Proceedings of the Third International Conference on Cognitive Modelling.

Device-Oriented and Task-Oriented Exploratory Learning of Interactive Devices

Anna Louise Cox (a.cox@herts.ac.uk)
Richard M. Young (r.m.young@herts.ac.uk)
University of Hertfordshire
College Lane, Hatfield
Herts. AL10 9AB, UK

Abstract

Our research examines the strategies employed by users during exploratory learning, and what they learn about how a device works and how to accomplish tasks using the device. Initial empirical and modelling work suggests that users typically conduct device-oriented exploration resulting in the acquisition of specifically device-oriented knowledge, but not directly of task-oriented knowledge. Although the device-oriented knowledge can serve as a basis for further problem-solving, it does not, in itself, result in the necessary skills to complete domain tasks.

Introduction

Exploratory Learning

People often learn a novel device or software application by actually trying to use it, drawing on a combination of prior knowledge, information from the interface itself, and problem solving skills. We describe this phenomenon as *exploratory learning*. Rieman (1996) showed that in a real world situation, people prefer to learn by exploration in the context of a real task they need to perform, rather than taking time out to experiment with it or work through the documentation in a task-independent manner. Exploratory learning also occurs in situations when training or documentation is not available and in the case of walk-up-and-use devices.

Although there has been some work on issues such as what makes a label on an interface item a good one or not (Soto 1999), there has generally been little consideration of what makes an interface explorable. Explorability must ultimately be a function of, among other things, good labels, well-structured menu systems, the lack of 'hidden' features and the ability of the user to experiment with the interface. Experimentation requires not only that the user constructs a hypothesis and then performs the appropriate action to test that hypothesis but also that the interface provides the necessary feedback to allow the user to conclude whether or not his hypothesis is correct.

Device Knowledge and Task Knowledge

In this research, we distinguish between *device* knowledge and *task* knowledge. Device knowledge usually consists of a collection of facts about what the device as a whole (or parts of it) do. When a user is trying to learn about a new device, one of his goals, for example, will be to find out what each button does. Task knowledge, on the other hand, is knowledge about how to complete a task using a particular device. When learning about a novel device, a user may need to acquire both device knowledge and task knowledge, e.g., 'what does the button labelled *on* do?' and also 'how do I set this video-recorder to record my favourite program while I'm out?'.

The Device

The study reported here concerns a common household device: a central heating timer. The simulation of the central heating timer was created in Allegro Common Lisp and implemented on a PC, based upon the Potterton Puma Electronic Timer. As can be seen in figure 1, the timer has four on-events and four off-events. Unlike the more commonly found central heating timer where the day is selected first, with this timer the user initially selects either an on- or an off-event. He can then choose either to set the time of this event using the *advance* and *back* buttons or to set the day or days that this event applies to. By pressing the *day* button, each day is selected individually in turn, followed by Monday to Friday, Saturday to Sunday, and finally, all seven days of the week.

