Roles and Goals: How Different Learning Strategies Impact What Is Learnt

April N Sinn

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

ACKNOWLEDGMENTS

First of all, I would like to thank my supervisor Anna Cox for her help with shaping the project and all of her guidance, support and patience. Special thanks go to Magda Osman for all her help with pointing me in the right direction, analysing the data and funding the study.

I would like to thank my 40 participants who volunteered to participate in my study. Without them, my thesis would not have been possible.

Lastly, I would like to thank all my friends and family; especially those who have become involved by trying to understand the research, supporting me and believing I can do anything I set my mind to.

ABSTRACT

One of the main applications of effective Human-Computer Interaction (HCI) is to develop systems that allow people to carry out their activities productively and safely. Therefore, exploring and discovering ways to enhance learning is vitally important in the field of HCI.

Previous studies have found that when learning to use a novel system, the type of goals (free vs. focused) impact the type of knowledge gained. Similarly, the difference in people's roles (controller vs. predictor) when learning to use an interactive system impacts the way people evaluate their progress of gaining knowledge, and in turn impacts the level of knowledge gained. By varying the type of exploration and the role of the participant, this study explored what happened when these goals and roles were considered simultaneously to understand the most effective ways of learning.

Forty participants were split equally into four different groups – free-controllers, free-predictors, focused-controllers, and focused-predictors. All groups learned to use an interactive program and then completed a series of tests to compare the knowledge they had gained of the system. The learning phase differed for each group of participants. The controller groups actively controlled the program while the predictor groups were instructed to predict the outcome based on the input values of the yoked controller. Free and focused groups were

given different instructions at the beginning of the task. Free groups were instructed to learn by exploration in an attempt to understand how the program works; whereas focused groups were instructed to keep the outcome at a certain level.

The main focus of this study was how the participants' acquisition of knowledge applies to tests. It was found that as a whole, predictors are able to gain the same knowledge of a system as the active controllers, as long as they are making predictions based on their evaluation of the outcome. This confirms the suggestion that people do not have to actively control a system to gain knowledge about using it effectively.

This study found that focused-controllers were initially superior to all other groups at performing task-specific tests. They were more experienced at controlling the system in this fashion and their focused goal to perform the task-specific tests gave them the advantage. It was also found that initially free-controllers and free-predictors were superior to focused-controllers and focused-predictors at transferring their knowledge to the device-oriented tests of the system because they had a better idea of the underlying concepts. This advantage stems from their ability to test their hypotheses whilst learning to use the system.

CONTENTS

CHAPTER 1. INTRODUCTION

1.1. Research motivation

Problem solving is an important cognitive activity that happens every day in normal life. The extent of the problem solving can range from solving normal everyday tasks to more important and more complex problems. While engaging in problem solving, people often make decisions with a particular goal in mind. However, it has been found that the decisions they make are often sub-optimal even when people are confident they understand the system they're using. So, we are left with the questions: why does this happen? How do they learn how to use things and are they learning in the right way to gain the correct information? The way people acquire knowledge of a system is an important part of problem solving. One of the core concepts in the field of Human-Computer Interaction (HCI) is to develop systems that meet the needs of the person using it in a way to allow them to complete their tasks productively, efficiently and safely. (Preece, Rogers, Sharp, Benton, Holland & Carey, 1994) In order to be able to do this, we need to know how people learn and acquire the knowledge to apply to the systems they are using.

A study by Cox and Young (2000, 2003) looked at the strategies that people used during exploratory learning and what impact these strategies had on what type of knowledge they learned. Cox and Young (2000) argued that users who gained device-oriented knowledge (about the rules of a system or what sections of a system do) were theoretically in a good position to use problem-solving skills to complete novel tasks, even though they had not necessarily learnt task-specific skills. However, the device-orientated knowledge alone was not enough to be able to complete domain tasks (how to complete a specific task using a system).

Osman (2008) conducted a study that looked at how the role of the participant impacted on information they learnt. They found that people who observed the actions of another person controlling a system, then predicted the outcome and acquired the same knowledge as those who actively controlled the device. Contrary to this study, in other studies (Berry, 1991; Berry & Broadbent, 1988) concluded that observing a system was too passive to be able to learn relevant knowledge as efficiently as someone who was learning by controlling the system. It would appear that by observing, and thinking about (predicting), the person doing the controlling is better than just observing the system.

This thesis attempts to understand and build on these studies by examining the learning strategies people adopt when learning to use interactive programs and investigating the effects of goals on what is learnt. It will vary the type of exploration (free or focused goals) and the role of the participant (whether they are an active controller or a predictor) when learning to use an interactive system. Then, it will investigate the effect the type of exploration and the role of the participant had on later performance tests. This will be measured by tracking error rates in ongoing performance.

At the same time as contributing to the HCI field of research, this study attempts to build on understanding situations in which the learning of an interactive program is the most effective and efficient. This in turn will help designers to build systems that allow people learn in a more enjoyable and natural way. Some examples of these systems where this can be applied are – walk-up-and-use systems, training for dynamic systems and simulation systems.

1.2. **Structure of this thesis**

This thesis starts with some background research (Chapter 2) in the area to be studied. This chapter will include previous research contrasting free and focused goals and research about controlling a system and predicting the outcome. The background chapter will build up to the aims of this study and the hypotheses that are being tested. Chapter 3 informs the reader how this study has been conducted; detailing the method, participants, materials, design and procedure. The results will be presented in Chapter 4 and will include the statistical tests and graphs showing the findings of the experiments. Chapter 5 then provides a discussion of the findings from the results of the study and what the findings suggest. Chapter 5 will also include the implications this study has to the field. Finally, the main conclusions from the discussion will be outlined in Chapter 6.

CHAPTER 2. BACKGROUND

This chapter introduces previous research of goals (free vs. focused) and roles (control vs. predict). Then it brings them together to form the aims and hypotheses of this study.

2.1. Free vs. Focused

When people are faced with a program they are unfamiliar with, they have shown a tendency to prefer to learn by exploration and avoid looking at the instructions unless they get stuck. Learning by exploration is essentially where people use a system and build up their own knowledge of the program by using trial and error without the benefit of instructions. Rieman (1996) stated that learning performance is best when people learn to complete tasks in the same context they will be using the tasks in practice. But this often means that their knowledge will be limited to particular system tasks they have learned rather than a concept of the system as a whole.

Burns & Vollmeyer (2002) compared free and focused goals when transferring knowledge and in problem solving skills. They suggest that free goals (where the goal is to find out what the underlying rules of a system are) lead to better transfer of knowledge and problem solving than focused goals (where the goal is to complete certain task or keep an output at a certain level). Many others (Geddes & Stevenson, 1997; Miller, Lehman, & Koedinger, 1999; Osman, 2008; Sweller, 1988;

Vollmeyer, Burns, & Holyoak, 1996) have shown similar results; that is, free goals lead to better understanding of a task when participants are learning by exploration in an attempt to discover the underlying rules. They suggested the reason behind this was that the free groups would engage in hypotheses-testing behaviour whereas the focused goals would show a more goal-oriented approach.

However, other studies have not shown the same results. Shrager & Klahr (1986) found that their participants, who were learning by exploration with free goals, had a tendency to only perform tasks that supported their hypotheses. This resulted in a lack of knowledge about the functionality of the device and caused them to form hypotheses on inadequate information. These studies have found that participants did not set themselves goals that were good enough to be able to transfer their knowledge, i.e. they did not test their hypotheses well. Payne and Howes (1992) discuss other limitations of exploratory learning; they highlight these limitations as “exploration traps”. They suggest in some cases, people accomplish the goal they intended to complete but then forget how they completed it. Or, people discover a non-ideal method of reaching a goal and then continue to use this imperfect method.

Studies (such as, Ascher, 1984; Lan, Bradley & Parr, 1993; Lavoie & Good, 1988; Pintrich & deGroot, 1990; Pressley & Ghatala, 1990; Schunk, 1983; Schunk & Rice, 1987) suggest there are limitations to exploratory learning and go on to identify reflection and self-monitoring as a key component that facilitates the necessary hypothesis testing to perform well on the tasks. Taatgen (1999) builds on this by suggesting that reflection causes people to deliberately acquire knowledge.

Trudel and Payne (1995) investigated exploratory learning through goal management and reflection. Using three different groups of participants, they tested performance whilst the participants were learning to complete tasks on a digital watch. Their findings from the study suggest that the group who performed the best was one that had a keystroke limitation, where the number of times a participant could press a key to perform a specific function was limited. They concluded this was because the keystroke limitation enforced the participants to reflect more fully on their actions because the actions were limited. These limits lead to a positive effect on the participants' acquired learning. Trudel and Payne concluded that participants in this study without the keystroke limitation, had "overacted", meaning they had made too many moves and not considered each move sufficiently enough to learn to use the system as well as those who had reflected. When people have too many options they tend to vary too many things at once and end up with no idea which action had what effect and which had no effect.

