). Following this approach, I synthesise a system that optimises interaction. After synthesising the system, the behaviour of this system is compared with the behaviour of participants in an experimental task. The synthesised system is usually defined on three levels: a goal level, a strategy level, and an implementation level (Marr, 1982). In this thesis, I define the goals and strategies of a self-regulatory system for the interaction between users and interactive applications. The goal of the self-regulatory system is to optimise the interaction between users and such applications. This optimisation is specified as being carried out by a feedback mechanism that monitors, evaluates and adjusts the interaction process (figure 2.1).

The monitoring mechanism selects the elements of the physical world that are relevant for the optimisation of interaction (Nijenhuis and Blommaert, 1997). The selected elements are perceived and transformed into a subjective representation of the state of the world, which is used as the input for the evaluation process. The evaluation mechanism determines the adequacy of the interaction by comparing the monitored interaction elements with reference values. The final satisfaction value that is generated by the evaluation mechanism can be interpreted as a combination of quantitative and qualitative information about the adequacy of the interaction. In this model, the quantitative information is interpreted as hedonic tones following Johnston (1999). A hedonic tone is negative when the evaluated interaction is worse than the reference value, and it is positive when the evaluated interaction is better than the reference value. The hedonic tones are aggregated into a single value to evaluate the adequacy of interaction; this value is referred to as pleasure (Cabanac, 1992). The qualitative information should give a direction for

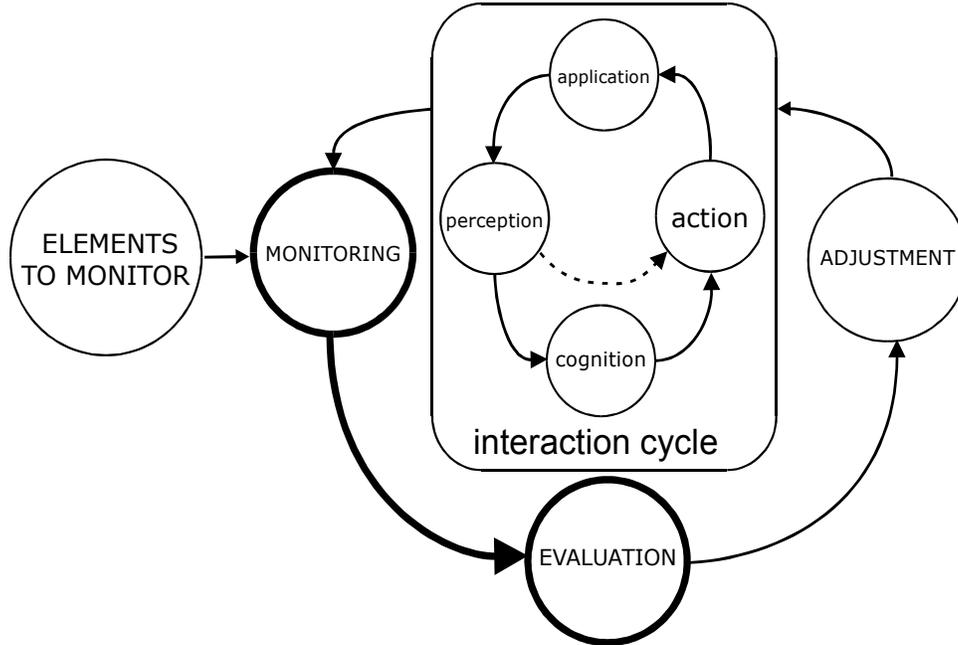


Figure 2.1: The control of interaction consisting of a monitoring, evaluation, and adjustment mechanism. The focus of this chapter is on the bold mechanisms.

the action adjustments. Such signals can be conveyed, at least partially, by emotions (Oatley and Johnson-Laird, 1987). Based on the evaluation, user actions are adjusted to improve the interaction.

In this chapter, I take a first step in validating this approach by investigating the first stages of the control of interaction, starting with the monitoring mechanism. Additionally, I investigate the evaluation mechanism, paying special attention to its relationship with the monitoring mechanism. In chapter 3 I will study the evaluation mechanism in more detail. Adjustments and the optimisation of interaction will be the focus of chapters 4 and 5.

2.1.1 Requirements of monitoring and evaluation mechanisms

The monitoring and the evaluation mechanisms are specified in more detail, to allow experimental investigation of these mechanisms. After specifying the

mechanisms in more detail, I make some predictions about the behaviour of the synthesised control system for interaction.

By adopting the intentional stance (Dennett, 1981), the monitoring mechanism is interpreted as a rational agent. The tasks of the monitoring mechanism are to select the relevant information about the interaction process and to transform this information into a format that can be used by the evaluation mechanism. The relevant information is not always the same for all interaction processes. For example, if I want to take the train to Gouda next weekend to go hiking in the wetlands, I first need the departure times of the trains. As I am trying to access the train timetables on the Internet between my other tasks, the speed of the connection is important for my interaction with the railway server. To determine whether the service is adequate, the self-regulatory system should therefore monitor the interaction time. On the other hand, once I am actually hiking, interaction time is no longer important. Hence, the monitoring mechanism should decide what elements are relevant and should therefore be recorded for different interaction processes.

The second task of the monitoring mechanism is to transform physical signals to provide the evaluation mechanism with information that it can evaluate. Once physical signals are suitably transformed, the relative importance of different occasions can be compared. Using the example of the trip to Gouda, such a transformation is useful to compare the required accuracy of the reviewed train timetables. If I were to decide to go right now, the decision of what train to take, would need to be specific within minutes. On the other hand, if I plan to go next Sunday, it does not matter if I take a train an hour earlier or later, as long as I keep Sunday morning free. The importance of a minute right now might be greater than that of an hour next week. A monitoring mechanism for time should therefore transform the physical occurrences of a phenomenon into different representations of that, in such a way that they can be compared as meaningful differences.

For the implementation of the monitoring mechanism, an obvious choice for investigation is the perception process. Not only are these the only processes that can receive information from the outside world, but perception processes are also found to provide both the functions that are part of the monitoring mechanism. The selection of relevant elements of interaction can be interpreted as the result of a perception process (Nijenhuis and Blommaert, 1997). There is also evidence that

perception processes transform physical data into subjective information according to power laws (Stevens, 1975), in a way that takes the relative importance into account (Janssen and Blommaert, 2000).

The evaluation mechanism's task is to use the information received from the monitoring mechanism to determine the adequacy of the interaction, and to give a signal to the adjustment mechanism to optimise that interaction. The evaluation mechanism knows, from experience or deliberation, which properties of the ongoing interaction process are of importance. In the example of finding a travel advice on the Internet, the shortest waiting time means that the other jobs get interrupted the least and that this site should be used in the future. On the other hand, during the hike in the wetlands, slow plodding through thick mud may be evaluated as more positive than a quick stroll on an asphalt road.

When the evaluation mechanism can compare the properties of different interaction processes directly, it can easily decide which is the better. However in daily choices which device to use, such face-to-face comparisons are not always available. Therefore, the evaluation mechanism should be able to determine the adequacy of interaction without comparing the different interaction processes (travel sites) face-to-face. If the evaluation mechanism is capable of making estimates of adequacy for an isolated interaction process, this limits the number of processes that have to be considered before finding a good enough interaction process. To achieve this, the evaluation mechanism should be able to compare an ongoing interaction to a reference or anticipated interaction, resulting in a single assessment value for the adequacy of interaction (Cabanac, 1992).

2.2 Experiments

To confirm that the proposed monitoring and evaluation mechanisms can determine the adequacy of interaction, an experiment was conducted. Adequacy of interaction depends on two elements. First, users have to achieve the goals of the ongoing interaction, otherwise no success is achieved and the interaction is not effective. Once the ongoing interaction is at least effective, a second property of interaction, efficiency, becomes important in the evaluation of the adequacy of interaction (ISO 9421-11, 1997). As a first-order effect, efficiency should be related to the efficiency of the interaction cycle, consisting of system output, user perception, cognition and action. A second-order effect on efficiency is the

efficiency of the control of interaction itself. The easier an interaction process can be regulated, the more efficient the combination of interaction and control of interaction will be.

To test these two types of efficiency, I investigated the influence of a typical aspect of the interactive application that should be monitored and evaluated. This was done by introducing time-delay to manipulate the efficiency of the interaction. Time-delay is an existing source of dissatisfaction in the evaluation of Internet services (Dellaert and Kahn, 1999) that still exists in spite of faster systems and networks (Dix, 1994). Time-delay is the time between the user's action and the output of the system; time-delay can be significant, especially for complex calculations (e.g. SPSS or MATLAB) or for distantly located, slow Internet servers. In some cases, time-delay might result in users deciding to stop the interaction or choosing another web site in the future.

The user cannot influence the system time-delay when interacting with applications. In an ongoing interaction the user will not adjust his or her actions, since this will not lead to any improvements in the interaction. The efficiency of interaction can therefore be influenced by manipulating the amount of time-delay.

When time-delay is not the same in a series of subsequent actions, it becomes harder for the control system to monitor the ongoing interaction. The less predictable the sequence of interaction cycles, the more attention a user has to invest in monitoring (Cohen, Atkin, and Hansen, 1994). We can therefore assume that the efficiency of the control mechanism can be influenced by manipulating the variability of time-delay over a sequence of interactions.

To investigate how the monitoring mechanism transforms physical variables of time-delay into subjective representations, I studied the relationship between the actual time-delay and the estimated duration of a task. Subjective duration assessment was earlier found to be relevant for the implicit evaluation of the efficiency of the task (Zakay and Shub, 1998).

The evaluation mechanism should compare the estimated duration with a reference value for duration, resulting in a representation of the efficiency of interaction as a single assessment value (Cabanac, 1992). Cabanac not only found that pleasure maximisation led to behavioural changes, but also found that participants were able to interpret physiological experience on consciously

reported scales. I therefore assume that it should be possible to measure pleasure as self-reported satisfaction.

2.2.1 Experiment 1

In the first experiment, I investigated the generation of an assessment of relevant elements of interaction (delay time) by the monitoring mechanisms. This was done by influencing system time-delay and recording how this influenced the estimated duration of the task. Second, I looked at whether the output of the monitoring mechanism could be the input of the evaluation mechanism. This was done by recording the self-reported satisfaction of an interaction process in which the system time-delay was manipulated. If the outcome of the monitoring mechanism is the input for the evaluation mechanism, a relationship between the monitoring and the evaluation mechanisms should be evident in the data. A third issue addressed in this experiment was whether and if so, how inconsistencies in sequences of interaction were monitored and evaluated by the control system. To do this I manipulated the variability of time-delay in a sequence of interaction cycles, leading to a single target.

2.2.2 Methods and Materials

Participants and design

Two groups of ten native Dutch-speaking participants, all students at Eindhoven University of Technology participated in the within-participant experiment. One group was asked to estimate task duration, in order to investigate the monitoring mechanism. The other group was asked to report their satisfaction for the same tasks, to investigate the relationship between the monitoring and the evaluation mechanism. The choice to use two groups was made, because participants should not realise that time and satisfaction were linked in this study. This linkage could also have been prevented by asking participants to scale satisfaction in a first series of trials and estimate task duration in a second series of trials. This option was not chosen in this first study to eliminate possible effects of practice. A linear relationship between the two groups was studied. If such a relationship can be found the outcome of the monitoring mechanism was the likely input for the evaluation mechanism.

The efficiency of a search task in an Internet database was manipulated by adding delay times. Each task consisted of finding a product in a hierarchical database. In each task, the participant had to make a total of five sequential selections for a category or product, after which the next category was shown. Before the next category appeared the system had an idle time in which the screen was blank. The properties of the completed sequence leading to product selection were varied in two ways. First to study the effect of changes in objective efficiency there were six different levels of average delay time ranging from 0.5 to 10.5 seconds (0.5, 2.5, 4.5, 6.5, 8.5, and 10.5). This manipulation changed the objective efficiency of interaction by inducing a longer time to achieve the same goals. Furthermore, to study the effect of inconsistency in delay time, variability was introduced in the occurrences of delay time within the same task. The variability was a fraction of the average delay time and was superimposed on the average delay time. The variability ranged from 0 (all values equal), through 0.5, 1 and 1.5 to 2 (where the longest time was twice as long as the average delay and the shortest time was instantaneous). For example, when the average delay time was 6.5 and the variability was 1, the instance of time-delay were: 3.25: $(1-1/2*1)*6.5$, 5.42: $(1-1/6*1)*6.5$, 7.8: $(1+1/6*1)*6.5$, and 9.75: $(1+1/2*1)*6.5$. The order of these instances was randomised. The full factorial combination of average time-delay and variability in time-delay makes up for the total of 30 stimuli that were presented to the participants in a randomised order.

Apparatus and task

The experiment was fully computerised. Participants were asked to find a product in a database environment, which was modelled in layout, contents, and structure, after an existing Internet database of a Dutch supermarket². The database consisted of over 7000 products and had a hierarchical tree structure of four levels followed by a confirmation screen. After each selection of a product or product-category, an instance of time-delay (Td) was generated before the next option became visible (figure 2.2). For example, one of the tasks was to find the product 'Swiss chocolate', where the participant had to choose subsequently: Snacks - Sweets - Chocolate - Swiss chocolate - Confirm Swiss chocolate. If a participant chose the incorrect branch of the database, for example, by choosing

² <http://www.ah.nl>, as it was available on the web in 2000

Biscuits instead of Chocolate after choosing Sweets, the participant had to backtrack and try again until the correct product was found.

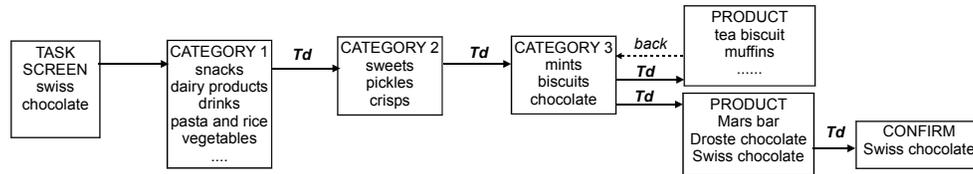


Figure 2.2: Structure of the product acquisition task.

Procedure

Participants were welcomed into a laboratory room containing a chair, a desk on which a PC with a 17" colour screen were located. The participants were randomly assigned to one of two groups. The participants were given on screen instructions, a printout of which was also available. They were then given eight practice tasks, after which they were asked to find 30 target products in the hierarchical database, as quickly and accurately as possible. Only when the correct product was found could a participant go on to the next task. After locating each product they were required to answer a single question. Participants in the first group were asked to estimate the duration of the task in seconds using an on-screen numeric pad, which is one way of acquiring the subjective duration of a task (Zakay, 1993). This group is referred to as the 'time-estimation group' in the remainder of this chapter. In the instruction given to the time-estimation group, task duration was defined as the total time from the moment that the search screen became visible until the participant confirmed the product in the final screen. Participants were given initial help to establish a better level of time estimation, to reduce variation in time estimation. This was done during the practise tasks, in which participants were shown the actual task time immediately after they had estimated task duration. This is one way of calibrating time estimation (Zakay, 1993). Participants in the second group were asked to score their satisfaction on a 10-point scale, according to the Dutch high-school grading system (where 10 is the highest mark). This group is referred to as the 'satisfaction group'. In the instruction participants were asked to take every conceivable part of the interaction into account when giving a satisfaction score. A completed experiment lasted about 50 minutes. The

participants were debriefed, thanked and paid 15 Dutch guilders (€ 6.80) for participating in the experiment.

Recorded variables and experimental hypotheses

The scores for satisfaction and estimated task duration were recorded. The objective interaction measures time-to-task-completion and the number of clicks needed to find a target (a measure for user-errors), were also recorded.

The amount of time users need to make a choice is assumed to be independent of system delay. Time-delay should therefore not influence the user time. If this is the case, longer time-delays should simply result in longer task duration, and an overall decrease in interaction efficiency.

To be able to determine the adequacy of interaction, users should notice the relevant properties of the interaction cycle. In this experiment, manipulation of time-delay resulted in variations in the time-to-task-completion, while the interaction remained the same. The main effect of this manipulation was a decrease in interaction efficiency, of which should be monitored. The time-to-task-completion in the experiment was somewhere between 10 seconds and 2 minutes. Over such a limited duration range, perception follows a monotonic relationship that can be described by a power law (Eisler and Eisler, 1992). The perception of time over long duration ranges can be interpreted as being made up of different subsequent power functions (Katar, Britton, and Nehaniv, 2000). The power relation between time-to-task-completion and estimated duration is therefore not expected to intersect the origin. Furthermore, comparisons of properties within a limited range should be fairly accurate, meaning that the relationship between estimated duration and time-to-task-completion should be a power relationship with an exponent slightly lower than 1. These requirements account for the relationship between time-to-task-completion (T_{act}) and estimated task duration (D_{est}).

Hypothesis 1: The relationship between time-to-task-completion and estimated duration is a power function which does not intersect the origin and has an exponent smaller than 1 (equation 2.1).

$$D_{\text{est}} = (T_{\text{act}} + x)^y \tag{2.1}$$

The adequacy of a task is related to the efficiency and effectiveness of the interaction. If the evaluation process gives a judgement about the different interaction processes based on adequacy, time-delay properties, which are typically related to efficiency, should have an effect on the evaluation. If satisfaction is a measure for the evaluation of efficiency in the interaction then there should at least be a monotonic relationship between time-to-task-completion and self-reported satisfaction.

Hypothesis 2: An increased time-to-task completion leads to lower satisfaction.

Time-to-task-completion can only have an effect on satisfaction through the perception of time-to-task-completion. There should therefore be a simple relationship between the estimated duration of a task and the self-reported satisfaction.

The variability of different instances of time-delay is a secondary effect of time-delay, which introduces additional attention requirements during the interaction. This additional attention results in reduced efficiency (Bobrow and Norman, 1975). If the addition of variability in interaction leads to a lower interaction efficiency, this should lead to a more negative evaluation of interaction.

Hypothesis 3: Increasing the variability of time-delay leads to lower satisfaction.

However, this effect depends on the interpretation of less efficient monitoring of the task, rather than that of a less efficient task. This effect is not an obvious indicator for task efficiency and is therefore likely to be considerably smaller than the effects of actual changes in time-delay. The evaluation generates a satisfaction value, which should allow the self-regulatory system to form a preference for the best interaction process. In practice this means that users will use the interface with the highest satisfaction score more often.

Finally, to investigate the effect of the effectiveness of the interaction, the influence of number of times users had to backtrack from an earlier choice on the evaluation mechanism is studied. The effectiveness of interaction is essential for the adequacy of an interaction process. Effectiveness was not manipulated in the laboratory-experiment however. To explore the influence of the lack of effectiveness on user satisfaction the cases with user errors are studied.

Hypothesis 4: User errors are a sign of lower effectiveness and should result in lower satisfaction.

2.2.3 Results of experiment 1

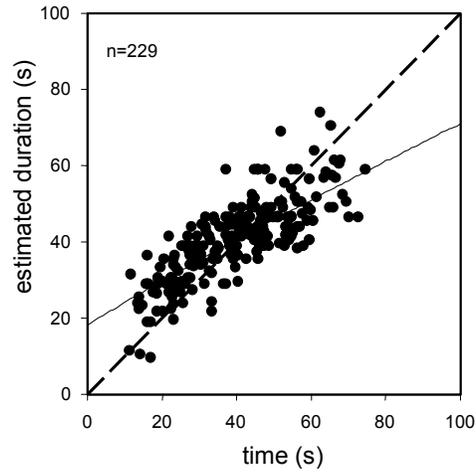
If the manipulations only influence the efficiency of interaction as measured by time-to-task-completion because of the additional task time generated by the application, there should be no difference in the time the participants take to pick an option. Therefore the relationship between time-delay and user time per interaction cycle was investigated. Time-delay did not influence the actions of the participants, neither for the magnitude of time-delay, $F(5,462)=1.63$, $p=0.15$, or for the variability of time-delay $F(4,463)=1.69$, $p=0.15$. This confirms that the user parts of the interaction are not adjusted when the delay time is longer, at least not as can be derived from observed interaction time.

In the time-estimation group, 71 of the 300 trials contained one or more user errors. In the satisfaction group, 61 out of the 300 trials contained one or more user errors. When user errors occur, the user time per cycle becomes longer than when no errors occur, $t(174.0)=5.61$, $p<0.01$, which indicates that users require longer interaction time when confronted with errors. To study the effect of the manipulations independently of these changes the cases in which no errors occurred were investigated separately. The effect of user errors will be dealt with in the final part of the results of this experiment. The z-scores of each participant's satisfaction scores were calculated, which means that the mean was set to 0 and the standard deviation to 1.

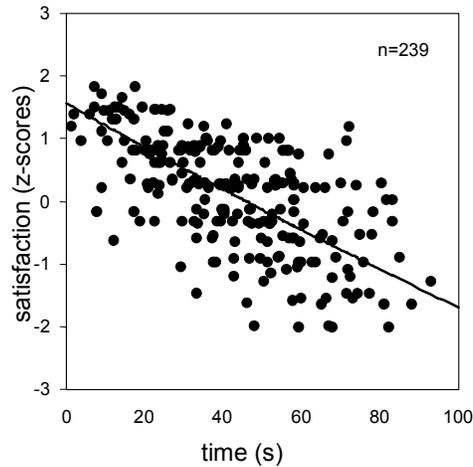
Effects of physical task time

The recorded values for the time-to-task-completion were between 10 and 80 seconds. The recorded values for time-to-task-completion were corrected for differences between participants. Two parameters of the power function actual time (T_{act}) and estimated duration (D_{est}) were determined (figure 2.3a; equation 2.2; $R^2=0.63$, $p<0.01$). The constant 0.88 results in a practically linear transformation of each additional second of objective time into perceived duration. The constant 27.1 to the power 0.88, makes up for the positive intercept at 18.2 seconds.

$$D_{\text{est}} = (T_{\text{act}} + 27.1)^{0.88} \quad (2.2)$$



a) effect of time on estimated duration



b) effect of time on satisfaction

Figure 2.3: Influence of time-to-task-completion on estimated duration and satisfaction. a) effect on estimated duration, the dashed line is a 1 to 1 relationship and the continuous line is the power function (equation 2.2). b) effect on user satisfaction a linear transformation of the power function shown in a) is used to describe satisfaction (equation 2.4)

Using the estimated duration (equation 2.2) as a linear regression term to predict scaled user satisfaction the relations between the estimated duration (D_{est}), time-to-task-completion (T_{act}) and satisfaction (S) could be determined (equations 2.3 and 2.4).

$$S = 2.7 - 0.062D_{\text{est}} \quad (2.3)$$

$$S = 2.7 - 0.062(T_{\text{act}} + 27.1)^{0.88} \quad (2.4)$$

Equation 2.4 gives a reasonable prediction of the actual relationship between physical time and the z-scores of satisfaction (figure 2.3b; $R^2=0.41$; $p<0.01$).

Effects of time-delay

The discrete effects of the manipulations on the estimated duration and satisfaction were studied separately to investigate the effect of the manipulations. The magnitude of time-delay had the greatest influence (on estimated duration: $R^2=0.49$, $p<0.01$ and on satisfaction $R^2=0.29$, $p<0.01$). The variability in time-delay

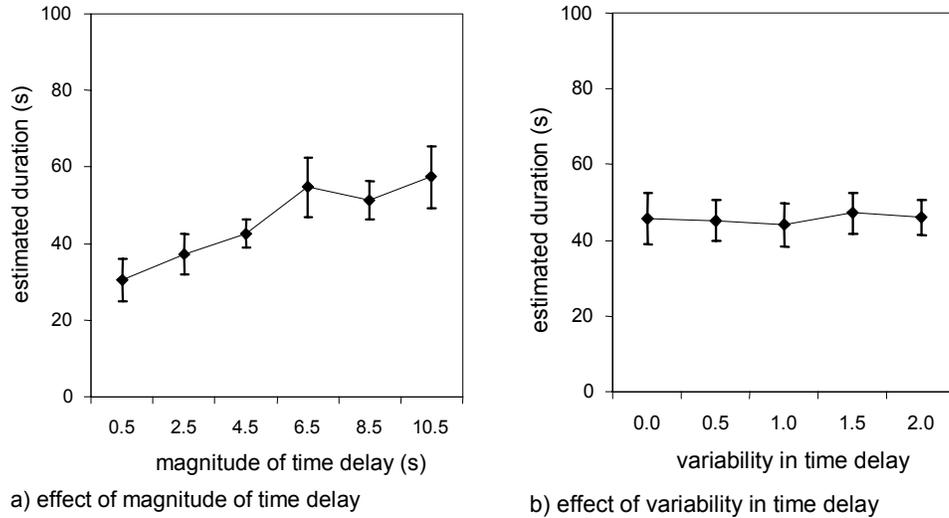


Figure 2.4: Effects on estimated duration of: a) time-delay, b) variability in time-delay for error free cases ($n=229$). The error bars represent significant differences at the 0.05 level.

had only a very small effect on satisfaction ($R^2=0.03$, $p=0.01$). Variability in time-delay showed no effect on estimated time ($p=0.85$; figure 2.4b). Based on these findings the separate effects of time-delay and variability of time-delay on estimated duration and satisfaction were investigated.

The influence of the average magnitude of time-delay (Td_M) on the estimated duration (D_{est}) was investigated in more detail (figure 2.4a; equation 2.5).

$$D_{est} = 28.7 + 0.55 \sum Td_M \quad (2.5)$$

The constant 28.7 accounts for the subjective duration of the total user time. The constant 0.55, multiplied by the sum of the delay time means that one second of delay time results in an increase of 0.55 seconds in estimated task duration.

The contributions of the magnitude of time-delay (figure 2.5a), and the small additional contribution of the difference in time-delay (figure 2.5b) to satisfaction (S) were modelled as equation 2.6. The constant 1.1 represents the extrapolated score of a zero time-delay situation. This equation reveals that the maximum effect

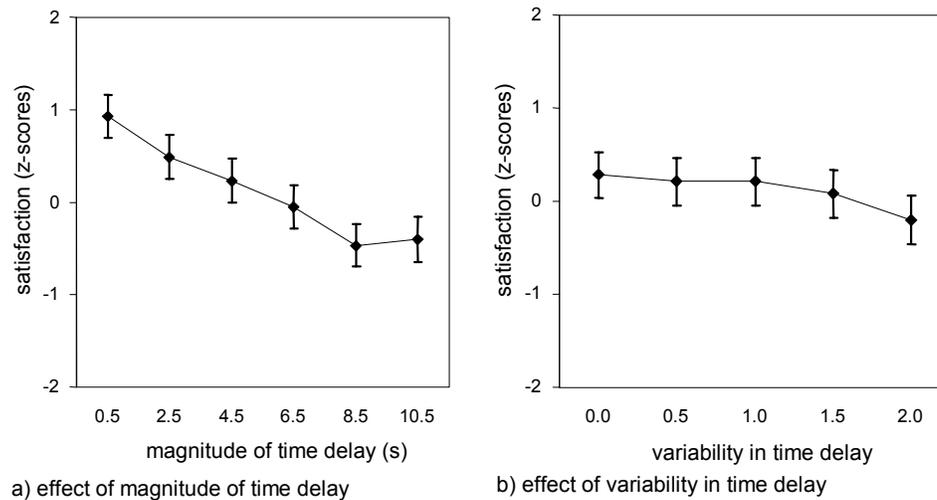


Figure 2.5: Effect on satisfaction of: a) magnitude of time-delay and b) variability in time-delay for error free cases ($n=239$). The error bars indicate significant differences at the 0.05 level.

of the variability in time-delay on satisfaction is equal to about 3 seconds of additional delay time between screens.

$$S = 1.1 - 0.035 \sum Td_M - 0.22Td_d \quad (2.6)$$

An alternative prediction for satisfaction (S) was derived by combining equations 2.5 and 2.4 (equation 2.7). If the average value for variability of time-delay (1) is taken into account, equation 2.6 is identical to equation 2.7.

$$S = 0.9 - 0.035 \sum Td_M \quad (2.7)$$

The task of the evaluation mechanism is to signal which interface should be selected. To investigate to what amount the objectively better interaction is preferred over a worse interaction, I compared differences in rank order of the estimated duration and scaled satisfaction scores. In doing so, the combined effects of the magnitude of time-delay and the variability of time-delay were investigated, by connecting points with the same rank numbers (figure 2.6). The almost horizontal lines in this representation show that magnitude of time-delay is the only predictor for estimated duration (figure 2.6a). For satisfaction, the influence of

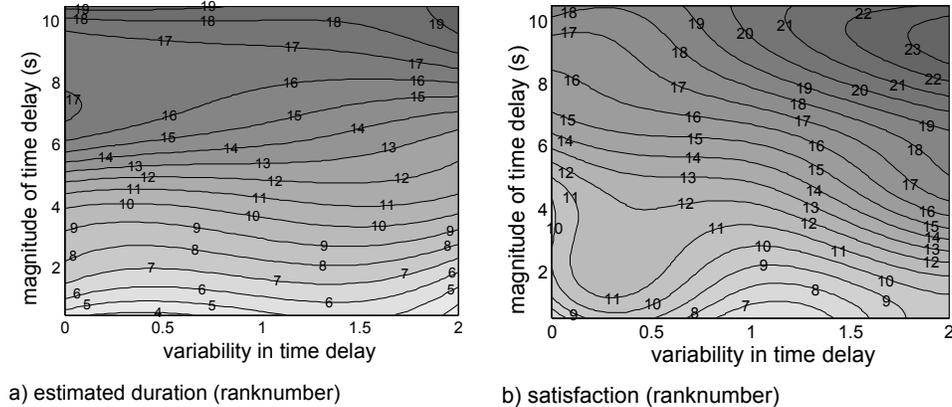


Figure 2.6: Lines connecting a smoothed surface of equal average rank numbers for a) estimated duration and b) satisfaction. Rank numbers are given on the lines.

variability in time-delay increases when magnitudes of time-delay get larger (figure 2.6b). This interpretation means that the increasing absolute variability of time-delay has a negative influence on satisfaction. This increase in absolute variability can only be recorded when larger magnitudes of time-delay occur.

The effect of user errors

To study the effect of user errors on satisfaction, I considered all cases (N=300). User errors (e) result in an increase of time-to-task completion with about 16 seconds for each error ($p < 0.01$, $R^2 = 0.54$), which should on average result in an additional estimated duration of about 8 seconds (equation 2.2). To study the influence of user errors on satisfaction, user errors are studied as an additional independent variable. User errors (e) turned out to contribute significantly to the estimated duration (D_{est} , $p < 0.01$, $R^2 = 0.70$; equation 2.8; figure 2.7a). The difference between the effect of errors can be explained by the differences between the estimation through the power function and the linear equation 2.8.

$$D_{est} = 26.8 + 0.65 \sum Td_M + 10.3e \tag{2.8}$$

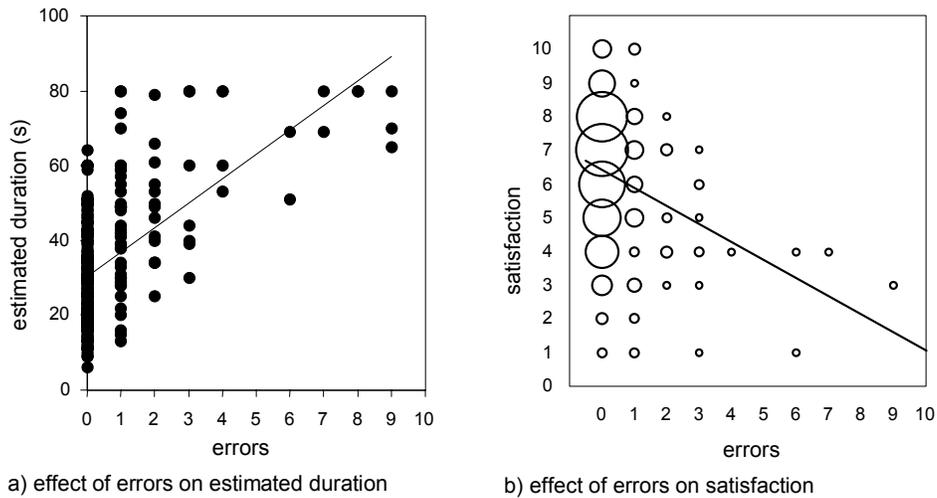


Figure 2.7: Effect of user errors on: a) estimated task duration and b) scaled user satisfaction. In b) the size of the circle indicates the frequency of the value.

Users errors (e) also contribute significantly to the self-reported satisfaction (S , $p < 0.01$, $R^2 = 0.38$; equation 2.9; figure 2.7b).

$$S = 1.1 - 0.038 \sum Td_M - 0.15Td_d - 0.34e \quad (2.9)$$

Assuming that the effect of user errors on satisfaction is evaluated as a decrease of efficiency, the estimated duration should be a predictor for satisfaction. By combining equations 2.8 and 2.4, I was indeed able to predict satisfaction in this way (equation 2.10). However the negative effect of errors (e) on satisfaction (S) in equation 2.10 is about twice as large (0.64), as in equation 2.9 (0.34) where it was derived from the data. Other mechanisms apparently play a role in the evaluation of user errors besides time-to-task-completion.

$$S = 1.0 - 0.040 \sum Td_M - 0.64e \quad (2.10)$$

2.2.4 Discussion of experiment 1

Experiment 1 shows that time-to-task-completion has a significant effect on estimated duration, which confirms the assumption that the monitoring mechanism records time related interaction properties. Estimated duration was also found to be influenced by the magnitude of time-delay and by user errors. These interaction properties are directly related to the physical task time. The relationship between this physical time and estimated duration follows a power function with the hypothesised constants. This indicates that the monitoring mechanism transforms the recorded interaction properties. The lack of effect of variability of time-delay on the estimated duration indicates that the monitoring mechanism regards the entire interaction process as a whole instead of monitoring the different instances of time-delay separately and combining them afterward.

This experiment also showed a significant relationship between physical task time and scaled satisfaction. This supports the idea that satisfaction is an adequate measure to evaluate interaction efficiency. More specifically, the linear relationship between the estimated duration and the user satisfaction was the same when derived in two ways. This relationship was made by comparing different stimuli over different groups of participants. This finding shows that the synthesized

transformation of physical time into estimated duration, results in an otherwise linear evaluation of the monitored physical variables. This supports the assumption that the estimated task duration is a suitable candidate for the input of the evaluator (Zakay and Shub, 1998).

The experiment also showed that the variability of time-delay had a significant, but very weak effect on satisfaction. The participants seemed to evaluate tasks in which time-delay varies as worse than tasks where the time-delay is constant. This does not influence the perception of task duration. The effect of the variability of time-delay must influence satisfaction through an intermediate variable or process other than the estimated duration. A possible explanation is that the variability in time-delay is monitored as an essentially different subjective representation that is only aggregated with time-delay properties in the evaluation mechanism. This argument is in line with theories that claim that the consistency of interaction in itself is monitored and evaluated as more positive than inconsistent interactions (see: Shepperd and McNulty, 2002). A substantially different explanation is that changes in the time-delay influence the evaluation mechanism. It was found that reference values for evaluation are continuously updated with experience (Mellers, 2000). This means, that a medium time-delay preceded by a long time-delay is evaluated positively, while the same medium time-delay is evaluated negatively if it follows a short time-delay. This mechanism of changing reference values has two advantages for the regulation of interaction. Firstly, it signals the direction of change in the interaction, enabling adjustments to prevent the emergence of bad interactions. Secondly, if reference values change, this means that when the world changes, the reference values of the evaluation mechanism follow and can find a new state. However, there is one problem with interpreting the data according to this mechanism, namely that the overall effects of the difference in time-delay should cancel each other out, thus resulting in a neutral score. The negative overall effect could have been caused by an asymmetry when interpreting negative and positive occurrences, where negative occurrences were judged more severely (Shafir and Tversky, 1995). The explanations of the differences in satisfaction involving either the monitoring of inconsistency or changing evaluation references are essentially different. The first mechanism focuses on the monitoring mechanism while the second focuses on the implementation of the evaluation mechanism. This could mean that each of the effects play an independent role, or

that the modelled distinction between monitoring and evaluation cannot be made. This issue was further investigated in the second experiment.