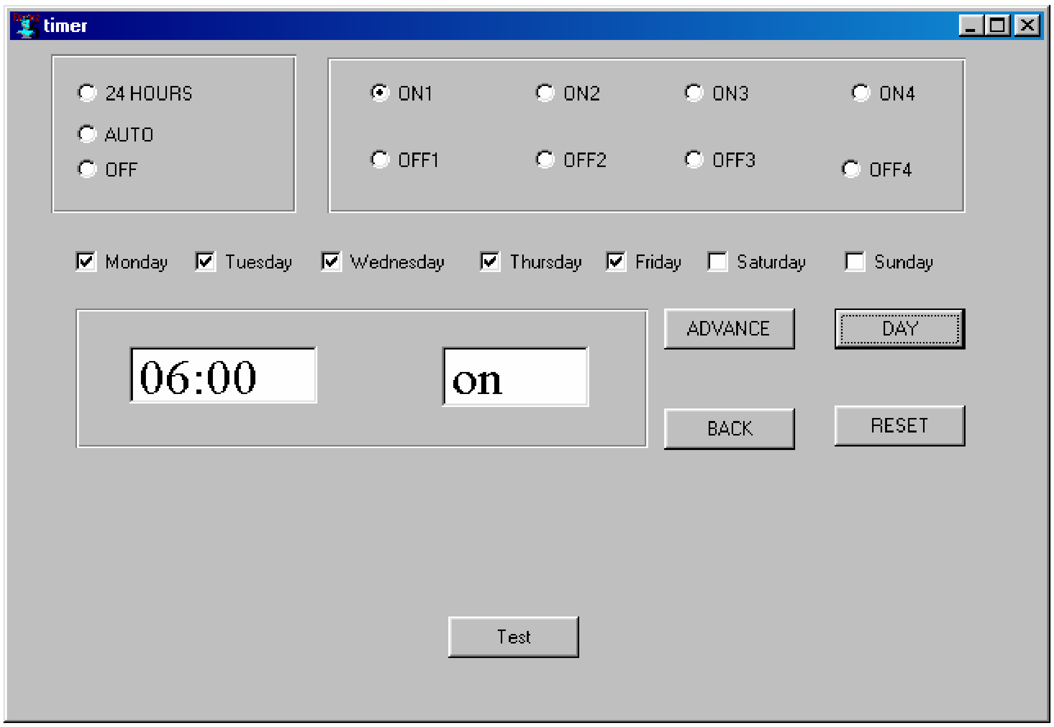

Figure 1: The simulated device used in the experimental work.

Once users have set both the time and day for an event, pressing the *Test* button moves them to the test screen (see figure 2) where they can 'probe' each day and see what the behaviour of the device will be over the week. The provision of this test screen allows the user to 'experiment' with the device, and overcomes the fact that although the device interface itself gives feedback about the settings, it does not give feedback about the behaviour of the central heating system.

The second difference between this device and the more conventional models is that the off1-event does not necessarily have to follow the on1-event. In fact, these two events could relate to different time periods on different days. One could be programmed for

Monday mornings while the other is set to Saturday and Sunday evenings. On standard timers an event called on1 would normally be paired with a following event called off1.

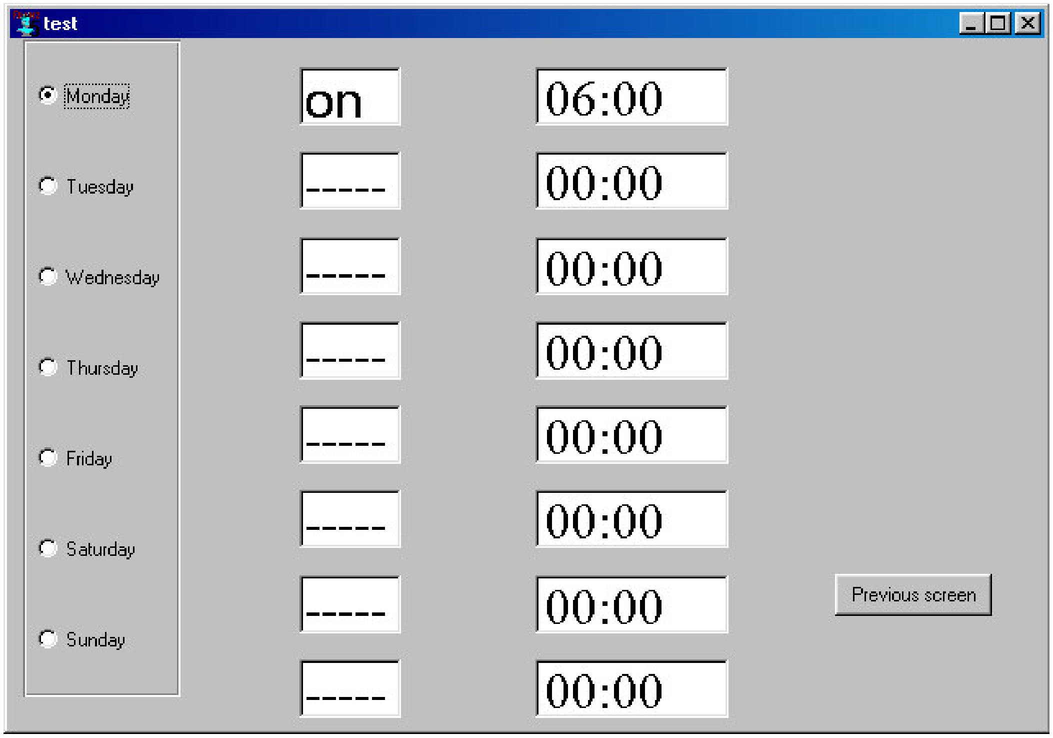

Figure 2: The test screen of the simulated device used in the experimental work.