Cox and Young (2000) explored this further; they proposed that participants who learn by free exploration tend to develop device-oriented knowledge. In this case, the device-specific knowledge is how the system works as a whole and the task-specific knowledge is how to complete specific tasks given. They found that the device-specific knowledge did not necessarily lead to people being able to successfully complete task-specific interactions with the system. Cox and Young's (2000) study was based on that of Trudel and Payne (1995) but their study did not obtain the same results. Instead, Cox and Young found participants in free exploration groups showed better performance on scenario-based tasks and on a questionnaire about what the device did. They concluded that the participants concentrated on gaining device-orientated knowledge, rather than domain-oriented knowledge. Those who had gained more device knowledge, tended to perform better on both domain- and device-oriented tasks. However, in contrast to Trudel and Payne's (1995) study, they did not find that the keystroke limit caused their participants to reflect on their actions. They concluded this might have been because the device used in their study was less modal than that of the Trudel and Payne study; this might have caused their participants to believe the device was simple making them overconfident before the learning phase was complete. Cox and Young also proposed that the keystroke imposition caused the participants to change their strategy, which could be to reflect more during the learning phase. A later study by Cox and Young (2003) focused more on the type of device and what was learned. They found that putting a limit on interactions during exploratory learning had a positive effect when the device was a device-oriented one (e.g. a digital watch or a computer game) rather than a domain-oriented device (e.g. a telephone or pocket calculator).

A similar outcome suggested the idea of the cost of performing an operation on the system and how the cost affected people's decisions. O'Hara & Payne (1998) argued that when the cost was higher, in terms of more effort involved, the participants were encouraged to reflect and plan more before making each interaction with the system. The outcome of their study was that the participants performed better when the cost was higher, and thus they were planning and reflecting. In support of this, Sweller (1988) proposed a cognitive load theory and suggested that the cognitive load was greater when learning to control a system with a focused goal rather than a free goal in mind. Sweller went on to suggest this higher cognitive load leads to poorer learning and understanding of a system. Sweller argued that in comparison with focused goals, the free goals allowed more room in working memory. This allows people to group system states together and develop schema of the tasks to build up their conceptual view of the system and how it works.

Based on this research about the effects of exploratory learning, it is apparent that free learning can be an effective way to gain in-depth knowledge of a device, but this is often only the case if people set themselves good goals and reflect on their actions.

2.2. **Controllers vs. Predictors**

The role a person is instructed to take when learning a novel system can have an impact in what they learn. Studies (Osman, 2008b & 2008c; Osman &

Speekenbrink, under review) have shown that people who are observing the actions of others can gain the same knowledge of a system as the active person who is controlling it. When driving a car and learning a new route, does the knowledge gained of the new route depend on whether the person is in the driver's seat or the passenger's seat? It could be that the driver is paying attention to the traffic rather than the new route they are taking, or the passenger could be aimlessly looking out the window with no reason to pay attention to where they are going. It is possible for either of these people to obtain better knowledge of the route in question; but for this to happen, at least a person needs to have a plan or a goal. When dealing with dynamic situations, people often build up their decisions based on interdependent decisions they have made previously in order to achieve their goals (Brehmer, 1992). Funke (2001) argues that most decisions people make are made based on previous decisions or plans. Dynamic situations mean that people have to make decisions on the fly which might not have the same outcomes as they have experienced in the past. It is also device-specific: a goal-oriented person learning to drive a car presumably performs better than the passenger watching the driver's action and reflecting on the outcomes than they would sitting working on a simulation, but the observer would learn more by sitting in the driver's seat and performing the actions. On the other hand a goal-oriented person learning to operate a driving simulator would presumably learn equally well observing and reflecting and predicting, as they would by operating the controls.

Osman and Speekenbrink (under review) used complex dynamic control (CDC) systems to investigate which people will perform better, someone actively controlling a system or someone who is predicting what will happen next based on the actions of others whilst learning to use a system. CDC tasks often involve a person interacting with various cues to change an outcome. For example, there might be a hormone cue (hormone A, hormone B, and hormone C) all of which have an effect on an outcome (athlete's performance level) depending what levels they are set to. The person interacting with the cues can control the level of the cue value to change the outcome. The system is dynamic and the interaction could be positive or negative depending in what was entered. The effects of the outcome build up gradually in successive actions to get closer to the goal.

CDC systems are used to transfer skill-based knowledge in a real-world environment. Funke (2001) suggested they are a good way to study problem solving behaviour in the real world because they can simulate simplified versions of real-world environments while still allowing them to be controllable, easily traceable and easy to analyse. For example, the CDC system used for this study is an incubator where air pressure, oxygen and humidity are cues that can be manipulated to control a baby's heart rate. Several studies (Berry & Broadbent, 1984, 1987, 1988; Broadbent, 1977; Broadbent & Ashton, 1978; Broadbent, Fitzgerald, & Broadbent, 1986) have successfully demonstrated the use of CDC systems in research where the participants have learned to perform task-specific tests.

In Osman and Speekenbrink's study (under review) they teamed the CDC tasks with multiple cue probability learning (MPCL) tasks to consider the effects of control and prediction on learning. They are similar to CDC tasks except the people interacting with the MCPL tasks are not interacting with the cues. Instead, MCPL tasks usually involve people learning the possible relationship of cues on an outcome. Taking the example above, the person predicts the performance level of the athlete based on various levels of hormones A, B and C. After each prediction of the outcome, feedback is given that shows the actual outcome.

It was proposed by Lagnado, Newell, Kahan & Shanks (2006) that an important cognitive processing skill is being able to accurately predict the future based on past experiences, and though it is virtually impossible to make all predictions correct on every occasion, learning from past experiences can be an effective way to learn to predict what is likely to happen. Lagnado et al (2006) have shown that with the introduction of prediction, the observers are not only watching what is happening; they are putting thought into what they think the outcome will be based on past trials. This self-insight about what they are doing causes them to develop strategies and in turn helps them to make accurate predictions about complex environments.

Previous research (Berry, 1991; Berry & Broadbent, 1988) has suggested that in order to learn relevant knowledge of a system, people need to be actively controlling it. These studies suggest that mere observation is too passive to acquire the same knowledge as an active controller. A study by Enkvist, Newell, Juslin & Olsson (2006) using a MCPL task found that control learning was superior to predictive learning. Enkvist et al (2006) claimed that in order to acquire knowledge about the tasks, it was critical to be directly controlling the environment. By contrast, several studies (Osman 2008b, 2008c; Osman & Speekenbrink, under review) do not show the same results. These studies added the element of predicting the outcome to a CDC task where both the controller of a system and a predictor's trials are identical so they can be compared later. Osman found that the person who was observing and predicting the outcome performed as well as or in some instances marginally better than the person controlling the system. Osman suggests this might be because the studies encouraged hypothesis testing where as Berry (1991) and Berry & Broadbent (1988) discouraged hypothesis testing.

Osman and Speekenbrink (under review) furthered the work of others (Sweller's, 1988 "Cognitive Load Theory" and Burns and Vollmeyer's, 2002 "Dual Space Hypothesis"), and argue that controllers and predictors learn through interacting with a system by their own evaluation of the outcome. The main difference between them in this case is the way they evaluate the outcome. Osman (2010) proposed a Monitoring and Control Framework (MC Framework) that states the differences between the evaluations of controllers and predictors. Controllers evaluate the outcome by the difference between the achieved outcome and the target outcome, while predictors evaluate the outcome by determining the relationship between the outcome they predicted and the actual outcome. This suggests that if controllers and predictors both are learning by evaluation the outcome they are therefore learning by the same underlying process and this should allow them to acquire the same knowledge.

Based on this research about the effects of the role of the participant, it is believed that observing and predicting the outcome as a basis for learning to use a system can be as effective as controlling the system.

2.3. **Aims and Implications**

This study was based on work completed by Osman and Speekenbrink (under review). Part of the study consisted of investigating the effect of noise levels when learning to use a system. Osman found that with low noise, controllers and predictors were able to acquire the knowledge to control the system correctly in about 10 trials. On the other hand, when the environment was slightly more noisy and unstable, this greatly increased the difficulty of learning for controllers and predictors to control the system effectively. Osman suggested this was because the participants in the noisier condition needed a longer period of familiarisation.

Based on Osman's study, when the noise was set to a low level, participants learned to operate the system effectively after only having made changes to the cue values 10 times. The learning phase consisted of 40 trials and the participants learned to operate the system after about 10 trials. This is believed to lead to the tests of this study being too easy to correctly complete because participants would figure out the interactions of the cues well before the tests were given. Increasing the noise in the environment will obscure the relationship between cues and outcomes, because when people manipulate the cues it is much less obvious how they affect the outcome. The extension of this study will set the noise to a high level. This means the variation in the outcome could change a great deal even when participants do not change the cue levels, so controllers and predictors will have to be very systematic and pay careful attention to be able to learn the relationships between the cues and the outcome. They will have to process statistical information mentally, for example, '4 times out of 5, increasing cue A had outcome X', requiring a high commitment to goals, as opposed to knowing the cue-outcome relationship more immediately in a low noise environment.

The aims of this study are to find out what strategies people adopt when they are learning to use a novel interactive system. If they are allowed to explore, does it lead to them having a better understanding of the program? Does it allow them to apply their knowledge to control the system more efficiently than those who are asked to learn by applying specific focused goals?

The literature for comparison of free and focused goals implies that free goals enable people to build device-oriented knowledge and focused goals build tasks-specific knowledge. The people who learn with free goals are more able to transfer their device-oriented knowledge to the task-specific tasks within the system because they have gained knowledge of the underlying concepts of the system.

But, this is only the case if they engage in reflection or hypothesis-testing behaviour.

The literature for comparison of controller and predictor roles has been less explored; some of the research implies that a person must actively participate in order to learn by engaging in procedures within the system, whereas others say that people are able to learn to use a system by observing so long as they are simulating hypothesis-testing behaviour. This research is also system-specific.

This study combines the goals (free vs. focused) and roles (controller vs. predictor) attempting to understand what happens when they are considered simultaneously. The type of learning will differ between participants but their tests, after learning to use the system, will be identical so as to compare the impact of the learning.