The combined effects of magnitude of time-delay and variability of time-delay determine the evaluation of an interface. When the rank numbers of interfaces are compared, the hypothesised distinctions becomes even clearer than when looking at the values of satisfaction alone. Users can apparently rank interaction processes for a combination of the magnitude of time-delay and the variability of time-delay. In practice this means that user's dissatisfaction can be partially compensated by limiting variability of time-delay.

If we define effectiveness as the achievement of goals (ISO 9421-11, 1997), the effect of user errors becomes an effectiveness issue, rather than an efficiency issue. Effectiveness, as defined here, is related to the goal of interaction (Norman, 1984), which means that changes in user strategy could be expected. The user time per interaction cycle increases when errors occur, indicating a change towards a slower, more deliberate search strategy (Teal and Rudnicky, 1992). This change in itself indicates that interaction is constantly regulated.

The influence of user errors on satisfaction as expressed by the estimated task time, is larger than the actually recorded effect of user errors on satisfaction. The increase in estimated duration is apparently not the only explanation for the differences in the self-reported satisfaction when errors occur. This difference cannot be due to a ceiling effect, however, since the lowest satisfaction score is hardly given at all, even when errors occur. I suggest another explanation, namely that the participants compensated their reported satisfaction. Participants were asked to interpret the experiment as a usability test. The participants may therefore have felt that making errors is a stupid mistake. They therefore scored a higher satisfaction than they actually felt because they did not want to 'blame' the computer, an effect that was encountered earlier in the evaluation of interactive devices (Reeves and Nass, 1996). The verbal remark of one of the participants during debriefing supports this suggestion. He stated that he compensated his satisfaction score somewhat whenever he made a mistake because it was not the computer's fault. These findings indicate that to interpret effectiveness issues understanding of the structure of control of interaction is needed, which is beyond the scope of this study.

2.2.5 Experiment 2

To understand the effect of variability of time-delay on satisfaction, a second experiment was conducted. The aim of this experiment was to determine which of the two explanations offered in the discussion of experiment 1 best describes the effect of variability in time-delay.

Variability in time-delay means that the user is less certain about how much time-delay there will be after each action, which makes the regulation of interaction less efficient. The self-regulatory system should select the sequences of interaction that can be most efficiently regulated. To do this, the evaluation mechanism should generate a lower satisfaction score for sequences of interactions that have more variability in time-delays. This will be particularly true when users want to use the delay time to perform a secondary task. The more variable the instances of time-delay, the less accurately these secondary tasks can be controlled, which should result in lower satisfaction. An example, where the waiting time can be used, is when I type in my password for Windows in the morning, I know that I just have time to get a cup of coffee before the computer has logged onto the network. If I make a (unanticipated) typing error, a failed password message appears within seconds. By then I have already left the room even though the computer is still waiting for the correct password. On returning I have to sit through the whole log-on procedure. I have lost the otherwise useful time during my computer start-up.

In the discussion of the first experiment I gave a first explanation how variability of time-delay influences the evaluation of interaction, namely by monitoring the inconsistency of time-delay and evaluating this inconsistency like other properties of the interaction process. The evaluation mechanism aggregates the separate outcomes of monitoring into a single evaluation value.

The second explanation of the influence of patterns of time-delay on satisfaction is related to the task of the self-regulatory system to improve interaction. To determine the direction of improvement, the evaluation mechanism compares the ongoing interaction with a reference value for that interaction. If this comparison is positive the interaction is better than anticipated and the ongoing interaction should be continued. To be able to regulate different interaction processes, and to continue improvement of interaction, the evaluation mechanism should be able to change the reference values of interaction. This can be done if the reference values

for evaluation are related to accumulating experience (Mellers, 2000). This means that the reference values are partly determined by all prior experiences and partly by the previous occurrence of time-delay. If reference values for evaluation are constantly updated the feedback control system continues to search for a better interaction process. This has the advantage that the control of interaction results in both a search for the optimal interaction and flexibility in the optimisation of interaction. In practice this means that a medium time-delay following a long time-delay would be considered positive, while the same medium time-delay, following a short delay, would be considered negative. To determine which of these explanations gives a better explanation of the observed differences in evaluation, I study the effect of patterns of time-delay, which was found to have an effect on satisfaction (Chow, 2001).

Although the focus of this experiment was on the effect of the variability of time-delay on satisfaction, this influence is probably only a second order indicator for the adequacy of interaction. Another property of the interaction, the number of user errors, influences the effectiveness of the interaction process itself. Although this property cannot be controlled, I nevertheless expect that this property has a far larger influence on satisfaction than the variability of time-delay.

2.2.6 Methods and Materials

Participants and design

Participants were two groups of 15 native Dutch-speaking university students, both male and female. Participants took part in a within-participant experiment in which variability of time-delay was manipulated. The variability of time-delay was offered to participants in five discrete patterns. These patterns were implemented by generating an instance of time-delay after selecting a product or product-category and before the next level appeared on the screen. The average magnitude of time-delay was 4.5 seconds for all conditions and the variability was 1; the shortest instance of time-delay was 2.25 seconds and the longest was 6.75 seconds. The occurrences of time-delay were ordered in five fairly stable patterns to amplify context effects: *slowest-fastest-fastest-slowest* (S-F-F-S) and *fastest-fast-slow-slowest* (F-f-s-S) which both had the longest instance of delay time last, and patterns: *slowest-slow-fast-fastest* (S-s-f-F) and *fastest-slowest-slowest-fastest* (F-S-S-F) which both had the shortest instance of delay times last. The fifth pattern consisted of four equal

instances of time-delay (neutral). Each pattern was assigned to each participant 3 times for a total of 15 tasks, in a randomised order.

Task and apparatus

The experiment was fully computerised, in which participants were asked to locate a product in a hierarchical database. The task was based on an Internet database with a hierarchical structure of four levels, which was a simpler version of the database used in experiment 1. The reduced database consisted of 625 products with 5 choices at each hierarchical level. In all other aspects the database and task structure were the same as those used in experiment 1.

The first group of participants were given a secondary task during the time-delay. In this secondary task, participants were shown a painting and an array of 20 painting fragments. They were asked to drag the fragment of the shown painting onto that painting. If the fragment was correctly matched, a new painting was shown. The second group of participants acted as a control group and were not given a secondary task. The participants in this group were given identical tasks but were shown a blank screen during the time-delay period.

Procedure

Participants were welcomed into a cognitive laboratory where they were assigned one of eight cabins. Each cabin contained a chair, a desk, and a PC with a 15" colour monitor. The participants used only the mouse as input device. They were asked to find 2 series of 15 target products in the database, as quickly and accurately as possible. Before each of the series was started, participants were given 5 practice tasks. Once a participant had found a target product, a single question was asked. In the first series of 15 tasks, participants were asked to report their satisfaction on a scale from 1 to 10 (10 being the highest level of satisfaction). In the instruction, the participants were told to take every conceivable impact on satisfaction into account, not just problems in finding items but also time-delay, erroneous choices and anything else that happened. In the second part of the experiment, the same participants were asked to estimate the duration of a completed task in seconds. In the practice tasks of this series, participants were told the time-to-task-completion immediately after the estimation to help them calibrate their estimation of task duration (Zakay, 1993). All the participants were asked to scale their satisfaction first to prevent a conscious link between interaction time

and satisfaction. The experiment lasted about 35 minutes. The participants were debriefed, thanked, and paid 10 Dutch guilders (€ 4.50) for participating in the experiment.

Recorded variables and experimental hypothesis

The reported variables, satisfaction, and estimated task duration were recorded. The time-to-task-completion was recorded along with the number of mouse-clicks, which is a measure for the number of errors participants made. The number of correctly chosen painting fragments per interval was stored as an indicator of the success of the secondary task. First, I will investigate the data to find confirmation for results from experiment 1. I did this by investigating whether a power function gives a good fit between the actual delay time and the estimated duration and whether a linear relationship can be defined to relate estimated duration to user satisfaction.

There are two possible effects of variability in time-delay on the self-regulatory system. First, the variability of time-delay makes the regulation more difficult because the self-regulatory system itself operates less efficiently. This means that the less variability influences the optimisation process the higher the satisfaction score for that the sequence of time-delays should be. The evaluation of the interaction should result in the highest satisfaction scored when there is no variability at all. Monotonic changes of interaction can also be controlled with relatively little effort, and should therefore receive intermediate satisfaction scores. Erratic changes of time-delay are the least predictable and should therefore get the lowest satisfaction scores.

Hypothesis 5: Sequences of delays that do not change are evaluated as being most satisfactory, followed by patterns that exhibit a monotonic change.

If, on the other hand, the difference of time-delay does influence user satisfaction because the evaluation reference is updated, this would mean that time-delay is evaluated by comparing it to a different reference value. The specific values are based on the case when the reference value for each subsequent instance of time-delay is the average of the previous reference value and time-delay. For example, if the reference value is 4.5 seconds and the actual time-delay was 3.5 seconds, a positive evaluation of 1 is registered and the reference value for the subsequent instance of time-delay becomes $(4.5+3.5)/2= 4$ seconds. This way of

combining experience with a reference value means that the reference values will slowly adjust to systematic deviations in the actual situations but will not be overly influenced by random fluctuations. In chapter 3 the updating of reference values for evaluation will be modelled in more detail.

This interpretation of the evaluation of a sequence of interaction means that the patterns ending in the shortest delay time (*slowest-slow-fast-fastest* and *fastest-slowest-slowest-fastest*) should be evaluated as being most satisfactory. To test this hypothesis, these patterns, are compared with the patterns ending in the longest delay time (*slowest-fastest-fastest-slowest* and *fastest-fast-slow-slowest*).

Hypothesis 6: Sequences of delays that have the shortest delay last, receive a better evaluation than sequences that have the longest delay last.

To understand how a satisfaction score for a sequence of interaction cycles is generated as the result of changing reference values, I introduced two ways in which the evaluation mechanism could generate such a satisfaction score. The evaluation mechanism could either take into account only the last instance of time-delay as indicator for the whole sequence, or alternatively it could use the evaluation of each instance of time-delay as a property of the whole sequence. By aggregating the evaluation scores of the separate instances, an overall score for the sequence is generated. These two ways of evaluating a sequence of interaction cycles do not result in different predictions for the main effect that fast-last should be best, however, they could result in slightly different predictions about the specific patterns.

In figure 2.8 the evaluation value of the different patterns of time-delay is predicted in both of these two ways. Since the expected effects of the manipulations are small, the differences between these two evaluation values will not be tested on the data from this experiment, they are used here only to illustrate the differences resulting from the possible implementations of the system.

To confirm that the success of the secondary task is indeed influenced by the instance of time-delay, I will determine the relationship between secondary task score and duration of the interval in which it was scored.

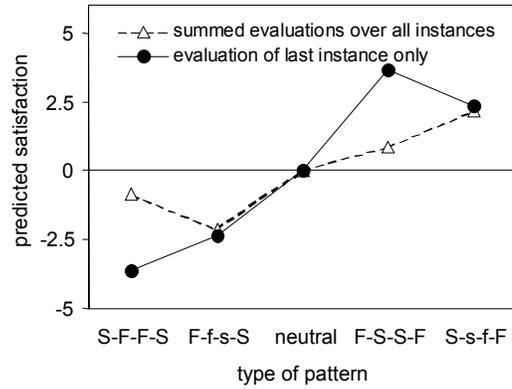


Figure 2.8: Two possible effects of patterns of time-delay on predicted satisfaction, under the assumption that the reference values are updated. Circles give a prediction if a sequence is evaluated taking into account the last instance of the sequence only, triangles show the outcome if the evaluation of all instances is added up.

2.2.7 Results of experiment 2

Effect of physical task time

The effect of time-to-task-completion on the estimated duration and on user satisfaction is established by estimating the parameters of a power function (equation 2.1). The results showed a significant relationship between actual time-to-task-completion (T_{act}) and estimated duration (D_{est}) on the one hand, and satisfaction (S) on the other hand. The relationships were similar to those found in experiment 1 and were found for both the participants without a secondary task (ns : $D_{est} R^2=0.28$, $p<0.01$; $S R^2=0.12$, $p<0.01$) and the participants with a secondary task (s : $D_{est} R^2=0.28$, $p<0.01$; $S R^2=0.13$, $p<0.01$; equations 2.11 to 2.14).

$$D_{est}(ns) = (5.9 + T_{act})^{0.94} \quad (2.11)$$

$$S(ns) = 2.5 - 0.08(5.9 + T_{act})^{0.94} \quad (2.12)$$

$$D_{est}(s) = (12.3 + T_{act})^{0.90} \quad (2.13)$$

$$S(s) = 3.3 - 0.11(12.3 + T_{act})^{0.90} \quad (2.14)$$

The effect of errors, size and pattern of differences in time-delay

The number of user errors had a large, significant effect on estimated duration and scaled satisfaction (table 2.1) for both groups. In the discussion of the first experiment, I argued that errors are related to effectiveness. Since the focus of this study is on the evaluation of patterns, the influence of errors is not further investigated. Z-scores of satisfaction were calculated for each participant. These values will be treated as independent cases for the remainder of this second experiment.

Table 2.1: Influence of user errors on estimated duration and satisfaction for participants with and without a secondary task. Each value is based on 225 observations.

	No secondary task (control)		Secondary task	
	<i>p</i>	<i>R</i> ²	<i>p</i>	<i>R</i> ²
Estimated duration	<0.01	0.51	<0.01	0.58
Satisfaction	<0.01	0.21	<0.01	0.40

An ANOVA was used to test the effect of patterns of variability in time-delay on satisfaction and estimated duration. The manipulation did not influence estimated duration in either group. The scaled satisfaction of the control group was also not affected. The pattern of variability had a marginally significant ($p=0.05$) effect of on satisfaction for the group with the secondary task (table 2.2; figure 2.9).

Table 2.2: Influence of pattern of variability in time-delay on estimated duration and satisfaction

	No secondary task (control)					Secondary task				
	<i>Hyp. Error</i>					<i>Hyp. Error</i>				
	<i>df</i>	<i>df</i>	<i>F</i>	<i>p</i>	<i>R</i> ²	<i>df</i>	<i>df</i>	<i>F</i>	<i>p</i>	<i>R</i> ²
Estimated duration	4	220	0.9	0.44	0.02	4	220	0.6	0.64	0.01
Satisfaction	4	220	0.1	0.77	0.02	4	220	2.3	0.05	0.04

To investigate hypothesis 5 the patterns were classified into the following three groups: no change in time-delay (neutral), patterns that change monotonically (*fastest-fast-slow-slowest* and *slowest-slow-fast-fastest*) and patterns that change in a non-monotonic way (*slowest-fastest-fastest-slowest* and *fastest-slowest-slowest-fastest*;

figure 2.10a). To investigate hypothesis 6, the patterns were classified into three different groups: no change in time-delay (neutral), slow-last (*slowest-fastest-fastest-slowest* and *fastest-fast-slow-slowest*) and fast-last (*fastest-slowest-slowest-fastest* and *slowest-slow-fast-fastest*; figure 2.10b).

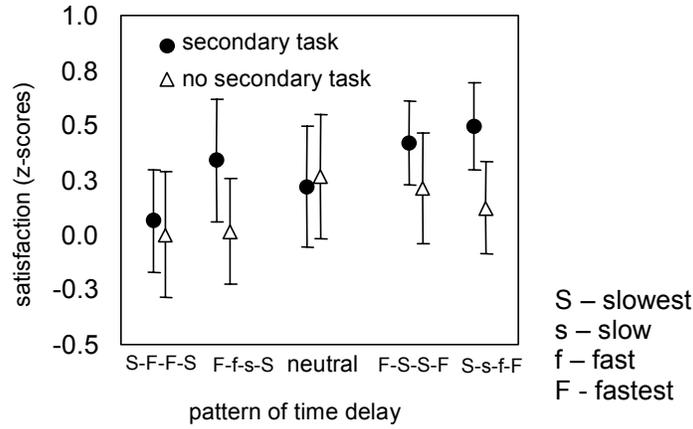


Figure 2.9: Effect of patterns of time-delay on satisfaction in a task with and without a secondary task. The error bars show the significance level for differences of mean at the 0.05 level.

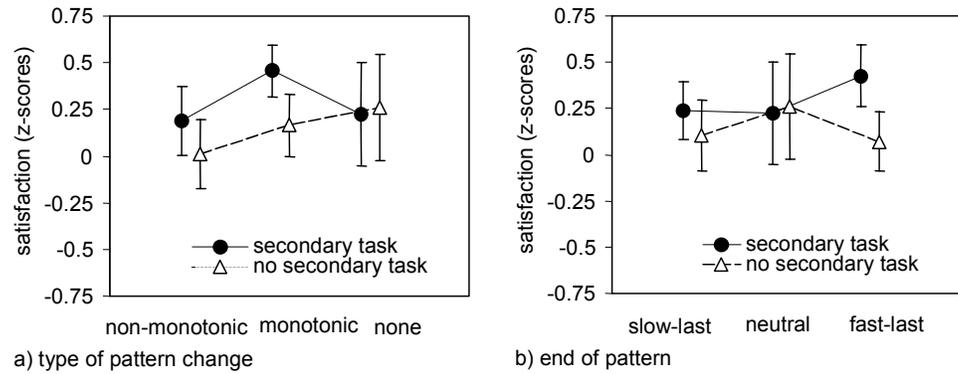


Figure 2.10: Effect of different properties of patterns of time-delay on satisfaction. a) the type of pattern change and b) the end of the pattern. The error bars show the significance level for differences of mean at the 0.05 level.

There were no significant differences for cases without a secondary task. For cases with a secondary task there seems to be a trend that monotonic changes in time-delay were most satisfactory ($F(2,177)=2.6, p=0.07$). By comparing the fast-last trials with the slow-last trials a marginal preference for the fast-last trials is found ($t(177)=1.6, p(\text{one-sided})=0.05$). The fact that this difference is so small is mainly due to the high score for the *fastest-fast-slow-slowest* pattern (F-f-s-S; figure 2.9a). The participants apparently evaluated this option as more positive, which might be the result of the increase in time available for the secondary task. This idea is supported by the significant positive influence of the length of the time-delay interval on the success of the secondary task during that interval ($p<0.01, R^2=0.13$).

2.2.8 Discussion of experiment 2

The results of this experiment support the main findings from experiment 1, namely that time-to-task-completion influences estimated duration and that there is a straightforward relationship between estimated duration and user-satisfaction. However, these relations accounted for less variation than in the first experiment: 0.28 versus 0.59 for estimated duration and 0.13 versus 0.39 for satisfaction. The main reason is likely to be the smaller variance in time-to-task-completion, since all the stimuli had the same total time-delay in this second experiment. On the other hand, the influence of user errors on satisfaction increased. The participants apparently looked for a measure to determine interaction adequacy of which errors were the most obvious. This indicates that whatever effects of variability in time-delay should be expected, they are probably very small.

The results showed a marginally significant difference between the fast-last and the slow-last patterns. This was only true in the cases in which a secondary task was given, which supports hypothesis 6 that the shape of the patterns influences satisfaction, and thus also supporting the idea that the reference values for the evaluation process are constantly updated (Mellers, 2000). A shorter delay time is evaluated as being more positive because it is compared with anticipated longer delay. The effect of variability in time-delay on satisfaction remains very small. To find support for the ideas developed here, in chapter 3, another interaction process in which the change of reference values of evaluation has a larger influence will be investigated.

The tasks in which the instances of time-delay were equal, were not evaluated as being more satisfactory than other cases. So, there is no evidence to support hypothesis 5, although there was a non-significant trend towards this effect in the participants who had no secondary task. A marginally significant positive influence of monotonic increase or decrease of time-delay was found. This is mainly due to the pattern '*fastest-fast-slow-slowest*'. This pattern did not result in the hypothesised lower satisfaction, but tended towards a higher satisfaction than predicted by either hypothesis. A possible explanation is that the satisfaction score is not only based on the control of the primary task to find a product, but also includes the secondary task of finding a match between a painting and a fragment. Whereas the satisfaction for the primary task is probably negatively influenced by this pattern, the satisfaction related to the secondary task will increase when the instances of time-delay increase monotonically, which is exactly the case in the pattern '*fastest-fast-slow-slowest*'. These contradictory effects could well have cancelled each other out in this case. To determine more specifically the influence of patterns of time-delay on satisfaction a secondary task for which longer time results in less adequate situations as was used by Ariely and Loewenstein (2000) might provide evidence for my hypothesis.

2.3 General discussion of chapter 2

The relationship between time-to-task-completion and estimated duration showed that a relevant property of interaction, namely system-time-delay, is monitored. This relationship is monotonic and almost linear, which confirms that monitoring can be executed by the perception system.

The results also showed a monotonic relationship between actual efficiency and self-reported satisfaction. This means that objectively better interactions are reported as more satisfactory, which supports the idea that self-reported satisfaction can be used as an indicator for the result of the evaluation process. The monotonic relationship between satisfaction and time-to-task-completion was related to a linear relationship between satisfaction and subjective task duration. The intentional description of a monitoring mechanism, that uses perceived duration as input for the evaluation of interaction, fits the observed actions of participants. These findings were confirmed by deriving the relations in different

ways and in three different experimental groups. It seems that estimated duration is indeed a candidate for predicting of the satisfaction in this type of interaction.

There are some limitations to this conclusion however. Firstly, in all studies the explained variance of the satisfaction score is not very large (about 30% in experiment 1 and about 15% in experiment 2). Apparently in a database search task, differences in time-delay only make up for part of the observed variance in satisfaction. This might be caused because other elements of the interaction process were also important for experienced satisfaction, such as the actual product that had to be found. A product that is less easily found in the database might be scored less satisfactory than more easily found products. It might even be the case that products that were more favoured by the participants (e.g. cake over toilet paper) received higher satisfaction scores. In future research these effects should be taken into account, for example by counterbalancing products over larger samples of participants. Another reason for the noisiness might be caused by the between participant design for estimating duration and scaling satisfaction. This was done so participants did not have a conscious relationship between time-delay and satisfaction, but this decision might have caused some noise. To have better control of this effect, the following experiments will be within participant designs when possible. A more fundamental issue is, that by asking participants to estimate duration, or by asking participants to report satisfaction, they are made aware of these concepts, which in itself could introduce systematic errors. Furthermore, in all the studies presented in this chapter, there was a significant positive correlation between the physical time and subjective duration. To confirm that the subjective duration is indeed the input signal for the evaluation process, I recommend additional experiments where the relationship between the physical time and the subjective duration is manipulated, for example by giving ongoing feedback about the waiting times, or by giving secondary tasks to reduce the experienced waiting times.

An alternative view at the rank numbers for the satisfaction also shows that the evaluation mechanism is capable of ordering the interaction processes based on the combination of two aspects from the interactive application: time-delay and variability in time-delay.

User errors are a measure of how well the interaction is executed. I therefore argued that this is an indicator that errors are related to interaction effectiveness.

The primary goal of the self-regulatory system is to achieve effective interaction. All related aspects of the interaction, including higher-levels at which strategy, and goals are controlled play a role in achieving this goal. This makes the straightforward interpretation of the effects of users errors on satisfaction hard to interpret without knowledge about the different levels of interaction control and their interactions.

Variability in time-delay over a sequence of interactions influences user satisfaction. I argue that the effects of variability indicate that the reference values for the evaluation mechanism change with experience. Support for this idea was found in experiment 2 where I found that interactions that end in the shortest time-delay were in general evaluated as being more satisfactory than interactions ending in the longest time-delay. However, my findings could not conclusively exclude that inconsistency is monitored as a negative trait of sequence of interaction. To explore this issue further I recommend future research into more complex combinations of interaction processes. The idea of changing reference values is taken up in more detail in the next chapter where I investigate the evaluation mechanism in more detail.

In practice, participants seem to be able to estimate duration based on the physical properties of the interaction process. Estimated duration of the task is one determinant of user satisfaction and is therefore an important measure for determining usability (Czerwinski, Horvitz, and Cutrell, 2001). However, although the effect is small, variability in time-delay influences satisfaction as well. When such effects are more fully understood, designers will be able to design interactions that partly compensate for the effects of time-delay on satisfaction.

Chapter 3: Evaluation of visual search tasks

Abstract

Behaviour of users can be studied as the outcome of a self-regulatory system (Carver and Scheier, 1998). In doing so, a self-regulatory system, or more specifically a feedback control system, is assumed to take care of the optimisation of user-system-interaction. The feedback control system consists of monitoring, evaluation, and adjustment mechanisms. The evaluation of the adequacy of interaction for a visual search task was investigated in an experimental study, with the hypothesis that the evaluator interprets information that is relevant for the optimisation of the interaction process. The evaluator should use this information to tell the adjustment mechanism how to improve the interaction. The result of the evaluation, self-reported satisfaction, is higher for search fields that have a better objective efficiency thus confirming that these signals are functional. Evidence was also found that the reference values of evaluation change as the feedback control system learns more about the specific task in hand. A possible mechanism for feedback control was studied by designing a detailed synthesis of the evaluator. High correlation between simulated and self-reported satisfaction of participants was found. To investigate whether the simulated evaluator initiates the best possible choice of actions, a strategy selector was specified that uses information from the evaluator to improve interaction. The proposed strategy selector is successful in selecting the most adequate of two interaction strategies for the task in hand. Finally, to understand how accumulating information can be used to learn more about the ongoing interaction, three ways of changing reference values for evaluation are compared.

3.1 Introduction

A great deal of information can be presented on a computer screen. Some of this information is usually irrelevant to the user. Internet pages in particular contain considerably varying amounts of useful information and irrelevant messages (such as advertisements). In some cases large amounts of information are presented on a

single page, for example in the so-called portals such as Yahoo. However, it seems as if the preference of users is currently moving towards simpler web design, such as the Google¹ search engine, which features a single text field and button.

Users work with both types of interface and compare them. To do this, they need a guidance system for interaction. Such a guidance mechanism can be called a self-regulatory mechanism, more specifically implemented as feedback control (Carver and Scheier, 1998). In this chapter, elements of a self-regulatory mechanism for the evaluation and adjustment of user-system interaction are studied. For the description of the feedback system, a reverse engineering approach is adopted (Dennett, 1994). According to this approach the goals of the self-regulatory mechanism and the strategies to achieve these goals are specified. The goal of interaction control is the optimisation of behaviour. The strategy to achieve this goal is through a feedback control mechanism that has monitoring, evaluation and action adjustment mechanisms. Users monitor the interaction by perceiving the ongoing interaction and forming subjective representations of the relevant elements. The output of the monitoring process is then compared with reference values for that interaction. The result of such an evaluation can be interpreted as a hedonic tone (Johnston, 1999). When the actual outcome of evaluation is better than a reference value, the target, the hedonic tone is positive. If the actual outcome is worse than the target, a negative hedonic tone is generated. The hedonic tones are aggregated into a single value called, pleasure or satisfaction, which is a value for the adequacy of the interaction (Cabanac, 1992). In this study I assume that the result of evaluation can be expressed as self-reported satisfaction. A high level of satisfaction signals adequate interaction, a low level of satisfaction signals that the interaction has to be adjusted. Action adjustments based on the evaluation results in a change of interaction towards an optimum (figure 3.1).

Humans, being multifunctional entities, are flexibly running a complex combination of different interaction processes (Dennett, 1995). To do this, they have to be able to determine the optimal interaction in different situations. This is

¹ The mentioned examples relate to the layout of Yahoo (<http://www.Yahoo.com>) and Google (<http://www.Google.com>) as they were available on the Internet early in 2003.

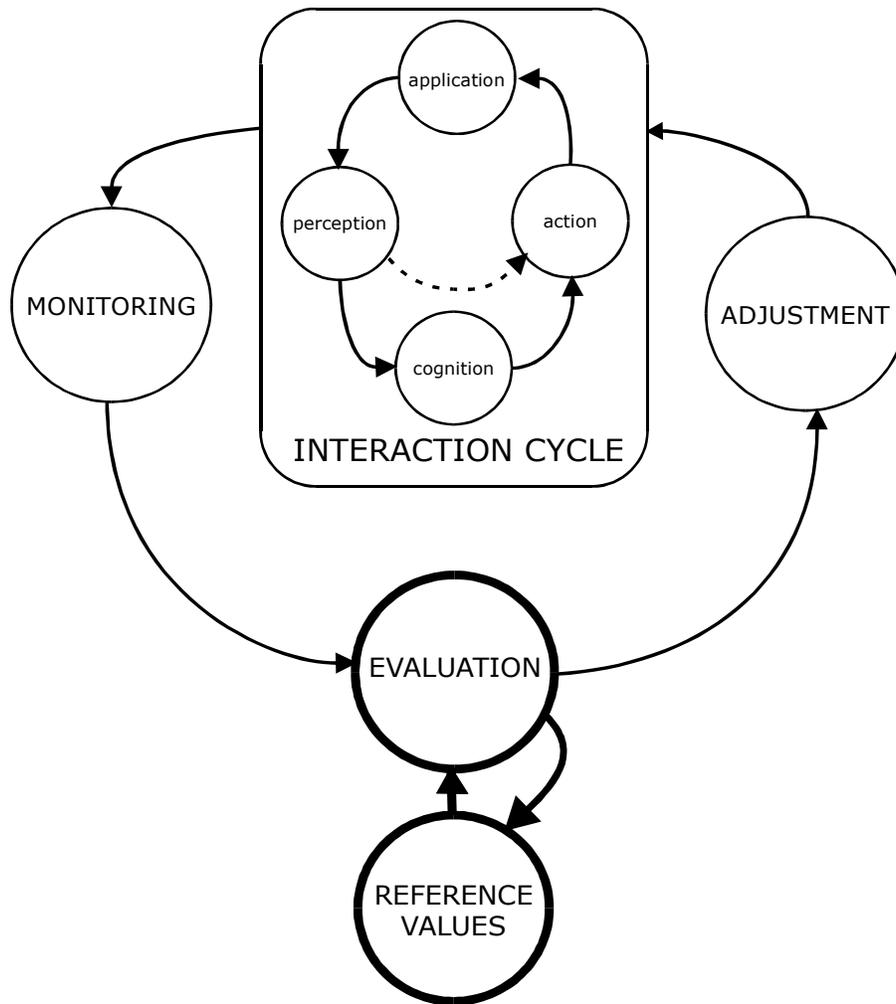


Figure 3.1: Feedback control of interaction. The whole interaction cycle is monitored. In addition to action, the cognitive or perceptive effort can also be adjusted when required. This chapter focuses on the investigation of the bold elements.

possible if reference values for evaluation are continuously changing based on new experience (Mellers, 2000).

An example of this type of optimum finding happened to me while hiking in the Scottish hills a couple of years ago. At one point I suddenly felt uncomfortable

without knowing why. I was getting more tired than I had anticipated, so I thought I was in a worse physical condition than expected. My step and my hiking pace decreased until it felt comfortable again. A while later the hike became much easier and more pleasurable once more. Only then did I realise that I had been climbing a hidden gradient. Climbing a hidden gradient gives no cues that the interaction has changed. However, the apparently, same interaction required more effort, and is therefore perceived as less efficient. This example illustrates that a change in the situation (start of climb) requiring additional effort can be signalled as negative. Functional action adjustments are made to overcome the problem (slower pace, shorter steps). This example also shows that the correct adjustments can be made, even without consciously knowing what the exact nature of the problem is. When the negative situation ends and the previous situation is restored this is experienced as positive.

To investigate the assumed self-regulatory system for optimisation of user-system interaction, the different elements of the feedback mechanism are sequentially investigated. In this chapter the behaviour of the evaluation mechanism is investigated. To do this the evaluation mechanism is first specified in more detail.

The task of the evaluation mechanism is to generate a signal to adjust actions, thus leading to improvement of interaction. To fulfil this task, the evaluation mechanism processes the information that was monitored about the interaction and determines the course of action that leads to improving the interaction. When confronted with two interfaces, such as Google and Yahoo, the evaluation mechanism should determine which is the better of these two interfaces. This comparison can be a face-to-face comparison. However, in many situations it is impossible to make face-to-face comparisons between all the alternatives. To determine whether an interaction process is adequate, the evaluation mechanism should therefore compare the interaction process with a stored reference value for interactions of a similar type, which results in a single evaluation value (Cabanac, 1992) that signals whether the interaction is better or worse than average. Based on this evaluation value, the adjustment mechanism determines whether the current action (for example use Google) is good enough or that adjustments must be made (look for other search engines such as Yahoo or AltaVista).

3.2 Experiment

The evaluation mechanism was investigated in an experimental study. The evaluation mechanism should generate a signal that allows the selection of the best interaction process. The evaluation mechanism does this by comparing the interaction process properties with reference values. These reference values are based on the available knowledge of the evaluator.

Although no adjustments could be made in the experimental interaction process, and the outcome of the evaluation could therefore not actually be used to improve interaction, nonetheless I assume that the evaluator runs its course and generates a judgement on the encountered interaction nonetheless.

To investigate whether the evaluation mechanism gives functional judgments about the adequacy of interaction, the efficiency of a visual search task is manipulated in two ways. Efficiency is recognised as one of the elements that contribute to the adequacy of interaction (ISO 9421-11, 1997). First the total amount of visual stimuli in a search field is varied. This follows the simple heuristic 'larger search fields are worse', which will often be successful in visual search tasks and is therefore likely to be one of the first heuristics to be applied by the evaluator. Second the number of targets is varied, which does not completely follow the available knowledge. Search tasks with both a large number of search fields and target fields might actually be more efficient than tasks with fewer search fields and targets fields.

3.2.1 Methods and Materials

Participants and design

Participants were 16 Dutch-speaking students from Eindhoven University of Technology, 9 were male and 7 were female. The average age was 20.5 (SD=2.2 yrs).

In a visual search task efficiency was manipulated. The main idea was that the evaluation mechanism should be able to give information about the adequacy of different combinations of target and nonsense fields in a search task. The fields in this task are boxes in which a list of options is given, therefore in this chapter I will from now on use boxes referring to the specific fields in the experimental task. A possible indicator for the outcome of the evaluation process, self-reported

satisfaction, was used to judge the ongoing interaction. The adequacy of the interaction process was determined using different combinations of the total number of boxes and the number of boxes containing database information (table 3.1). The first manipulation can be evaluated using a heuristic that can be referred to as ‘larger is worse’. The second manipulation influenced the objective efficiency independently of this heuristic. The total number of search fields on the screen ranged from 1 to 16 boxes (total); the number of boxes containing database information was 1, 2 or 3 (database). Of the possible 45 combinations, 24 were selected (table 3.1). The selected stimuli were presented to each of the participants in a randomised order.