It is due to these two peculiarities that we predict that users will have difficulty learning to use this device by exploration.

The Experiment

This section first describes the experimental work carried out by Trudel & Payne (1995) regarding reflection in exploratory learning, which serves as a starting point for our research. We then describe an experiment using the device described above, and the modelling work we have done in order to help us explain both how people behave in an exploratory learning situation and what they learn.

Trudel and Payne 1995

Trudel & Payne (1995) carried out a series of investigations into reflection and goal management in exploratory learning. In their initial study, they compared the performance of three groups of participants on tests of declarative and procedural knowledge. They presented their participants with a simulation of a digital watch and asked them to explore it with a view to being tested afterwards. The participants either had 20 minutes in which to do this, or a limit of 250 keystrokes, or were given a list of 7 tasks to complete. Their manipulations yielded dramatic results. They found that the keystroke-limited group did significantly better than the other two groups on the tests for declarative knowledge and procedural knowledge, despite actually spending less time exploring it. They concluded that the imposition of the keystroke limit had made the keystrokes a precious resource and had forced the participants to reflect more fully on each action.

Following the work of Trudel & Payne (1995) it was decided to conduct a partial replication of their first experiment using a different, less-moded device. The central heating timer described above was used to see whether their results were replicable and generalisable to other devices.

Method

The 18 volunteers in the study consisted of 11 males (5 in the unstructured exploration (UE) group and 6 in the mouseclick limited (ML) group) and 7 females (6 in the UE group and 1 in the ML group). They ranged in age from 21 to 68 (UE group 21 to 55; ML group 23 to 68) with a mean age of 33 years (UE group 31; ML group 36). All but 2 of the participants had programmed a central heating timer at least once prior to the study. The 2 who had not were both in the UE group. Most participants (15) said that they had programmed a central heating either a few times or many times. With regard to previous computer experience, 1 participant had between 12 months and four years experience while the other 17 participants had 4 years experience or more.

In the UE condition, participants were allowed to explore the simulated device using the mouse for 15 minutes without any supporting literature. In the ML condition, participants were allowed to explore the simulated device using the mouse for 15 minutes or 100 ‘meaningful’ mouseclicks, whichever finished sooner. A ‘meaningful’ mouseclick was one that involved any of the buttons or switches on the device except for ‘advance’ and ‘back’. The limit of 100 mouseclicks was chosen following the pilot study in which it was noted that the UE group used an average of 200.

After the exploration phase, participants completed a pencil and paper questionnaire that consisted of 11 questions about the timer and the effect of the buttons on it. Following the questionnaire, each participant completed a test of 5 cumulative scenario-based tasks using the simulated device. Each task had a separate time limit in which it had to be completed in order to limit the amount of new exploration that could occur during the test. The time limits were set to be twice the amount of time it took an expert to complete the task, rounded up to the nearest half minute. The wording of the tasks refers to a real-world situation which is elaborated cumulatively. Thus, following the first two questions which set up a pattern of on- and off-times for all seven days of the week, the third question reads: *“Sarah and Frank are finding it a bit chilly at the weekends. Change the timer so that, at the weekend, the heating comes on at 7am and goes off at 11am and then comes back on at 3pm until 11pm. Leave the existing settings for the other 5 days of the week. Program the timer and then check that you have programmed it correctly. (4 minutes).”*

Results

Examination of the time-stamped keystroke log file revealed that neither group used their full 15 minutes of exploration time (UE group = 12mins 17secs; ML group = 8mins 28secs). Participants in the ML group also chose not to use up their full allocation of 100 meaningful mouseclicks using a mean of only 81 compared to the mean of 219 used by the UE group.

The data shows that not only did the UE group use over twice as many mouseclicks overall than the ML group (595 compared to 279), they also used over twice as many 'meaningful' mouseclicks. The UE group also spent almost half as much time again exploring as the ML group. An independent samples t-test[1] was conducted on the mean number of 'meaningful' mouseclicks. This showed a highly significant difference between the two groups (t (12.4) = 6.045, p<0.001)

The questionnaires were scored by assigning one point per correct answer for each of the eleven questions. The mean scores for the two groups show that the UE group scored higher on the test of device oriented knowledge (70%) than the ML group (46%). An independent samples t-test showed that the UE group remembered significantly more device oriented knowledge than the ML group (t (16) = 2.734, p<0.016).