This research aims to contribute to the body of academic knowledge in HCI, and also explore useful implications for HCI designers. The importance of understanding, with a view to enhance, the way people learn to control a system effectively can help designers to build better systems that support peoples' needs. This will be discussed further in Chapter 5.

2.4. **Hypotheses**

The literature on distinctions between prediction and control suggests that predictors are able to perform the same or better than the controllers. On the other hand, the literature on distinctions between free and focused suggests that free groups perform better than the focused groups on tests determining the underlying workings of a system and sometimes on task-specific learning when they set themselves good goals while learning to use a device.

In the free-predict group, there is no cost or consequence to the actions, but they are expected to engage in hypotheses testing behaviour. Based on previous research, this suggests they might be the best at performing tests of device-oriented knowledge. The literature suggests they might perform the worst with test of task-specific knowledge. However, the literature also suggests that they may be able to transfer their device-oriented knowledge to the task-specific tests if they are able to set themselves goals that support hypothesis-testing behaviour.

The focused-controllers will have a cost to their actions but they will not be expected to engage in hypotheses testing behaviour. Based on research, the expectation is that this group will perform best at task-specific knowledge but they will not be able to transfer this knowledge to device-oriented tests.

The full complement of this study is an exploratory one, specifically testing to see whether there is an interaction between all groups together, and if so which group performs the best.

CHAPTER 3. METHOD

This study will compare both the type of exploration (free or focused goals) and the role of the participant (whether they are an active controller of the system or if they were observing and then predicting the outcome) when learning to use an interactive device. The knowledge gained from a learning phase will be tested by comparing how quickly the different groups of participants can gain understanding of the device, by comparing how far away the participants were from the target levels.

The aim is to investigate what happens when free and focused goals are compared directly alongside controller and predictor roles.

3.1. **Participants**

Forty people participated in this study for a £5 incentive. The participants were current students of UCL (graduate and undergraduates) with an age range of 19- to 46-years-old. The participants' genders were almost evenly split 21 females and 19 males. They were selected on the basis that they were able to participate within the time frame for this study and were recruited from the Psychology pool or fellow classmates.

3.2. Design

The study used a between-subjects yoked design which consisted of two between-subject independent variables contrasting exploratory (free) learning and goal orientated (focused) learning, and comparing process control learning and predictive learning (controllers and predictors). There were four groups in total – free-controllers, free-predictors, focused-controllers, and focused-predictors.

	Free	Focused
Controllers	Group 1	Group 2
Predictors	Group 3	Group 4

The dependent variable was percentage-correct at test. This was measured at four points during the experiment: two measures of participants' ability to control the system (control test 1 and control test 2) and two measures of the participants' ability to predict the performance of the system (predict test 1 and predict test 2).

3.3. Materials

Participants interacted with a control task during both the learning and test phases of the experiment. The control task was presented on a 15-inch screen which was attached to a standard personal computer and used a software program written in C#. The screen layout, main instructions and cover-story were the same for both predictors and controllers. Participants were informed that an incubator has been developed for newborn babies with irregular heart rates. This incubator enables medical staff to monitor the heart rate and adjust 3 variables (cues) – air pressure, oxygen and humidity to attempt to maintain a steady heart rate.

The interface of the task displayed 3 cues (air pressure, oxygen and humidity) which the participant could manipulate to produce different outcomes: one positively affected the outcome, one negatively affected the outcome and the third was unrelated to the outcome (null cue). The 3 cues were varied randomly for each participant. For example, the cues could be as follows - positive cue (air pressure) negative cue (oxygen) and null cue (humidity).

For participants in the control groups, and for the control tests, the screen displayed three cue values with a sliding scale ranging from 0-100 (shown in figure 1). All cue values were set to 0 at the start. The first outcome value was set to 178. The target value throughout the studies was set to 62. Participants had a visual representation of the values on a graph. The graph showed the heart rate on the y axis and the trial number was displayed on the x axis. As trials were completed, the previous values were displayed on the graph; it also displayed a green line encompassing values 52-72 (target value +/- 10) which were termed the 'safe' values. (The green line encompassing the 'safe' values is shown in figures 1 and 2).

For each trial, participants could change any, all or none of the cue values and then click 'submit'. For participants in the predictor groups, the screen layout was almost identical. It displayed the same 3 cues at the same point in the system as their yoked controller had input. The predictor was able to see the cue values entered by the yoked controller and predicted where they believed the outcome would be after the values had been submitted. The predictor would do this by moving the sliding scale next to the outcome value and click 'submit'. Figure 2 shows an example screenshot of the display when a predictor has clicked submit. The dotted line shows the predicted estimate of the outcome and the solid line shows the actual outcome. Once they were ready to move on to the next trial, they clicked 'continue'.

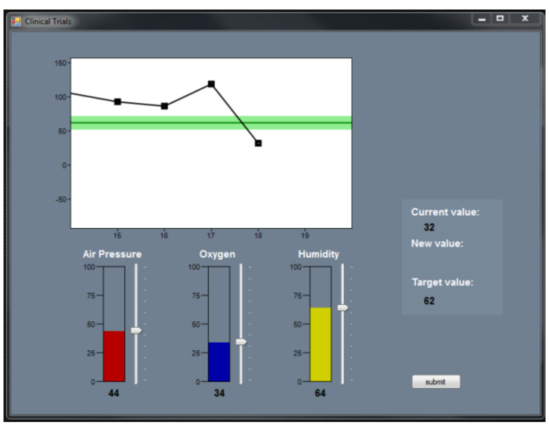

Figure 1: Screenshot during the learning phase for a controller.

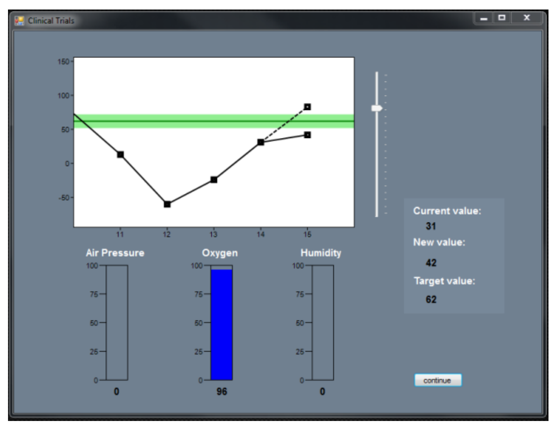

Figure 2: Screenshot during the learning phase for a predictor.

The structure of the system is as follows:

$y(t) = y(t-1) + b_1 x_1(t) + b_2 x_2(t) + e_t$

Positive cue = $b_1 = 0.65$

Negative cue = $b_2 = -0.65$

Random noise[3] = e_t

Outcome value = $y(t)$

Previous outcome value = $y(t-1)$

e_t = random noise component, with a mean value of '0', it was normally distributed with a standard deviation of 32.

The raw data was recorded by the program (i.e. starting values, predicted cue input levels, predicted outcome levels, and actual outcome levels).

3.4. **Procedure**

Each participant was asked to complete a task which consisted of 112 trials. The trials were divided into two phases – a learning phase and a testing phase. All participants completed a learning phase of 40 trials, then a control test of 20 trials, a prediction test of 16 trials, a further control test of 20 trials and finally a prediction test with 16 trials.

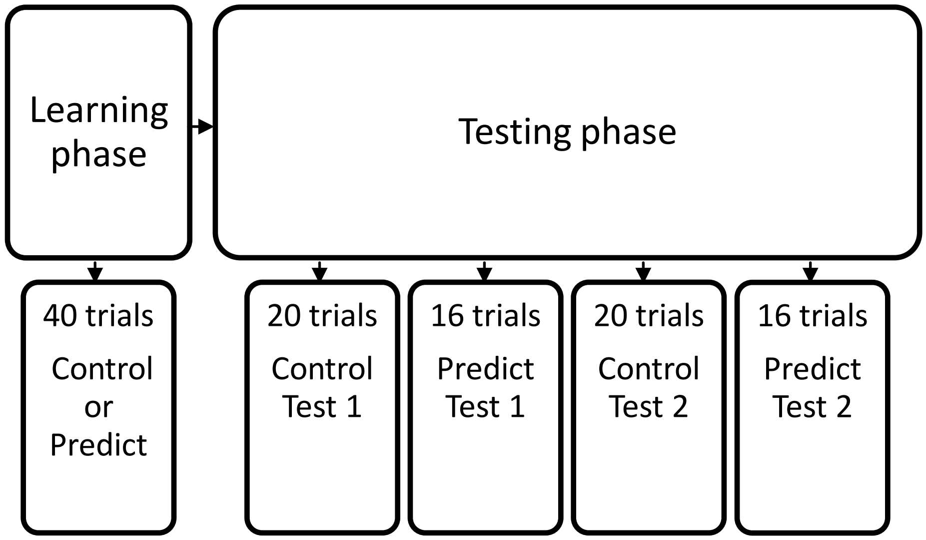

Figure 3: Learning and testing phases showing the number of trials in each phase and the order in which they are experienced.

All participants were given the same generic instructions to introduce the task as follows:

"A recently patented Baby Incubator (Patent:: ISO 9012:2008, ISO 13485:2007) has been developed specifically for newly born babies that have an irregular heart rate. The Incubator is a clean, secure, and regulated environment which enables the medical staff to monitor a baby's heart rate, and adjust its conditions to maintain steady heart rates.