The average number of boxes that had to be investigated to find a box containing database information depended on both of the manipulations. In the stimulus set, equal objective efficiency was not always monotonically related to the total number of search fields. Consider, for example, the case when 1 database box is embedded in a total of 5 boxes (fraction database information = 0.2), which has the same objective efficiency as the case of 2 database boxes among 10 boxes or 3 database boxes were amidst 15 boxes (bold O’s in table 3.1).

Table 3.1: The 24 stimuli used. (The bold O’s are the examples from the text)

Database \ Total	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	X	X	X	X	O	X	X	X		X		X				X
2	-	X		X		X		X		O		X		X		X
3	-	-	X			X			X			X				O
Average efficiency (D/T)	1.00	0.75	0.67	0.38	0.20	0.33	0.14	0.19	0.33	0.15	-	0.17	-	0.14	0.20	0.09

In general, when one box with information was hidden among a number of nonsense boxes, the average number of boxes that had to be considered before a target was found equalled 0.5xtotal field size (top straight line in figure 3.2).

If two or three boxes containing information were hidden among a number of boxes this average number was lower (0.25 or 0.17; lower straight lines in figure 3.2). The stimuli that were used (table 3.1), are different combinations of the total size of search field and number of boxes containing information. The average number of boxes that have to be considered before a target was found is the

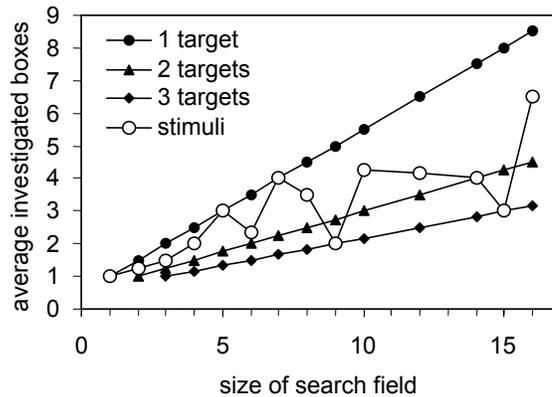


Figure 3.2: Efficiency for search fields of different sizes with 1, 2, or 3 boxes containing database information (targets) and the average efficiency of the stimuli. The search task with 1 target can be interpreted as the heuristic ‘larger is worse’.

average of the presented stimuli with the same total size of search field (bottom row in table 3.1; open markers in figure 3.2). Since not all possible combinations were made, the sequence of presented stimuli is not monotonically related to the total size of the search field.

Apparatus and task

An interactive experiment was programmed using Visual Basic 5.0. Participants received an on-screen database search task in which they had to find a product at the lowest level of a hierarchical database. A single stimulus, a combination of database boxes and nonsense boxes, was given during each task. The database consisted of four levels and was based on the product list of an online supermarket². Each level consisted of a choice between five products or product categories, which were listed in a box. The participant had to confirm the choice to finish a task. For example if the target product was Gouda Cheese, the participant had to select the following items subsequently: (1) Dairy products from the options: Vegetables / Dairy products / Baking products / Drinks / Personal Care

² The database was the same simplified version of the database taken from the website: <http://www.ah.nl> in 2000, as used in experiment 2 of chapter 2

- (2) Cheese from the options: Milk / Deserts / Butter / Cheese / Eggs - (3) Dutch cheese from the options: French cheese / Foreign Cheese / White crust / Blue veined / Dutch Cheese - (4) Gouda Cheese: from the options Edam / Gouda Cheese / Mustard Cheese / Maaslander / Westzaanse smoked cheese - and finally (5) Confirm Gouda Cheese.

The experiment was implemented on a computer. The selected products were presented to the participants in random order. The list boxes contained either a list of 5 database items, which was always identical in all those boxes, or a list of 5 nonsensical letter-combinations in a similar format (string length, number of strings, capitalisation, etc.) to the boxes containing database information. The screen was divided into a 4x4 array of box positions. On each individual screen, boxes were randomly positioned at one of the 16 positions (figure 3.3). The colours and positions of the boxes, and the applied font types were varied over the various levels of the database search.

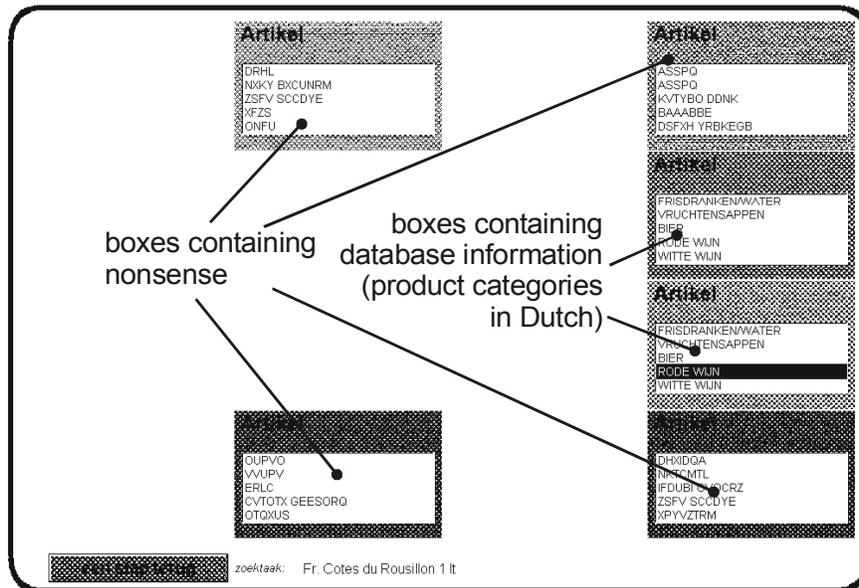


Figure 3.3: Stimulus example with 2 boxes containing (identical) database information and 4 boxes containing nonsense letter combinations. The boxes are placed at a randomly selected position. The colours of the boxes are varied randomly.

Procedure

Participants entered a cognitive laboratory, in which a cabin was assigned to them. Each cabin contained a PC with a 15" screen and a mouse. The participants did not have to use the keyboard. Participants were given on-screen instructions asking them find each product as quickly and accurately as possible. The instructions were followed by 8 practise tasks, after which the 24 experimental tasks were given. For each of the 24 tasks, one of the stimulus combinations, consisting of the total size of search field and number of boxes that contained database information was selected at random. If the participants selected an incorrect product or category they had to continue the task until they found the correct product. The trial was completed after the participant confirmed the correct product. After each product had been found, the program asked participants to scale their satisfaction on a 10-point scale according to the Dutch high-school grading system (where 10 is the highest mark). The experiment lasted about 30 minutes, after which the participants were debriefed, thanked, and paid 10 Dutch guilders (€ 4.50).

Recorded variables and experimental hypotheses

During the experiment, self-reported satisfaction, the number of mouse clicks before task completion, and the time-to-task-completion were recorded. The manipulated variables, i.e. the number of target boxes and the number of boxes containing database information were stored in the data file. A search field could contain one, two, or three targets for the same size of search field.

Visual search can be assumed to be a practised task, which means that participants know about this type of interaction. The evaluation of the search task should result in a judgment that can be used to select the most adequate interaction. In the first instance this choice should be related to the participant's existing reference values, in other words to their 'world knowledge'. This world knowledge is assumed to be something like 'larger is worse', which is a search heuristic for the most common search of one target among many non-targets (top line in figure 3.2). Even when the actual interaction is occasionally better than expected, the evaluator should know that this might be due to chance. This means that the evaluator should follow the world knowledge unless there is a reason to suspect that it does not suffice.

Hypothesis 1: A monotonic relationship initially exists between the size of the search field and satisfaction.

The ratio between the target boxes and the total number of boxes is a better indicator of the actual search efficiency than the total size of the search field. The target box is found more efficiently when fewer boxes have to be scanned. With the same user effort this would mean that there is a relationship between the ratio of target boxes compared to the total number of search boxes and the time-to-task completion. Both the number of considered boxes and the time-to-task completion can be indicators for the evaluation mechanism that the objective efficiency of the interaction process is not varying according to the world knowledge. The evaluator should start to notice that there are consistent discrepancies between the heuristic 'larger is worse' and the executed interaction. If this happens, a second interaction property, the number of target boxes, will be included in the evaluation of the search task.

Hypothesis 2: The evaluator notices that with a constant size of the total search field, those trials with more target boxes are more efficient, which is reported as higher satisfaction.

With this additional, specific, knowledge about the task in hand, the evaluator can learn to generate judgements that are more closely related to objective efficiency. This means that the evaluation mechanisms should adapt its judgment when the objective interaction efficiency deviates from the world knowledge.

Hypothesis 3: With accumulating task experience the evaluator learns to base its estimation of interaction efficiency less on the 'large is worse' heuristic and more on actual efficiency.

3.2.2 Results

As a preparatory operation, z-scores of the reported satisfaction measures were calculated to compensate for the different base levels and variances between the participants. In this operation the mean for each participant became 0 and the standard deviation of each subject was made equal to 1. After these transformations, trials were treated as independent cases.

About 15% (58) of the trials contained at least one instance in which participants had to backtrack in the database. These errors has a large effect on time-to-task

completion, $R^2=0.79$, $p<0.01$. The number of errors also had a small effect on reported satisfaction was found, $R^2=0.05$, $p<0.01$, which indicates that the evaluator noticed these failures to achieve the highest possible task effectiveness, but that this did not weigh heavily in the determination of which interaction process was better. The effect of these additional clicks was studied in more detail in chapter 2.

Figure 3.4 shows a significant relation between the ratio of database boxes to the total field size and the time-to-task-completion ($R^2=0.10$, $p<0.01$). The total number of boxes did not contribute significantly to the time-to-task-completion. The assumed measure for objective efficiency, the ratio of database boxes to all fields is related to time-to-task-completion, which can also be interpreted as a measure for objective efficiency. The results also show a relationship between time-to-task-completion and satisfaction ($R^2=0.16$, $p<0.01$). The evaluator apparently takes information about the objective efficiency into account.

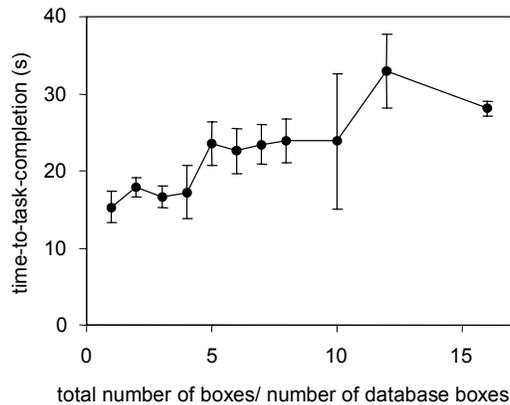


Figure 3.4: The influence of objective efficiency on time-to-task-completion as determined by the total number of boxes divided by the number of database boxes. Two extreme values (time > 100 seconds) were removed from this graph. The noise is partly caused by the occurrence of user-errors. The error bars indicate the significant difference at the 0.05 level.

The total number of boxes influences satisfaction, $R^2=0.28$, $p<0.01$ (figure 3.5a). There is an additional, significant effect of the number of database fields in a search task ($R^2=0.03$, $p<0.01$). Tasks with more database fields for the same number of total fields were scaled as more satisfactory than tasks with only one database field

(figure 3.5b). This was the case for all but two of the total field sizes (open markers in figure 3.5b). Participants appear to use the information that more database fields mean that the interaction is more efficient.

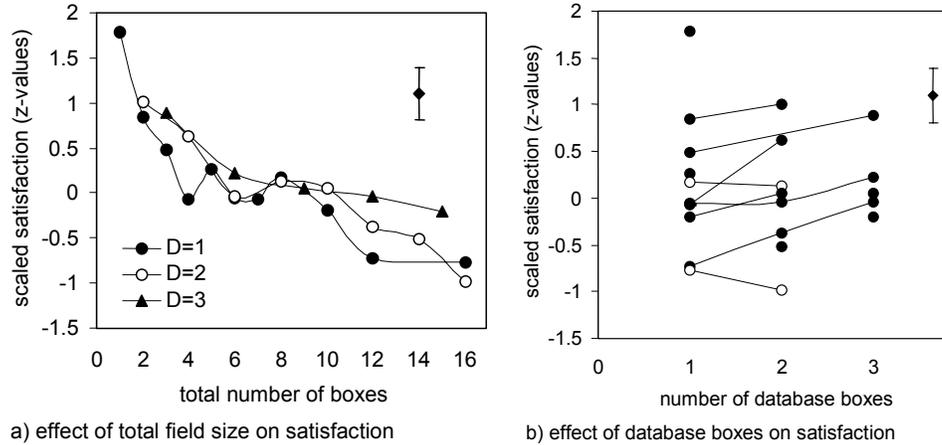


Figure 3.5: Mean of satisfaction in the 24 conditions. a) The effect of the total number of boxes on satisfaction. The lines are the different levels of number of database boxes. b) The effect of the number of database boxes on satisfaction. The different lines are the different levels of the total number of database boxes. Error bars indicate the significant distance between points at the 0.05 level.

If the evaluation mechanism uses the knowledge it gathers to generate a more accurate evaluation, the influence of the different properties of the search field should change during the course of the experiment. To study this change of evaluation in the course of the experiment, trials 1 to 12 are compared with trials 13 to 24. The total field size had a significant influence on the reported satisfaction in trials 1-12 ($R^2=0.33$, $p<0.01$, figure 3.6). However, the number of boxes containing database information had no influence on the reported satisfaction in these trials ($p>0.10$). This indicates that the heuristic 'large is worse' is the only property of the visual search field that influences satisfaction scores for the first trials in the experiment. In trials 13 to 24, however, the number of database boxes did significantly contribute to the satisfaction score, $R^2=0.04$, $p<0.01$, which indicates that the evaluation mechanism learns to take this information into account.

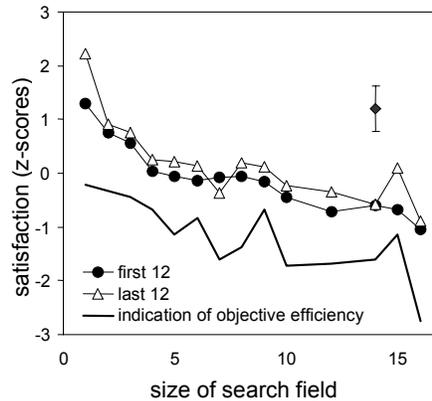


Figure 3.6: Comparison of the mean satisfaction for the 14 sizes of the total field. In the plot comparisons between trials 1 to 12 and trials 13 to 24 show that participants learn to evaluate stimuli closer to objective efficiency of which an indication is plotted. The error bar indicates the significant difference between observations at the 0.05 level.

If the evaluator learns the actual efficiency, with accumulated experience it should generate satisfaction scores that are more closely related to the actual efficiency of the interaction. To investigate this change of evaluation in more detail, the effects of total size of the search field and the objective efficiency of the search task on satisfaction were separately studied for trials 1 to 12 and 13 to 24. To compensate for the correlation between the independent variables the structural equation program LISREL was used. For the trials 1 to 12 I only found a significant effect of the total size of the search field $t(192)=9.66$, $p<0.01$. The effect of objective efficiency did not contribute significantly to the reported satisfaction $t(192)=1.21$, $p>0.10$. For trials 13 to 24 both the total size of the search field $t(192)=2.38$, $p<0.02$ and the objective efficiency $t(192)=2.37$, $p<0.02$ contributed to the reported satisfaction, in approximately equal amounts, $\beta_{\text{total}}=-0.22$, $\beta_{\text{efficiency}}=-0.21$. The participants apparently learn to use objective efficiency to evaluate the interaction.

An alternative way of studying the results is by focussing on the task of the evaluation mechanism. The task of the evaluation mechanism is to signal which interface should be selected. By ranking the stimuli on objective efficiency and comparing them with the ranked satisfaction scores an average difference in rank number of 1.57 is found for the first 12 trials and an average difference in rank

number of 1.28 for the trials between 13 and 24. The participants learn to order the stimuli better (22%) according to their objective efficiency. When the difference in the satisfaction rank-number was compared with the difference in stimuli rank-number I found using a LISREL analysis, that in trials 1 to 12 the total field size is the most important determinant for preference ($\beta_{\text{total}}=-0.63$, $\beta_{\text{efficiency}}=-0.34$), while objective efficiency has the most influence in trials 13 to 24 ($\beta_{\text{total}}=-0.30$, $\beta_{\text{efficiency}}=-0.60$).

3.2.3 Discussion of the experimental data

The significant relationship between the time-to-task-completion, the ratio of the number of information boxes to the total boxes, and satisfaction confirms that there is a relationship between objective efficiency and self-reported satisfaction. However, as in the experiment that was reported in chapter 2, the explained variance is limited. This might be caused by additional factors such as the change of appearance of the boxes, the random position of the boxes on the screen or the actual products that had to be found.

The total size of the search field mainly influenced user satisfaction in the first 12 trials. The participants initially ranked the stimuli based on the total size of the search field. For the action adjustment mechanism this would mean that when confronted with a choice, the interface with the fewest fields would be selected. This would be a good choice for most visual search tasks and can be attributed to the knowledge of the evaluation mechanism prior to the experiment. In several conditions in this experiment, however, this simple heuristic did not apply because more information fields were present. The evaluation mechanism did indeed discover that the simple 'larger is worse' evaluation is not the best, as it learns to recognise that in some cases the total size of the search field does not give all the information about the efficiency of interaction. A positive effect of the number of redundant database boxes on satisfaction becomes significant.

This adaptation resulted in a change from evaluation that only uses the total size of the search field, towards the use of the information about objective efficiency occurs. This is interpreted as a change of the reference values for interaction (Mellers, 2000), in which the evaluator notices the mismatch between the reference values and the executed interaction process and reduces it.

3.3 Simulation

In the first section of this chapter I specified the evaluation mechanism for the feedback control of user-system interaction as an intentional system or rational agent. At a lower level of specification, rationality and intentionality no longer describe these processes (Dennett, 1994). Mechanisms at the physical level are executed by simple functions. A synthesised mechanical evaluation mechanism that exhibits the same behaviour as the intentional system would support the proposed control system. The developed mechanisms give additional tools for the interpretation of the adjustment of interaction, which will be the focus of the next chapter.

In this section I will explore whether such a simple, non-rational, function can exhibit the same behaviour as the evaluation mechanism that is defined as an intentional system. To do this a mechanism is specified that evaluates the ongoing visual search task. Different requirements for the evaluation mechanisms will be studied sequentially.

The first step is to specify a basic evaluator. The evaluator receives information about the interaction, which is compared with a reference value and produces an estimate about the quality of the interaction. The behaviour of this synthesised evaluator is compared with the experimental results to confirm that the approach is in principle sound. The outcome of this system should give information about the adjustment of the interaction. After the evaluator has been specified, a system that actually selects different actions for different interactions is specified and its behaviour is observed. Finally, as the experiment shows that the evaluation mechanism uses a mixture of past knowledge and recent observations, some possibilities for updating of the reference values are investigated.

3.3.1 Simulation of the evaluation mechanism

The evaluation of the visual search task is based on a positive experience when finding a target and a negative experience when finding a non-target. In the simulated evaluator for the search task; the evaluation values for the individual trials are recorded as a simulated experience, the 'hedonic tone' (Johnston, 1999). The reference value for the evaluation depends on the function of the evaluation. If the evaluation mechanism has to determine which search field is preferable (e.g., to use Yahoo or Google) interactions have to be compared with the same reference

value. This reference values is assumed to be based on experience with all search engines. For the experimental data this means that the same reference value is used to evaluate all search fields. The simulated interaction process, that executes the tasks described in the experiment, searches for the target box by starting at the first box and considering the boxes one by one. If the considered box is not a target box the experience of the interaction is negatively adjusted. This is repeated until a box that contains database information is found. The experience value at that moment is the result of the evaluation (figure 3.7). A negative outcome signals that the interaction is worse than the reference for this type of interaction, a positive hedonic tone means that the target was found faster than expected.

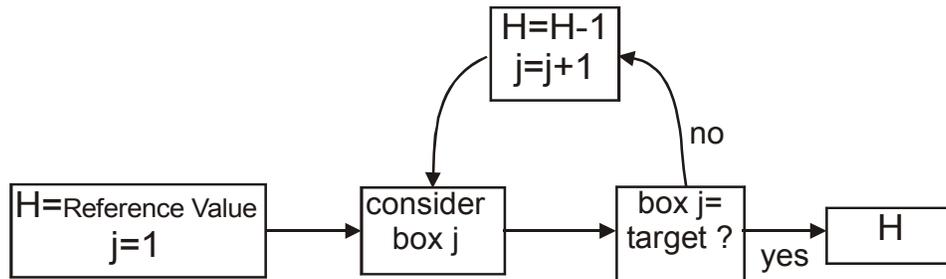


Figure 3.7: The simple version of the hedonic tone simulator

To investigate whether such a simple mechanisms exhibits a behaviour similar to that of the participants, 80 simulated subjects were given all 45 possible stimulus combinations of 1, 2 or 3 boxes containing the database information within a total of 1 to 16 total boxes. The position of the targets was randomised. The synthesised evaluation mechanism shows a lower evaluation for larger search fields. The mechanism also showed a better evaluation for search fields in which more target boxes were present. In that respect, the evaluation values generated by the simulation are highly correlated to the empirically derived values of satisfaction ($R^2=0.90$; $p<0.01$). A mechanism that simply counts the boxes that have to be considered before a target seems to give a reasonable description of the actual evaluation process for this kind of task. The next step is to look at the function of the evaluation mechanism as an indicator for action adjustment.

3.3.2 Different search strategies

A main assumption in this thesis is that the evaluation of interaction has a function in the optimisation of interaction, more explicitly in action adjustments. To allow adjustment, two strategies for visual search are synthesised. If there are two strategies, then the possible adjustment is choosing the more efficient of these two strategies. Therefore, I chose to simulate two known strategies for visual search: 'random search' and 'systematic search' (Hornof and Kieras, 1997). Using random search, the user skips through the search field erratically until the target is located. In systematic search the user starts at a single location and systematically considers each target (similar to reading a book). Random search is usually preferred for unstructured or small search fields, while systematic search is most suitable for large structured search fields. The problem of random search strategies is that the same non-target can be considered more than once. It is highly unlikely that this will occur immediately; a visual memory is therefore implemented to prevent the last few fields being looked at again (Shore and Klein, 2000). Systematic search on the other hand requires more psychological effort (Wolfe, Alvarez, and Horowitz, 2000). This difference can be simulated by making the consideration of each box more costly, e.g. more time consuming or requiring more effort, in systematic search than in random search.

3.3.3 Simulated strategy selector

The next step is to specify a heuristic strategy-selector. This selector has to choose the best available action in a range of situations. To do this it has only access to the results of the evaluation process. Two different strategies, random and systematic, are compared with different effort and memory parameters. The effort was set to 1.5 for systematic search and to 1 for random search. Four occurrences of the random search were held in memory. The random strategy should therefore result in better interaction as long as a target was found within 4 search cycles after registering the first failed target. The specified evaluation-by-counting mechanism generates a value for the simulated search strategies. The most recent outcome of the evaluation, a hedonic tone (H), was stored in a different memory for each of the strategies.

The difference between the two memorised hedonic tones was calculated as input for the strategy selector (figure 3.8). When there was no preference for an

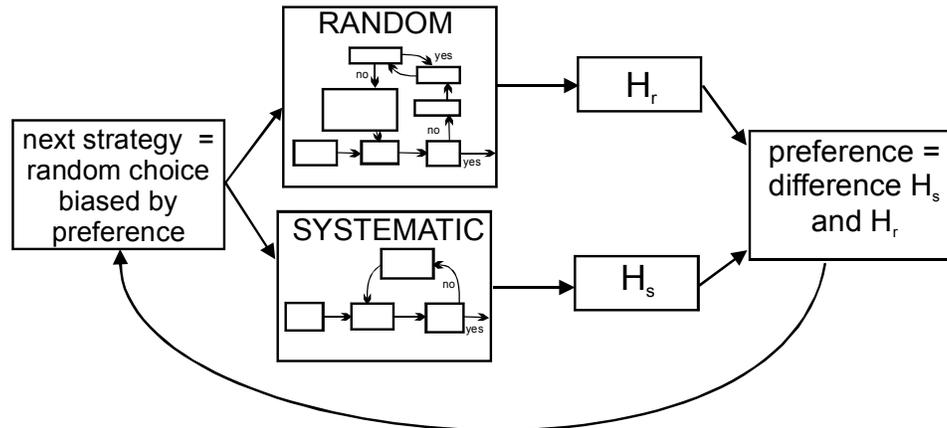


Figure 3.8: Selection system to choose between random and systematic search

interaction the decision was random. The larger the preference, the more the decision is based on the preference alone. In this simulation, this effect was modelled by modifying a random number with the preference, resulting in a biased chance. The biased chance determines the strategy for the next execution of the task. The task, executed using the selected strategy, results in a new hedonic tone for that strategy, which replaces its previous value. The process is repeated with the new hedonic tone. For the same stimulus the most successful strategy will, in general, have the best hedonic tone. The biased chance in the selection prompts this strategy to be chosen more often when the process is repeated.

Results

All 45 possible combinations of 1, 2, or 3 targets from a total of 1 to 16 fields were simulated. In each case the same stimulus was given 80 times. An initial hedonic tone was set at an arbitrary high value for both strategies to make sure each strategy was chosen at least once. The choice selector could either choose one of the two strategies or could remain in an undecided state. I considered a series of more than 20 consecutive choices for the same strategy to be a stable selection strategy (figure 3.9).

Where the simulator made a stable choice, this was generally a choice that was well suited to the applied search field (figure 3.10). The random search was generally favoured for small fields or fields with many targets. When between five

and seven boxes had to be considered the choice became undecided, after which a consistent preference for the systematic search strategy emerged.

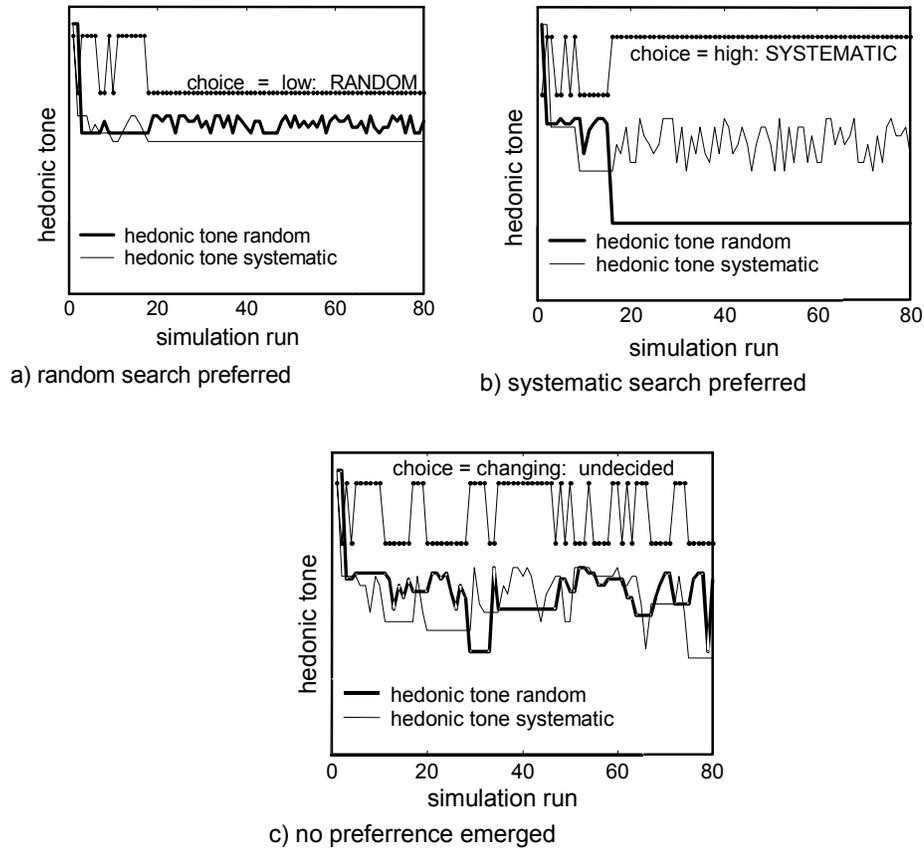


Figure 3.9: Examples of decisions made by the evaluator a) random search, b) systematic search, c) no preference for either strategy.

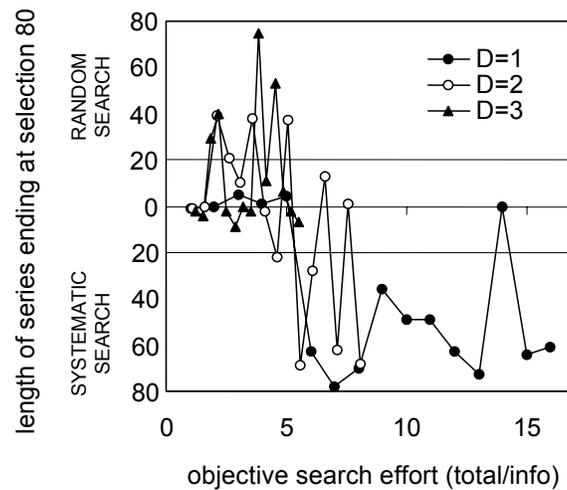


Figure 3.10: Selection of the two strategies at different levels of objective search effort as determined by the ratio total search field to info fields. The number of consecutive choices for that same strategy is counted starting from the strategy that was chosen for the last search task and working backwards. If the same strategy is chosen more than 20 times consecutively this is interpreted as a stable choice.

Discussion of the strategy choice

The strategy selector showed a preference for random search when a target was encountered relatively quickly, while for larger fields systematic search was consistently chosen. This selection was solely based on the stored evaluation values for the different strategies. This means that the strategy selector made the best possible choice, and that the information provided by the evaluator enables the optimisation of the interaction. However, the strategy selector made an occasional sub-optimal choice. These unexpected values were probably the result of the randomness of the stimuli, which can occasionally lead to a lucky series in a single strategy, or alternatively to an early unlucky result in a strategy leading to the permanent rejection of that strategy.

The reference for selection was established by memorising the previous result of each strategy. Even with this simple rule for generating a reference value for the

interaction, the best strategy was often selected, even a simple strategy selector appears to be capable of making reasonable choices. Longer series of trials or better use of all past experience would probably increase the success of the selector since these factors should suppress the effect of random occurrences.

3.3.4 Updating reference values for evaluation

I concluded the previous section with the remark that to improve the operation of the evaluator, past experience should be used more extensively. The gradual change in the evaluation of the stimuli in the experimental study indicates that past and present experiences are integrated in one way or another. In a computational model of the feedback control system this can be implemented by feeding both the experienced hedonic tone and the memorised reference value for evaluation into a function that determines the reference value (RV) for the next evaluation (figure 3.11).

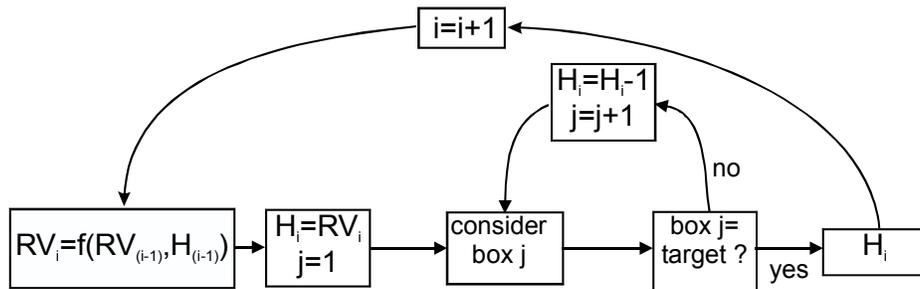


Figure 3.11: A mechanism for updating anticipation.

Similar mechanisms for change in anticipation have been suggested for human actions in the ill-defined cases of gamble-games (Damasio, 1994) or decision processes (Mellers, Schwarz and Ritov, 1999), and were implicitly used in the analysis of the results of the second experiment in chapter 2. The function that combines past reference values and hedonic tones into a new reference value determines how fast and how well people discover the underlying mechanism of the interaction.

A first requirement of such a function is that it takes accumulating experience into account. Further requirements are that it should react as quickly as possible to a new experience and at the same time be fairly insensitive to random occurrences.

The simplest way to update anticipation based on a new experience is simply to replace the memorised reference value (RV_i) with the most recent experience ($H_{(i-1)}$; equation 3.1).

$$RV_i = H_{(i-1)} \quad (3.1)$$

The reference values determined by this function follow experience (figure 3.12a). This does not suffice as a stable reference frame for interaction evaluation, since the reference values are as noisy as the experiences themselves. For more stable evaluations to emerge, prior experience should be used more when determining the reference for interaction. This is the case when a reference value depends not only on the most recent experience but also in part on past reference values ($RV_{(i-1)}$; equation 3.2). A stability index (S) is introduced to determine how large the influence of past reference values is compared to the most recent experience.

$$RV_i = \frac{S \cdot RV_{(i-1)} + H_{(i-1)}}{S + 1} \quad (3.2)$$

Although equation 3.2 results in a stable change of the reference value towards the average of the experienced interaction (figure 3.12b), such a function is either slow in adapting (large values of S) or sensitive to noise (small values of S). It is initially difficult for people to determine which part of the signal is noise and which part of the signal is systematic. To compensate for these problems the influence of the past should count less at first, allowing for quick adaptation, but should weigh more heavily later on, to guarantee a stable reference. This is the case when the stability index has an initial low value, which increases as the number of interaction processes increases (i ; equation 3.3).

$$RV_i = \frac{S(i) \cdot RV_{(i-1)} + H_{(i-1)}}{S(i) + 1} \quad (3.3)$$

For example, if a linear increase of the stability index ($S(i)=c \cdot i$) is used, the reference values converge rapidly upon the average of experience and become stable after a while (figure 3.12c).