An efficiency score was calculated for each subject on each test question of the scenario tasks test, by dividing the ideal number of mouseclicks by the actual number of mouseclicks and multiplying by 100.

The mean efficiency scores on the scenario tasks test showed that the UE group scored higher than the ML group (Table 1). These scores were found to be significantly different by an independent samples t-test[1] (t (12.8) =3.349, p<0.006).

As it was possible that these scores could have been reduced by the fact that participants had not completed a task within the allotted time and were therefore given a score of 0%, the mean efficiency scores of completed questions were also calculated and are displayed in Table 1. These scores showed that for those tasks completed, neither group out-performed the other.

	Unstructured Exploration	Mouseclick Limited
Mean efficiency scores	47%	20%
Mean efficiency scores for completed questions	79%	72%

Table 1: The mean efficiency scores for the two groups on the scenario tasks test.

Analysis was also carried out on the number of tasks in the scenario tasks test that were completed successfully. The UE group completed a mean of 2.9 tasks while the ML group completed a mean of 1.4 tasks out of 5. This difference was shown to be statistically significant by an independent samples t-test[1] (t (14) = 3.210,p<0.007).

Discussion

Surprisingly, these results have not shown a replication of the effect found by Trudel and Payne (1995). The participants in the unstructured exploration group out performed those in the mouseclick-limited group on both the questionnaire (device oriented knowledge) and the scenario tasks test (task-oriented knowledge). These results therefore suggest that although neither group acquired a significant amount of task oriented knowledge to enable them to complete tasks using the device without further problem

[1] Welch's t-test used due to unequal variances of the samples.

solving, the UE group did acquire sufficient device oriented knowledge to enable them to answer questions about the device.

The low scores on the post exploration tests, coupled with the overconfidence exhibited by the participants during the exploration phase that resulted in their choosing not to use their full 15 minutes or 100 'meaningful' mouseclicks, suggest that the participants failed to realise the mismatch between their prior knowledge of how this class of devices works and the conceptual model of this particular device.

In their 1995 study, Trudel and Payne distinguished between the participants' focus on acquiring device-oriented knowledge as against task-oriented knowledge. They showed that participants spent more time acquiring the former and almost no time on the latter (a mean of 15.1 events acquiring device-oriented knowledge and a mean of 0.7 events acquiring task-oriented knowledge). The high scores on the questionnaire and the low scores on the scenario tasks test found in this study, suggest that this finding has been replicated and that participants failed to set themselves 'realistic' tasks when exploring the interface, but concentrated on acquiring device oriented knowledge.

Taatgen (1999) has remarked that "reflection corresponds to the use of explicit learning strategies" suggesting that reflection is a deliberate attempt to learn a piece of knowledge. This suggests that the high performance of Trudel and Payne's keystroke-limited group can be explained by an increase in the number and/or quality of their deliberate attempts to learn new pieces of knowledge about the device. It would appear that in this study, however, the imposition of the mouseclick limit has not encouraged participants to reflect more fully during their exploration of the device.

There are two main differences between the study reported here and the one carried out by Trudel and Payne, which, we believe, interact to account for the non-replication of their findings. The first is that the two studies employed different devices and that the less-moded central heating timer may encourage overconfidence in participants in a way that the digital watch would not. The second difference is that none of Trudel and Payne's participants had had experience of using a digital watch before whereas all but two of the participants in this study had used a central heating at least once before. We suggest that this prior mental model of how this class of devices works resulted in the participants thinking that they only needed to learn which button did what in order for them to use the device. This would account for, not only their overconfidence in the exploration phase, but also why the participants only focused on device-oriented exploration when this was an inappropriate strategy.

The Model

The initial modelling of the empirical data obtained in this research was concerned with two issues. What do people do during exploration, i.e. what is their behaviour? And, what do people get from exploration, i.e. what do they learn? We have argued that while exploring, people are focused on acquiring device knowledge and therefore try to find out what each button on the interface does and that device knowledge is exactly what they learn.

With this in mind, a simple initial model has been built in ACT-R (Anderson & Lebiere, 1998) and interacts with a pared down version of the interface that was used in the empirical study. This pared down version consists of two buttons (advance and back) and the time display.