The conditions of the Incubator can be changed in three ways, by modifying air pressure the amount of oxygen released, and the humidity, all of which can affect heart rate activity in different ways either by increasing or decreasing it."

All participants were asked to imagine they were a trainee nurse in a hospital who will to be performing manipulations on three different cues (air pressure, oxygen and humidity) in an attempt to change the outcome (heart rate) of incubator. Thereafter, the instructions were slightly adjusted from the same template to direct participants to learn to either control or predict the behaviour of the system in either a focused or exploratory situation. Complete copies of the instruction scripts for all groups can be found in Appendix A. Here we provide a summary of the differences in instructions provided to each of the groups:

Free-controllers: control a simulation incubator, familiarise themselves with the system, and try to find out what effect the cues have.

Free-predictors: monitor the outcomes of the simulator and predict the outcome based on the values shown on the screen.

Focused-controllers: control an incubator with a baby where the heart rate should be kept as close to the 'safe' levels as possible.

Focused-predictors: monitor the heart rate and predict the outcome of the heart rate based on the values shown on the screen.

Learning phase:

All controllers began with the learning phase in which they undertook 40 trials of controlling the system.

All predictors were asked to predict the outcome based on the levels displayed on the screen. The levels displayed were taken from the respective yoked controller's 40 trial learning phase.

Testing Phase:

All participants had an identical testing phase. Their learning history was the only thing that differed among the participants. The testing phase consisted of 20 control test trials, 16 prediction test trials, 20 control test trials and 16 prediction test trials.

Control tests were designed to tests the task-specific knowledge of the environment.

The prediction tests were designed as structured tests that investigated device-oriented knowledge of the environment. The participants were either required to predict a cue value based on the values of the other cue values and the outcome presented to them (for example, predicting the negative cue value based on the displayed values of the positive, null and outcome values) or they were required to predict the effect of the outcome based on the values of the three presented cue values.

Feedback was not given for either prediction test. The prediction tests consisted of a mixture of eight old tests that participants had seen during the learning phase and eight tests that were randomly generated that they had not previously experienced. Participants were not told that any of the trials were ones they had already seen. For controllers, these trials were randomly selected ones which they had generated themselves in the learning trials. For predictors, these previously seen trials were the same as those that they had previously predicted during the learning phase. Old and new trials followed the same structure, 2x positive cue value, 2x negative cue value, 2x null cue value, 2x outcome value. The presentation of the prediction test trial was randomised. The predictive value was recorded for each trial.

Scoring = error scores

The error scores for predictors are calculated by subtracting their prediction of the outcome, with the actual outcome for that trial. The error scores for controllers are calculated by subtracting their achieved outcome for a trial from the target value they were designed to achieve.

CHAPTER 4. RESULTS

This section contains a comparison of the error scores for the control tests and the predict tests. The control tests 1 and 2 have been compared together. The predict tests have been separated to compare the error scores for old trials and new trials and consider the prediction scores for the cues – positive, negative and outcome.

Comparison of the learning phase is not a fair representation of the results as the free and focused groups had different aims and therefore the error scores for the learning phase did not need to be considered.

4.1. Summary of control test phase

The control test phases contain 20 trials each. The scores for the control test phase were separated into two blocks, each containing 10 trials each. The error scores were averaged across each block for all participants.

Control error scores for test 1 and 2

Figure 4 and 5 (below) show the control error scores for controllers and predictors in the free and focused groups for test 1 and 2 respectively. The ANOVA was conducted on control error scores and showed there were no significant main effects. The analysis did reveal a 4-way interaction between test (control test 1, control test 2) x block (block 1, block 2) x role of participant (predict, control) x type of exploration (free, focused), $F(1, 36) = 4.72$, $p < .05$. In order to identify the location of the interaction, further tests were conducted. These revealed that, in control test 1, for those participants in the free condition, control performance was stable, but for the focused conditions controllers improved across block 1 and 2 $t(9) = 1.84$, $p = .09$. No other analyses were significant.

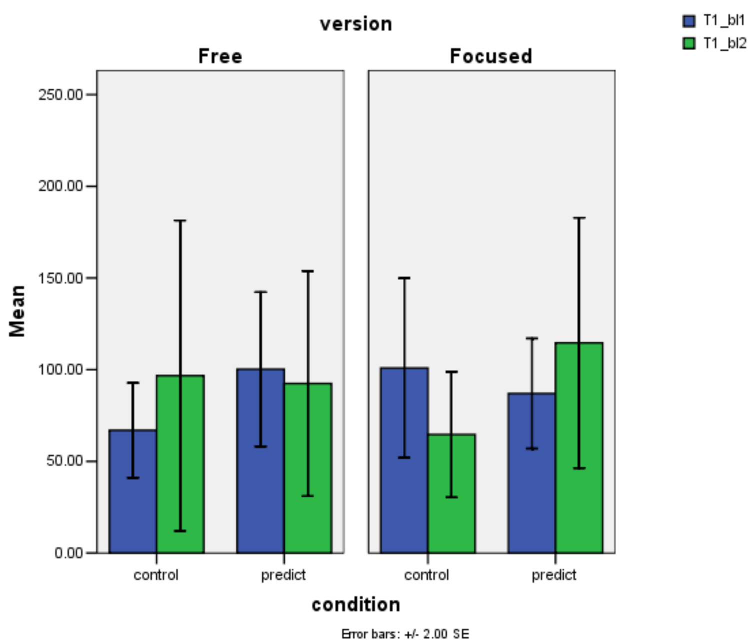

Figure 4: Control test 1 - control error scores of the control vs. predict groups by block of 10 for the free and focused conditions during control test 1.

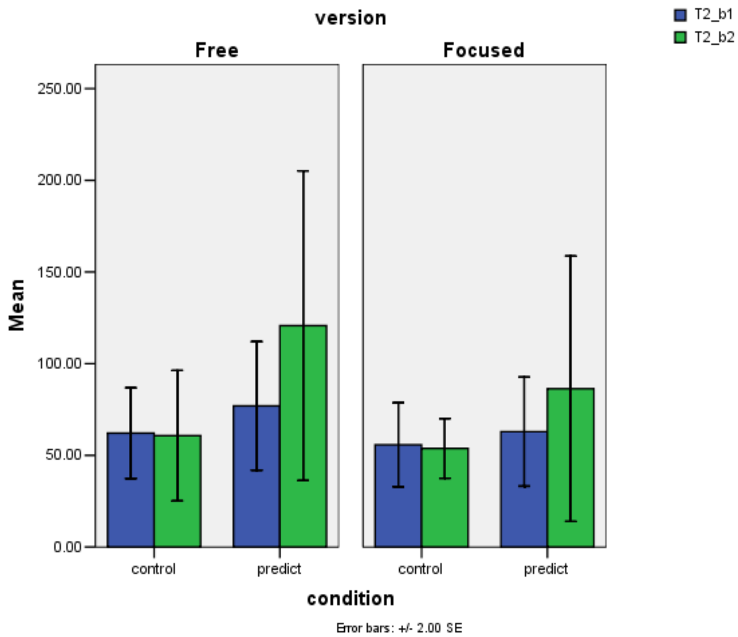

Figure 5: Control test 2 - control error scores of the control vs. predict groups by block of 10 for the free and focused conditions.

4.2. Summary of predictive test phase

Predictive tests contain 16 trials each so the scores were separated into two blocks containing 8 trials each. The error scores were averaged across each block for all participants. First the error scores for both tests will be considered.

This section compares the error scores for the old and new prediction tests in predictive tests 1 and 2. The old tests are ones that the participants had encountered previously from their learning phase and the new tests are ones that were randomly generated. Old and new tests were identical for yoked pairs.

Test 1

Figures 6 and 7 (below) show the prediction error scores for the old (figure 6) and new (figure 7) prediction tests. The error scores were calculated by using a repeated measures ANOVA which included the following within subject factors: test (old, new), cue (positive, negative, outcome), the between factors: role of the participant (control, predict) and type of exploration (free, focused). Comparisons for error scores shows there was a cue (positive, negative, outcome) x test (old, new) interaction, $F(2, 72) = 19.82$, $p < .005$. Further analyses revealed that predictions for the positive cue ($t (39) = 4.39$, $p < .005$), the negative cue ($t (39) = 2.47$, $p < .05$) were significantly more accurate for old trials compared to new trials. But for the outcome ($t (39) = 3.37$, $p < .005$) overall people were more accurate for new trials compared to old trials. No other main effects were significant.

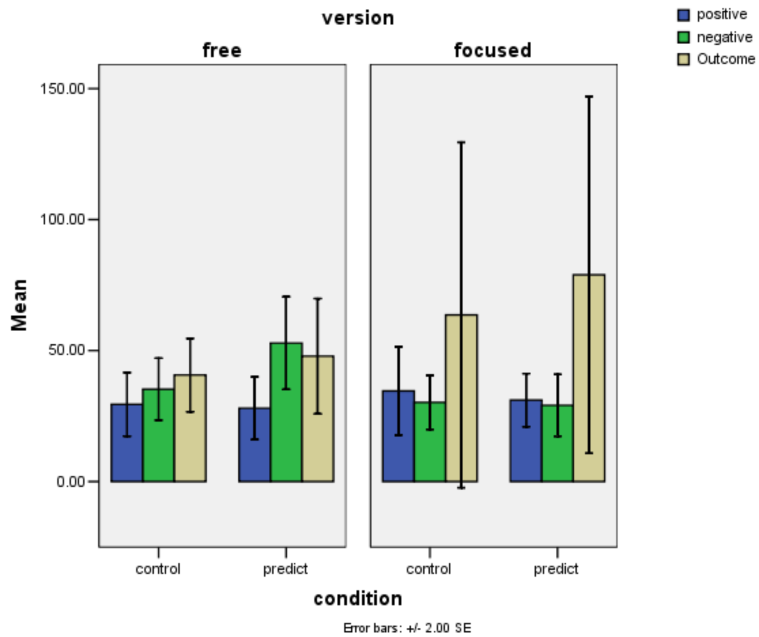

Figure 6: Predictive test 1 (Old trials). Prediction scores of the control vs. predict groups by cue for the free and focused conditions during old predictive trials.