This way of integrating recent experience with past reference values meets all of the requirements. Recent experiences are integrated, so reference values converge to actual experience. The function for reference values converges rapidly towards a first estimation for a completely new experience, which is increasingly but slowly

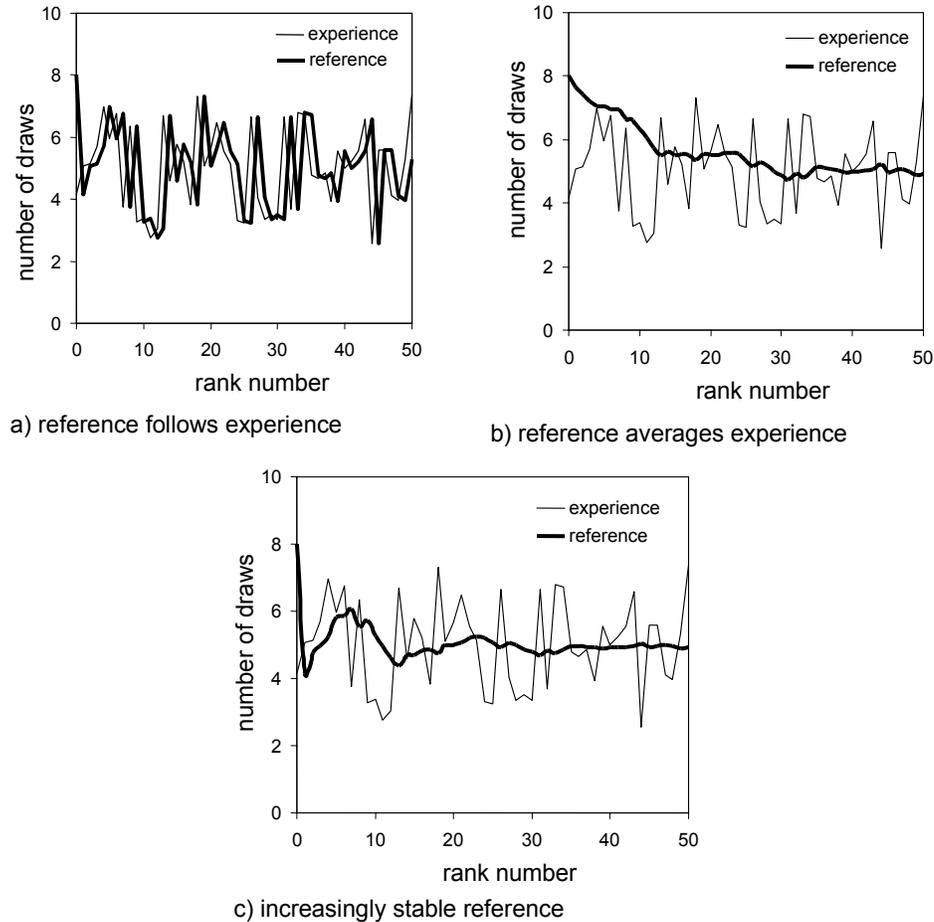


Figure 3.12: Convergence of the reference value of evaluation (thick line) and experience (thin line) based on a) equation 3.1, b) equation 3.2, and c) equation 3.3. The depicted example is the case of 3 targets in a total of 15 boxes.

refined when the experience is repeated. This type of integration of experience is observed when people learn new tasks. New experience counts more heavily in new tasks than in repeated tasks (Davidson, 1995). In the control of tasks that have become skilled, references can be assumed to hardly change anymore (Rasmussen, 1983). A drawback is that with increasing stability the reference values hardly react to new experience anymore, which could lead to the wrong choice of actions, skill type errors (Reason, 1990).

3.4 General discussion of chapter 3

The results of the experiment confirm that users can evaluate ongoing interaction by comparing that interaction with reference values that define the anticipated interaction. The relationship between objective efficiency and satisfaction supports the claim that evaluation gives information about the adequacy of the interaction.

Initially, this evaluation can best be described as being based on the existing knowledge of the evaluation mechanism. Later, as the knowledge about specific properties of the experiment accumulates, the evaluation mechanism can determine satisfaction that is more closely related to objective efficiency. This indicates that the reference for the interaction is updated as more experience about the task at hand is gathered. There are two advantages in modelling an evaluation mechanism that at first estimates the adequacy of interaction based on prior knowledge. First, this means that the mechanism is fairly insensitive to random fluctuations. Second, the use of a memorised reference frame, means that the evaluation can be fast, for example through heuristics. One drawback however is that when prior knowledge does not match the actual situation the correct evaluation of the interaction is not generated straight away. A practical implication is that if a highly adequate application is designed with an interface that leads the users to anticipate inefficiency, they will probably evaluate the quality of the interaction as being poor. With experience, anticipated adequateness will be adjusted towards a more realistic value. However this might never happen if users have decided not to use the interface based on the initial experience.

To learn more about the evaluation of interaction, in this chapter I developed and studied a computer simulation. Evaluation scores derived from the simulation, which counted the number of search fields before a target is found are highly

correlated with the reported satisfaction scores. This finding provides further support for the existence of a functional evaluation mechanism. The results of the simulation were then used as the input for a strategy selection function. When the different strategies had positive and negative elements, the selection mechanism was able to choose between strategies. Even in early choices the synthesised mechanism usually selected the objectively best strategy. To refine the applicability of the evaluation mechanism, three functions were compared that could integrate new experience into the reference values for evaluation. The simplest of these replaces anticipation by new experience. Both other functions integrate prior experience more extensively. The first of these functions reacts slowly on changing environments, the second reacts quickly at first but then stabilises. This last function is probably closest to the actual mechanism of human behaviour. For example, the occurrence of skill-based errors in highly practised interactions can be explained using this integration function. The conclusion that a simple model exhibits functionality similar to empirically observed interaction evaluation and subsequent control, supports the proposed self-regulatory system for user-system-interaction.

Chapter 4: Action adjustments in an ongoing gambling task

Abstract

People assumedly learn to operate an application even when they have little knowledge about its working. The development of actions in such ill-defined situations is interpreted as the outcome of a system for the self-regulation of behaviour, specified as a feedback mechanism. When confronted with a new situation the self-regulatory system should be able to learn the new task with experience. To do this, the system needs to gather information to improve the adjustment of interaction towards an optimum while simultaneously improving the ongoing interaction. In an experimental gambling task, the participants tended to choose a strategy before sufficient information about the game was gathered, which shows that the self-regulatory systems can operate in such ill-defined situations. Emotions were assumed to play an important role as signals for action adjustment in behaviour. The mood of the participants was manipulated to investigate the influence of emotions in interaction control. Those participants that listened to hard rock music took more risks than the other participants. A simulated selection mechanism produced parameter values that were interpreted as being similar to the outcomes of the experimental study.

4.1 Introduction

When I first installed Microsoft Word on my computer, I could only use its simple functions, such as typing plain text. After a while, I discovered more about the rationale and functionality of the application, which allowed me to make better use of the program's possibilities. These days, I can generate a table of contents and I even made the layout for this thesis.

This example indicates that interaction processes change with experience. These changes can be explained if some kind of guidance system is assumed that optimises the interaction process (Carver and Scheier, 1998). This guidance system

allows the users to interact in the best possible way, based on the available knowledge of the interaction process. The guidance system accumulates more information about an ongoing interaction as it progresses. With this increased knowledge, future interaction processes can be improved to meet a higher standard.

To study the guidance system, I adopted the reverse engineering approach (Dennett, 1994). Following this approach, the investigated system is usually specified at three levels (Marr, 1982; Newell, 1990). At the highest description level the guidance system is interpreted as an intentional system (Dennett, 1981). The goal of this system is to optimise the interaction process. To achieve this optimisation at a lower level, the self-regulatory guidance system (Carver and Scheier, 1998) for user-system interaction is specified as a feedback mechanism consisting of three components: a monitoring module, an evaluation module, and

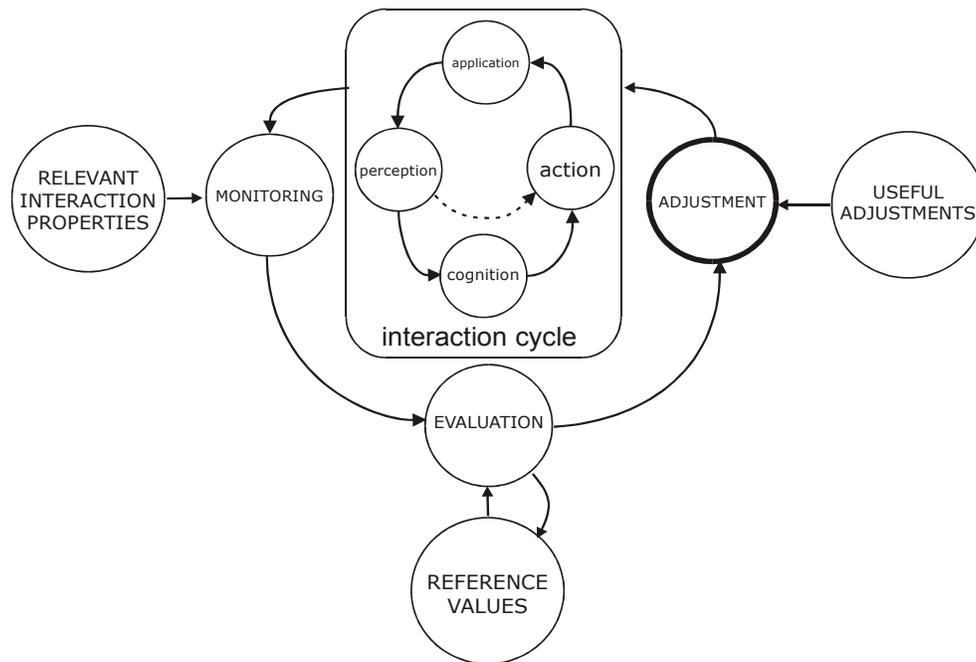


Figure 4.1: A computational model for the regulation of user system-interaction. In addition to regulating the interaction cycle, the reference values for evaluation and the relevant elements for monitoring also have to be determined.

an adjustment module (figure 4.1).

The monitoring mechanism records the ongoing interaction. This is followed by an evaluation mechanism, which determines the quality of the ongoing interaction by comparing the output of the monitoring mechanism with a reference for similar interaction processes. The outcome of this evaluation can be interpreted as a hedonic tone (Johnston, 1999) that is positive if the ongoing interaction is more optimal than the reference value, and negative if the ongoing interaction is worse than the reference value. Based on the evaluation of different interaction processes, the self-regulatory system chooses the interaction process with the highest hedonic tone. In practice this means that the adjustment mechanism will continue the ongoing interaction if it has a positive hedonic tone, but will choose another interaction process if the ongoing interaction has received a negative hedonic tone.

In the study of human-computer-interaction, three types of interaction processes are often assumed (e.g., Norman, 1984). The goals of interaction are established at the highest level. At a lower level a strategy to achieve these goals is specified, which is executed as an action sequence at the lowest level. For example, sometimes on a Sundays I decide to go hiking in the wetlands near my apartment (goal). Before I leave, I plan where my hike will take me. After I have determined this strategy, the hike itself is executed by making a long sequence of steps. Each of these levels of the interaction process is optimised by a, similar, independent guidance system (Dennett, 1994).

The primary aim of my research is to understand the optimisation of interaction at the action execution level. Action sequences can be interpreted as consisting of skilled actions (Norman, 1984). These skilled actions have emerged from earlier strategies that led, through practise, to actions that no longer require conscious control (Rasmussen, 1983). To understand the emerged automated actions, in this chapter I investigate the process leading to automated actions, namely the optimisation of the strategy for a new interaction task.

When optimising the strategy for a new task, each of the mechanisms of interaction control should learn what interaction properties to focus on. The monitoring mechanism should find out what the relevant properties of the interaction process are. The evaluation mechanism should learn the correct reference values for the interaction process, and the adjustment mechanism should know what elements can be adjusted and what the effects of these adjustments will

be on the interaction process. Only when these settings are known, can the interaction be optimal. However, since it is not certain whether the optimal settings will be found within the allowed time, the guidance system has to put effort into optimising the interaction immediately using whatever information is available (Newell and Simon, 1972). If too much effort is spent in working out the exact properties of the interaction, the benefits of the application may not be high enough to warrant the costs of learning how to use them. On the other hand, major benefits could be missed if too little time is spent understanding an application. In the example of hiking in the wetlands, if I spend too much time planning my hike using maps, the day may be half over before I start to walk. On the other hand, if I plan not at all, I might get stuck on paved major roads and miss all those interesting small trails that I would have found if I had studied the map better.

These two requirements to achieve optimal interaction can be executed best if the self-regulatory system is able to initiate the best possible action strategy at any time, even if it does not yet have sufficient information. The self-regulatory system can do this by initially applying simple heuristics. Only at a later stage does the gathered knowledge allow more elaborate processing leading to understanding the problem in more detail. In the hiking example, a heuristic could be that every time I see a small trail I turn off the road.

Different types of heuristics can be applied in a search task where there is not enough information to make a rational decision. Two heuristics search methods are associated with problem solving tasks where no information about the task is available (Winston, 1992). The first is the depth-first search, which is a method to continue the search along the lines of the best-known possibilities. If this heuristic finds a successful strategy it continues following that strategy without ever exploring other possibilities. The drawback is, that the depth-first heuristic might find an unsuccessful strategy and after its discontinuation would still have gathered little knowledge about the search space. The second method is the breadth-first search heuristic, which is the selection of different options leading to the most knowledge about the search space. The breadth-first heuristic explores a large proportion of the possibilities, which can mean that it takes longer to find an optimal solution. The self-regulatory system determines the success of the chosen heuristic in optimising the task and, if necessary, can change it (Cohen, 1975).

To be able to control the application of these heuristics, the self-regulatory mechanism must be able to determine what an improvement of a solution is. In chapters 2 and 3 I have shown that the self-regulatory system in well-practised interaction tasks can rank different interactions on their adequacy. For more complex tasks, Damasio (1994) argued that the anticipated feeling (somatic marker) is an indicator for the adequacy of the selected option. By comparing neurological patients with a control group, it was found that emotional impairment rather than rational impairment negatively influenced a gambling task (Bechara, Damasio, Damasio, and Anderson, 1994).

There are several ideas that emotional experience is not only of importance in the evaluation of the different options. Specific emotions also convey qualitative information about how to adjust the ongoing interaction. Each affective experience has its own specific interaction control function (Oatley and Johnson-Laird, 1987). For example, anger is a signal that the chosen interaction strategy does not result in the desired outcome and that additional effort has to be invested to overcome obstacles (Oatley and Jenkins, 1996). This function of emotions can be interpreted as a heuristic (Epstein, 1994).

To indicate how the affect heuristic plays a role in the different hierarchic levels of interaction control, emotions are interpreted as having three levels (Sloman, 1999). Moods are closely related to goals, medium length emotions are related to strategies, and short-lived affective experiences are related to actions. As an example of this independent control of the hierarchic levels of interaction, consider the sunny Sunday when my mood made me shift my goals from typing this thesis, sitting in the grey concrete university buildings, to taking a walk in the wetlands near my apartment. Happy with my choice, I adopted the strategy of seeking out the sunniest patches of grassland and avoiding the damp gloomy forest edges. When I hit my foot on a log during the hike, I felt a touch of anger and my actions changed, so that for a while I lifted my feet higher as I walked. However, I was still happy to be on the sunny grassland and glad that I had decided not to stay indoors. This interpretation of the signals of moods and emotions gives a tool to understand how sub-conscious regulation of interaction is implemented.

4.2 Experiment

To test the system for the subconscious optimisation of a new task, an experiment is conducted in which participants had to work out the rules of a card game and make a profit. This experiment is a variation on the experiment by Bechara et al., (1994), in which they showed that intuitive, emotional control plays an essential role in the optimisation of a task. The main goal of this instance of the experiment was to investigate in more detail how the accumulation of knowledge in a search task can lead, in a non-rational way, to the convergence of behaviour. Investigating this effect in a new search task, rather than in a practised task, allowed the study of the process of optimisation. In the next chapter automated behaviour will be interpreted using the understanding of the process of convergence of behaviour towards stability.

When the game started, all properties of the interaction could be of importance for example the colour of the drawn card, the deck of the card, the sequence of wins or losses, etc. The only adjustment available was the selection of cards. In addition, because the participants did not know how many cards would be drawn, they had to find a balance between striving for a quick profit and certainty about the best strategy.

The exhibited behaviour should be the result of a trade-off between spending a lot of effort in trying to figure out how the game works, after which informed, profitable choices can be made, and taking gambles from the start. Damasio (1994) argued that the gambles are made based on the emotional anticipation of the outcome of the action. This can initially result in very successful, or in very unsuccessful interactions. If a participant discovers a successful strategy, this should be evaluated as being positive and this strategy should be continued. The control system should interrupt un-successful interaction sequences.

To investigate how emotions influence their actions, the mood of participants in the experiment is manipulated independently of the context of the task. Mood manipulation is assumed to be related to the interaction goals (Sloman, 1999) and, through it, the interaction strategy.

Finally, to explore how the selection mechanism might be implemented, the results of the study will be interpreted using the simulations that were introduced in chapter 3.

4.2.1 Methods and Materials

Participants and design

Forty-eight participants, all TU/e students, took part in the experiment. Sixteen of the participants were female and thirty-two were male, with an average age of 21.6 ($SD=2.2$). Mood was manipulated between participants, by exposing participants to a music sequence, meant to induce one of three different moods, i.e. angry, happy, and sad (Lewis, Dember, Schefft, and Radenhausen, 1995). After listening to the music, each participant played a game in which maximum profit had to be achieved by successively selecting a card from one of four decks of cards with 100 cards in total.

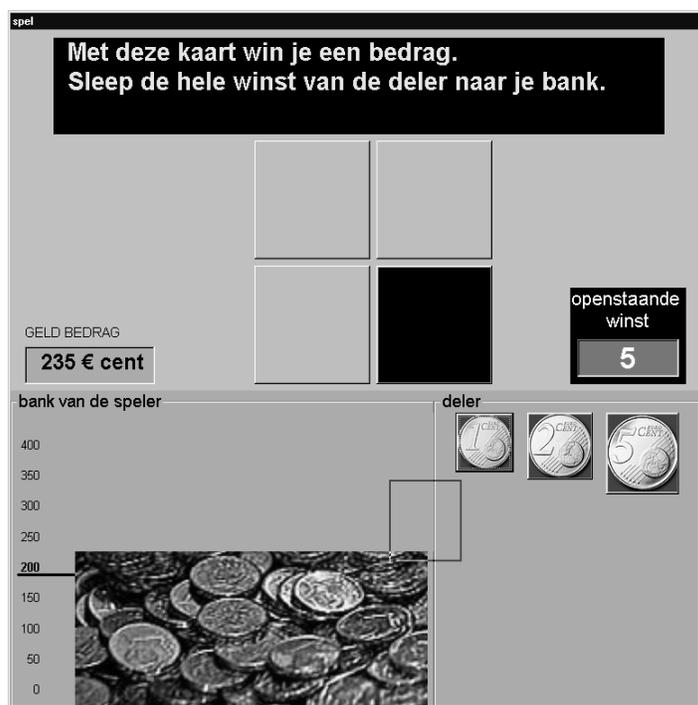


Figure 4.2: The card game interface. A dialogue box is positioned at the top, above the four decks of cards. The bank balance is shown on the left, both numerically and graphically. The dealer is represented in the lower right quadrant, with an indicator of the unpaid wins / losses above.

Interface and rules of the game

The experiment was fully computerised. A graphical user interface was built for the card game using Visual Basic 6.0 (figure 4.2).

The four decks of cards were positioned in a square on the interface. The card game started with an animated card shuffle, after which all decks showed the blank side of a card. A dialogue box at the top of the screen instructed the participants to select a card. After the participant had selected a card by clicking it, the card was revealed to show either a black or a red face. The dialogue box told the participants whether a profit or loss was made. The participants then had to move all profit from the dealer's account to their bank by dragging the images of Euro coins. Losses had to be dragged from the participant's bank to the dealer's account. The participant's bank was displayed simultaneously by a counter showing the balance (centre left of the screen in figure 4.2), and by a pile of coins that grew or shrunk depending on the balance (bottom left of the screen). After dragging all of the required coins, the card faces once again became blank and the participant was asked to draw a new card.

The participants were given an initial balance of € 2. The wins and losses depended on which of the four decks was selected, and were randomly distributed within blocks of ten cards. Table 4.1 shows the different yields of the four decks. As the participants were paid their earnings from the game, the monetary values were adapted from the original experiment, in which (non-paid) yields were between 50 \$ and 100 \$ (Bechara et al., 1994).

Table 4.1: Deck pay-offs per block of 10 cards.

A	B	C	D
5 x 10 € cents	9 x 10 € cents	5 x 5 € cents	9 x 5 € cents
5 x -15 € cents	1 x -115 € cents	1 x 2 € cents	1 x -20 € cents
		1 x 1 € cent	
		3 x -1 € cent	
Net -25 € cents	Net -25 € cents	Net +25 € cents	Net +25 € cents

In the original experiment participants always received some money for a card, but sometimes had to pay a loss as well. This rule was modified so that, instead of separately paying winnings and asking for losses, if applicable, a net outcome was

generated. The reason for this was because it became apparent in a pilot study that participants developed a dislike of the stacks with frequent losses (A and C), because these stacks required the most mouse operations to separately pay the losses and gains.

To summarise, decks A and B were ‘losing’ decks that resulted in a net loss of 25 euro cents per 10 cards; decks C and D were ‘winning’ decks that resulted in a net profit of 25 euro cents per 10 cards. A completed card game consisted of a total of 100 drawn cards.

Procedure

The participants were welcomed into a cognitive laboratory, where they were assigned one of eight cabins with a PC, a 15" display, a standard Microsoft mouse, and stereo headphones. Participants were played about 12 minutes of music to induce different moods (Lewis et al., 1995). The music, which was rock music¹ in all conditions, was assigned to participants based on the order of arrival at the laboratory. The types of rock music differed and could be described as hardcore, meant to induce anger or aggression (anger condition), up-tempo Caribbean and beach music, meant to induce happiness (happy condition) and slow, depressive music meant to induce sadness (sad condition). The experimenter gave a short verbal introduction, in which participants were told that they would receive a reward of € 2 for filling out a questionnaire on their music experience, which would be increased by however much money was in the bank at the end of a subsequent card game. The participants listened to the assigned music and then filled out a Dutch translation² of the PANAS mood scale (Watson, Clark, and Tellegen, 1988) on the computer. The PANAS scale is a two-dimensional mood scale that measures positive affect (PA) and negative affect (NA) using 10 items each. The dimensions are not correlated. After the participants had filled out the scales, they were given on-screen instructions for the card game. The instructions told the participants that they had to draw cards from four decks, that the cards were randomly shuffled before the experiment, that each card could generate either a profit or a loss, that the game had strict rules, and that they could make a profit if they could work out these rules. No further information was given. The

¹ see appendix A for the titles and performing artists of the songs.

² See appendix B for the Dutch translation used.

participants were not told how many cards were available. After these instructions the user interface was displayed and the game began. The order of the decks on the screen was balanced between the participants so that each deck occurred equally often at each position in each condition. After drawing the hundredth card, the participants were asked to fill in the PANAS mood scale once more, after which they were thanked, debriefed, and paid an amount modified by the final bank balance, a total between € 3 and 6. The entire experiment lasted about 35 minutes.

Recorded variables and experimental hypothesis

The items of the PANAS scale (pa1..10, na1..10) were recorded for each participant, both after the music and at the end of the card game. The overall measures of the scales (PA, NA) were calculated. The deck and rank number of each card was also recorded. These variables were used to determine the length of all series, defined as the number of consecutive cards drawn from the same deck.

The main task of the feedback mechanism is to optimise the interaction. To do this, it should be able to determine the success of the current interaction, which in this case is simply the monetary reward. Pleasure can be seen as the outcome of the evaluation of the interaction (Cabanac, 1992), therefore affective experience should be more positive if the interaction is more adequate.

Hypothesis 1: The direction of mood change during the experiment is positively correlated with profit.

Hypothesis 1 states that the evaluation mechanisms of the self-regulatory system can, in principle, determine the success of the ongoing interaction sequence of drawing cards is by relating profit to pleasure. In the specification of the evaluation mechanism, a hedonic tone is assigned to each interaction process. In this experiment this would mean that each deck of cards is given its own hedonic tone, i.e. the anticipation for that deck. Cards should be selected based on this anticipation. In other words, as long as the experience with a certain deck of cards is better than anticipated result of the other decks, the self-regulatory system should ensure that the current deck continues to be selected. The more obviously profitable the selected deck is compared to the other decks, the more stable the selections from that deck should become. Such an ongoing series can be interpreted as the application of a successful depth-first search.

Hypothesis 2: Choices from profitable decks are continued.

Additionally, the combination of hypotheses 1 and 2 means that if both the length of series and change of mood are positively correlated with profit then the length of series should also be positively correlated with mood change. This co-occurrence of a stable action pattern and positive mood change has the same effect as happiness, signalling and stabilising adequate interaction (Oatley and Jenkins, 1996).

However, there is a chance that a losing action sequence is selected based on initial positive experiences. The outcome of the evaluation mechanism should lower the anticipated profit of the losing deck as the losses accumulate. The self-regulatory system should interrupt any ongoing series from a losing deck (stop depth-first exploration) and re-initiate a broader search phase among the options.

Hypothesis 3: Series of losing decks are discontinued.

There are several ways by which the control mechanism can learn to improve the interaction, either through a lasting period of trial and error, by choosing a good action sequence and sticking to it, or by choosing a bad action sequence first and abandoning it. The accumulated information increases in the course of the experiment, which means that later card choices should be more profitable.

Hypothesis 4: There is a positive relationship between the rank number of the card drawn and the number of profitable cards.

Since there is initially little information about the interaction, early selections will probably rely more on the uninformed choice of a heuristic. I assume that by inducing moods, heuristic signals are generated that will influence the initial interaction control. An aggressive mood should initiate strategies to overcome obstacles (Oatley and Johnson-Laird, 1987). This can be interpreted as a signal to explore short-term wins and to be insensitive to losses and past experiences that limit this exploration (Evans, 2001). In the context of this game this means that the participants with an aggressive mood would tend to draw from the losing decks because these appear to have the highest wins. However the losing deck A only generates a gain five out of ten times, and results in high losses in the other five cases. Whichever heuristic is being used, this deck will soon be identified as a losing deck. Deck B, on the other hand, exhibits frequent high gains and only

sporadic losses. The participants that followed the aggressive, gain maximising and loss ignoring, strategy should therefore favour deck B.

Hypothesis 5: The participants who listened to aggressive music will initially prefer deck B, although this is a losing deck.

Happiness signals the continuation of the ongoing state of interaction (Oatley and Johnson-Laird, 1987). In problem solving this leads to careful, risk evasive behaviour if there is a chance that the interaction will diminish the happiness (Isen, 2000). Happiness should therefore lead participants to draw cards from the 'low risk-low gain' decks, which is the profitable interaction since these are the winning decks.

Hypothesis 6: The participants who listened to happy music will prefer low risk-low gain decks.

Sadness is an emotion that signals that the interaction is frustrated, but that no solution is readily available; the current goal hierarchy should be reconsidered and a new approach to goals and strategies should be initiated (Oatley and Johnson-Laird, 1987). This would mean that sadness might lead to two types of behaviour, i.e. a low intensity goal reprioritisation, or an aggressive initial search for a new strategy. I argue that the specific effect of sadness is therefore hard to predict in detail, and the effect of this mood will only be explored.

4.2.2 Results

The Dutch translations of both the PA and the NA scales were found to be reliable, PA $\alpha=0.85$, NA $\alpha=0.83$, and not significantly correlated ($p=0.12$).

On average participants made a small loss in the game, ($M=12$ euro cents loss; $SD=102$ euro cents). The greatest loss was € 1.86; the highest profit was € 2.44.

To give an idea of how the cards were sequentially selected, the choices of three different participants are shown (figure 4.3). One participant found a successful deck after an initial period of trial and error and subsequently stuck with that successful strategy (figure 4.3a). A second participant never got further than trial and error and showed a random card drawing pattern, which can be interpreted as an extreme case of exploring the problem space (figure 4.3b). A third participant had an initial preference for a losing deck but changed to a period of trial and error

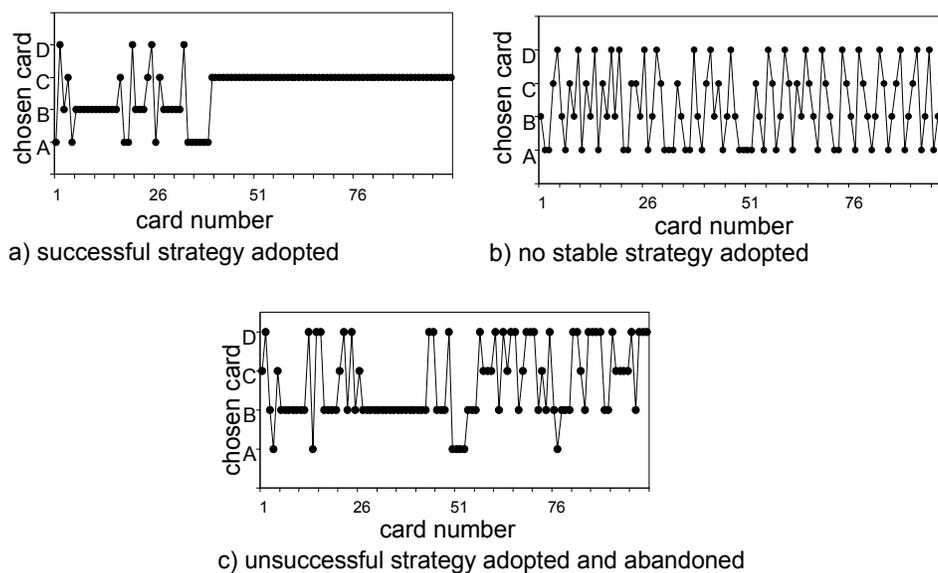


Figure 4.3: Typical examples of different states of interaction. a) the control system found a successful strategy and stabilised this strategy b) the control system did not (yet) find a correct strategy and continued its explorative behaviour, and c) the control system found a losing strategy, abandoned this strategy, and went back to explorative behaviour, ignoring deck A.

that lasted until the last card (figure 4.3c). Note that this last participant had correctly decided that deck A should be disregarded as a search option.

A MANOVA shows that the mood of participants significantly changes during the experiment, $F(2,44)=4.9$, $p=0.01$. The overall mood-change is negative, but this effect is mitigated by a positive correlation between the final bank balance and the change in the PA scale, $r=0.44$, $p<0.01$ (figure 4.4). The participants who develop winning strategies do apparently evaluate this as positive, thus confirming the first hypothesis.

Since profits depended strictly on the choice of deck, the numbers of cards from the different decks are compared. There was a difference between the number of cards drawn from each deck, $\chi^2(3, N=4800)=339$, $p<0.01$. For each deck it was tested whether the number of cards drawn from that deck differed from a random selection (25 out of 100). The losing deck A was chosen significantly below chance

level, $M=14.9$, $t(47)=10.1$, $p<0.01$. The losing deck B was chosen more often than the chance level, $M=33.1$, $t(47)=3.6$, $p<0.01$. The winning decks C, $M=28.1$, $t(47)=1.0$, $p=0.30$, and D, $M=24.0$, $t(47)=0.5$, $p=0.65$, did not significantly deviate from the chance level. The participants appeared to notice that they should stay away from deck A, but were attracted to deck B.

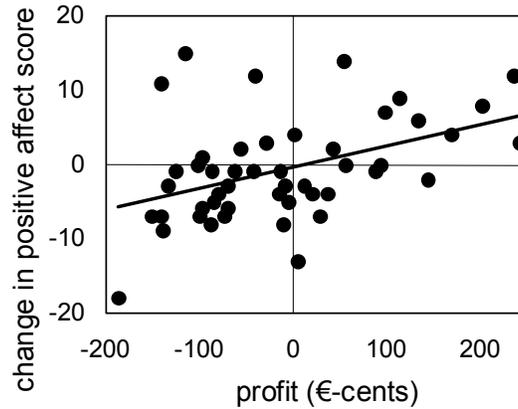


Figure 4.4: Relationship between profit and the change of positive affect.

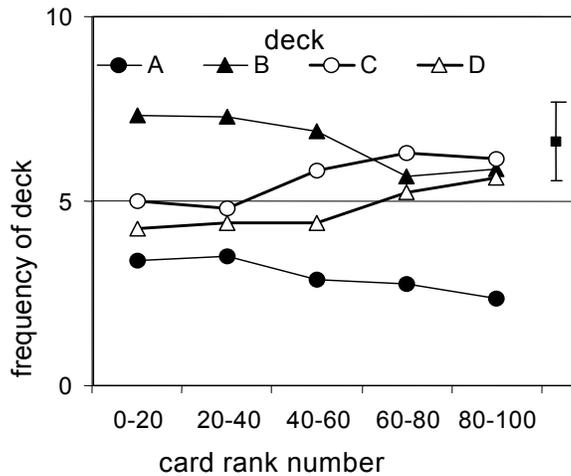


Figure 4.5: Selected decks for blocks of 20 cards. The error bar indicates the significant difference between observations at the 0.05 level.

The development of choice is investigated by looking at blocks of 20 cards (figure 4.5). This showed an overall decrease of the losing decks A, $F(1,45)=9.6$, $p<0.01$, and B, $F(1,45)=4.6$, $p=0.04$, as well as an increase of the winning deck D, $F(1,45)=6.9$, $p=0.01$. A trend towards an increase of cards drawn from deck C was also found, $F(1,45)=3.0$, $p=0.09$. All these changes led to more profitable card selections. This confirms hypothesis 4 that, on average, the participants develop more profitable selections in the course of the experiment.

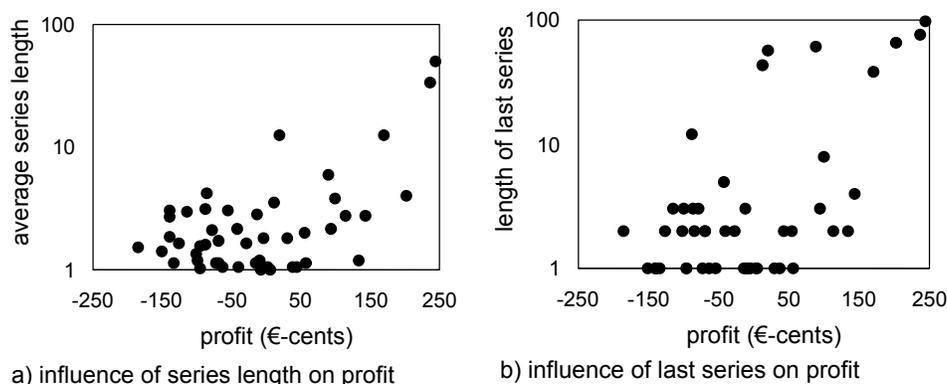


Figure 4.6: Relationship between profit and length of series plotted on a logarithmic scale. a) average series length b) the length of the series ending with the last card is considered. Note that because of average the longest series in panel a) are necessarily shorter than in panel b). Very long series occur only in profitable situations. All other situations are undecided.

To determine whether a successful strategy was discovered, the length of the series of consecutive cards drawn from the same deck is investigated. The average length of a series per participant varied between 1 and 33. On average, the participants changed decks 57 times, $SD=30$. Two participants changed decks 99 times, which meant that no two consecutive cards were drawn from a single deck throughout the entire experiment. Most participants changed decks less frequently. One participant only changed decks for the first two cards and then started drawing from one (profitable) deck, which he continued doing for the remaining 98 cards of the game. There was a significant positive correlation between the average length of a series and the profit at the end of the game, $r=0.56$, $p<0.01$. This

correlation was slightly higher if the last card in the series was also the last of all of the cards, $r=0.65$, $p<0.01$. Although these correlations are mostly due to the extreme values (figure 4.6), these findings confirm hypothesis 2 that long series generally only occur if they are profitable. The higher correlation towards the end indicates that more profitable series are drawn later in the experiment. When comparing the rank number of the card at the end of a series with the length of that series, I found that the length of series increased during the course of the experiment, Kruskal-Wallis $H(99)=181$, $p<0.01$. This, together with an increase in profitable cards indicates that participants were starting to converge on a profitable strategy.