The model chooses an interface item (a button) and its goal is to predict what happens to the display of the device after pressing the item. First, the model searches its memory to see if it has already acquired a piece of knowledge that encodes this information, i.e. to see if it already 'knows' what will happen next. If the model can find such information then it outputs the new value of the display. If the model does not already have this information then it sets about acquiring it by pressing the button and reading the new display value. It associates the button it has pressed, the new value of the display and the previous value of the display, and learns this as a piece of declarative knowledge.

Future developments of the model include introducing additional strategies which the model can use to predict the outcome of an action on an interface item, based on its interactions with that item in other situations or on other similar items, and capturing the distinction between device and task-oriented knowledge.

Implications of the Model

Although the model is very simple, just thinking in terms of the model and what will be needed for its future development, has already proved valuable in offering interpretations of the experimental findings and suggesting what is happening during this kind of exploration.

It is obvious that the kind of learning done by the model, i.e. focusing on individual components of the device, concerns the functional relationship between the component and the device. The model might learn, for example, that pressing the button 'advance' leads to changing the setting of the time, or, that to change the time involves pressing the button 'advance'.

Churchill & Young (1991) contrasted the effects of giving Soar models 'how-to-do-the-task' instructions and 'how-the-device-works' instructions. Their results showed that giving a model the 'how-to-do-the-task' instructions (i.e. purely task-oriented procedures) resulted in a situation where the model was able to perform tasks from start to finish but was unable to complete a partially solved task. The model that received the 'how-the-device-works' instructions (i.e. device-oriented), on the other hand, acquired pieces of knowledge that referred to the state of the device rather than which action has just been completed, and allowed the model both to perform tasks from start to finish and also to achieve the final goal from any intermediate state of the problem. This ability to perform tasks, however, depended upon the model's having acquired explicit problem-solving ("task-action mapping") knowledge, and this in turn depended upon its having been trained on a succession of increasingly complex tasks. Although the how-it-works knowledge provides, in principle, a sufficient basis for performing all the tasks, the problem solving required to apply it to realistic tasks *ab initio* would have been impracticably difficult.

The same general argument would appear to hold in the case of our central heating timer. In order to solve the tasks given in the post exploration test, participants need knowledge about the relationship between the scenario situations and the actions on the device (e.g. what to set up, what buttons to press). In Soar and ACT-R, such knowledge can be learned only in the context of realistic tasks. It follows that if, during exploration, participants focus primarily upon device goals (e.g. achieving some setting of the device), they will fail to acquire the kind of task knowledge needed to perform the scenario tests. To do that, they would have to set themselves more realistic task goals (e.g., getting the central heating to come on according to a desired pattern of times), which participants seem not to do spontaneously.

Conclusions

Initial consideration of the results of the experimental work shows that the imposition of the keystroke and mouseclick limits during exploration do not automatically encourage reflection. In fact, in the particular circumstances described in this investigation, the mouseclick limit impeded the acquisition of knowledge by the participants. However, the results of this study also confirm our prediction that people do find this device difficult to learn.

Further analysis of the pattern of results on the post exploration tests shows that people tend to perform device-oriented exploration rather than task-oriented exploration. Those with higher levels of device-oriented knowledge, such as the UE group, are able to complete device-oriented and domain-oriented tasks more successfully than those with lower levels.

Consideration of the experimental data in conjunction with the model suggests that device-oriented exploration by a model or user leads to the acquisition of device-oriented knowledge and does not, in itself, equip the model or user with procedures that enable it/him to complete domain tasks. Instead, from this position the user or model can undertake further problem solving that can lead to task completion. However, in the absence of prior training, that problem solving can be too difficult to perform on-the-fly in the context of unfamiliar, time-pressured tasks. In order to be able to use the device effectively in realistic tasks, users need to have set themselves appropriate task-oriented goals during exploration, or to have tackled an appropriate sequence of training examples.

We can therefore suggest a new gloss on Rieman's (1996) observation that users prefer to learn a novel device by exploration in the context of real tasks. On the basis of experimentation and arguments based upon cognitive models, we suggest that users *need* to undertake at least some of their learning with realistic tasks, in order to acquire the task-oriented knowledge they will require to perform tasks successfully in the future.

References

Anderson, J. R., & Lebiere, C. (1999). *The Atomic Components of Thought*. Mahwah, New Jersey: Lawrence Erlbaum.

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*.

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

Soto, R. (1999). Learning and performing by exploration: label quality measured by latent semantic analysis. *Proceedings of the ACM CHI 99 Conference on Human Factors in Computing Systems*.

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.