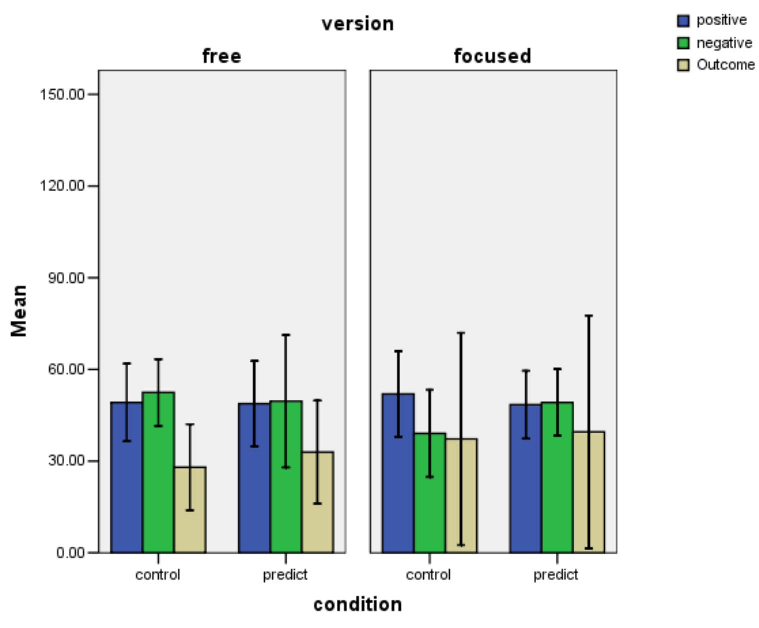

Figure 7: Predictive test 1 (New trials). Prediction scores of the control vs. predict groups by cue for the free and focused conditions during new predictive trials.

Test 2

Figures 8 and 9 (below) show the prediction error scores for the old (figure 8) and new (figure 9) prediction tests. The error scores were calculated by using a repeated measures ANOVA which included the following within subject factors: test (old, new), cue (positive, negative, outcome), the between factors: role of the participant (control, predict) and type of exploration (free, focused). Comparison of prediction error scores shows there is a main effects of cue $F(2, 74) = 13.62$, $p < .0005$. This suggests that overall predictions for the positive cue were more accurate than the outcome ($t(39) = 3.83$, $p < .005$), and similarly, predictions for the negative cue were more accurate than the outcome ($t(39) = 4.27$, $p < .005$). No further analyses were significant.

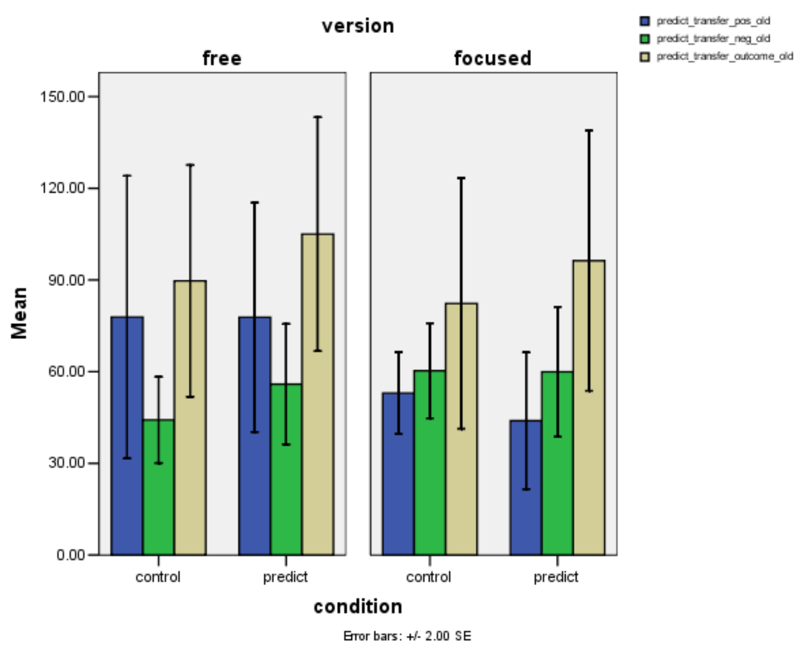

Figure 8: Predictive test 2 (Old trials). Prediction scores of the control vs. predict groups by cue for the free and focused conditions during old predictive trials.

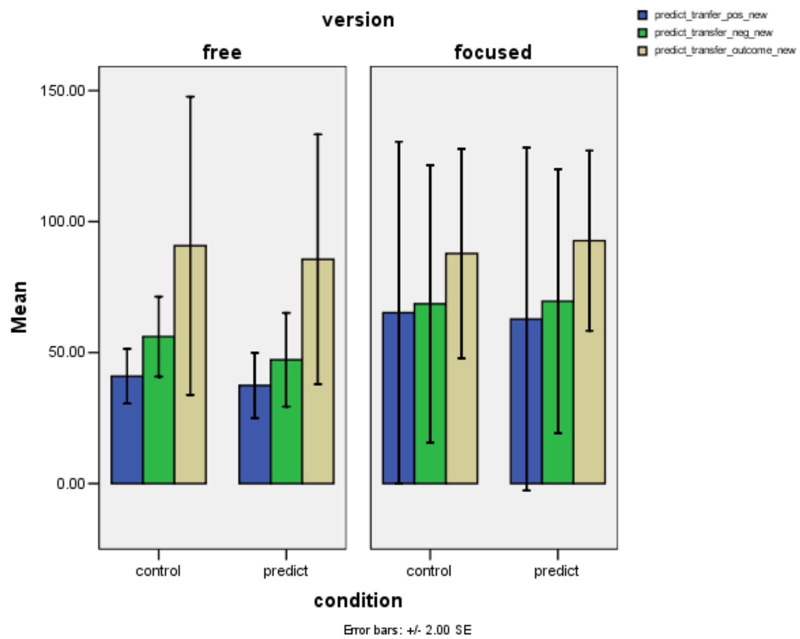

Figure 9: Predictive test 2 (New trials). Prediction scores of the control vs. predict groups by cue for the free and focused conditions during new predictive trials.

Summary of results

Control test phase 1 and 2: The results showed a 4 way interaction between the conditions. After the learning phase, all participants were instructed to control the system. At first, the error scores were comparable but after the first block, the participants in focused groups performed better.

Predictive test 1 – old and new: The results showed a cue interaction and a test interaction. The positive cue and negative cue were significantly more accurate for old trials than new trials. However, the predictions for the outcome were significantly more accurate for the new tests than the old tests.

Predictive test 2 – old and new: The results showed a cue interaction. Positive cue and negative cue were more accurate than the outcome for both old and new trials.

CHAPTER 5. GENERAL DISCUSSION

The aim of this study was to explore how the type of learning impacted what was learned and how this knowledge was applied to the tests. The acquisition of knowledge was a factor that was varied amongst the conditions but it was not a factor that was explicitly analysed in the results; the focus of the study is how the conditions applied to tests. This chapter discusses the findings from the control tests and predictive tests in relation to the literature in Chapter 2. Next it suggests implications of the research and finally, limitations and directions for further research.

5.1. Expected findings

It was expected the results would show that free participants would have gained better knowledge of the underlying concepts of the system and therefore they would perform better on the predictive tests (where they would be transferring their knowledge) than focused participants. However, the focused participants were expected to perform better on the control tests as the design for the control tests were the same as the way they learned to control the system. Participants in the predictor groups were expected to perform equally as well as the controllers at being able to transfer knowledge of the underlying concepts of the system to the predictive tests and in the control tests.

More specifically, the focused-controllers were expected to perform the best at the control tests as they were focusing on task-specific tests with a focused goal. The free-predictors were expected to perform the best at predictive tests as the free goals would allow them to obtain device-oriented knowledge of the underlying concepts of the system by exploring and testing hypotheses. They would also be reflecting on their actions to build up device-oriented knowledge.

It was not known how the full complement of groups would perform when tested all together.

5.2. **Findings**

The most interesting results of this study were uncovered in the first control test: the results showed a 4-way interaction between the conditions. Immediately after the learning phase, all participants were instructed to complete the same task of controlling the system for 20 trials. The initial control test was the first opportunity for the free-predictors and focused-predictors to control the system. For the free-predictors and free-controllers, this was the first time they were instructed to use the cues to try to reach a specific outcome. The focused-controllers' instructions for this test were the same as their learning phase. They were the only group whose goal was to control the system in this way from the beginning. Both of the control tests were task-specific tests of the system which determine whether participants had gained task-specific knowledge. The first block of trials showed the error scores for all four groups were stable. The second block showed that the performance of the participants in the focused-controller groups improved and thus they performed better than all other groups. In line with studies comparing controllers and predictors in a similar way (Osman 2008b & 2008c; Osman and Speekenbrink, under review), it was expected that focused-controllers' and focused-predictors' performances would be the same. However, there are a couple of possible reasons for the difference here. The focused-controllers were the most used to controlling the system in exactly this way; they had been focused on the goal of reaching and maintaining the set outcome (i.e. the green line) for the entire time. This is likely to have been the reason they performed the best in this test. The focused-predictors on the other hand, might have become overconfident in their knowledge after the first block of the control test and then unable to match the focused-controllers during the second block. The focused-predictors did perform as well as the free-controllers and free-predictors in the second block. These results support CDC studies that found focused goals successfully allowed participants to perform task specific tests of a system (Berry & Broadbent, 1984, 1987, 1988; Broadbent, 1977; Broadbent & Ashton, 1978; Broadbent, Fitzgerald, & Broadbent, 1986).