The typical properties of longer series were studied in more detail by only considering series of 10 cards or more. The series with at least 10 cards are longer than average, and these series are singled out because they span a complete combination of losses and wins for a deck. There were 34 series of at least 10 cards. Nineteen of these series were drawn from the losing deck B, and 15 were either from winning deck C or D. When comparing the actual length of the series containing at least 10 cards, long series from the winning decks were significantly longer, $M=41$, than those from the losing deck $M=16$, Mann-Whitney $U(n=34)=47.5$, $p<0.01$ (figure 4.7).

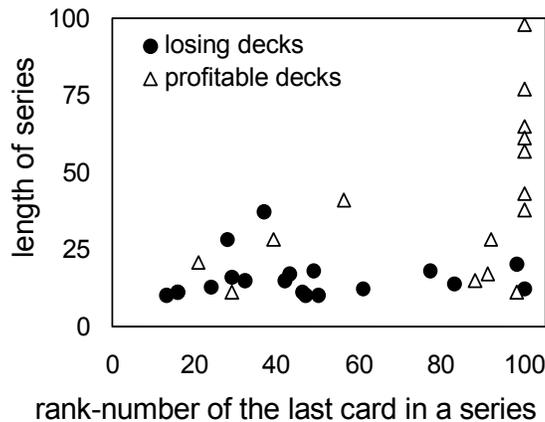


Figure 4.7: Rank number of the last cards of the series. Only series of at least 10 consecutive cards from the same deck are considered. Most losing series ended early in the experiment, while profitable series continued.

This confirms that winning series are continued, but losing series are discontinued. Further evidence for this statement is found by investigating the rank-number at which these series ended. Of the series from the losing deck B, only 5 out of 19 ended after card 50, which was significantly lower than the 12 series out of 15 that ended after card 50 for the winning decks C and D, $\chi^2(1)=4.8$, $p<0.05$. The 6 longest series ended at card 100; each of these series was drawn from one of the winning decks. This confirms that although the participants generate long series of cards from the losing deck B, they discover that drawing long series from losing decks was not a successful strategy, while drawing long series from winning decks was.

Mood manipulation

Using a MANOVA a significant difference was found in mood between the participants after listening to different types of music, $F(2,45)=3.7$, $p<0.03$. This difference was based on the participants who listened to sad music. No difference on the PANAS scales between happy and angry music could be found (figure 4.8).

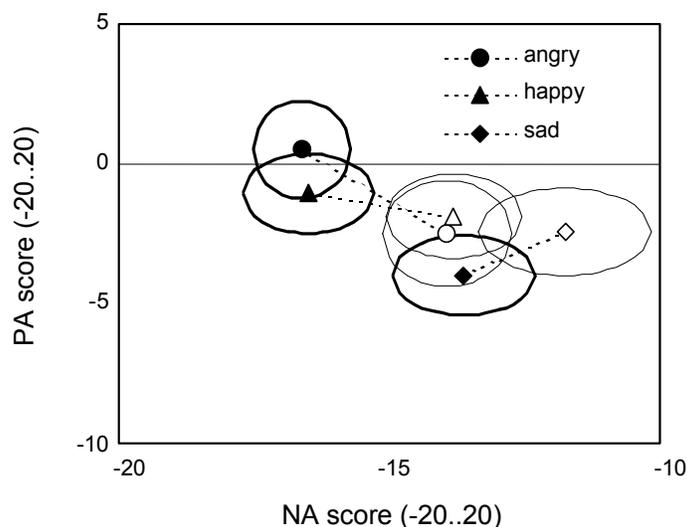


Figure 4.8: PA and NA scores immediately after listening to music (closed markers) and after finishing the game (open markers), for the three music files. The angry music scored higher than expected on the PA scale. The ellipses depict 1 standard error of mean; when ellipses are not touching there is a significant difference at the 0.05 level.

To investigate whether there were any differences between the happy and the angry mood, the different items of the scale were reviewed separately. Contrasts to compare the scores on single items of happy and angry music indicate that the participants who listened to angry music were more excited, $t(2,45)=2.4$, $p=0.02$ (pa4), more active, $t(2,45)=2.1$, $p=0.04$ (pa7), and more hostile, $t(2,45)=2.2$, $p=0.03$ (na10), than the participants who listened to happy music. This failure to thoroughly check the manipulation might be due to a liking for the 'cool' aggressive music (as remarked by a participant during the debriefing), that became reflected in the mood scores, or it might indicate that the PANAS scale is not the best scale to measure differences between emotions with similar arousal (Green and Salovey, 1999). When the assumed emotions are interpreted as locations in the two-dimensional emotion space with the dimensions activation and pleasantness (Russell and Lemay, 2000) the PANAS scale only differentiates the effect of music on the activation dimension, where happiness and anger are close. Therefore it cannot conclusively be stated that the aggressive music indeed induced anger, or that the happy music indeed induced happiness.

After the game there was no longer a difference in mood between the participants who listened to different types of music, $F(2,45)=0.9$, $p=0.43$.

A trend is found that music influences the overall yield, $F(2,45)=2.6$, $p=0.09$. The differences in the final loss and profit varied slightly between the conditions: angry ($M=57$ cents loss, $SD=77$), happy ($M=24$ cents profit, $SD=116$) and sad ($M=2$ cents loss, $SD=109$). If the bank balance of participants is studied throughout the experiment, it is found that on average, the participants in the angry condition had a lower momentary yield throughout the 100 cards of the experiment, (repeated measure ANOVA) $F(2,45)=5.0$, $p=0.01$. This means that the participants who listened to aggressive music were worse off than the participants who listened to the other music types, but only in the course of the experiment. Later on in the experiment, they partly compensated for their earlier losses.

To study the difference between moods, the number of cards from the different decks is compared. The number of cards drawn from deck B was significantly higher for the angry condition, $\chi^2(2, n=1585)=49$, $p<0.01$. The participants in the angry condition also encountered more losses in deck B ($M=4.4$), than those in either the happy, $M=2.8$ or the sad music condition, $M=3.4$, $\chi^2(2, n=170)=6.2$, $p<0.05$. This difference occurred mainly in the first 20 cards (figure 4.9). The

participants who had heard aggressive music initially chose deck B more often, but in doing so learnt to avoid this deck. This finding confirms hypothesis 5, with the reservation that it was not possible to determine conclusively that the participants were actually put in an aggressive mood by the music.

The participants who listened to sad music preferred cards from deck C to cards from deck D, $\chi^2(1, n=858)=24, p<0.01$. The participants in neither other condition exhibited a preference for one profitable deck over another.

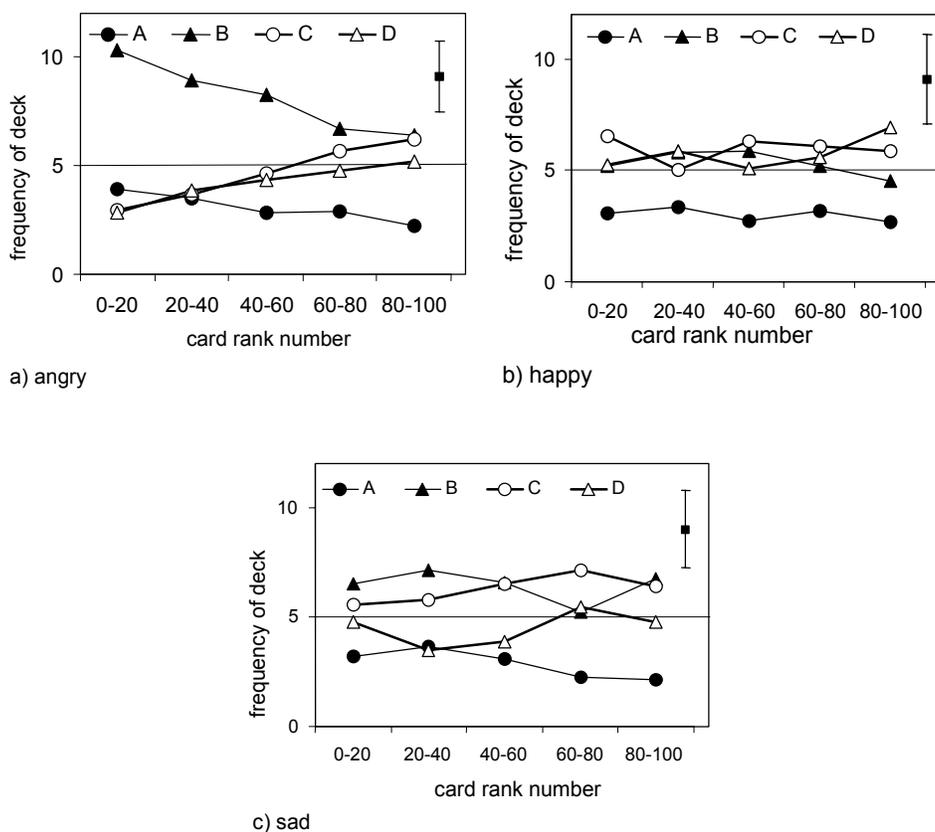


Figure 4.9: Deck selection for blocks of 20 cards for a) angry music, b) happy music, c) sad music. The legends indicate the marker for each of the decks. The error bars indicate the significant difference between observations at the 0.05 level.

4.2.3 Discussion of the experimental results

To optimise their behaviour, the participants have to be able to determine the adequacy of the undertaken interaction. The positive correlation between profit, an indicator for interaction adequacy, and the change in the mood scores, confirms hypothesis 1, i.e. that participants can evaluate their strategy. The fact that this change was measured on a mood scale indicates that the evaluation is at least partially affective. The participants seem to have access to information about the adequacy of their interaction. However, it is not clear precisely how wins and losses are evaluated. This evaluation is probably not linear, and negative outcomes are likely to be evaluated as larger than similar positive outcomes (Shafir and Tversky, 1995). There were indications that the participants used the information they gathered during the game to improve their interaction. The participants who won the most did so by making the correct choice at once, when they could have had only very little information about the rules of the game. Such a long series of profitable cards can be considered to be the result of the choice of a depth-first exploration of the possibilities. Even without being able to work out the possibilities, the accumulation of profit proved to be enough information to stay with the chosen strategy, confirming hypothesis 2 that profitable series are continued. On the other hand, the participants who made an early choice for a losing deck worked out that ongoing selections from that deck were generating a consistent loss and they stopped drawing from that decks (hypothesis 3). The self-regulatory-mechanism apparently selects a strategy of only drawing continuing series of cards from a single deck when that strategy results in long-term profit. If losses accumulate, the selection is expanded to include the other options. Across the whole group, a positive correlation between the rank number of the drawn card and the series length was found. When interpreting the co-occurrence of these findings, it becomes apparent that longer and more profitable series were emerging, which is the start of convergence towards the more stable and profitable strategy from selecting long series of the profitable decks.

Many of the participants, however, could not find the correct strategy to make profit within the 100 cards of the game. This was mainly because they favoured the losing deck B. A possible explanation is the make-up of deck B, which had a high frequency of high profits (90% of the cards). On average the participants

encountered only 3.5 losses in deck B compared to 29.5 wins. The participants seemed to need a lot of information before discarding deck B.

The induced moods did influence the number of cards drawn from the different decks. Anger is a signal to find a new way to start optimising behaviour; the control system should go for easy profits and not be put off by initial losses. Hypothesis 5 was confirmed by the results of the participants who listened to the aggressive music who had a clear preference for the deck that yielded the most frequent, high, wins (B). Participants that listened to happy music exhibited careful risk-evading actions, by either sticking with proven profitable decks or with low risk random drawing. These results, however, have to be taken with some reservations since the applied mood scale could not distinguish between participants who listened to happy or aggressive music. This might be due to an inherent shortcoming in the scale used, or by the fact that the participant cognitively changed their mood reported for the hard-rock music towards being more positive than they actually experienced the music at a sub-conscious level. The differences found might also be related to a completely other property of the music, such as the speed. To substantiate the claims about the emotions, research should be carried out to find out whether different 'high-arousal' types of music have a different influence on the mood of participants.

Another effect of the mood is found on participants who were exposed to sad music. These participants chose significantly more often from deck C than from the winning deck D. The difference between the decks is that deck C produced only low yields (5, 2, or 1 euro cents) and very low losses (1 euro cent). A possible explanation is that the participants in the sad condition preferred low-intensity experiences while changing their goal priority (Oatley and Johnson-Laird, 1987) and maintained a low mental activity (Izard and Ackerman, 2000). On the other hand, the participants in the sad condition also favoured deck B. A possible explanation is that sadness leads to the search for new goals, which can result in a period of aggression (Blumberg and Izard, 1986) to overcome initial obstacles. This explanation indicates that the effect of sadness on behaviour is complex.

A limitation, for the behavioural measures reported in this chapter, is that less clear results were found than in an earlier version of the same experiment (Bechara et al., 1994). The participants in the version of the experiment reported in this chapter, in particular, exhibited a preference for the losing deck B compared to the

results of Bechara et al.. There are several factors that may have contributed to this difference. One explanation for this observed difference is that the participants in my experiment were students from a university of technology. The technological environment puts a heavy emphasis on rational and logical reasoning, while it was argued that intuitive decision-making is faster in this type of task (Bechara, Damasio, Tranel, and Damasio, 1997). The rational thinking of participants might have been increased by the fact that the experiment was fully computerised, since a computer could give an additional logical framing to the experiment. This interpretation is supported by remarks of some of the participants during the debriefing that they were looking for complex combinations of the number of consecutive wins, combined with cards of the same colour and their relationship to the size of the loss and profit. A related point is that my instructions did not indicate which elements of the game participants should focus on; which might have distracted the participants. Schmitt, Brinkley and Newman (1999) encountered similar problems and found that the final instructions of an earlier instance of the experiment (Bechara et al. 1997) included hints to look specifically at wins and losses. Another explanation for the difference in card choice could be a motivational issue. Whereas a human experimenter conducted the original experiment, my experiment was fully computerised. This could result in a lower level of social compliance to achieve the aims of the experiment and therefore to a lower motivation in the participants. A lower motivation can result in less clear relationships between the available information and the behaviour (Petty and Cacioppo, 1981). In an attempt to motivate the participants, they were paid according to the outcome of the game. The monetary amount of the reward was therefore adjusted. This might have influenced the experience of wins and losses. Another adjustment to the game might also have played a role. In the study of Bechara et al. (1994) each deck yielded the income determined for that deck at every card. Losses had to be paid separately only when they occurred. This means that the positive gain of a deck might be more easily discovered, after which only the rules for losses had to be found. In the reported experiment yields and losses were not handled separately but as a net value, which might have caused a less clear interpretation of the structure of pay-offs of the decks.

Nevertheless, the mechanism of card selections was well on its way to discover more profitable strategies, although it was not there yet. Furthermore, anticipated

effects of anger, leading to high-risk choices, and happiness, resulting in risk avoidance, were found. This supports the idea that affect heuristics are involved.

4.3 Simulation

In this section, the results of the experimental study will be interpreted as the outcome of a selection mechanism, which enabled me to investigate to extent to which the selection of cards can be described by a mechanism similar to the one was theoretically specified in chapter 3 (figure 4.10).

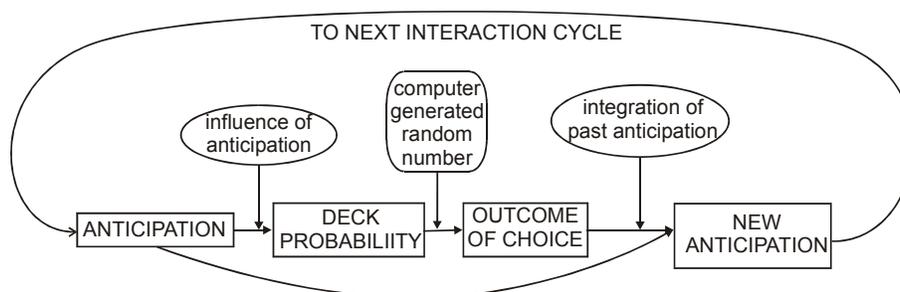


Figure 4.10: The simulated mechanism with two parameters to describe the influence of experience.

The mechanism selects an action, in this case the deck from which a card is drawn. The mechanism needs some knowledge of the game to select the most profitable card. The mechanism therefore has a memory for the anticipated outcome of each of the decks, called a hedonic tone. This anticipation is based on previous experiences as specified in equation 3.2. The difference in anticipation is modified by the first parameter, influence of anticipation, which determines how much the differences in anticipation influence the selection process, this is a parameter for the effect of the random value that was used in the selection mechanism as specified in paragraph 3.3.3. Based on the differences in anticipation between the decks, and the parameter for influence, the probability that each deck will be selected is calculated. Decks with a better anticipation will have a higher chance of being selected, the larger this, the higher the chance. A random number generated by the computer, determines which deck is actually chosen based on the probabilities of the decks. The selected deck generates an outcome, similar to the one in the experimental task (see table 4.1). The outcome is integrated in the

anticipation of the outcome of that deck. A second parameter, the stability parameter (S) of equation 3.2, models strong the effect of past occurrences is compared to a new experience. If this parameter is small, past occurrences will be readily forgotten. If the outcome of the deck is better than expected, the anticipation for that deck will increase. A positive experience, therefore, means that the deck has a higher chance of being selected during the next selection. A negative experience, on the other hand, results in a smaller chance of the same deck being chosen again. The anticipated outcomes of the three decks that are not selected are also updated, as if they have generated a zero outcome. This means that decks that are not chosen revert back to a zero anticipation level over time. In this way a decaying memory for past occurrences is integrated into the simulation. The process starts over again with the new anticipation values, so new selection can be made. An example of the selection of a single card and the integration of this experience into the anticipation for the next draw is computed in table 4.2.

Table 4.2: Example computation of a single card selection, and the subsequent updating of anticipation.

ANTICIPATION	$A = 5 / B = 0 / C = -5 / D = 0$
profit / deck:	
PARAMETER 1	$= 1$
Influence of anticipation	
DECK PROBABILITY (%)	$A = 25 + 1 \cdot 5 = 30 / B = 25 + 1 \cdot 0 = 25 / C = 25 + 1 \cdot -5 = 20 / D = 25 + 1 \cdot 0 = 25$
RANDOM NUMBER (0..100)	$= 42 \Rightarrow$ DECK B SELECTED
OUTCOME	$YIELD = 10$
OF CHOICE	
PARAMETER 2	$= 4$
integration of past anticipation	
NEW ANTICIPATION	$A = (4 \cdot 5 + 0) / (4 + 1) = 3 / B = (4 \cdot 0 + 10) / (4 + 1) = 2 /$
Profit /deck	$C = (4 \cdot -5 + 0) / (4 + 1) = -3 / D = (4 \cdot 0 + 0) / (4 + 1) = 0$

The task of the self-regulatory mechanism is to control how the selections are undertaken in a way that leads to the most profitable interaction. To do this, the self-regulatory system can control the two parameters in the simulation. The first parameter that can be adjusted is the influence that the anticipated pleasure has on the selection of the decks. The more important the ongoing task is for the user's current well being, the higher the value of this parameter should be. The second parameter is the influence of past experiences on the anticipation of future actions.

If this influence is high, the past occurrences are taken strongly into account, making a stable, but not very flexible setting.

The self-regulatory system should define a different way of selecting cards in different situations, in this case induced by mood differences. To test the specified card selection simulation, the two parameters were estimated that fit the different induced moods the closest. Epstein (1994) argued that affect can be interpreted as a heuristic that determines the adjustments of action (Oatley and Johnson-Laird, 1987). Mood induction should therefore result in different parameters, set by the self-regulatory system, based on the interpretation of the induced mood. If the function of anger is considered, this would mean that past experience should not be taken into account in defining a new anticipation of the task. This means that the parameter for integration of past occurrences should be the lowest when fitting the parameters to the participants who listened to aggressive music.

4.3.1 Outcome and discussion of the simulation

The simulated mechanism had to select 100 cards, repeated 48 times. This resulted in a number of cards being drawn from each deck. The outcome of the simulation was compared with the number of cards drawn from each deck that were recorded in the experiment. The parameters were estimated by iteratively minimising the squared difference between the empirical and the simulated data. The simulation rendered overall results that fit the overall outcome of the data fairly well (table 4.3), although the data remained inherently noisy.

An idea of how the simulated selection mechanism chooses cards is given in figure 4.11. The anticipated outcome for each deck at any moment is plotted in the

Table 4.3: Fit parameters and variance of the experimental data that is explained by the parameters compared to random card selections.

DATA	Explained variance (%)	Influence of anticipation	Integration of past experience
All experimental	94%	0.15	2.9
Angry music	96%	0.11	1.1
Happy music	97%	0.17	5.3
Sad music	75%	0.14	4.9

left hand figure. The decaying memory is clearly visible in the change of the negative values towards the neutral over time. The result of this decay is that after a while an option is chosen again, even if a very negative experience was recorded. This gradual improvement of negatively anticipated decks prevented the generation of long sequences of cards.

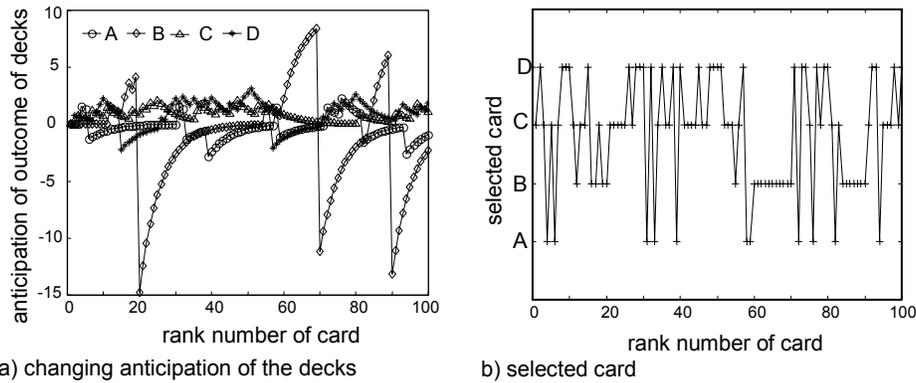


Figure 4.11: Example of the simulation. a) the anticipation for each of the decks. b) selected card. The parameters were those found for the happy music condition.

The participants who listened to aggressive music had the lowest parameter for the integration of the past experience. The influence of the anticipation was also lowest for this group. This outcome indicates that accumulating experience has little predictive value for the behaviour of aggressive people, which agrees with the functional idea of aggression, that by ignoring past experience it helps to generate behaviour that is needed to break through barriers (Evans, 2001).

However, there were limitations to the simulation that prevented a more detailed interpretation of the parameters. First, no long series were generated, while they were observed in the experiment. This is probably because the simulation did not differentiate between gambling early on and going into a prolonged period of trial and error. The problem is probably that in reality two distinct modes of selection alternate (Sloman, 1996), whereas the simulation only had one parameter set. To overcome this problem, the simulator could be nested in a secondary control system to qualitatively determine how different heuristic or deliberate selections should be made. This adjustments is put off to future work

because a secondary control system would further complicate the simulation and would rapidly lead to more complex strategy generation (Dennett, 1995); the data set in this experiment is not extensive enough to fit reliable parameters for such a complicated mechanism.

Another limitation of the simulation is that it could not distinguish between the two objectively equally good options C and D, which received different scores for the sad music (32% versus 22%). The effect of sadness on choices was argued to be a combination of two different strategies, i.e. the preference for low intensity stimuli or generating anger to start a new sequence. This problem can probably also be solved by adding a secondary control system to determine whether low intensity or aggressive solutions are favoured.

A final limitation is related to the evaluation of the outcome of the game. In this simulation it was assumed that the monetary yields were straightforwardly evaluated. However, there is evidence that negative outcomes should be perceived as more significant than corresponding positive outcomes and that extreme values are basically perceived as 'very big' (Shafir and Tversky, 1995), which complicates the generation of the anticipation for the next action.

A more complete mechanism for strategy selections can be designed if these problems are dealt with.

4.4 General discussion of chapter 4

In an ill-defined problem-solving task people optimise their behaviour in a way that can be explained by a self-regulatory system that sets parameters that determine how the problem-solving task should be executed. To do this, the self-regulatory system selects strategies for the interaction, such as trying to make a profit as soon as possible, or alternatively working out the rules of a game to be able to make a profit later on. The participants showed high-risk early choices and more informed choices later on. This indicates that the psychological mechanism that performs the same tasks as the synthesised system probably acts on a mix of heuristic and more deliberate solutions. This can be interpreted as a dual-process, which allows the system to select reasonable actions at any moment (Slooman, 1996). The ongoing control ensures the adjustment of heuristics that result in bad choices and the confirmation of those that results in successful choices. This means that good interactions are enforced by positive feedback while bad interactions are

interrupted by negative feedback. With increasing information about the task, a more stable action selection emerges. Not all of the participants reached the winning stage. There are indications that this might be due to motivational problems or the complexity of the task.

A simulation generated results that were on average the same as those of the participants. However, the simulation was not able to exhibit the same behaviour as the individual participants. This is attributed to a limitation of the simulated mechanism, which only uses one mode of processing, whereas participants possibly use both intuitive, heuristic and rational processes (Sloman, 1996; Bechara et al., 1997).

A possible mechanism that plays a role in the self-regulatory system is the affect heuristic. The induced change in mood resulted in different strategies, which indicates that the outcome of the evaluation is at least partially affective and that mood inductions influence strategic decisions. The assumed mood induction did not show the anticipated differentiation on the used mood scale, however. I suggest that this is a shortcoming of that mood scale. This idea should be substantiated by future research into the influence of music on self-reported mood and the exact measurement of this influence. Mood induction is most important at the start of the experimental task, when the simple affect heuristics is the only available cue and is therefore apparently used as the starting point for the search process. Later, as more information has been gathered, the effect of mood induction diminishes, the chosen actions are based on strategies that have been generated by the self-regulatory system based on accumulated experience.

Chapter 5: Regulation of mouse-cursor-movements in different contexts

Abstract

Behaviour can be interpreted as the outcome of a self-regulatory system for interaction optimisation. Such a mechanism should lead to effective and efficient settings of interaction process parameters. This mechanism was tested in an experiment, in which participants showed that they were able to execute three different versions of a mouse-cursor target-acquisition task. The participants were shown to have stereotypical combinations of movement time and movement precision, indicating that interaction adjustment is based on varying a personal action pattern. By computing the force and co-contraction and comparing these parameters for effort, it is illustrated that a changed interaction results in a different distribution of cost factors. In a second experiment mouse-cursor-movements were measured but the participants were not given an explicit target-acquisition-task. Even without a specific target-acquisition-task, cursor-movements exhibit properties of converged action programs, thus confirming the control of these interactions. The role of emotions as heuristic signals for the adjustment of action programs for mouse-cursor-movements was studied by inducing different moods, which resulted in predictable changes of movement time and precision.

5.1 Introduction

If a user wants to make drawings using programs like Microsoft Paint or CorelDraw, the user needs to be able to work very accurately. When the same user starts playing games, such as Pong, hunting, or racing games high-speed interaction is required. All these tasks will probably be executed using a computer mouse. To be able to meet the specific demands for these tasks, users have to adapt

their actions to fit the required interaction. Users are therefore assumed to have a system that regulates behaviour.

In this chapter the regulation of automated processes in different contexts is studied by assuming an intentional description (Dennett, 1981) of a self-regulatory guidance system for behaviour (Carver and Scheier, 1998). I used the top-down approach of reverse engineering (Dennett, 1994) to study the self-regulatory guidance system for the interaction between users and computer applications. Following this approach an intentional system is synthesised that optimises interaction. The behaviour of this system is then compared with the behaviour exhibited by users, to determine the extent to which the synthesised mechanism describes the actual control of user-system interaction.

The self-regulation of the interaction between users and applications is implemented as a feedback control loop, which is specified using three mechanisms: monitoring of ongoing interaction, evaluation of adequacy of interaction and adjustment of interaction (figure 5.1). The monitoring mechanism perceives the relevant elements of the interaction (Nijenhuis and Blommaert, 1997) and passes them on to the evaluation mechanism (chapter 2; Fischer and Blommaert, 2001). The outcome of the monitoring mechanism is compared with reference values in an evaluation process (see chapter 3). The reference values of evaluation are either determined by prior experience (Mellers, 2000), which allows a flexible adaptation of action sequences, or are determined by changes in goals or strategies. The results of these comparisons are measures for the adequacy of different elements of interaction. These measures can be interpreted as hedonic tones (Johnston, 1999). The hedonic tones are aggregated and form a single quantitative evaluation value for the interaction, referred to as: pleasure (Cabanac, 1992). Besides quantitative information on how good the interaction is, the evaluation mechanism also gives signals indicating how to improve the ongoing interaction, which can be interpreted as a heuristic function of emotions (Epstein, 1994; Oatley and Johnson-Laird, 1987). A positive evaluation, accompanied by happiness, signals that the ongoing process should be continued, negative feelings, accompanied by anger or sadness, give a signal to change the ongoing interaction in specific ways (Oatley and Johnson-Laird, 1995).

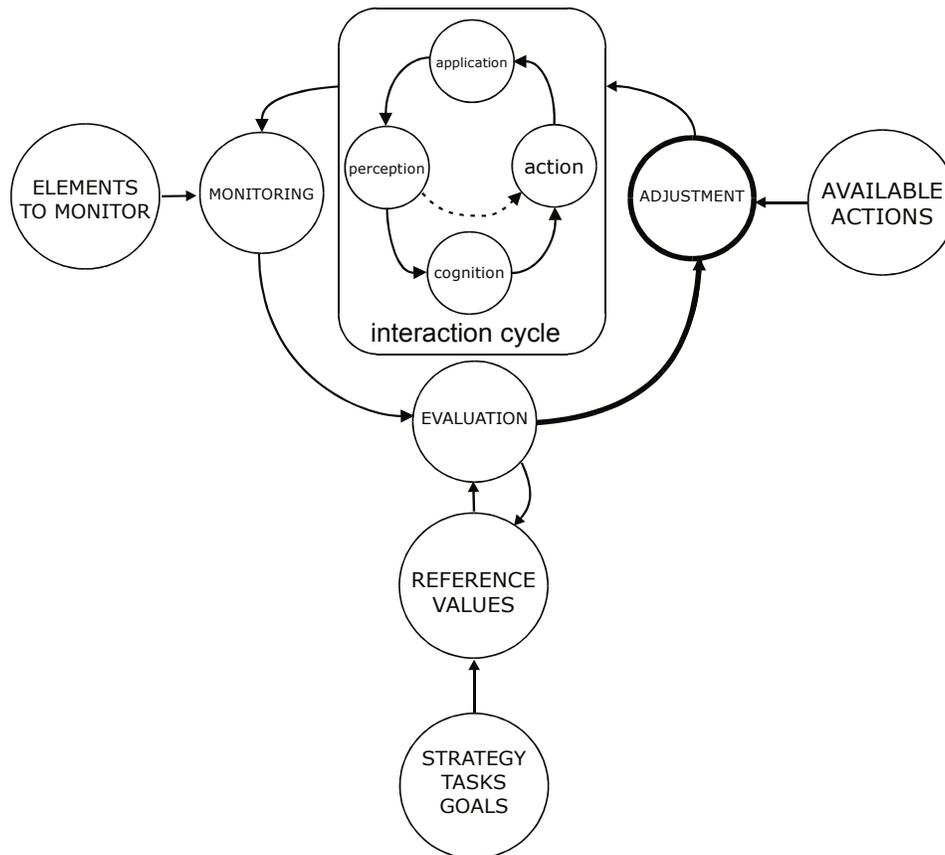


Figure 5.1: A feedback mechanism for the control of interaction. Monitoring of relevant interaction properties gives information to an evaluation mechanism. The evaluation mechanism compares the outcome of the monitoring with reference values that depend on goals and tasks. The evaluation mechanism generates signals on how to adjust actions (**bold**), the research of which is the main focus of this chapter.

The self-regulatory mechanism uses the information that is generated by the evaluation mechanism to adjust actions for the better. In chapter 4, I showed that this mechanism can improve the process of finding a good strategy for a new task. If such a new task is often repeated, its operation converges onto a stable situation (Rasmussen, 1983). In this chapter I investigate how the synthesised mechanism for

interaction control manages such converged action patterns. To do this, adjustments of a typical example of skilled behaviour, mouse-cursor-movements, are studied in different tasks and in a field study.

The action program should be adjusted when the constraints of the interaction change. This can be illustrated using the example of cycling to work in the morning (as many Dutch people do), compared to cycling as a sport during the weekend. I have decided that I should not get overly tired on the way to work and that I should not arrive at the university dripping in sweat. Based on these constraints on the cycling task, I choose not to use too much movement force. Sometimes, if it is warmer than I expected, I start to sweat anyway. In such a case the perception of sweat becomes an important negative element in the monitoring and evaluation of the interaction. Without further thought an action adjustment is made to reduce my cycling velocity and with it the amount of sweat. On the other hand, when I go cycling in the woods near Eindhoven during the weekend, I like to exercise by going as fast as possible. Sweating is no problem because I will anyway take a shower when I return home. I might even consider sweating to be positive since it confirms that I am making an effort.

5.1.1 Motor action regulation

In this chapter the regulation of skilled actions is investigated. Skilled actions can be considered as the outcome of a motor program in which the quantity of two parameters: movement force and co-contraction forces, can be adjusted (Van Galen and Schomaker, 1992). Movement force controls movement velocity, while co-contraction forces are applied to control the noise of the movement. Co-contraction is a combination of forces that are applied to opposing muscle groups to control joint stability. An increase in co-contraction leads to greater joint stability and reduces movement noise. Larger movement forces result in more motor noise. When the movement force is increased, velocity increases but accuracy decreases.

In the intentional description of the self-regulatory system it is the task of the self-regulatory system to determine the parameters of the motor program. Motor actions in user-system interaction can be controlled by applying different combinations of force and co-contraction. The self-regulatory system can set the movement force and co-contraction in a number of ways within a possible movement space (figure 5.2). However, once the possible actions are limited by the

requirements of the task, this defines specific solution spaces within the movement space. When the task requirements have determined the space of movements that fit the task, the self-regulatory system should change the combination of movement force and co-contraction in such a way that the movement falls within the applicable solution space. Even with such limitations there are an unlimited number of ways in which an action can be effectively executed (Wolpert, 1997). When using these motor skills, people are found to exhibit behaviour that is highly predictable. Fitts' law, an often-replicated finding, refers to the high correlation between that movement time to a target and the difficulty of the task, where the difficulty of a task is based on the distance to the target and the size of the target (Fitts, 1954). Such a convergence of behaviour indicates that the motor actions have reached a sufficiently optimal state and have become stable and stereotypical (Bernstein, 1967).

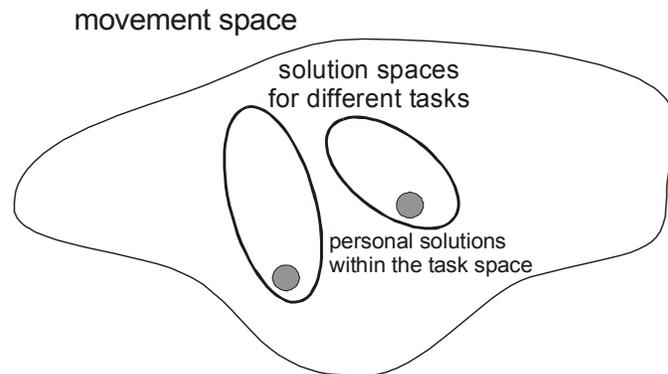


Figure 5.2: *The space for all possible movements. Within this space specific areas give the boundaries of effective movements to complete a task. Within those areas action patterns converge based on personal experience of efficiency. When a task changes, the action pattern is expected to change to the solution space for the new task.*

To regulate motor behaviour, the self-regulatory system should find the best possible combination of the results of the actions and the invested effort (Todorov and Jordan, 2002). In chapters 2 and 3, I argued that the evaluation mechanism of the self-regulatory system can do this by monitoring properties of the interaction. These properties are then compared with reference values, resulting in assessments of the different properties of interaction as a value on a single scale, referred to as

hedonic tones (Johnston, 1999). The outcome of the evaluation of the different properties of interaction depends on the reference values. In chapter 4, I have studied how such reference values can emerge from the personal experience of a single person. In the emergence of these reference values, minimal differences in the user's physical makeup, but also the personal history and chance occurrences in the person's past or even character traits may have influenced the resulting reference values. This would mean, that within certain limits, what is good for one person is not necessarily good for everyone, and that the self-regulatory system acknowledges this. This means that interpersonal differences can be expected, that show a consistent distinction between users over different situations since the reference frame differs. The aggregation of the different hedonic tones gives a value for pleasure (Cabanac, 1992). The self-regulatory mechanism aims to increase pleasure as a way of improving interaction.

The adjustments made to maximise pleasure are probably heuristic rather than rational. The reason is that, even with only a few different experiences, the search for an optimal interaction has to be executed in a complex, multidimensional space, in which it becomes infeasible to find a rational solution rapidly. It is therefore likely that such tasks are controlled heuristically (Reason, 1990; Sloman, 1996). Affect is one such a heuristic (Epstein, 1994). A positive outcome of the evaluation will result in the highest level of pleasure. The user might become happy, which signals that the ongoing interaction is going well and that no further adjustments are needed. A negative outcome of the evaluation, leading to anger, is a signal that the goal has not been achieved and that more effort, and especially more force, should be invested in the interaction (Oatley and Jenkins, 1996). Another negative outcome, sadness, signals there is a more fundamental problem with the interaction. Sadness indicates that action adjustments have to be undertaken that might involve changing the operational goals (Oatley and Jenkins, 1996). Using these heuristic signals for interaction regulation, the self-regulatory system can regulate interaction in fixed situations. When the movement task changes, the reference values for an adequate interaction also change. These new reference values indicate a different, fixed situation. Actions are adjusted to regulate an interaction for the new situation by changing the parameters of motor programs. The new parameters result in ballistic movements aimed at achieving the target in a single non-adjusted action.

5.2 Experiments

In the introduction of this chapter it was assumed that the control system for interaction relies on experience and heuristic adjustments. The first task of the proposed regulatory system is to adjust automated actions for the task in hand. Once the actions have been adjusted effectively, they should be adjusted to the most efficient combination of action and invested effort. To investigate this idea, two experiments were carried out to investigate the control of a highly trained motor-control task, mouse-cursor-movements. In both experiments observable actions of participants are measured (movement time and path). To establish how the interaction is adjusted, the participants in the first experiment were given three different target-acquisition-tasks. In the first task participants were asked to determine freely how to move the cursor to the target (free condition) in the other conditions they were required to achieve the target within either a certain maximum time (time condition) or with a certain precision (precision condition). Effective action adjustments should result in the successful completion of the tasks. However, even within the range of effective actions still an infinite number of actions are possible. In general, the regulation towards more efficient interaction will further limit the possible movements. The physical action sequence is the final stage in a cascade of goals and strategies (Norman, 1984). The regulation of such a system requires each separate stage to be as efficient as possible to achieve the overall goal as efficiently as possible (Dennett, 1994). This means that mouse-cursor-movements should be regulated even when the task is not explicitly related to mouse-cursor-movements. In the second experiment, mouse-cursor-movements were recorded although the participants were not given an explicit target-acquisition-task. A second issue of investigation, addressed in this experiment, is whether affect is indeed an adjustment signal in the regulation of a highly practised task.

5.2.1 Methods and materials for experiment 1

Participants and design

Ten participants, 4 female and 6 male, volunteered for the experiment. The participants were aged between 27 and 31 years, and worked at the TU/e. All participants were right handed when writing and when using a mouse. The

experiment was a within participant design with three task conditions (free, time, and precision). In each condition, each participant executed a target-acquisition-task of moving the cursor to an on-screen target for 88 times after practising with 16 trials. The 88 trials were made up of 11 targets in each of the 8 directions.

Apparatus

An interface was created using Visual Basic 5.0. A round start button, 36 pixels in diameter, was shown at the centre of a 17" screen that had a resolution of 1024 by 768 pixels. When the start button was clicked it disappeared and a target, a black circle 30 pixels in diameter, appeared. There were 8 targets. The centre of each target was at a distance of 250 pixels (about 8 cm on the screen) from the centre of the start button. The targets were located around the start button at 45° intervals. In figure 5.3 an example of a stimulus from one of the conditions is shown.

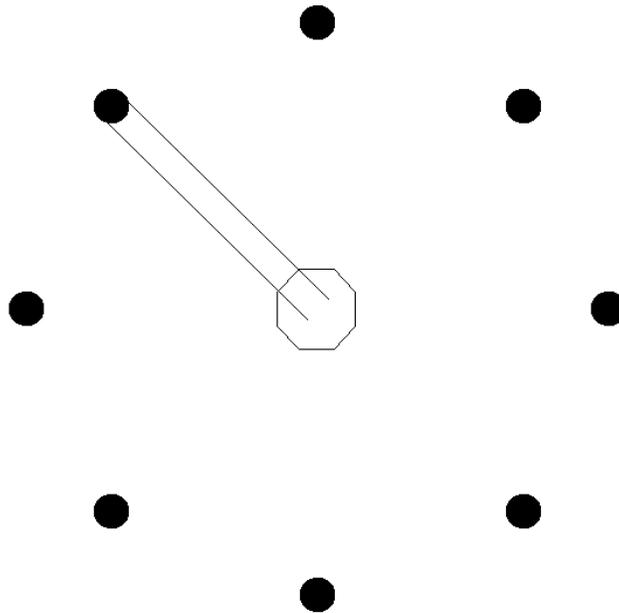


Figure 5.3: Example stimulus of the precision condition showing all eight targets. In this condition participants had to move to the target while keeping the cursor between the lines that determined the 'safe' area.

The experiment was conducted using a standard Microsoft mechanical mouse. The gain was set to the medium level of about 0.23, which means that every centimetre movement of the mouse resulted in a cursor-movement of about 4 centimetres. The cursor was a hairline cross, instead of the more common arrow. The centre of the cross was the recorded cursor position. The X- and Y- coordinates of the cursor position, in pixels, and the system time, in milliseconds, were sampled at a frequency of 50 Hz (20 ms).

Procedure

The experimenter welcomed the participants into the laboratory and gave a verbal introduction. After the introduction they sat at a distance of approximately 50 cm from the computer screen. The participants had to move the cursor to the start button and click on it. After the start button was clicked, it disappeared and one of the targets, which was randomly selected, appeared on the screen. The participant had to move the cursor to the target. A hit was registered when the cursor-position was within the target range for a continuous period of 100 ms. After a hit was registered the target disappeared and the start button reappeared. All participants had to complete 88 trials of successful hits for each of the task conditions, 11 in each of the 8 target directions. The task ended after all 88 hits had been recorded.

There were three conditions, a condition where the participants were asked to move to the targets as they liked (free condition), a condition where the participants had to reach the target within one second (time condition) and a condition where the mouse-cursor had to be kept between two lines 25 pixels apart when moving to the target (precision condition). The time limit of one second and the spacing of 25 pixels between the lines were about one standard deviation from the minimal time and the minimal path curvature recorded in a similar experimental task by Keuning-van Oirschot and Houtsmá (2001). If a participant failed to achieve the task, because the target was not reached within a second or because the cursor moved outside the lines, the screen was coloured red and a miss was recorded. Missed targets were repeated until the full set of 88 successful hits was completed.

All participants were given the free condition first. This was done to record actions that were as close as possible to normal mouse operations. Half of the

participants were given the time condition next; while the other half were given the precision tasks as the second condition. Prior to each experimental condition, participants were given 16 targets (2 in each direction) to practise on. After the practice and halfway through each condition (after approximately 45 targets) the experimenter suggested that the participant took a few moments of rest. The completed experiment lasted about 40 minutes, after which the participants were debriefed and thanked for their cooperation.

Recorded variables and experimental hypothesis

The overall movement time was recorded from the actual start of a movement, which occurred just after clicking the start button, to the recorded hit. The actual start of the movement was registered when the first instance of a movement of more than two pixels per sample occurred. A hit was recorded when the cursor was in the target area for at least 100 ms. At every sample moment, the X- and Y-coordinates were recorded. The total path length was approximated by the summation of the (straight) distances between consecutive sample points. The minimum possible path length between the start and the end of the movement was calculated as the (straight) distance between the location of the start and end of the movement. The difference between the minimum possible path length and the total path length was called the additional path length, and was taken as a first-order measure for lack of accuracy (Zhai and Milgram, 1998). The average movement velocity (\bar{v}) was calculated by dividing total path length by movement time. The maximum velocity was the largest distance moved in a single sample instance (20 ms). If the initial movement did not take the cursor into the target area, corrections had to be made. Each correction was interpreted as an additional sub-movement. The number of sub-movements was approximated by counting the number of times the sample velocity dropped below 10% of the maximum velocity (Keuning, 2003). If the task was not achieved within the required time or precision, a miss was recorded, and the other variables of that movement were not taken into account for further analysis.

The first task of the control system is to adjust actions to make interaction effective or, in terms of a solution space (figure 5.2), to make sure the executed actions fall within the solution space of a task. If this is the case the system will change the behaviour of participants to meet task demands.

Hypothesis 1: The self-regulatory system adjusts the outcome of the participant's actions to meet the task requirements. In other words, participants complete the task faster if there is a time limit, participants move more precisely if there is a precision demand.

If the self-regulatory system undertakes these action adjustments, the task will be effectively performed. Once this has been established, the interaction should be improved by optimising the efficiency of the ongoing interaction. This means that the self-regulatory system should ensure the convergence of the interaction to the most efficiently evaluated combination of invested resources (force, co-contraction) and motor action results (task time, path length). This would result in the participants exhibiting converged behaviour in each condition. In the free condition the participants will try to find the best subjective balance between the invested effort and actions. This will probably mean that less invested effort leads to a relatively large additional path length and movement time. In the other conditions the task requirements enforce a deviation from this balance of interaction properties.

Hypothesis 2: In the free condition the participants will find the most optimal interaction. This is achieved by maximising the subjective experiences of the outcomes of the interaction (additional path length, movement time) compared to the subjective experiences of invested effort (co-contraction and movement force).

As argued in 5.1.1, the translation of the physical elements of interaction into subjective experience is based on personal aspects of the participant. If a task changes, the self-regulatory system should adapt the movement parameters. So, the self-regulatory system determines how to adjust the movement parameters in relation to personal reference values for interaction. This means that a personal preference for investing large forces but high imprecision will remain for different situations, as long as this is possible within the task requirements. In other words, the behaviour in skilled actions (Rasmussen, 1983) converges to stereotypical motor actions (Wolpert, 1997), as long as those stereotypical motor programs allow the achievement of the task.

Hypothesis 3: Personal stereotypes will remain intact throughout different tasks, if the task-constraints allow this.

5.2.2 Results of experiment 1

The time lapse between clicking the start button and the start of the movement was about 200 ms. In 36 of the 2640 trials the program did not record the moment that a hit occurred reliably, these trials were removed from the analysis. The trial-number had no significant influence on either the movement time or the additional path length. This shows that the participants no longer needed much training and that interaction had indeed converged to a stable action pattern. As the data had no normal distribution, non-parametric rank number tests were used.

To confirm the first experimental hypothesis, that participants adjust their action to meet task demands, the measures that were relevant for the given tasks (movement time and additional path length) are compared for the three conditions. The type of task in the conditions (free, time, precision) significantly influenced the additional path length, Kruskal-Wallis $H \chi^2(2)=278, p<0.01$, and movement time, K-W $H \chi^2(2)=1646, p<0.01$. The additional path length was longest in the free condition (median (*Mdn*)=15 pixels) and shortest in the precision condition (*Mdn*=7 pixels), which confirms that the participants make more precise movements if this is required. The movement time was shortest in the time condition (*Mdn*=0.5 s) and longest in the precision condition (*Mdn*=1.8 s), which confirms that the participants move fast when required to do so, and that they sacrifice speed for precision. These findings confirm that participants, in principle, made the necessary action adjustments to execute the task effectively.

The chosen solutions for the three conditions can be plotted as areas in the possible space of additional path length and movement time (figure 5.4). In this figure it is shown that the space determined by the time condition consists of a subset of the free condition. The subjects were apparently easily able to adapt in order to comply with the time condition. However, the precision condition falls beyond the interactions chosen in the free condition. This indicates that movements had to be made in the precision condition that are not normally part of the interaction, which can be interpreted as an indicator that the precision condition is the most difficult to achieve, and might even require the use of a different motor control strategy than is used when the mouse is handled freely or under (moderate) time pressure.

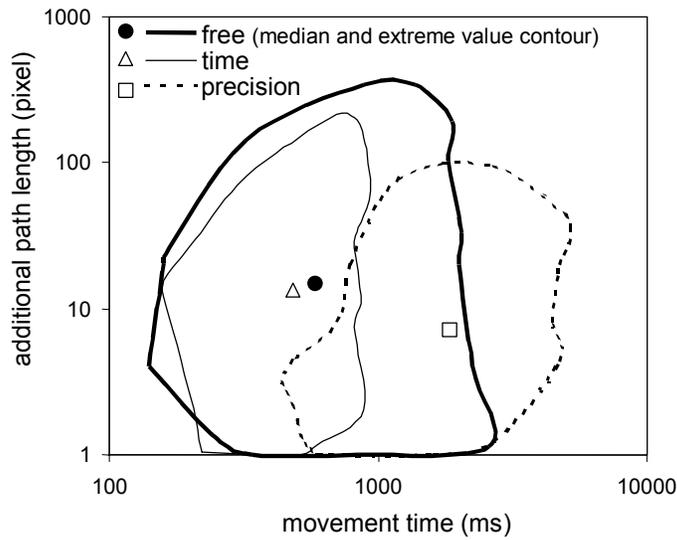


Figure 5.4: Combination of movement time and additional path length used in the three conditions. The movements of the time condition are a subset of those in the free condition, whereas the movements of the precision condition fall outside the range of movements of the free condition.

An interpretation of one of the indicators for the adequacy of the interaction, effectiveness, is related to the number of times the target was missed (see chapter 2). In the time condition, 59 misses (7%) were recorded before all the targets were completed. In the precision condition, 163 misses (18%) were recorded, which indicates that the precise movement was the most difficult task.

A single participant (P8) accounted for 48 (29%) of all misses in the precision condition. On the other hand, in this condition, P8 was both more precise and faster than the other participants (additional path length Mann-Whitney ($N=867$) $z=3.2$, $p<0.01$; time M-W($N=867$) $z=9.9$, $p<0.01$). This indicates that P8 tried to make single ballistic movements and simply refused to add sub-movements, even when this resulted in less effective movements. This finding offers the first evidence for hypothesis 3, that participants adjust the properties of the interaction process differently.

The number of sub-movements can also be interpreted as an indicator for the effectiveness of the movement. If only one sub-movement is made, the target is

acquired by executing the motor program with the initial parameters. If the target is missed, additional movements have to be executed. The most sub-movements were made in the precision condition, Kruskal-Wallis $H \chi^2(2)=193$; $p<0.01$. P8 scored significantly fewer sub-movements in the precision condition than the other participants, M-W ($N=867$) $z=4.9$, $p<0.01$, which confirms the idea that P8 preferred single ballistic movements over many sub-movements, even when this meant he had to accept more failures.

To find further evidence for hypothesis 3, that personal stereotypes of interaction can be observed in different conditions, differences between participants were studied. In general, differences between the 10 participants in the additional path length, K-W $H \chi^2(9)=196$, $p<0.01$, and the movement time, K-W $H \chi^2(9)=125$, $p<0.01$, were found. To investigate whether these differences are related to typical properties of the participants, the participants were classified based on the additional path length.

The first class consisted of the participants whose precision significantly improved (decrease of additional path length) when they had to complete the target acquisition in tasks within a limited amount of time (participants 4, 5, 7, and 8). This group of participants exhibited the least precise movement in the free condition and is referred to as “quick and dirty” (figure 5.5a). The second group consisted of the participants whose precision did not increase in the time condition, but did for the precision condition (participant 1, 6, and 10). This group is referred to as intermediate (figure 5.5b). The final group consisted of the participants whose additional path length did not significantly decrease for the precision condition (participants 2, 3, and 9): the precise group (figure 5.5c).

The relationship between the movement time and additional path length in figure 5.5 can be interpreted as the result of a cost-benefit trade-off and a speed-accuracy trade-off. The more effort (cost) invested, the better the interaction should be. A first order measure for the amount of effort, or the quality of the outcome of the interaction is the product of the movement time and the additional path length. The larger this product the less optimal the combination of time and additional path length will be. The ratio between the time and the additional path length is a second descriptor of participant behaviour (table 5.1).

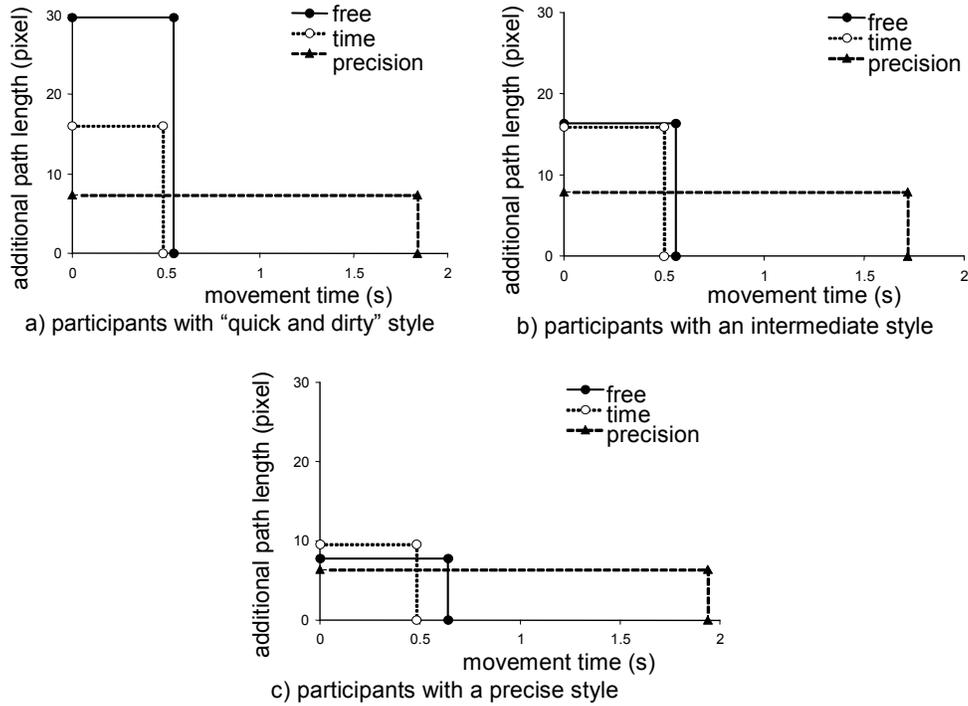


Figure 5.5: Medians of additional path length and movement time for the 3 groups of participants: a) "quick and dirty" b) intermediate and c) precise.

Table 5.1: Medians for effort invested in the interaction.

	Free condition		Time condition		Precision condition	
	<i>n</i>	Median	<i>n</i>	Median	<i>n</i>	Median
	effort style		effort style		effort style	
"quick and dirty"	348	16 18	348	7.3 30	348	13 250
Intermediate	260	9.5 34	260	7.1 27	259	14 248
Precise	261	8.6 80	260	3.8 46	260	13 298

effort = $T \cdot A_{add}$, where a larger value indicates less effort.

style = T / A_{add} , where a larger value indicates a higher regard for accuracy.

Note that the values cannot be straightforwardly compared over the conditions.

The smaller the ratio time divided by additional path length, the more a participant favours speed over accuracy. Such a preference can be interpreted as a personal style. The different types of participants are compared using these measures for effort ($T \cdot A_{\text{add}}$) and the level of movement style (T/A_{add}) (table 5.2). However, the linear metrics that are used in this first approximation might not be the most suitable for comparing the different tasks. This is indicated by the effort value for the precision movements, for example, which indicates that the precise movements require the least effort while other measures indicated that this condition required the least effort.

Table 5.2: Comparisons on z-scores, derived from Mann-Whitney U, between types of participants

	Free condition		Time condition		Precision condition	
	effort	style	effort	style	effort	style
“quick and dirty” – intermediate	3.6**	5.7**	0.4	0.0	0.4	0.4
“quick and dirty” – precise	9.6**	14.2**	5.3**	5.7**	1.1	4.5**
Intermediate – precise	6.1**	9.9**	5.1**	5.3**	0.8	4.4**

** significant at the 0.01 level

The “quick and dirty” participants had the lowest value for the time/additional path length ratio, indicating that they exhibited the most explicit quick-and-dirty style. On the other hand the product of time and additional path-length indicated that the “quick and dirty” participants also invested the least resources in the interaction. As soon as the requirements on movements increased, the “quick and dirty” participants had to increase the resources invested in the interaction, and had to change their style towards a more precise type of movement. These changes made them indistinguishable from the intermediate participants. In all three conditions the precise participants exhibited a significantly more precise style than both the “quick and dirty” and the intermediate participants. The high precision of these participants was achieved at the cost of more invested resources, at least in the free and the time condition. In the precision condition all the participants had the same product of time and additional path length. Together with the large

number of errors, this indicates that this condition required a specific amount of effort and did not allow many differences in regard to the available effective solutions.

The relationship between invested effort and the movement parameters is studied to learn more about the settings of the motor program. To do this movement force and co-contraction have to be estimated first.

To approximate movement force (\mathbf{F}) Newton's law ($\mathbf{F} = m \cdot a$) was used. A constant (c_m) was defined as substitute for mass to account for unknown limb and device parameters, such as inertia, mass, the ratio between the mouse-movement and the cursor movement, and static friction forces. The acceleration of movement (\mathbf{a}) is $\mathbf{a} = \mathbf{F}/c_m$. The assumption of practically constant acceleration and deceleration for ballistic movements (Harris, 1998) leads to a first order approximation of the velocity-time diagram (figure 5.6), allowing the estimation of movement force based on the experimentally recorded variables (equation 5.1).

$$\mathbf{F} = c_m \cdot \frac{4\bar{v}}{T} \quad (5.1)$$

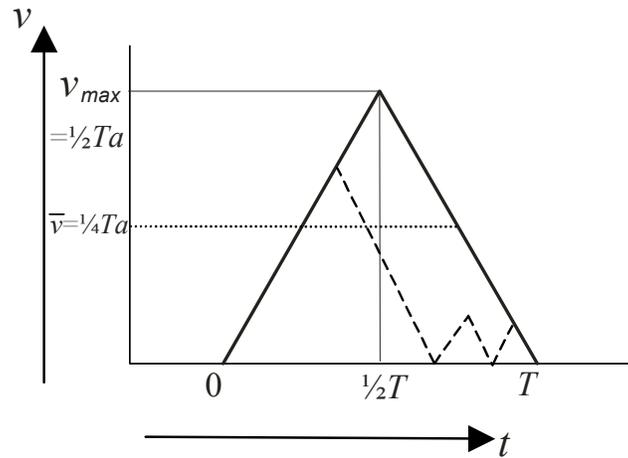


Figure 5.6: Velocity-time diagram with constant and equal acceleration and deceleration (equal force and co-contraction). The maximum velocity will be lower for movements with multiple sub-movements (dashed line).

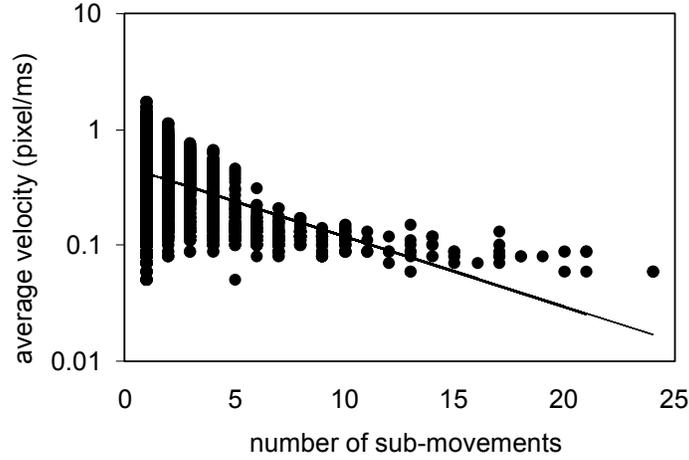


Figure 5.7: Influence of the number of sub-movements on the average movement velocity (plotted on a logarithmic scale).

When additional sub-movements (s) are made, the average velocity (\bar{v}) is lower than would be expected for movements executed with the same movement force (dashed line in figure 5.6). In equation 5.1, the average velocity determines the estimation of force. Since the average velocity for multiple sub-movements is lower than would be the case for a single sub-movement, the estimate for movement force will also be too low. To compensate for this effect introduced by sub-movements, an exponential relation between the observed average velocity (\bar{v}_{obs}) and sub-movements (s) was derived from the data ($\bar{v}_{\text{obs}} = 0.42 \cdot e^{-0.14 \cdot (s-1)}$; $r=0.45$; $p<0.01$; figure 5.7).

The equation was fitted using the $(s-1)$ instead of (s) to make the exponential term equal to 1 in the case of a single sub-movement. Equation 5.2 specifies a virtual average velocity (\bar{v}_{vir}) for the approximations of force.

$$\bar{v}_{\text{vir}} = \bar{v}_{\text{obs}} \cdot e^{0.14 \cdot (s-1)} \quad (5.2)$$

In cases with multiple sub-movements, the movement is assumed to be made up of a sequence of ballistic movements (dashed lines in figure 5.6). To compensate for the resulting lower average velocity, the empirically measured average velocity is replaced by equation 5.2. With a low mass constant ($c_m=1$) and a correction for the number of sub-movements, the movement force is approximated by equation 5.3.

$$\mathbf{F} = \frac{\bar{\mathbf{v}} \cdot e^{0.14(s-1)}}{T} \quad (5.3)$$

With a constant distance to the target (A), there should be a relationship between the movement time and the average velocity $A/T = \bar{\mathbf{v}}$. This would mean that an increase of force is related to the inversed squared decrease of movement time: $\mathbf{F} = c_m \cdot 4 \cdot A/T^2$. However, Fitts' law (equation 5.4; 1954) gives empirical evidence that the movement time (T) of a goal directed movement is strictly determined by the relationship of the distance to the target (A) and the width of the target (W) and not by movement force.

$$T = c_1 + c_2 \log_2 \left(\frac{2A}{W} \right) \quad (5.4)$$

The positive effect of higher force in the form of an increased velocity is apparently nullified by another movement property. Neuromotor noise was suggested as this property (Van Galen and Schomaker, 1992), where an increased force leads to more neuromotor noise. Increased neuromotor noise, in turn, results in a higher probability of missing the target, which takes time to correct. To optimise interaction, neuromotor noise should therefore be controlled. There are two ways to compensate for this noise. First, by increasing pressure on the input device the static friction increases and as the stiffness of the interaction system increases, the effects of noise decrease. This means that the constant c_m in equation 5.1 increases. If neuromotor noise is controlled this way additional force no longer simply corresponds to an increase in velocity. A second way to control neuromotor noise is to apply co-contraction forces, thus increasing the stiffness of the limb of the user (Van Galen and Schomaker, 1992).

By interpreting the additional path length as a measure for the endpoint error, the neuromotor noise and the path precision can be related. When the endpoint error of the first sub-movement is interpreted as the target size of the initial ballistic motor program, the relationship between the endpoint error and the neuromotor noise can be related to Fitts' law (equation 5.4; Van Galen and De Jong, 1995). The simplest implementation of these relationships is to substitute the width of the target (W) by the additional path length (A_{add}). The larger the co-contraction (F_c), the smaller the end-point-error (A_{add}) becomes (Van Galen and Schomaker, 1992). With an arbitrary, initial neuromotor-noise level ($\alpha_0=9$), a first order approximation of co-contraction can be derived from Fitts' law (equation 5.5).

$$F_c = 9 + \log_2 \left(\frac{2A}{A_{\text{add}}} \right) \quad (5.5)$$

Differences in the approximated force and co-contraction were found between the conditions and the participants (table 5.3). Co-contraction was lowest in the free condition and highest in the precision condition. This shows that when high precision is required, co-contraction is increased to suppress motor noise. Force was highest in the time condition. Shorter movement time could only be achieved at the cost of more force and more co-contraction compared to free movements. Force was lowest in the precision condition. An increase in force results in higher velocity but also in more noise. Since noise suppression was an essential element of the precision condition, these lower forces were to be expected.

With these approximations of force and co-contraction, we now have four values for interaction process parameters. With the two effort parameters (force and co-contraction) every possible combination of two parameters for the outcome of

Table 5.3: Effect of conditions (free, time or precision) and participants (Kruskal-Wallis H) on co-contraction forces and movement force.

	df	Co-contraction		Force	
		χ^2	p	χ^2	p
Condition	2	265	<0.01	1671	<0.01
Participant	9	259	<0.01	60	<0.01

actions (movement time and additional path length) can be specified. The self-regulatory system should be capable of deciding the subjective adequacy of the different task conditions using these four properties in combination with the personal reference values, indicating the importance and the base value of each of these properties. Hypothesis 2 states that the self-regulatory system finds the best balance of the ongoing interaction by combining the outcome of the evaluation of the separate interaction properties: additional path length, movement time, co-contraction, and force. To compare the different interaction properties they need to be separately evaluated as a measure of adequacy (hedonic tone). One way of doing this is by comparing the scores of each property with the overall median. To make further comparisons possible, the variance of the different interaction properties should be the same, which can be achieved by standardising the variances of the different properties. This results in measures for the adequacy of interaction properties, relative to those of the other tasks (figure 5.8).

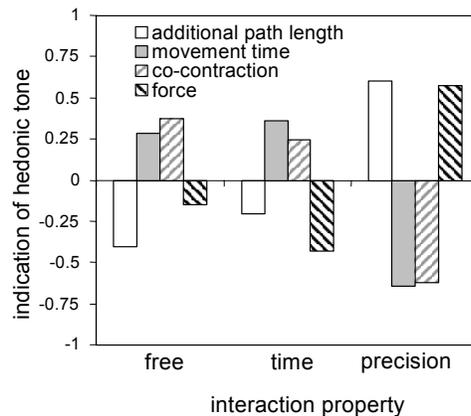


Figure 5.8: Interpretation of the influence of context as hedonic tones for path-precision, movement time, co-contraction forces, and force.

In the free condition the participants scored lower on movement time and path precision, but invested less resources (force and co-contraction). In both of the other conditions, the observable actions were improved, but at the cost of effort. In the precision condition in particular, the path precision was highly improved at the cost of co-contraction. As a first order estimate for the overall adequacy (pleasure) of the different task conditions, the relative adequacy measures are simply added

up. The free condition has the highest, most optimal value (0.11). The time condition comes second (-0.02), with the precision condition being least optimal (-0.09). This finding is in line with the earlier finding about the number of failures, where the precision condition scored the highest number, indicating this was the least adequate interaction.

5.2.3 Discussion of experiment 1

The findings of this experiment confirm the first experimental hypothesis, that users can adapt mouse-cursor-movements to meet task demands. The relevant interaction property, time or precision, is adjusted. To do this, they have to accept a lower efficiency for the remaining interaction properties. This is most evident for the increase in movement time when high precision is demanded. When force and co-contraction were estimated, it was shown again that when task demands increase, i.e. more effort has to be invested to complete the task within the required limits. When a shorter time is demanded, the participants both increase their precision and decrease their movement time, at the cost of force and co-contraction. The self-regulatory system appears to use two ways of improving movement time: an increase of velocity (movement force), and a decrease of path length for a constant velocity (co-contraction). To confirm this idea, force and co-contraction were approximated based on the data. Tasks with higher demands on the user are more difficult, as was confirmed by the finding that in these conditions the number of misses and sub-movements increased, which are indications of lower effectiveness. The precision condition proved to be especially difficult, as can be observed in the number of misses and in the observed movement parameters beyond those recorded for the free movement. The applied forces and co-contraction were increased to overcome the additional difficulties, which led to effective, although less efficient, movements. The number of failures revealed that participants used different cost/benefit trade-offs. One participant did not deviate from the ballistic movement strategy, which is the optimal strategy in free or time limited target-acquisition-tasks (Harris and Wolpert, 1998), and subsequently registered the most errors, but also the fastest movement time in the precision condition. This participant apparently interpreted the cost of a larger number of misses as lower than the cost of changing his strategy, which is an indication of interpersonal differences as predicted in hypothesis 3. This strategy of accepting

additional misses can be interpreted as a way of optimising the overall efficiency of the complete sequence of tasks, rather than optimising each task separately. When considering all of the tasks, P8 had to make 136 movements in the precise condition, which he did in 144 seconds. On average, the other participants made only 100 movements, which required 196 seconds. The time for restarting a trial is not taken into account, but an estimate of somewhere in the region of 50 seconds for these 36 additional moves will probably be close to the mark. This indicates that P8 executed the tasks in the precision condition in about the same time as the other participants. This was not the only confirmation of the idea of personal stereotypes that were only changed when needed. Another way of looking at personal stereotypes was through the product of the outcomes of the motor program: movement time and additional path length. This product was taken as a first order measure for the quality of this outcome, or in other words, the amount of effort invested to achieve this outcome. The participants who invested the most in interaction can be interpreted as the most motivated to perform the task. When explicit tasks were set, these participants only have to apply limited additional effort. On the other hand, the “quick and dirty” participants had to increase the amount of invested effort considerably. These participants keep exhibiting the lowest possible amount of invested effort, and a high level of quick-and-dirty movements when adapting the interaction to meet the task demands. This shows that personal stereotypes remained largely unchanged throughout the various conditions, which once again shows that action adjustments in skilled tasks are variations of well-established stereotypes (Rasmussen, 1983) rather than rational determination of action.

The approximated values for force and co-contraction, together with the measured movement time and additional path length, should be enough to determine the optimal balance for an action. To test hypothesis 2, it was explored whether the self-regulatory system could determine the adequacy through the assumed system of adding hedonic tones for the different interaction properties. Hedonic tones are assumed to be the results of the evaluation, i.e. the comparison with the observed interaction property and a reference value of that property. The simplest way to determine such hedonic tones is to use standardised values of the four relevant interaction properties. An increase of the two properties for effort (force and co-contraction) was applied in both conditions with increased task

demands. This confirms that the participants increased the use of resources to be able to meet the task demands. The overall sum of standardised interaction values is highest for the free condition, which indicates that, when allowed to choose freely, participants adjusted their actions towards the best possible combination of interaction properties. A limitation of this conclusion is that the measures for relative adequacy are artificially derived from the data. To determine whether these measures are the actual input values for the optimisation of the actions, empirical data about the evaluation of the mouse-cursor movements should be gathered.