By the time the participants reached the second control test, their error scores were stable across the four groups showing no significant differences between them. This showed that the participants had successfully learned to control the system to the same degree across the four groups. This suggests that the focused-predictors might have needed slightly more trials to be able to match the performance of the focused-controllers when completing task-based control tests.

Based on previous research (Shrager & Klahr, 1986; Cox & Young, 2000), the free-predictors might have been the worst at this test because they had not had experience of task-specific tests where they were completing the learning phase, and they were also relying on their ability to form and test their own hypotheses with no cost to their interactions. Although there was no significant difference between the groups, there was a trend that the free-predictor group performed slightly worse than the other groups in the second control test. This could be because their experience of controlling the system was more limited or they had a lack of task-specific goals to reach (i.e. reach and maintain a set level for the outcome) which had caused their performance level to be slightly worse when controlling the system. Their knowledge of the underlying concepts of the system at this point was also not as great as the other groups to be able to transfer their knowledge to task-specific tests. This finding is in line with the findings of Shrager & Klahr (1986) and Cox & Young (2000) where the free-predictor group may have set themselves inadequate hypotheses that did not allow them to acquire knowledge at the same level as the other groups. In saying that, this is only a trend and not a statistically significant result.

Overall, the results also support the research surrounding the evaluation of the outcome (Osman 2010; Osman & Speekenbrink, under review). In this case the controllers are evaluating the outcome by the difference between the achieved outcome and the target goal that they were instructed to reach and maintain. The predictors were evaluating the outcome by determining the relationship between the predicted outcome and the actual outcome. These results show that the participants in all groups have displayed an ability to evaluate the outcome effectively in order to gain task-specific knowledge and display the same level of performance.

The results for the first predictive test showed an interaction between the positive, negative and outcome cues and between the two tests. The positive cue and negative cue were significantly more accurate for old trials than new trials. The prediction of the cue values specifically tested the participants' ability to transfer device-oriented conceptual knowledge of the system because this type of interaction with the system had not been attempted by any of the participants previously; they had only been able to change the cue levels, not attempt to predict the levels based on the outcome or other cue values. The superior level of accuracy for the positive and negative cues in old trials might have simply been the participants recognising the old trials' values from the learning phase. They might not have been as able to transfer their knowledge to device-oriented tests at this stage as they are at recognising previous actions in predicting the outcome. The prediction of the outcome value was significantly more accurate for the new trials than the old trials. It can be seen in figure 6 that this is because the focused-controllers and focused-predictors were worse at predicting the outcome for the new trials. The free-controllers and free-predictors were stable in both old and new trials. This suggests the free-controllers and free-predictors were better at transferring their knowledge to device-oriented tasks. These results are in line with previous literature that free goals are better than focused goals for gaining knowledge that can be transferred to device-oriented tasks (Burns & Vollmeyer (2002); Geddes & Stevenson, 1997; Miller, Lehman, & Koedinger, 1999; Sweller, 1988; Vollmeyer, Burns, & Holyoak, 1996).

The results for the second predictive test showed that the positive cue and negative cue were more accurate than the outcome for both old and new tests. This suggests the participants gained knowledge of how the system works and the underlying concepts rather than merely being able to accurately predict the outcome. There were no significant differences between the groups when comparing their overall scores against each other. This shows that by the end of the second predictive test, the controllers and predictors were displaying the same level of accuracy across all four conditions. Although there were no significant differences between the groups, figure 9 shows there was a trend that free-controllers and free-predictors performed to a higher degree of accuracy than the focused-controllers and focused-predictors for the new trials for the positive and negative cue values. This suggests their transfer skills were marginally better than the focused-controllers and focused-predictors and shows that the free-predictors were able to acquire the same level of transfer skills as the free-controllers.

The overall results for the second predictive test were worse than the first predictive test. There are a couple of suggested reasons for this difference; it is possible the wild fluctuations of the noise levels were too high to accurately predict the outcome for the second test. The noise levels made it impossible to gain completely accurate results. This was designed to encourage participants to pay more attention and reflect more to gain knowledge of the underlying rules of the system but it could have also had a negative effect. Participants might have realised that it is impossible to be give accurate predictions and may have gotten frustrated or bored and stopped trying hard to make a good prediction.

By the end of the study, all participants showed the same knowledge of the system in terms of control skills (task-oriented knowledge) and in terms of transfer (device-oriented knowledge). Osman and Speekenbrink (under review) found that the predictors had performed better than controllers at gaining device-oriented knowledge of the system however; this finding was not significantly replicated. This study found trends of both free-predictors and focused-predictors being able to transfer their device-oriented knowledge more accurately than the free-controllers and focused-controllers but these results were not statistically significant. This may have been because the noise levels were too high and fluctuating too wildly for all participants to be able to get a solid understanding of the system.

The final control test and predict test show that the findings support previous research that found reflection (in this case through predicting the outcome) is an important part of acquiring device-oriented knowledge and to be able to transfer this to task-oriented tests (Schunk, 1983; Ascher, 1984; Schunk & Rice, 1987; Lavoie & Good, 1988; Pintrich & deGroot, 1990; Pressley & Ghatala, 1990; Lan, Bradley & Parr, 1993, Taatgen, 1999; Osman, 2008b & 2008c; Osman and Speekenbrink, under review).

5.3. **Implications**

As mentioned previously, enhancing learning and studying the way people learn are a vitally important part of HCI. Understanding the most effective methods for people to learn to acquire specific knowledge of systems can help lead to the development of systems that enable and support people performing at their best. The implications of this study are to inform and provide designers in HCI direction about the effective learning methods for different systems. This can be used in training for CDC systems and Effective training of systems is important when they don't want to read the instructions. The systems that are designed need to support the way people learn to enable them to enhance the interaction and make it more pleasurable to use.

Osman (2010) discusses the growing complexity of interactive systems (such as computers, hand held devices, phones, automatic driving systems, automatically and safety systems) in daily life puts more importance on research into how people learn to interact and control these dynamic environments.

5.4. **Limitations of this research and directions for further research**

After the experiment, in debriefing sessions, the participants were informally asked what their goals were and how they were trying to achieve them. Many had issues with the scenario. Almost all of the participants laughed and made a comment that they had ‘killed the baby’. The heart rate levels are not realistic as these include negative values; there is no logical definition of a negative heart rate, and no human would survive long with zero heartbeat! Other features of the scenario may have distracted participants from learning to use the system effectively, were their preconceived ideas about how the cue values affected heart rate in the real world. One participant said ‘it wouldn’t be right to not put any oxygen into the incubator.’ Another commented that they did not want to raise the humidity very much because the baby would not want to be hot and sweaty. Most participants knew what oxygen does to the human heart rate and often would at least put some oxygen into the incubator at every trial. A suggestion for future research would be to change the scenario to something else that realistically can go below 0 without causing death.

There are some interesting directions for extensions. The design of this study tested how accurately the controllers and the predictors learning skills applied to later tests. An interesting extension would be to study the difference between a predictor observing a controller who is controlling the system with a great deal of accuracy and to one observing a controller who is performing very badly.

It would be interesting to measure the outcome in a low noise condition. Osman and Speekenbrink (under review) found that the level of noise did not show a significant difference when varying the noise of the tests. However, in this study, the noise was set far higher than high noise condition in their study so it would be possible for this to show a difference if this study was replicated in a low noise condition. In high noise situations, it may be hard to tell what the associations between the cue and the outcome were. Many participants at the end of the study commented on how there was one cue value they did not understand at all and they often did not realise how much the values could fluctuate even if they performed no manipulations. None of the participants finished the study believing they had mastered the system. Lower noise could allow them to control the system more efficiently and make it less frustrating and less likely for them to give up on learning to use the system effectively.

This study consisted of only 40 participants; this means that only 10 participants were divided between the 4 groups. Given more time, a study with a greater sample size could show more statistically significant results.

CHAPTER 6. CONCLUSION

This study aimed to combine free and focused goals with controller and predictors roles in an effort to discover the effect on the knowledge acquired by the type of learning and how it applied to later tests on an interactive device. The results of this study found that by the end of the testing phase participants in all groups showed the same level of task-oriented and device-oriented knowledge with no significant differences between them. This concludes that all groups were able to perform to the same level eventually when they had enough time to learn the interactions.

The beginning of the testing phase showed the differences between the acquired knowledge of the groups. The focused controllers were the best at initially gaining the task-specific knowledge of the system. It is believed this is because they had trained specifically on the system in exactly this way. The focused-predictors were not able to match their knowledge initially but had gained the same level of knowledge after a further 10 trials. The free-predictors showed a tendency to perform slightly worse at performing task-specific tests. It is suggested that this result occurred because they had not set themselves realistic goals to perform at the same level as the other groups.