There seems to be a contradiction in the fact that the self-regulatory system adjusts interaction towards an optimum, while different participants exhibit stereotypical behaviour. As I mentioned in the introduction of this chapter, this difference between optimal interaction and personal preference can be partially explained by differences in personal character traits. This means that when the difference in optimality is small these personal attributes, such as physical strength might explain the observed differences. There might also be other character traits such as preference for risk over certainty, that lead to separate similarly optimal styles of interaction patterns at a higher level, that in turn give rise to the observed differences in motor control. Following this argument, optimal interaction must be interpreted as an optimal interaction in the context of a single participant. Optimal interaction for a single participant should not only take into account an objective interaction optimum, but is also partially determined by the participant's physical makeup, personality and personal history.

Another limitation is the use of the approximated values of force and co-contraction, which were also derived theoretically with the assumption that movements were ballistic. In the precision condition, in particular, force was fairly low compared to the other conditions. Precision tasks should probably be considered as a continuous target-acquisition-task (Accot and Zhai, 2001), rather than a ballistic movement. This is confirmed in the velocity-time diagrams for the different conditions (figure 5.9), both the free and the time condition show a rapid increase to the maximum velocity followed by a subsequent decrease of ballistic movements (Harris and Wolpert, 1998). In the precision condition the velocity-time has a distinctly different shape, indicating that the movements are not ballistic.

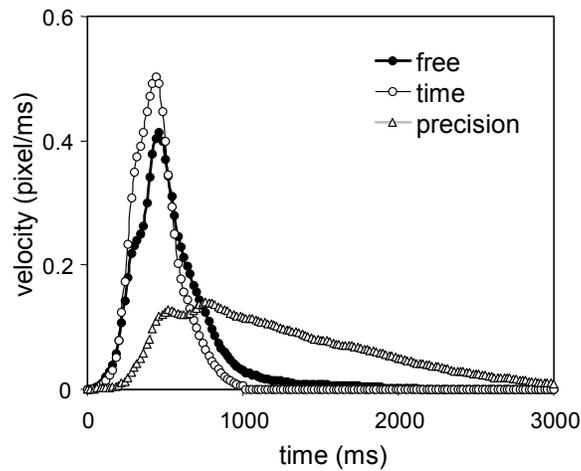


Figure 5.9: Averaged velocity-time diagrams for the three conditions. The sharp acceleration-deceleration of the free and time conditions indicates ballistic movements.

The velocity-time relationship of the precision condition shows a low maximum velocity and a long period of relatively high velocities, compared to this maximum velocity. This means that estimating sub-movements at 10% of the maximum velocity will probably result in too few recorded sub-movements for the precision condition, which means that the estimate for force is probably too low. To make more detailed claims about the regulation of precise movements, better estimates of movement force and co-contraction should therefore be generated. On the other hand, is the estimation of force is too low, the precision condition, which was already evaluated as the worst, would be even less adequate, which supports the other measures for adequacy of this condition. So, although the approximation of force limits quantitative conclusions, it does not influence the qualitative conclusion that the participants adjusted their mouse-cursor-movements to the most optimal interaction pattern when they were able to choose freely how to move, and that they invested the additional resources when effective interaction achievement required them to do so.

5.2.3 Experiment 2

Even when mouse-cursor-movements are not the primary task, some control should be applied to these interaction processes. This is true if we assume that interaction strategies are optimised (Dennett, 1994), which can only be the case if the separate elements of a strategy are controlled. To investigate this, in a second experiment mouse-cursor-movements were recorded while the participants executed a gambling task. By not telling the participants beforehand that their movements were being tracked, I was able to investigate how movements are regulated when the motor task is not explicit. Since mouse-cursor-movements are part of daily human-computer-interaction, these movements should be controlled. This control is probably neither conscious nor rational, but is likely to depend on heuristics such as affect (Petty and Wegener, 1999). To understand the influence of affect on interaction regulation, the emergence of emotions is interpreted as a process (Frijda, 1986), similar to the stages of the feedback mechanism from the studied self-regulatory system. This process consists of a stage where the ongoing process is monitored, followed by an evaluation stage that initiates changes in the action plan. The whole process leads to an emotional experience. A negative emotion is a signal that the situation is unsatisfactory and initiates improvement of the situation by making specific changes in the action plan. The influence of affective control on interaction optimisation was investigated by manipulating the mood of the participants. Mood is an affective phenomenon that is argued to influence the goals of the interaction processes (Sloman, 1999). The mood will therefore generate a signal about all ongoing interactions. This signal will be spread to all processes that regulate the various ongoing tasks. Although the influence of mood on the evaluation has nothing to do with the actual control of mouse-cursor-movements, this would mean that the self-regulatory system evaluates the interaction as being less adequate when a bad mood is induced. The self-regulatory systems of the participants who are put into a bad mood, will then generate signals that the ongoing interaction should be improved by investing more effort. More explicitly, anger indicates that the interaction should be improved by applying more force (Oatley and Johnson-Laird, 1987). The same functional signals of emotions indicate that a sad mood means a lack of satisfaction with the current situation and leads to the current goals being abandoned. Conversely, people in a happy mood get the signal that the interaction is good so

no action adjustments are made, even when the interaction is perhaps not as efficient as it should be.

5.2.4 Methods and materials experiment 2

Participants and design

Thirty students from the TU/e and Fontys hogeschool participated in the experiment. There were 12 female and 18 male participants, ranging in age between 19 and 26. Each participant was played one of three music files, compiled to manipulate mood¹ i.e., angry, happy, and sad (Lewis, Dember, Scheft, and Radenhausen, 1995). The music files were randomly assigned to participants, balanced for gender. The effect of the music on participant mood was determined in an earlier study (see chapter 4). For each of the participants, 100 mouse-cursor paths were recorded. These paths were recorded while participants played a card game. The card game was a pilot study for the experiment reported in chapter 4.

Apparatus

A fully computerised experiment was created with Visual Basic 5.0. The computer first played about 12 minutes of music on a pair of headphones and then showed the interface of a card game. The interface was shown on a 15" computer screen with a resolution of 1024 by 768 pixels. In the game, participants had to select 1 out of 4 decks of cards for a total of 100 times. Each deck was 130 by 130 pixels (about 4x4 cm) in size. After a card was drawn, losses and gains were paid in a bank area 250 by 250 pixels (about 7.5x7.5 cm) in size. This bank area was the starting area for the cursor-movements that resulted in the drawing of a card. The centre of the bank area was located at 425 pixels (about 12.5 cm) from the centres of the cards (figure 5.10). During the movement the X- and Y- coordinates of the cursor position were sampled at a frequency of 50 Hz (20 ms). The experiment was carried out using a standard Microsoft mechanical mouse. The cursor shape was an arrow, which is standard in most applications.

¹ See Appendix A for titles and performing artists of the songs.

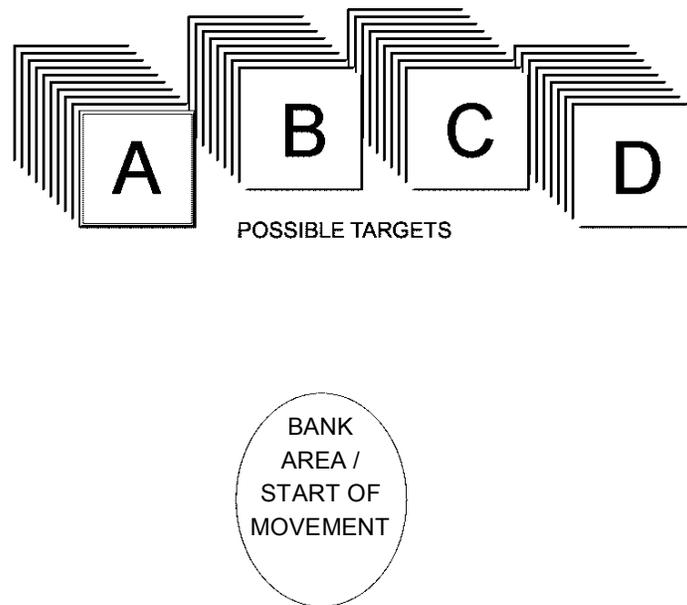


Figure 5.10: The screen in which each participant had to select 100 times from 1 of 4 decks (A, B, C, or D). After each card selection, a task had to be executed in the bank area at the bottom of the screen.

Procedure

The participants were welcomed into a cognitive laboratory and were each assigned to one of eight cabins. They sat in front of a computer screen and were told that the aim of the experiment was twofold, first to ask them some questions about music, and second to test their skills in a card game. In the first stage of the experiment the participants heard one of 3 music files lasting 12 minutes on a pair of headphones. When the music had finished, the participants were asked to fill out the PANAS² scale (Watson, Clark, and Tellegen, 1988). They were then given on-screen instructions for the card game, but were not told that mouse-cursor-

² See appendix B for the used Dutch translation.

movements would be recorded. The card game consisted of drawing 100 cards. A completed experiment lasted about 45 minutes, after which participants were debriefed, thanked, and paid between € 3.00 and € 5.00 depending on the success of their card game.

5.2.5 *Recorded variables and experimental hypotheses*

The cursor-movements towards the decks of cards were recorded from the actual start of a movement until the click on a card. The start of a movement was judged to be either of two occurrences. A click on an interactive control immediately followed by a continuous movement ending in a card click counted as the start of a movement. If the click on a control was followed by a period in which no movement occurred, the end of the last recorded interval of 2 seconds (2000 ms), in which the cursor moved less than a total of 2 pixels, preceding a continuous movement ending in a card click counted as the start of a movement. The X- and Y-coordinates were recorded at every sample instance, i.e. every 20 ms. The total path length was approximated by summing the straight distances between consecutive sample points. The straight distance between the start and the end of movements was calculated as the minimal possible path length. The case was removed if the minimum possible path length was 0, because immediately before clicking a card no movement occurred for 2 seconds (151 out of 3000 cases). The additional path length, a measure for the imprecision of movement, was calculated as the difference between the minimum path length and the actual path length. The average movement velocity was calculated by dividing the path length by the movement time. Movement force and co-contraction were approximated by using equations 5.1 and 5.5.

Clicking on cards is considered to be an action sequence (Norman, 1984) consisting of a movement to a deck and a click on a card. If mouse-cursor interaction is regulated, the combination of these actions should be controlled.

Optimal movement to a deck can be characterised by a high initial acceleration of movement velocity followed by a longer deceleration and readjustment phase (Harris and Wolpert, 1998). At the end of the movement the next action should be clicking the card, which should follow immediately after recognising the end of the movements. GOMS, a cognitive model for highly practised tasks, divides interaction into perception, cognitive and motor processes (Card, Moran, and

Newell, 1983). A quantitative prediction of interaction time can be made by modelling the studied interaction in GOMS. The time lapse between ending the movement and clicking a card should involve one perception process (observing end of movement on card), one cognitive cycle (deciding to make a click), and the start of a motor program. GOMS predicts that these cycles together should take on average 240 ms (105~475 ms).

Hypothesis 4: A regulated card clicking process consists of an optimised movement, followed by an immediate click, together adding up to about 240 ms.

Following the self-regulatory system that is investigated in this thesis, the regulation of the card clicking process is assumed to be heuristically regulated. If this is the case, mood manipulation should influence action adjustment. The induced mood was argued in the introduction of this experiment to generate specific signals to the action adjustment mechanism that permeates through all the levels of interaction control. When investigating the specific signals given by the three induced moods, I expect that the ongoing interaction in an angry mood is evaluated to be unsatisfactory and that a specific action signal is generated to aggress, that is, to apply more movement force. The participants in a happy mood are expected to invest few additional resources in improving the movement; a happy mood signals that the interaction is good. The participants in a sad mood should reprioritise goals, which can lead to slower and more deliberate actions.

Hypothesis 5: Different moods result in different signals about the adjustment of the interaction, inducing participants to invest more force (anger), to invest more effort in a general way (sadness) or to maintain the current situation (happy).

To conclude the analysis of the second experiment the results of the second experiment are compared with those of the first experiment. Since the mouse-cursor-movements were not the explicit task in the second experiment, less stringent control of the interaction is expected in this experiment. This difference could be observed as a lower degree of convergence of the actions.

The differences between the two experiments will complicate the straightforward comparison of the results. The number of trials in the two experiments was comparable and is of no consequence for the comparison: 2604 in experiment 1 and 2849 in experiment 2. Important differences were, that in

experiment 1 the mouse-cursor-movement task was the goal of the interaction process, whereas in experiment 2 mouse-cursor-movements were a means to achieve another goal. Another important difference is the distance between the start of the movement and the target. Where the average distance was defined as being 250 pixels in experiment 1, in the natural interaction tasks of experiment 2, this distance was less clearly defined and was larger (425 pixels). The targets of experiment 2 were also larger (130 pixels) than those of experiment 1 (30 pixels). The second term of Fitts' law ($\log_2(2A/W)$; equation 5.4) is called the index of difficulty (ID) of a task. The index of difficulty for experiment 1 is 4.1 where the index of difficulty of experiment 2 is 2.7. If we assume that the tasks are similar and that the same amount of effort is invested, this difference in difficulty indices means that movement times in experiment 2 should be shorter than the movement times in experiment 1.

Hypothesis 6: If an explicit task is regulated in the same way as a non-explicit task, the task with the lowest difficulty index should be executed in the shortest movement time.

When interpreting the regulation of the motor-program by applying different levels of force and co-contraction, the difficulty index from Fitts' law means that with the same amount of effort invested in the task, more of this effort can be allocated to movement force because less co-contraction is required.

Hypothesis 7: If an explicit task is regulated in the same way as a non-explicit task, the task with the lowest difficulty index should have the highest acceleration.

Some differences, which will not be taken further investigated here, were that the movements in experiment 2 ended by clicking a card, where in experiment 1 a period of time in the target area was counted as a hit. There was also a difference in the interfaces used in the experiments, namely that the task in experiment 1 was executed on a grey, almost empty, 17" screen, whereas the interface of experiment 2 was colourful, contained multiple interaction elements and was presented on a 15" screen.

5.2.6 Results of experiment 2

Examples of recorded cursor paths shown in figure 5.11 give an immediate idea that mouse-cursor-movements are goal directed.

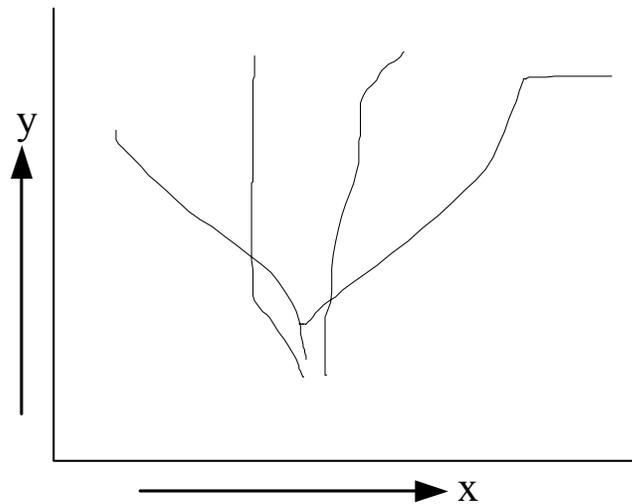


Figure 5.11: Example of cursor paths (x, y coordinates) for a choice of each of the four cards.

When reviewing the velocity-time diagram (figure 5.12a), average velocity of the movements can be seen to increase from the start with a high acceleration. After the maximum velocity is reached and the cursor approaches the chosen card the movement becomes more precise and slower. A linear increase ($v=5.9T$) followed by a linear decrease ($v=c-0.9T$) describes the velocity profile well ($R^2=0.92$), and is similar to the description for converged movement presented by Harris and Wolpert (1998). Beyond 1 second, there is a lingering period of low velocities caused by additional corrections in the movement (figure 5.12).

There is a time lapse between the end of a movement and the clicking on a card, which is most frequently between 170 and 230 ms ($sd=170$ ms, figure 5.12b). There are however some extreme values close to 0 ms, probably because the click was initiated before the movement ended, and beyond 1000 ms, probably because the participant reconsidered his or her card selection before confirming it by a click. The most common click time around 200 ms is even a bit shorter than the expected time needed for an automated process to execute a command (Card, et al., 1983). Therefore, action leading to drawing a card can indeed be interpreted as a goal directed process, with an action sequence consisting of two automated actions, an optimised movement towards the card followed by a click on the card.

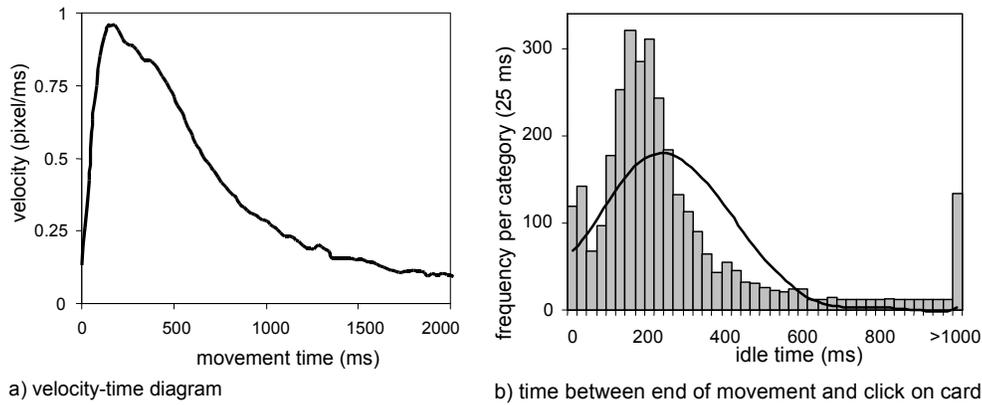


Figure 5.12: Indicators for action sequence. a) velocity-time diagram. b) distribution of the time between the end of a movement and the subsequent click on a card. The line is the normal curve with the same mean and sd as the data.

Before going into the results of the mood manipulation it should be noted that, using the PANAS mood scale, it was not possible to make a clear-cut distinction between the participants who listened to aggressive music and those who listened happy music. The interpretation of differences between these two conditions should therefore be taken with some reservations. See chapter 4 for a more complete discussion of the effect of music on mood.

Table 5.4: Effect of moods on additional path length, movement time, approximated co-contraction and movement force. Also shown are the indicators for effort (T^*A_{add}) and style (T/A_{add}).

Mood	n	Additional path length (pixels)	Time (s)	Co-contraction	Force	effort	style
		Mdn	Mdn	Mdn	Mdn	Mdn	Mdn
Angry	958	76	1.7	1.9	173	14	26
Happy	944	78	2.0	1.8	137	16	26
Sad	957	62	1.8	2.0	174	13	30

Mood induction influenced the additional path length, K-W $H \chi^2(2)=7.1$, $p=0.03$, movement time, K-W $H \chi^2(2)=18.1$, $p<0.01$, and the approximations for co-contraction, K-W $H \chi^2(2)=6.8$, $p=0.03$, and movement force, K-W $H \chi^2(2)=15.8$, $p<0.01$. Significant differences between the types of music were found. When an

angry mood was induced, the movement time decreased and the additional path length increased. When people had listened to happy music, their movement time and additional path length were longer than in other conditions. The participants who listened to sad music achieved a shorter time and the smallest additional path length (tables 5.4 and 5.5).

Table 5.5: Effect of moods on additional path length, movement time, approximated co-contraction and movement force; z-scores based on Mann-Whitney U

Mood	Additional path length z	Time z	Co-contraction z	Force z	effort z	style z
Angry – Happy	0.5	3.2**	1.8	3.3**	1.8	1.0
Angry – Sad	2.0*	1.0	0.7	0.5	1.6	2.1*
Happy – Sad	2.5**	4.0**	2.6**	3.6**	3.4**	1.0

* significant at the 0.05 level

** significant at the 0.01 level

When comparing the amount of invested effort by considering the product of time and additional path length, I found that the participants who listened to sad music invested significantly more effort in the interaction than the participants who listened to happy music (tables 5.4 and 5.5). This difference is marginally significant between the participants who listened to happy and angry music ($p=0.07$). The movement style of participants in the sad condition was significantly more precise than that of the participants who listened to aggressive music. This confirms the hypothesis that negative emotions signal that more effort should be invested in the ongoing interaction, and that aggression leads to the preference of force over precision.

In the discussion of the first experiment, I argued that in regulating interaction, the self-regulatory system can subconsciously adjust actions. These adjustments are made based on signals about the adequacy of the interaction. Typical emotions give signals that the interaction receives too little force (aggression), is good as it is (happiness), or is not good at all (sadness). To investigate how these signals influence the effort invested in mouse cursor movements, the different properties of the interaction are compared. I did this by individually transforming the different interaction properties into a measure of adequacy (hedonic tone), in a

similar way as in the first experiment, i.e. by comparing the scores of each property with the overall median for that property. To make further comparison possible the variance of the different interaction properties should be the same, which can be achieved by standardising the variances of the different properties. This results in measures for the adequacy of interaction properties, relative to those of the other tasks (figure 5.13). Note that the effects are about five times smaller than in experiment 1. Simple summation of these hedonic tones indicates that the aggressive interaction is the least satisfactory (-0.064) while happy (-0.028) and sad (-0.025) mood result in roughly equally adequate interaction. The main difference between the happy and the sad interaction is that in the happy condition, the participants invest the least effort, while in the sad condition the movement results are best.

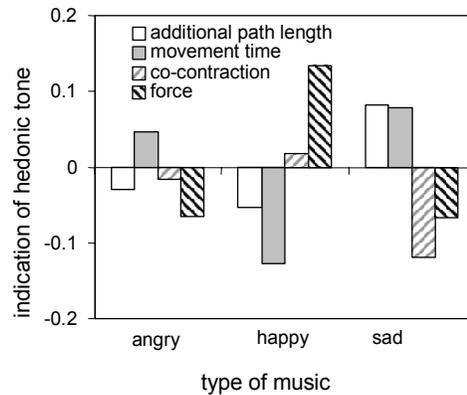


Figure 5.13: Influence of mood on the hedonic tones for different properties of mouse-cursor-movements.

5.2.7 Comparison between explicit and implicit task regulation

In the first experiment, the regulation of three versions of an explicit target acquisition task was studied. In the second experiment the regulation of a non-explicit mouse-cursor operation task was studied. To explore the similarities and differences between these two types of mouse cursor movements, the results are compared.

With the limitations with regard to the similarity of the data as sketched in section 5.2.5, some aspects of the two experiments are compared. First the contours defined by all the mouse-cursor-movements of the experiments are compared (figure 5.14). This comparison can give information about the utilised movement space. The most obvious difference between the experimental conditions is that there was more variance in the task without an explicit mouse cursor task (experiment 2).

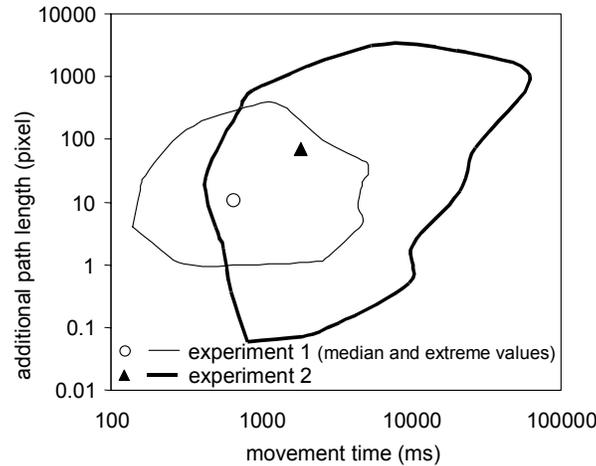


Figure 5.14: Areas of movement made in tasks when mouse-cursor-movement tasks were the explicit goal of the task (experiment 1) compared to tasks when the movements were not the explicitly goal (experiment 2).

A second difference is that, in general, the movement time of experiment 2 was longer than that in experiment 1, although an equally strictly regulated interaction process predicts that the movement time should be shorter in experiment 2 (hypothesis 6). There turned out to be a considerable overlap, which indicates that the mouse-cursor-movements in both types of tasks are executed in a similar space of possible movements. The larger area occupied by movements from experiment 2 is interpreted as an indication that the actions converged to a lesser degree than in the case where an explicit task was given. In fact, the additional path length in experiment 2 had a larger range than the additional path lengths in experiment 1, with a lower minimum value and higher maximum value. The shorter additional

path length was probably caused by the relative ease of the mouse cursor task, due to the large size of the targets.

A second comparison was made by investigating the velocity time diagrams. To do this I studied only the two conditions of experiment 1, free and time, that could be interpreted as the result of ballistic movements. If the velocity time diagrams are compared (figure 5.15), the acceleration in experiment 2 is shown to be higher than in experiment 1. This indicates that a relatively large part of the invested effort in experiment 2 was force, as was predicted in hypothesis 7. A second issue in the velocity time diagram, namely that the surface under the curve of experiment 2 is larger than expected based on the difference in the movement distance alone. This is in line with the observed differences in the additional path length, which were about 4% of the movement distance in experiment 1 and about 17% of the movement distance in experiment 2. This observation gives support for the argument that the actions in the second experiment converged to a lower degree than in the first experiment.

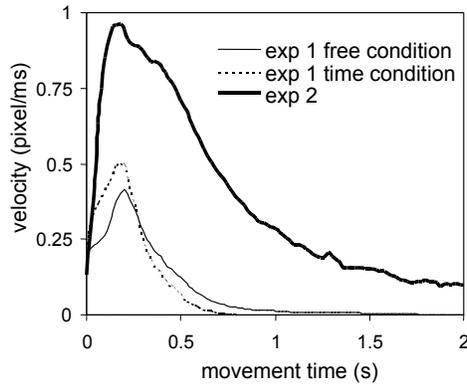


Figure 5.15: Velocity time diagrams for two tasks when mouse-cursor-movement tasks were the explicit goal of the task (experiment 1) compared to tasks when the movements were not the explicitly goal (experiment 2).

5.2.8 Discussion of experiment 2

Even when the participants were not told to concentrate on mouse-cursor-movements, goal directed action sequences were observed. This confirms that actions that are outside the focus of attention are also regulated to some extent as

well. Context effects, such as music-induced mood differences, partly determine the outcome of the regulatory process. These emotional effects on the adjustment of action can be interpreted as cognitive functions of emotions (Oatley and Johnson-Laird, 1987). The participants who listened to happy music exhibited slow and imprecise movements. This is consistent with the hypothesised signal of happiness that ongoing interaction is good enough, and that no additional effort should be invested in optimising the movement. When the participants listened to angry music, the movement time decreased but the imprecision remained high. This implies that a signal was given indicating that the interaction was not good enough and that more force should be generated to improve the interaction, i.e. the signal for anger. This conclusion should be taken with some caution, however, since I could not conclusively distinguish between the effect of the happy and the aggressive music on the participant's mood by using the PANAS mood scale (see chapter 4). The participants who listened to sad music improved both their movement time and movement precision. As seen in experiment 1 the improvement of these properties goes together. This finding can be interpreted as showing that sadness induces an explorative period of low-intensity experiences while changing goal priorities (Oatley and Johnson-Laird, 1987) resulting in low activity (Izard and Ackerman, 2000) or an aggressive exploration of new interactions (Blumberg and Izard, 1986). This is not a straightforward explanation of the effect of sadness, and points at a possible underlying problem that has been voiced by Barr-Zisowitz (2000) as the question whether or not sadness exists at all as a stable emotional state.

The interpretation of the relevant interaction properties as hedonic tones shows that angry and sad music only result in better observable behaviour at the cost of less satisfactory use of resources. This supports the idea that negative moods signal sub-optimal interaction and initiate action adjustments by increasing effort to optimise interaction.

There is, however, another possible explanation for the influence of music on the observed differences in performance. If we interpret the influence of music as an effect on participant's arousal, this would explain difference in invested resources, i.e. applied force and co-contraction forces (Van Galen, Müller, Meulenbroek, and Van Gemmert, 2002). A similar explanation was given for the effect of the speed of music on car-drivers (Brodsky, 2002). Although these explanations can also

partially explain the effects observed in this experiment, I nevertheless argue that there is a qualitative effect of mood, following the assumed functions of emotions. Support for this argument is found in the differences in observed behaviour between the happy and the angry music conditions, where no differences in arousal were measured, and where the music did not differ much in speed. With the data presented in this chapter no definitive conclusion can be made, however. This issue should therefore be researched to a greater extent in future experiments.

A comparison of the results of the first experiment and the second experiment indicate that the interaction is more regulated when the interaction task is an explicit task. This means that more resources are employed for the interaction process. We can make sense of this by considering the task of the self-regulatory systems that regulate each of the two situations. In the first experiment many resources could be used for the regulation of mouse-cursor-interaction. In the second experiment, at the goal level, a self-regulatory system had to decide how much resources to give to the two separate regulatory-systems, for mouse-cursor-regulation and for the regulation of the card game. By dividing the resources both systems will receive less resources and will be less able to generate the best action control. This shows that optimisation at a higher level may be at the expense of optimality on each of the lower levels.

5.3 General discussion of chapter 5

The results of both experiments confirm that user adaptation of interactions can be predicted as the outcome of a self-regulatory system that adjusts movements to meet contextual demands. Specific task related demands are achieved by changing movement interaction properties. In the first experiment, stereotypical action adjustments for participants as well as tasks were observed. This shows that individual action adjustments had converged on stable individual and task effective action patterns. Two elements of observable efficiency in a task, movement time and path precision, increased when the external demands on the participants were increased. However, the number of failures to achieve the task indicates that these tasks are more difficult and that a lower adequacy has to be expected. This lower adequacy was distributed over the remaining interaction properties and the number of failures to achieve the task; adjustments were made at a cost. The fact that one participant accepted failures rather than adjust the

interaction parameters is an indication that individually different, higher-level considerations (effectiveness) are made when demands on the execution of a skilled task increase. When low-level interaction processes are frustrated, the task at a higher strategic level is reconsidered in an effort to optimise the whole interaction structure. Conversely, it is concluded from the results of the second experiment that higher-level interaction optimisation influences all the elements of an ongoing interaction. This control includes the optimisation of lower levels such as mouse-cursor-movements, even when this is not a specific task.

A limitation to these conclusions is that the effort of the interaction (force and co-contraction) had to be approximated. Although the approximations of force and co-contraction forces give insight into the cost-benefit trade-off in the optimisation of behaviour, neither force nor co-contraction forces were measured directly. The approximation of force and co-contraction force, especially for high accuracy tasks, should be validated, particularly as the assumption that movement is ballistic does not hold for all the studied tasks.

Interpreting the data as hedonic tones, an affective experience, allows straightforward comparison between the different relevant elements of interaction. Such a straightforward comparison is a tool for the synthesised self-regulatory mechanism in determining the overall quality of interaction. All properties can be compared in this way leading to a single, affective estimation of adequacy referred to as pleasure (Cabanac, 1992). Maximising pleasure results in improvements in the interaction. The accompanying qualitative signals about the direction of action adjustments were assumed to be conveyed by specific emotions. In the second experiment, I found evidence that emotion plays a functional role in signalling action adjustments. Anger leads to an increased effort with little precision, sadness to reconsidering the importance of individual goals and happiness to a state of continuation of the current situation.

To summarise the findings of the experiments, I can state that users regulate movement in a way that results optimal interaction parameters. Increasing task demands result in the re-distribution of interaction properties for the better, based on a variation of an existing stereotype. Furthermore, I have found indications that users regulate mouse-cursor-movements even when this is not an explicit task, and that emotions play an important role in this regulation process.

Chapter 6: General discussion

The combined outcomes of the studies presented in this thesis show that the adaptation of user behaviour in simple interaction tasks can be described as the stable result of a self-regulatory system. More precisely, this self-regulatory system can be specified as a subconscious feedback mechanism made up of a monitoring, an evaluation and an action adjustment mechanism. This conclusion supports the main hypotheses that were drafted in chapter 1. I have also shown that in the implementation of the self-regulatory system, affect related phenomena can be interpreted as functional signals for the determination of the adequacy of interaction, as well as specific signals for the action adjustment mechanism. Mechanisms for combining and storing experience to provide a reference for the evaluation of interaction can also be specified. These mechanisms for storing and using experience explain the observed differences between users as the result of prior experience and deviations in personal properties rather than a deviation from the assumption of optimisation. The described system as a whole is capable of flexibly selecting adequate interactions and steering user actions towards an improvement in a wide range of user-system interaction tasks in changing contexts.

Designers can use the results of this work to get an idea of how an interaction process will be controlled by users, based on experience, emotions and the information that is acquired from the ongoing interaction process. This knowledge will enable them to design interaction processes that are as closely related as possible to the regulation of interaction. If interactions are designed in this way, users will be able to adapt better to different interactions. This easier adaptation could be especially advantageous when learning new applications, operating

applications that are for one-time-use only, or in operating changing applications or applications in a changing context. In the specified self-regulatory system, the emergence of reference values for the interaction process properties plays a central role. Further study of the process of emerging reference values might lead to the design of software that allows a gradual and consistent accumulation of user experience. Such a gradual change could be used to train users to adapt to the new application, without them being aware that training is occurring.

The main conclusion of this thesis is that users can adapt their behaviour while interacting with interactive applications. The research was carried out using the top-down reverse engineering approach. In this approach it is assumed that human behaviour can be regarded as functional and goal-directed. This assumption implies that interaction has to be controlled. To study the control of interaction, a self-regulatory system for interaction improvement was specified. At a high description level this self-regulatory system was specified as a feedback control system. To be able to improve interaction, such a feedback mechanism should be able to monitor and evaluate ongoing interaction, and to use the information gathered in this way to adjust actions. Based on these specifications, the behaviour of the system was predicted in different situations and compared with the behaviour of participants in experimental tasks. The following recapitulation of the findings from the empirical chapters shows how support for the functionality of the self-regulatory system was gathered by relating the different elements of the specified system to a range of interaction tasks.

The monitoring mechanism was studied in chapter 2. Here I confirmed that perception processes can execute the tasks of the monitoring mechanism. This conclusion was based on the observation that the participants keep track of interaction time, at least when it is an important interaction property.

In chapter 2 the main functionality of the evaluation of interaction was also studied by asking the participants to report their satisfaction. The negative influence of user errors on satisfaction indicates that the evaluation mechanism can correctly interpret and report errors as an indicator for a lower adequacy of the ongoing interaction. I argued that errors can be interpreted as mainly related to one indicator for interaction adequacy, namely effectiveness. A second indicator for the adequacy of interaction, efficiency, was also shown to be important for the

evaluation of interaction. By comparing the outcome of the monitoring mechanism with the outcome of the evaluation mechanism, I made plausible that the evaluation mechanism determines the adequacy of interaction by using the output signals from the monitoring mechanism as input.

In chapter 3, the functionality of the evaluation mechanism was studied in more detail. The additional findings focussed on the ability of the evaluation mechanism to flexibly determine the efficiency of an ongoing interaction process. The evaluation mechanism generates the most positive scores for the most efficient interaction processes. This confirms that the evaluation mechanism is successful in determining the adequacy of the interaction. The evaluation mechanism was also shown to be flexible in adapting to new or changed situations. This was concluded from the finding that in cases where the interaction deviates from a well-known standard, the evaluation mechanism gradually learns to interpret the deviating interaction correctly. Using these gradual changes, evaluation is interpreted as a process of comparing the ongoing interaction with reference values for that interaction. These reference values are determined by growing experience, rather than through rational deliberation. Although memory was not explicitly studied, these results suggest that reference values for typical interaction patterns are stored as an associatively accessed emotion (such as somatic markers, Damasio 1994). The argument that reference values, which emerged from experience, rather than rationally constructed reference values determine the evaluation of interaction suggests that a strict hierarchical chain of sub-goals and action strategies (as described in for example Newell, 1990; Norman, 1984; Powers, 1973) cannot always describe human interaction control. An alternative, psychological, explanation that accounts for such deviations from rationality can be found in dual-process theories (e.g., Chen and Chaiken, 1999; Petty and Wegener, 1999; Sloman, 1996). These theories introduce associative, intuitive or heuristic decisions in situations where, for one reason or another, the required rational effort cannot or should not be generated. When applied to practised human-computer-interaction processes, heuristic control of interaction uses fewer resources than rational deliberation. The findings that interaction processes are monitored and evaluated based on the relevance for the interaction task in hand, combined with the interpretation of the outcome of the evaluation process as emotion leads to the interpretation of the self-

regulation of interaction as a process of emergent emotions (Frijda, 1986), in which action adjustments are initiated without the necessity of conscious control.

To understand how such a sub-conscious, intuitive, evaluation of interaction as studied in chapter 2 and 3 can lead to the improvement of interaction processes, actual changes of the interaction were studied in the chapters 4 and 5. This was done by attributing heuristic signals to the synthesized self-regulatory system. Following the interpretation of interaction control along the lines of emergent emotions, it was also investigated how emotional influences are related to interaction control.

In chapter 4, an experiment was reported that investigated how actions are adjusted in a completely new, ill-defined, situation. The gradual accumulation of new knowledge allowed the participants to choose better options in the course of the experiment. This finding shows that the interaction control system can learn to recognise the better options and act upon this observation. The gradual change suggests that the underlying process is associative and probably not fully conscious. In this same experiment the ideas about the influence of emotions on the adaptation of user-system-interaction were more explicitly studied by changing the mood of the participants. Changes in mood influence the affective interaction-control, which confirms that new interactions have at least an initial period of heuristic or intuitive control. In this case emotions are interpreted as useful heuristics that allow users to make reasonable adjustments in ill-defined situations, especially when rational control cannot, or cannot yet, generate optimal behaviour. The intuitive, emotional judgement is interpreted as being based on accumulated experience (see for example Sloman, 1996). This means that the intuitive evaluation of different interactions will not be the same for different users, since the reference values for the evaluation have emerged from different personal experiences. This makes sense when considering the fact that there might be different but equally profitable interactions possible (such as decks C and D from the experiment in chapter 4). Past personal experiences leading to stable reference values might therefore also lead to lasting interpersonal differences, as long as the solutions are approximately equally adequate.

In automated action patterns, differences in the set of reference values for the evaluation mechanism can be observed as personal stereotypes. In chapter 5 such stereotypical interaction patterns were found, for a highly practised computer task,

i.e. mouse-cursor-movements. When the boundary conditions of the interaction tasks changed, the action pattern changed, but the participants still exhibited their own stereotype. I argue that this is an indication that the reference values of the different participants have emerged from personal attributes such as physical makeup, personal history or experience and possibly even character traits. In the final experimental study, I found further support for the idea that interaction processes are regulated by showing functional and converged interactions even when there was no explicit task associated with that interaction. This can be understood by assuming that each action sequence has to be regulated to achieve a goal. Although the automated actions that were the focus of this thesis were shown to be adjusted for goal achievement, it is highly unlikely that each action is consciously and deliberately regulated. The possibility that the control of such skilled tasks is at least partly affective was shown by manipulating mood of participants, which resulted in changes in the interaction patterns.

The research presented in this thesis does have some limitations and there are some open issues left that could be resolved in future research. If these issues are solved, further understanding of the regulation of user-system-interaction processes could be gained.

A first limitation is related to the measure for the outcome of the evaluation mechanism. In the studies reported in chapters 2 and 3, the relationship between the objective efficiency and the assumed measure for the outcome of the evaluation of interaction, self-reported satisfaction, was consistently found. This was interpreted as a confirmation of the idea that the evaluation mechanism integrates judgements about ongoing interaction into a single assessment value, and that this value can be reported. However, the relationship between objective efficiency and self-reported satisfaction was rather noisy. This noisiness can be partially explained by the fact that in asking participants to report their satisfaction, the participants had to translate the subconscious evaluation into a reported score. Such a transformation is likely to create noise, as well as a cognitive reinterpretation of the subconscious evaluation. In chapter 5, the outcome of the evaluation was estimated as the standardized difference between the effort and outcome of the tasks, in which the objectively better descriptions of tasks received the best (estimated) evaluation scores. However, there was no empirical evidence

to relate the estimated evaluation scores to the actual experience of the participants. To overcome these problems, I recommend additional experiments that allow better ways to measure self-reported satisfaction, possibly supported by physiological measures (e.g., galvanic skin resistance, heart beat variability).

At a more fundamental level, a future direction for an improved understanding of the subconscious regulation of interaction control is an investigation into the influence of emotions. Emotional control aspects were only superficially integrated in the synthesised system. As noted earlier, the relationship between a self-regulatory feedback control system and emotions has similarities with the process of emergent emotions as proposed by Frijda (1986). Further research should be carried out into the influence of affects on interaction control according to this theory of emotions, which implies that a situation generates a feeling that is responsible for action adjustments. A problem encountered in the presented investigation into the typical effects of emotions, was that the effect of the mood manipulation in chapter 4 and 5 could not be established in detail with the used measurement. Although the experimental results shows consistent differences according to the hypotheses, the mood measure could not distinguish between two types of music. To be able to make more explicit claims about the effects of specific emotions, effort should be invested in determining the exact influence of emotions on mood and how to measure this influence. At a high level, however, the interpretation of the data when using functions of emotions in the self-regulatory mechanism resulted in a consistent and fruitful interpretation of the control of different interactions processes. I therefore recommend that the role of emotions is further explored in future research, to gain a better understanding of the subconscious adaptation of user-system interaction.

A second fundamental issue applies to the applicability of the chosen method, that of the intentional top-down reverse engineering approach. An observation that already reveals the limitation of this cognitive model is related to the emergence of reference values through experience. This interpretation of reference values as emergent already suggests that there are not only intentional properties of the regulatory system that play a role. Many of these emergent phenomena can be described '*as if*' executed by rational agents. This interpretation is essential for humans to be able to interact with complex, otherwise un-intelligible devices, such as chess computers (Dennett, 1981). The question is, however, whether this

approach is always justified. At the higher specification levels of the control system (strategy and goals of the control system), the intentional stance successfully predicts the behaviour of many systems. However, in the final stage of the specification of a computational system, that of the implementation in terms of physical mechanisms (Marr, 1982), the intentional, reverse engineering, approach encounters problems. At the physical level the implementation of the functions should be seen as the results of the blind process of evolution, which led to the physics of all biological systems. With this process of evolution also physical and mental limitations to the biological systems evolved that limit the amount of leeway the rational agent has. Furthermore, the reverse-engineering approach requires the designer to specify the different elements of the agent beforehand and in isolation to prevent unforeseen side effects. With that limitation there are only scarce possibilities for the emergence of complex integrated and multifunctional functions. Bottom-up emergence of properties, for example optimised through natural selection, is not hindered by these limitations (Dennett, 1994) and might, in retrospect, lead to unforeseen levels of optimality beyond the grasp of a rational engineer. It is at this level that the actual and detailed understanding of human behaviour could be studied further. Studies into such emergent systems are a way to achieve better understanding of complex behaviour. This type of investigation is executed at a fundamental level in the field of artificial life (Langton, 1989). I therefore recommend that in the synthesis of mechanisms that mimic human behaviour both evolutionary and rational design processes should be investigated.

Based on one of the results of chapter 2, I argued that the full hierarchic chain of actions, actions sequences and strategies, is optimised. This means that the evaluator does not only determine the efficiency of the ongoing interaction, but also takes into account the effort needed to control the ongoing interaction sequence. The effect of the effort needed for control in interaction was hard to establish however, and was an order of magnitude smaller than the effect of the actual interaction efficiency. The difference in size of effect can be understood by looking at the implicit task of the control system, which is to optimise the interaction. To do this, the effort spent in controlling interaction should be smaller than the effort saved by the control process. The effort used to control interaction should be far smaller than the effort concerned in the ongoing interaction process. This effect was not investigated in detail because the efficiency of the self-

regulatory mechanism is mainly involved in the control of sequences of interaction, which are related to optimisation of action strategies. Additionally it was shown in chapter 5 that when an action sequence is not a conscious task, some effort is still invested in the optimisation of the action sequence. The control of the different levels of interaction can apparently be interpreted as an aggregated optimisation, i.e. not only on each individual level, but also accumulated across the levels. If this effect, and the suggested explanations, is investigated in more detail, more complete understanding of the control of interaction processes and its relationship to the optimisation of action strategies can be achieved. This can be done by using the reverse engineering approach to interpret the control of interaction at the level of strategic or goal-directed considerations. The simultaneous control of the different tasks in the second experiment of chapter 5 could be considered as the optimisation of the interaction at a strategic level. The heuristic way, in which ill-defined problems are treated, is another optimisation process for a strategy of interaction, consisting of the choice between different actions, rather than the optimisation of the actions themselves. The experiment in chapter 4 provides a stepping-stone for such research. If this effort is undertaken and the regulation of interaction at all levels is studied in more detail, a better understanding of how users integrate different goals, strategies and actions in their daily interaction with applications can be achieved. Such insights might lead to a better understanding of the context in which users interact with technology.

In sum, there were some limits to the approach that have to be solved in future research. These limits focus on the understanding of emergent properties and the general application of the ideas for the regulation of the strategic and goal-directed interpretation of interaction processes. Outside these limits the proposed self-regulatory system is successful at predicting user-adaptation in a wide range of different, simple, user-system-interaction processes.

References

- Accot, J., & Zhai, S. (2001) Scale effects in steering law tasks. In J. Jacko & A. Sears (Eds.) *Proceedings of SIGCHI conference on human factors in computing*, (pp. 1-8), New York: ACM Press.
- Ariely, D., & Loewenstein, G. (2000) When does duration matter in judgement and decision-making? *Journal of experimental psychology: general*, 129, 508-523.
- Barr-Zisowitz, C. (2000) "Sadness" – is there such a thing? In M. Lewis & J. M. Haviland-Jones (Eds.) *Handbook of emotions 2nd edition*, (pp 607-622), London: Guilford Press.
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. (1994) Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50, 7-12.
- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997) Deciding advantageously before knowing the advantageous strategy. *Science*, 275, 1293-1295.
- Bernstein, N. (1967) *The co-ordination and regulation of movements*. Oxford UK: Pergamon Press.
- Blommaert, F. J. J., & Janssen, T. J. W. M. (1999) The satisfied user. *IPO annual progress report*, 34, 54-62.
- Blumberg, S. H., & Izard, C. E. (1986) Discriminating patterns of emotions in 10- and 11- year-old children's anxiety and depression. *Journal of personality and social psychology*, 51, 852-857.
- Braitenberg, V. (1984) *Vehicles: Experiments in synthetic psychology*. Cambridge MA: The MIT Press.
- Brinkman, W. P. (2003) *Is usability compositional?* Doctoral dissertation, Eindhoven University of Technology, Eindhoven, The Netherlands.

- Brodsky, W. (2002) The effects of music tempo on simulated driving performance and vehicular control. *Transportation Research Part F*, 4, 219-241.
- Cabanac, M. (1992) Pleasure: the common currency. *Journal of theoretical biology*, 155, 173-200.
- Card, S. K., Moran, T. P., & Newell, A. (1983) *The psychology of human-computer interaction*. London: Lawrence Erlbaum.
- Carver, C. S., & Scheier, M. F. (1998) *On the self regulation of behavior*. Cambridge UK: Cambridge University Press.
- Chen, S., & Chaiken, S. (1999) The heuristic-systematic model in its broader context. In S. Chaiken & Y. Trope (Eds.) *Dual-process theories in social psychology*, (pp. 73-96), New York: Guilford Press.
- Cherniak, C. (1999) Rational agency. In R. A. Wilson & F. C. Keil (Eds.) *The MIT Encyclopedia of the cognitive sciences*, (pp. 698-699), Cambridge, MA: The MIT Press.
- Chow, D. S. L. (2001) The effects of time delay in electronic commerce. In J. Jacko & A. Sears (Eds.) *CHI 2001, extended abstracts*, (pp. 387-388), New York: ACM.
- Cohen, H. A. (1975) *The art of snaring dragons*. Massachusetts Institute of Technology, A. I. Memo 338. Retrieved June 14 2003 from the world wide web <ftp://publications.ai.mit.edu/ai-publications/pdf/AIM-338.pdf>.
- Cohen, P. R., Atkin, M. S., & Hansen, E. A. (1994) *The interval reduction strategy for monitoring cupcake problems*. University of Massachusetts Amherst, Computer science technical report 94-09. Retrieved May 16 2003 from the world wide web http://www.cs.umass.edu/Dienst/UI/2.0/Describe/ncstrl.umassa_cs%-2FUM-CS-1994-009.
- Czerwinski, M., Horvitz, E., & Cutrell, E. (2001) Subjective duration assessment: an implicit probe for software usability. *Proceedings of IHM-HCI 2001 Conference, Volume 2*, (p. 167).
- Damasio, A. R. (1994) *Descartes Error, emotion, reason, and the human brain*. New York: Putnam.
- Davidson, A. R. (1995) From attitude to action to attitude change: the effects of amount and accuracy of information. In R. E. Petty and J. A. Krosnick (Ed.) *Attitude strength: Antecedents and consequences*, (pp. 315-336), Hillsdale NJ: Lawrence Erlbaum.
- Dawkins, R. (1976) *The selfish gene*. Oxford UK: Oxford University Press.
- Dawkins, R. (1986) *The blind watchmaker*. London: Penguin.

- Dellaert, B. G. C., & Kahn, B. E. (1999) How tolerable is delay? Customer's evaluations of internet web sites after waiting. *Journal of interactive marketing*, 13, 41-54.
- Dennett, D. C. (1981) Intentional systems. In J. Haugeland (Ed.) *Mind design*, (pp. 220-242), Cambridge MA: MIT press.
- Dennett, D. C. (1994) Cognitive science as reverse engineering: several meanings of "top-down" and "bottom-up". In D. Prawitz, B. Skyrms, & D. Westerståhl (Eds.) *Logic, methodology and philosophy of science IX, a companion to the philosophy of mind*, (pp. 236-244), Oxford UK: Blackwell Press.
- Dennett, D. C. (1995) *Darwin's dangerous idea: evolution and the meaning of life*. New York: Simon and Schuster.
- Dennett, D. C. (1999) Intentional stance. In R. A. Wilson & F. C. Keil (Eds.) *The MIT Encyclopedia of the cognitive sciences*, (pp. 412-413), Cambridge MA: The MIT Press.
- Dix, A. (1994) *Seven Years on, the Myth Continues*. (Research Report RR9405). Huddersfield, UK: University of Huddersfield, School of Computing and Mathematics. Retrieved September 1 2002 from the world wide web <http://www.comp.lancs.ac.uk/computing/users/dixa/papers/myth95/7-years-on.html>.
- Eisler, H., & Eisler, A. D. (1991) A mathematical model for perception with experimentally obtained subjective time scales for humans and rats. *Chronobiologia*, 18, 79-88.
- Epstein, S. (1994) Integration of the cognitive and psychodynamic unconscious. *American Psychologist*, 49, 709-724.
- Evans, D. (2001) *Emotion: the science of sentiment*. Oxford UK: Oxford University Press.
- Fischer, A. R. H., & Blommaert, F. J. J. (2001) Effects of time delay on user satisfaction. In M. G. Helander, H. M. Khalid, & M. P. Tham (Eds.) *Proceedings of the International Conference on Affective Human Factors Design*, (pp. 407-414), London: Asean Academic Press.
- Fitts, P. M. (1954) The information capacity in the human motor system in controlling the amplitude of movement. *Journal of experimental psychology*, 47, 381-391.
- Frijda, N. H. (1986) *The emotions*. Cambridge UK: Cambridge University Press.

- Green, D. P., & Salovey, P. (1999) In what sense are positive and negative affect independent? A reply to Tellegen, Watson, and Clark. *Psychological science*, 10, 304-306.
- Harris, C. H. (1998) On the optimal control of behaviour: a stochastic perspective. *Journal of neuroscience methods*, 83, 73-88.
- Harris, C. H., & Wolpert D. M. (1998) Signal-dependent noise determines motor planning. *Nature*, 394, 780-784.
- Hornof, A. J., & Kieras, D. E. (1997) Cognitive modelling reveals menu search is both random and systematic. In: *Proceedings of SIGCHI conference on human factors in computing*, (pp. 107-114), New York: ACM.
- Isen, A. M. (2000) Positive affect and decision making. In M. Lewis & J. M. Haviland-Jones (Eds.) *Handbook of emotions 2nd edition*, (pp 417-435), London: Guilford Press.
- ISO 9421 - 11 (1997) part 11: Guidance on the usability specification of measures, part of ISO 9241, *Ergonomic requirements for office work with visual display units (VDT's)*.
- Izard, C. E., & Ackerman, B. P. (2000) Motivational, organizational, and regulatory functions of discrete emotions. In M. Lewis & J. M. Haviland-Jones (Eds.) *Handbook of emotions 2nd edition*, (pp 253-264), London: Guilford Press.
- Janssen, T. J. W. M., & Blommaert, F. J. J. (2000) Visual metrics: discriminative power through flexibility, *Perception*, 29, 965-980.
- Johnston, V. S. (1999) *Why we feel, the science of human emotions*. Cambridge MA: Perseus.
- Karoly, P. (1993) Mechanisms of self-regulation: a systems view. *Annual review of psychology*, 44, 23-52.
- Katar, M., Britton, C., & Nehaniv, C. (2000) Specifying multiple time granularities in interactive systems. In P. Palanque & F. Paternò (Eds.) *ISCE 2000 workshop on: design, specification and verification of interactive system*, (pp. 49-61), New York: ACM.
- Keuning, H. (2003) *Augmented force feedback to facilitate target acquisition in human-computer interaction*. Doctoral dissertation, Eindhoven University of Technology, Eindhoven, The Netherlands.
- Keuning-van Oirschot, H., & Houtsma, A. J. M. (2001) Cursor trajectory analysis. In S. Brewster & R. Murray-Smith (Eds.) *Haptic HCI, LNCS 2058*, (pp. 127-134), Berlin: Springer.

- Kotler P. (1991) *Marketing management seventh edition*. London: Prentice Hall.
- Langton, C. G. (1989) Artificial life. In: C. G. Langton (Ed.) *Artificial life*, (pp. 1-49) Reading, MA: Addison Wesley.
- LeDoux, J. (1996) *The emotional brain*. New York: Simon and Schuster.
- Lewis, L. M., Dember, W. N., Schefft, B. K., & Radenhausen, R. A. (1995) Can experimentally induced mood affect optimism and pessimism scores? *Current psychology: development learning personality social*, 14, 29-41.
- Marr, D. (1982) *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.
- Mellers, B. A. (2000) Choice and the relative pleasure of consequences. *Psychological bulletin*, 126, 910-924.
- Mellers, B. A., Schwartz, A., & Ritov, I. (1999) Emotion-based choice. *Journal of Experimental Psychology*, 128, 332-345.
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960) *Plans and the structure of behavior*. New York: Holt, Rhinehart, & Winston.
- Newell, A. (1990) *Unified theories of cognition*. Cambridge MA: Harvard University Press.
- Newell, A., & Simon, H. A. (1972) *Human problem solving*. Englewood Cliffs NJ: Prentice Hall.
- Nijenhuis, M. R. M., & Blommaert, F. J. J. (1997) Perceptual-error measure and its application to sampled and interpolated single-edged images. *Journal of the optical society of America A*, 14, 2111-2127.
- Norman, D. A. (1984) Stages and levels in human-machine interaction. *International journal of man-machine studies*, 21, 365-375.
- Norman, D. A. (1988) *The psychology of everyday things*. New York: Basic Books.
- Norman, D. A., & Bobrow, D. G. (1975) On data-limited and resource-limited processes. *Cognitive psychology*, 7, 44-64.
- Oatley, K., & Jenkins, J. M. (1996) *Understanding emotions*. Oxford UK: Blackwell Publishers.
- Oatley, K., & Johnson-Laird, P. N. (1987) Towards a cognitive theory of emotions. *Cognition and emotion*, 1, 29-50.

- Oatley, K., & Johnson-Laird, P. N. (1995) The communicative theory of emotions: empirical tests, mental models, and implication for social interaction. In L. L. Martin & A. Tesser (Eds.) *Striving and feeling: interaction between goals and affect*, (pp. 363-393), Hillsdale NJ: Erlbaum.
- Petty, R. E., & Cacioppo, J. T. (1981) *Attitudes and persuasion: classic and contemporary approaches*. Oxford UK: Westview.
- Petty, R. E., & Wegener, D. T. (1999) The elaboration likelihood model: current status and controversies. In S. Chaiken & Y. Trope (Eds.) *Dual-process theories in social psychology*, (pp. 41-72), New York: Guilford Press.
- Powers, W. T. (1973) *Behaviour: the control of perception*. New York: Aldine de Gruyter.
- Rasmussen, J. (1983) Skills, rules and knowledge; signals, signs and symbols, and other distinctions in human performance models. *IEEE transactions on systems, man and cybernetics*, 13, 257-266.
- Reason, J. (1990) *Human error*. Cambridge UK: Cambridge University Press.
- Reeves, B., & Nass, C. (1996) *The media equation: how people treat computers, television, and new media like real people and places*. Cambridge UK: Cambridge University Press.
- Russell, J. A., & Lemay, G. (2000) Emotion concepts. In M. Lewis & J. M. Haviland-Jones (Eds.) *Handbook of emotions 2nd edition*, (pp 491-503), London: Guilford Press.
- Schmitt, W. A., Brinkley, C. A., & Newman, J. P. (1999) Testing Damasio's somatic marker hypothesis with psychopathic individuals: risk takers or risk averse? *Journal of abnormal psychology*, 108, 538-543.
- Shafir, E., & Tversky, A. (1995) Decision making. In E. E. Smith & D. N. Osherson (Eds.) *An invitation to cognitive science 2nd edition volume 3: Thinking*, (pp. 77-100), Cambridge MA: MIT Press.
- Shepperd, J. A., & McNulty, J. K. (2002) The affective consequences of expected and unexpected outcomes. *Psychological Science*, 13, 85-88.
- Shore, D. I., & Klein R. M. (2000) On the manifestations of memory in visual search. *Spatial Vision*, 14, 59-75.
- Simon, H. A. (1967) Motivational and emotional controls of cognition. *Psychological review*, 74, 29-39.
- Simon, H. A. (1990) Invariants of human behavior, *Annual Review of Psychology*, 41, 1-19.

- Sloman, A. (1999) Beyond shallow models of emotion. *Presented at the 13 spring days workshop on behaviour planning for life-like characters and avatars*, retrieved March 25 2001 from the World Wide Web: <http://citeseer.nj.nec.com/253617.html>.
- Sloman, S. A. (1996) The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.
- Stevens, S. S. (1975) *Psychophysics: Introduction to its perceptual, neural, and social prospects*. Oxford UK: John Wiley and Sons.
- Teal, S. L., & Rudnicky, A. I. (1992) A performance model of system delay and user strategy selection. In P. Bauersfield, J. Bennett & G. Lynch (Eds.) *Proceedings of SIGCHI conference on human factors in computing*, (pp. 295-305), New York: ACM.
- Todorov, E., & Jordan, M. I. (2002) Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5, 1226-1235.
- Van der Aalst, W. M. P., & Van Hee K. M. (1997) *Workflow management: modellen, methoden en systemen*. Schoonhoven The Netherlands: Academic Service.
- Van Galen, G. P., & De Jong W. P. (1995) Fitts' law as the outcome of a dynamic noise filtering model of motor control. *Human movement science*, 14, 539-571.
- Van Galen, G. P., Müller, M. L. T. M., Meulenbroek, R. G. J., & Van Gemmert, A. W. A. (2002) Forearm EMG response activity during motor performance in individuals prone to increased stress reactivity. *American journal of industrial medicine*, 41, 406-419.
- Van Galen, G. P., & Schomaker, L. R. B. (1992) Fitts' law as a low-pass filter effect of muscle stiffness. *Human movement science*, 11, 11-22.
- Watson, D., Clark, L. A., & Tellegen, A. (1988) Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology*, 54, 1063-1070.
- Winston, P. H. (1992) *Artificial intelligence 3rd edition*. Amsterdam: Addison-Wesley.
- Wolfe, J. M., Alvarez, G. A., & Horowitz, T. S. (2000) Attention is fast but volition is slow. *Nature*, 406, 691.
- Wolpert, D. M. (1997) Computational approaches to motor control. *Trends in cognitive sciences*, 1, 209-216.
- Zakay, D. (1993) Time estimation methods - do they influence prospective duration estimates? *Perception*, 22, 91-101.

- Zakay, D., & Shub, J. (1998) Concurrent duration production as a workload measure. *Ergonomics*, 41, 1115-1128.
- Zhai, S., & Milgram, P. (1998) Quantifying coordination in multiple DOF movement and its application to evaluating 6 DOF input devices. *Proceedings of SIGCHI conference on human factors in computing*, (pp. 320-327), New York: ACM Press.

Appendices

Appendix A: Music files

Three music files were compiled to induce different moods. The amplitudes of loudness of the different tracks were equalised so all tracks were of approximately equal loudness.

Table A.1: Music tracks that were used to manipulate participant mood in chapters 4 and 5.

	Condition	Length	Tracks	Band	Year
1	Angry	12.42	Bombtrack	Rage against the Machine	1992
			The Future of War	Atari Teenage Riot	1997
			Killing in the Name	Rage against the Machine	1992
2	Happy	13.17	Surfin' U.S.A.	Beach Boys	1963
			I Shot the Sheriff	Bob Marley	1973
			Night Boat to Cairo	Madness	1982
			Brazilian Love Song	Steelbands of the Caribbean	1996
3	Sad	12.48	Fatal Wound	A Minor Forest	1999
			The Eternal	Joy Division	1980

Appendix B: PANAS scale

In this appendix the translations of the instruction, items and selection options of the PANAS scale are presented. The original English texts by Watson, Clark, and Tellegen (1988) are given for comparison.

Dutch instruction

“Dit invulformulier bestaat uit een aantal woorden die verschillende emoties en gevoelens beschrijven. Lees elk woord en geef dan achter het woord het antwoord aan dat van toepassing is. Geef aan in hoeverre je je nu, op dit moment, voelt volgens het gegeven woord.”

Original English instructions

“This scale consists of a number of words that describe different feelings and emotions. Read each item and mark the appropriate answer in the space next to that word. Indicate to what extent you feel this way right now, that is, at the present moment.”

Table A.2: Answer options and coding values for each of the PANAS items.

Dutch	English	Score
heel weinig of helemaal niet	very slightly or not at all	-2
een klein beetje	a little	-1
middelmatig	moderately	0
behoorlijk	quite a bit	1
heel erg	extremely	2

Table A.3: PANAS items.

Dutch word	(original) English word	Item code
enthousiast	enthusiastic	pa1
betrokken	interested	pa2
vastbesloten	determined	pa3
opgewonden	excited	pa4
geïnspireerd	inspired	pa5
alert	alert	pa6
actief	active	pa7
zelfverzekerd	strong	pa8
trots	proud	pa9
aandachtig	attentive	pa10
angstig	afraid	na1
bang	scared	na2
overstuur	upset	na3
bedroefd	distressed	na4
schrikachtig	jittery	na5
nerveus	nervous	na6
beschaamd	ashamed	na7
schuldbewust	guilty	na8
prikkelbaar	irritable	na9
vijandig	hostile	na10

Samenvatting

(Dutch summary)

In dit proefschrift wordt onderzocht hoe mensen de interactie met computerapplicaties reguleren. Hiervoor is de ‘reverse engineering’ aanpak gebruikt. In deze aanpak wordt een reguleringsysteem voor interactie ontworpen dat dezelfde interactieprocessen reguleert als bij daadwerkelijke gebruikers het geval is. Het gedrag van dit ontworpen reguleringsysteem wordt dan vergeleken met het werkelijke gedrag van proefpersonen.

Als wordt aangenomen dat mensen doelgericht gedrag zo adequaat mogelijk uitvoeren, dan heeft het ontworpen zelfregulerende systeem als doel er voor te zorgen dat interactietaken zo goed mogelijk uitgevoerd worden. Het onderzochte zelfregulerende systeem is zo gespecificeerd dat het functioneert op basis van de terugkoppeling van informatie over het lopende interactieproces. De terugkoppeling bestaat uit drie onderdelen: een waarnemingsmechanisme, een evaluatiemechanisme, en een actie-aanpassingsmechanisme. Deze drie onderdelen zijn uitgewerkt met hulp van bestaande psychologische en HCI theorieën. Het waarnemingsmechanisme is gespecificeerd als een perceptieproces, dat informatie filtert en transformeert. De resultaten van het waarnemingsmechanisme worden door het evaluatiemechanisme vergeleken met referentiewaarden voor interactieprocessen van een zelfde type. Voor het beschrijven van het evaluatieproces zijn inzichten toegepast over de invloed van ervaring op referentiewaarden, en over emoties als heuristische signalen aangaande de kwaliteit van het interactieproces. De uitkomst van de evaluatie geeft informatie over de kwaliteit van de interactie en een signaal welke aanpassingen gemaakt moeten worden om de interactie te verbeteren. Tenslotte wordt deze aanpassing uitgevoerd.

Het zelfregulerende systeem is bestudeerd in een aantal experimentele studies. In een eerste experiment bleken proefpersonen in staat de efficiëntie van een taak (tijdsduur), en de invloed van experimentele manipulaties op de tijdsduur in te

schatten. Deze inschatting kon vervolgens worden gebruikt om een oordeel over de kwaliteit van de interactie te vormen. In een volgend experiment werd bevestigd dat de gerapporteerde kwaliteit van interactie is gerelateerd aan objectieve efficiëntie. Tevens werd aangetoond dat met toenemende ervaring, het evaluatiemechanisme de specifieke eigenschappen van de experimentele stimuli beter leert te gebruiken. Om te onderzoeken hoe acties worden aangepast op basis van de evaluatie, werd een experiment uitgevoerd waarin proefpersonen winst moesten maken in een gokspel. Proefpersonen vertoonden een geleidelijke overgang naar meer winstgevende keuzes. Een aantal proefpersonen gokte reeds zeer snel op een winstgevende strategie en haalde hoge scores zonder het hele spel te verkennen. Dit toont aan dat behalve rationele, ook intuïtieve processen bepalen welke acties worden uitgevoerd. Door met muziek de emoties te manipuleren werden deze non-rationele processen beïnvloed. Met name agressieve muziek leidde tot het nemen van grote risico's. In een laatste serie experimenten is specifiek gekeken naar regulatie van motorische taken, meer specifiek cursor manipulatie met een computer muis. Proefpersonen varieerden een persoonlijk stereotypisch actiepatroon om aan de voorwaarden van verschillende experimentele taken te voldoen. In een laatste studie werd bevestigd dat ook wanneer bewegingscontrole geen expliciete taak is, muisbewegingen gecontroleerd worden. Als proefpersonen voorafgaand aan het experiment luisterden naar muziek met verschillende emotionele lading, leidde dit tot verschillen in muisbewegingen volgens verwachte patronen, zoals snel en niet precies bij agressieve muziek. Dit bevestigt dat intuïtieve regulatie ook op aanpassingen van geoefend gedrag invloed heeft.

De uitkomsten van deze studies leiden tot de conclusie dat adaptaties van gebruikers aan simpele mens-machine interactieprocessen kunnen worden beschreven als een terugkoppelingsmechanisme waarin op een onderbewuste manier, op basis van opgedane ervaringen en emotionele ervaringen objectieve verbeteringen van interactie in gang gezet worden. Deze kennis kan praktisch worden gebruikt om interactieprocessen te ontwerpen die de regulering van het gedrag zo goed mogelijk ondersteunen.

Summary

The topic of this thesis is the way in which users regulate their interaction with computer applications. The research was carried out by adopting the reverse engineering approach. According to this approach a self-regulatory system was designed that regulates the same interaction processes as those carried out by users. The behaviour of the self-regulatory system was then compared to the behaviour of actual users.

If it is assumed that users execute goal-directed tasks as optimally as possible, the task of the self-regulatory system is to optimise ongoing interaction processes. The specified self-regulatory system is based on the feedback of information about the ongoing interaction process. The feedback loop is specified using three modules: a monitor, an evaluator and an action adjuster. These modules are specified using existing psychological and human-computer interaction theories. The monitor is specified as a perception process that filters and transforms physical information. The outcomes of the monitor are compared with reference values for similar interaction processes by the evaluator. In the specification of the evaluator, knowledge about changing reference values under the influence of experience, and emotions as heuristic signals about the quality of interaction processes are applied. The result of the evaluator is a value for the adequacy of interaction and a signal to the action adjuster indicating the changes that have to be made to improve the ongoing interaction.

This self-regulatory system was studied in a series of experiments. In the first two experiments evidence was found that the interaction is monitored through the perception process. The participants were able to estimate the efficiency of a task (time-to-task-completion) and the influence of experimental manipulations on this efficiency. This estimation can be used to evaluate the adequacy of the interaction. In a subsequent experiment, it was confirmed that reported adequacy is related to

the objective efficiency of interaction. It was also shown that with increasing experience, the evaluator learns to use specific characteristics of the experimental stimuli in to estimate interaction adequacy. To investigate how actions are adjusted based on the outcome of evaluation, the participants were asked to try to make a profit in an experimental card game. In the course of this experiment participants started to make more profitable selections. Some participants chose profitable strategies right from the start and scored high profits without ever exploring the game. This finding shows that in addition to rational processes, intuition also determines the course of actions. These non-rational processes were influenced by influencing participant mood through music. Aggressive music in particular prompted the participants to take larger risks, thus confirming the influence of associations. In a final series of experiments I focussed on action adjustments of skilled interactions, or specifically on cursor operations using the computer mouse. Participants varied a personal stereotype to match the given tasks. In a final study it was confirmed that the participants also regulated mouse-cursor movements when this was not an explicit task. Music-induced mood differences resulted in differences in mouse-cursor movements, which followed the expected patterns (for example, aggression resulted in quick but imprecise movements). This finding confirms that skilled behaviour is associatively adjusted.

The results of all these studies lead to the conclusion that the improvement of simple user-system-interaction processes can be described as the outcome of a feedback mechanism in which objective improvements of interactions are initiated subconsciously, based on experience, associations and affective experiences. Practically, this knowledge can be used to design interactions that optimally support the regulation of human-system-interaction.

Curriculum Vitae

- 12th of July 1971 Born in Rotterdam, the Netherlands
- 1983-1989 Preparatory scientific school, Gymnasium Erasmianum, Rotterdam
- 1989-1996 Delft University of Technology; M. Sc. Industrial Design Engineering received December 12th 1996.
- 1999-2003 Ph. D. research at the Eindhoven University of Technology – J. F. Schouten school for User-System Interaction (1999-2001 at the IPO – center for user-system interaction research, continued at the Human-Technology-Interaction group of the department of Technology Management)