In conclusion, the free-controllers and free-predictors were best at transferring their knowledge to device-oriented tests. This suggests not only that the free goals supported the participants in gaining device-oriented knowledge, but also that the free-predictors were able to acquire the same level of knowledge as the free-controllers. This study strengthens the support for literature that concludes free goals are an effective way to gain device-oriented knowledge of a system and also provides support for that focused goals are a good method for obtaining task-specific knowledge.

This study reinforces and strengthens the support for prediction being an effective way means of learning both task-oriented and device-oriented knowledge by effectively evaluating the outcome to learn the cue-associations. The knowledge was gained by predictors more quickly when the goals were focused.

Together with providing a contribution to academic research in the HCI field, this study also demonstrates direction of interaction design in complex systems. The importance of supporting people's interactions with complex devices has been highlighted as an important criterion to consider when developing systems. Gaining greater understanding of how people learn will help provide guidance to how this can be achieved.

REFERENCES

Berry, D. (1991). The role of action in implicit learning. *Quarterly Journal of Experimental Psychology*, 43, 881-906.

Berry, D., & Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology*, 36, 209-231.

Berry, D., & Broadbent, D. E. (1987). The combination of implicit and explicit knowledge in task control. *Psychological Research*, 49, 7-15.

Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology*, 79, 251-272.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, 81, 211-241.

Broadbent, D. E., (1977). Levels, hierarchies and the locus of control. *Quarterly Journal of Experimental Psychology*, 32, 109–118.

Broadbent, D., & Ashton, B. (1978). Human control of a simulated economic system. *Ergonomics*, 78, 1035-1043.

Broadbent, D., Fitzgerald, P., & Broadbent, M. H. P. (1986). Implicit and explicit knowledge in the control of complex systems. *British Journal of Psychology*, 77, 33-50.

Burns, B. D., & Vollmeyer, R. (2002). Goal specificity effects on hypothesis testing in problem solving. *The Quarterly Journal of Experimental Psychology,* Section A, 55: 1, 241 — 261 DOI: 10.1080/02724980143000262

Churchill, E. F., & Young, R. M. (1991). Modelling representations of device knowledge in Soar. Proceedings of *AISB-91: Conference of the Society for the Study of Artificial Intelligence and Simulation of Behaviour*.

Cox, A. L., & Young, R. M. (2000). Device-Oriented and Task-Oriented Exploratory Learning of Interactive Devices. Proceedings of the *Third International Conference on Cognitive Modeling*, University of Groningen, Groningen, Netherlands.

Cox, A. L., & Young, R. M. (2003) Evidence for the Domainy-Device and Devicey-Device Distinction. *Eurocogsci 03: The European Cognitive Science*

Conference 2003, Schmalhofer, F., Young, R., Katz, G. (eds) Lawrence Erlbaum Assoc Inc, 85-90

Enkvist, T., Newell, B., Juslin, P., & Olsson, H. (2006). On the role of causal intervention in multiple-cue judgment: Positive and Negative effects on learning. *Journal of Experimental psychology Learning, Memory, & Cognition*, 32, 163-179.

Funke, J. (2001). Dynamic systems as tools for analysing human judgement. *Thinking and Reasoning*, 7(1), 69–89.

Kieras, D. E., & Bovair, S. (1984). The role of the mental model in learning to operate a device. *Cognitive Science*, 8, 255-273.

Lagnado, D.A., Newell, B.R., Kahan, S., & Shanks, D.R. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology General*, 135 (2). pp. 162-183. ISSN 00963445

Lan, W. Y., Bradley, L., & Parr, G. (1993). The effects of a self-monitoring process on college students' learning in an introductory statistics course. *Journal of Experimental Education*, 62, 26 – 40.

Lavoie, D. R. & Good, R. (1998). The nature and use of prediction skills in a biological computer simulation. *Journal of Research in Science Teaching*, 25, 335 – 369.

O'Hara, K. P., & Payne, S. J. (1998). The effects of operator implementation cost on planfulness of problem solving and learning. *Cognitive Psychology*, 35, 34-70.

Osman, M. (2008a). Evidence for positive transfer and negative transfer/Anti-learning of problem solving skills. *Journal of Experimental Psychology: General*, 137, 97-115.

Osman, M. (2008b). Observation can be as effective as action in problem solving. *Cognitive Science*, 32, 162-183.

Osman, M. (2008c). Seeing is as good as doing. *Journal of Problem Solving*, 2.1.

Osman, M. (2010). Controlling uncertainty: a review of human behaviour in complex dynamic environments. *Psychological Bulletin 2010*, Vol. 136, No. 1, 65–86 DOI: 10.1037/a0017815

Osman, M., & Speekenbrink, M. (under review). Learning to predict vs. Learning to control a dynamic environment. *Journal of Experimental Psychology, Learning, Memory and Cognition*.

Pinrich, P. R., & Degroot, E.V. (1990). Motivational and self-regulated learning components of classroom academic performance. *Journal of Educational Psychology*, 82, 33 – 40.

Preece, J., Rogers, Y., Sharp, H., Benson, D., Holland, S., & Carey, T. (1994). Human-Computer Interaction. Pearson Education.

Pressley, M., & Ghatala, E.S. (1990). Self-regulated learning: monitoring leaning from text. *Educational*, 25, 19 – 33.

Rieman, J. (1996). A field study of exploratory learning strategies. *ACM Transactions on Computer-Human Interaction*, 3(3), 189-218.

Shrager, J., & Klahr, D. (1986) Instructionless learning about a complex device: the paradigm and observations. *International Journal of Man-Machine Studies* Volume 25, Issue 2, August 1986, Pages 153-189. doi:10.1016/j.physletb.2003.10.071

Schunk, D. H. (1983). Progress self-monitoring: Effects on children's self-efficacy and achievement. *Journal of Experimental Education*, 51 , 89 – 93 .

Schunk, D. H., & Rice, J. H. (1987). Enhancing comprehension skill and self-efficacy with strategy value information. *Journal of Reading Behaviour*, 3, 285 – 302

Svendsen, G. B. (1991). The influence of interface style on problem solving. *International Journal of Man-Machine Studies*, 35, 379-397.

Sweller, J. (1988). Cognitive load during problem solving: Effects of learning. *Cognitive Science*, 12, 257-285.

Trudel, C.-I., & Payne, S. J. (1995). Reflection and goal management in exploratory learning. *International Journal of Human-Computer Studies*, 42, 307-339.

Taatgen, N. A. (1999). Learning Without Limits: From Problem Solving towards a Unified Theory of Learning. Unpublished PhD, Rijksuniversiteit Groningen, Netherlands.

APPENDIX A

Instruction Scripts

Control (free)

Main

A recently patented Baby Incubator (Patent:: ISO 9012:2008, ISO 13485:2007) has been developed specifically for newly born babies that have an irregular heart rate. The Incubator is a clean, secure, and regulated environment which enables the medical staff to monitor a baby's heart rate, and adjust its conditions to maintain steady heart rates.

The conditions of the Incubator can be changed in three ways, by modifying air pressure the amount of oxygen released, and the humidity, all of which can affect heart rate activity in different ways either by increasing or decreasing it.

Learning phase

As a trainee maternity nurse you have been called in to monitor your neighbour's newborn daughter Molly. This is particularly hard for you because though you are new to the job you feel a strong responsibility to take extra special care of Molly because your neighbour is a very good friend of yours.

Molly is ill and has an irregular heart rate which can go up as well as down, and so will need to be regularly monitored in the Incubator. Your job will be to adjust the conditions of the Incubator to the right levels in order to maintain Molly's heart rate at a stable level.

Any increase in heart rate above the normal levels is dangerous and indicates over absorption of oxygen. Any decrease in heart rate levels is dangerous and indicates lack of oxygen. On each of these occasions you can adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator. Once you make any change a signal is sent to the monitor that will immediate let you know what effect this had on the baby's heart rate. It is **ABSOLUTELY CRUCIAL** that you pay attention to this information because it will indicate the effects on the baby's heart rate.

Molly will be transferred soon and you have some time to use a simulation system to adjust the conditions of the incubator before Molly arrives. Use this time to make yourself familiar with the system and to try to find out what the controls do. You will have time to make 40 regular adjustments to the conditions of the Incubator.

Control test 1

Molly has now been transferred to the incubator. Now is the time to use the knowledge you have gained from the simulator to operate the system. You will be expected to adjust the (A) air pressure (B) Oxygen, (C) humidity in the Incubator yourself.

The monitor of the baby's heart rate indicates the safe heart rate levels for the baby. The nurse must keep within these levels otherwise she risks seriously endangering the baby's health. The medical staff have asked you do make adjustments for a short period of 20 intervals.

Predict test 1

From your experience of the adjustments made to the incubator and the effects on heart rate, you will have made some assumptions regarding the effects of air pressure, oxygen levels, and humidity on heart rate.

As part of your training you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the "submit" button when you are ready.

Control test 1

You have found out that Molly's condition has stabilised and she has been moved to the normal baby ward instead of being brought to the Incubator. However, Paul a newly born baby boy has now show signs of an irregular heart rate. Given your familiarity with the Incubator the medical staff have asked that for a short critical period of 20 intervals should adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator in order to keep Paul's heart rate at safe levels.

Baby Paul has a more severe condition, and so the heart rate levels that have to be kept are more stringent. Therefore, you need to try and maintain heart rate as close to the “safe” level as possible.

Predict test 2

As earlier, you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the “submit” button when you are ready.

Predict (free)

Main

A recently patented Baby Incubator (Patent:: ISO 9012:2008, ISO 13485:2007) has been developed specifically for newly born babies that have an irregular heart rate. The Incubator is a clean, secure, and regulated environment which enables the medical staff to monitor a baby's heart rate, and adjust its conditions to maintain steady heart rates.

The conditions of the Incubator can be changed in three ways, by modifying air pressure the amount of oxygen released, and the humidity, all of which can affect heart rate activity in different ways either by increasing or decreasing it.

Learning phase

As a trainee maternity nurse you have been called in to monitor your neighbour's newborn daughter Molly. This is particularly hard for you because though you are new to the job you feel a strong responsibility to take extra special care of Molly because your neighbour is a very good friend of yours.

Molly is ill and has an irregular heart rate which can go up as well as down, and so will need to be regularly monitored in the Incubator. Your job will be to monitor the changes the nurse makes to the conditions of the Incubator and predict the effects on the baby's heart rate.

Because you are in training it is important that you pay careful attention in order to learn to make accurate predictions. The nurse may adjust (A) air pressure (B) Oxygen, (C) humidity and will do so to try and maintain Molly's heart rate at a safe level. Any increase in heart rate above the normal levels is dangerous and indicates over absorption of oxygen. Any decrease in heart rate levels is dangerous and indicates lack of oxygen. Once you see the adjustments that the nurse has made, you will be expected to predict the resulting change in heart rate. Once you make a prediction this will be sent to the monitor. It is **ABSOLUTELY CRUCIAL** that you pay attention to this information because it will indicate the accuracy of your prediction.

Molly will be transferred soon and you have some time to use a simulation system to predict the conditions of the incubator before Molly arrives. Use this time to make yourself familiar with the system and to try to learn as much as you can about the effect of the adjustments on the Incubator. You will have time to make 40 regular predictions in order to learn about the conditions of the Incubator simulation and the effects on heart rate.

Control test 1

Molly has now been transferred to the incubator. Now is the time to use the knowledge you have gained from the simulator to operate the system. You will be expected to adjust the (A) air pressure (B) Oxygen, (C) humidity in the Incubator yourself.

The monitor of the baby's heart rate indicates the safe heart rate levels for the baby. The nurse must keep within these levels otherwise she risks seriously endangering the baby's health. The medical staff have asked you to make adjustments for a short period of 20 intervals.

Predict test 1

From your experience of the adjustments made to the incubator and the effects on heart rate, you will have made some assumptions regarding the effects of air pressure, oxygen levels, and humidity on heart rate.

As part of your training you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the "submit" button when you are ready.

Control test 2

You have found out that Molly's condition has stabilised and she has been moved to the normal baby ward instead of being brought to the Incubator. However, Paul a newly born baby boy has now show signs of an irregular heart rate. Given your familiarity with the Incubator the medical staff have asked that for a short critical period of 20 intervals should adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator in order to keep Paul's heart rate at safe levels.

Baby Paul has a more severe condition, and so the heart rate levels that have to be kept are more stringent. Therefore, you need to try and maintain heart rate as close to the "safe" level as possible.

Predict test 2

As earlier, you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the "submit" button when you are ready.

Control (Focused)

Main

A recently patented Baby Incubator (Patent:: ISO 9012:2008, ISO 13485:2007) has been developed specifically for newly born babies that have an irregular heart rate. The Incubator is a clean, secure, and regulated environment which enables the medical staff to monitor a baby's heart rate, and adjust its conditions to maintain steady heart rates.

The conditions of the Incubator can be changed in three ways, by modifying air pressure the amount of oxygen released, and the humidity, all of which can affect heart rate activity in different ways either by increasing or decreasing it.

Learning phase

As a trainee maternity nurse you have been called in to monitor your neighbour's newborn daughter Molly. This is particularly hard for you because though you are new to the job you feel a strong responsibility to take extra special care of Molly because your neighbour is a very good friend of yours.

Molly is ill and has an irregular heart rate which can go up as well as down, and so will need to be regularly monitored in the Incubator. Your job will be to adjust the conditions of the Incubator to the right levels in order to maintain Molly's heart rate at a stable level.

Any increase in heart rate above the normal levels is dangerous and indicates over absorption of oxygen. Any decrease in heart rate levels is dangerous and indicates lack of oxygen. On each of these occasions you can adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator. Once you make any change a signal is sent to the monitor that will immediate let you know what effect this had on the baby's heart rate. It is **ABSOLUTELY CRUCIAL** that you pay attention to this information because it will indicate the effects on the baby's heart rate.

The monitor of the baby's heart rate also indicates the safe heart rate levels for the baby. You have to keep within these levels otherwise you risk seriously endangering the baby's health. Your work-shift will start soon and you have to make 40 regular adjustments to the conditions of the Incubator.

Control test 1

After a short period in which Molly started to make progress and her heart rate returned to normal levels outside of the incubator, she suffered a relapse and had to be returned to the Incubator.

Because you've already been modulating Molly's heart rate levels while she's been in the incubator, the medical staff have asked that you resume doing this for another short period of 20 intervals in which you will be expected to adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator yourself in order to maintain Molly's heart rate at safe levels.

Predict test 1

From your experience of the adjustments made to the incubator and the effects on heart rate, you will have made some assumptions regarding the effects of air pressure, oxygen levels, and humidity on heart rate.

As part of your training you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the "submit" button when you are ready.

Control test 1

Molly's recovery has been good and so she has now been moved to normal baby ward. However, Paul a newly born baby boy has now show signs of an irregular heart rate. Given your familiarity with the Incubator the medical staff have asked that for a short critical period of 20 intervals should adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator in order to keep Paul's heart rate at safe levels.

Baby Paul has a more severe condition, and so the heart rate levels that have to be kept are more stringent. Therefore, you need to try and maintain heart rate as close to the "safe" level as possible.

Predict test 2

As earlier, you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the "submit" button when you are ready.

Predict (focused)

Main

A recently patented Baby Incubator (Patent:: ISO 9012:2008, ISO 13485:2007) has been developed specifically for newly born babies that have an irregular heart rate. The Incubator is a clean, secure, and regulated environment which enables the medical staff to monitor a baby's heart rate, and adjust its conditions to maintain steady heart rates.

The conditions of the Incubator can be changed in three ways, by modifying air pressure the amount of oxygen released, and the humidity, all of which can affect heart rate activity in different ways either by increasing or decreasing it.

Learning phase

As a trainee maternity nurse you have been called in to monitor your neighbour's newborn daughter Molly. This is particularly hard for you because though you are new to the job you feel a strong responsibility to take extra special care of Molly because your neighbour is a very good friend of yours.

Molly is ill and has an irregular heart rate which can go up as well as down, and so will need to be regularly monitored in the Incubator. Your job will be to monitor the changes the nurse makes to the conditions of the Incubator and predict the effects on the baby's heart rate.

Because you are in training it is important that you pay careful attention in order to learn to make accurate predictions. The nurse may adjust (A) air pressure (B) Oxygen, (C) humidity and will do so to try and maintain Molly's heart rate at a safe level. Any increase in heart rate above the normal levels is dangerous and indicates over absorption of oxygen. Any decrease in heart rate levels is dangerous and indicates lack of oxygen. Once you see the adjustments that the nurse has made, you will be expected to predict the resulting change in heart rate. Once you make a prediction this will be sent to the monitor. It is **ABSOLUTELY CRUCIAL** that you pay attention to this information because it will indicate the accuracy of your prediction.

The monitor of the baby's heart rate also indicates the safe heart rate levels for the baby. The nurse must keep within these levels otherwise she risks seriously endangering the baby's health. Your work-shift will start soon and you have to make 40 regular predictions in order to learn about the conditions of the incubator and their effects on heart rate.

Control test 1

After a short period in which Molly started to make progress and her heart rate returned to normal levels outside of the incubator, she suffered a relapse and had to be returned to the Incubator.

Because you've already been monitoring Molly's heart rate levels in the incubator, the medical staff have asked that you now take over and for a short period of 20 intervals you will now be expected to adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator yourself in order to maintain Molly's heart rate at safe levels.

Predict test 1

From your experience of the adjustments made to the incubator and the effects on heart rate, you will have made some assumptions regarding the effects of air pressure, oxygen levels, and humidity on heart rate.

As part of your training you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the "submit" button when you are ready.

Control test 2

Molly's recovery has been good and so she has now been moved to normal baby ward. However, Paul a newly born baby boy has now show signs of an irregular heart rate. Given your familiarity with the Incubator the medical staff have asked that for a short critical period of 20 intervals should adjust (A) air pressure (B) Oxygen, (C) humidity in the incubator in order to keep Paul's heart rate at safe levels.

Baby Paul has a more severe condition, and so the heart rate levels that have to be kept are more stringent. Therefore, you need to try and maintain heart rate as close to the "safe" level as possible.

Predict test 2

As earlier, you will need to answer 16 questions in which you will have the opportunity to apply your understanding of the relationship between the conditions of the Incubator and heart rate. Try to do this as accurately as possible; your responses are particularly important as they will help you in your future career as a maternity nurse.

In each of the questions, you will need to predict a value. In some cases you will be given the heart rate and based on this you will be asked to predict the corresponding values of one or several of the three conditions (air pressure, oxygen, humidity) of the Incubator. Or, you will be asked to predict the heart rate based on the corresponding values of the conditions of the Incubator. You will not receive feedback on the accuracy of your predictions, but try to think carefully about your answers, if you are unsure, then guess.

Move the corresponding slider to generate your answer, then click on the "submit" button when you are ready.