5. Experiments and Evaluation

This chapter describes experiments that were conducted using the application prototype and implementation outlined in Chapter 4. Firstly, each experiment has an aim and motivation. Secondly, the experiment's procedure and conditions of each experiment are explained. Finally, the results from the experiments are presented with a discussion and analysis. For the purposes of all the experiments conducted, it should be noted that each friendly name assigned to a device is characterised by what typically might be discovered in a specific context. Therefore, some of the friendly names assigned to a device are not actual items of technology but rather objects one may discover in a particular context. However, it is possible to imagine a world in which all objects identify themselves via RFID or Bluetooth. The actual devices used in all the experiments have Bluetooth support, and they are:

- 3 x Nokia N80 mobile phones
- 2 x IBM T60 ThinkPad's
- 1 x Palm Zire 72s
- 1 x HP Photosmart C7180 All-in-One Printer
- 1 x Motorola RAZR V3xx mobile phone
- 1 x Sony Ericsson W950i mobile phone
- 1 x DELL desktop computer
- 1 x Belkin USB Bluetooth Dongle

Section 5.5 evaluates the overall approach of this model and application from the experiments conducted.

Please refer to Appendix B for instructions for viewing the data collected and recorded from the experiments and application interactions.

5.1 Experiment One: Validation of Application and Accuracy

5.1.1 Aim and Motivation

The implementation of the application prototype needs to operate correctly to conduct any experiments. This experiment was motivated by a requirement to prove the accuracy and validation of the back-end software code. The aim of this experiment was to demonstrate that the software code programmed for the application is able to accurately identify devices in a selected context (known output), given the inputs into the application are known. The implemented application would be able to be validated based on the results of the experiment.

5.1.2 Experiment Procedure and Conditions

To validate the software code and prove the accuracy of discovering devices, the setup of the experiment involved assigning friendly names to Bluetooth devices. The context that was selected for this experiment is an Office context. In this experiment, an Office context is characterised by the following friendly names:

- Desk
- Bookshelf
- Filing Cabinet
- DELL Desktop
- Printer Multifunction
- Web Camera
- Palm PDA
- Nokia N80

A total of eight devices were used in this experiment (see above) and the friendly names for the Office context were assigned to these devices.

This experiment was tested for ten experiences and shows the accuracy of discovering a potential total of eight devices over the course of the experiences. All eight devices were located in the same room and each device was positioned in the same place in the room for all of the ten experiences that were conducted. Each application experience conducted was spaced in 10 minute intervals. Given the inputs (device-friendly names) were known and what the selected output (Office context) was going to be, the experiment setup was quite straightforward.

5.1.3 Results and Analysis

Figure 5.1 illustrates a graph of the experiment results for accuracy of device discovery in an Office context. The first experience was considered a training experience in discovering the devices for the first time in the Office context. Therefore, the constructive memory-match is 0%, as the application had not seen any of the devices before, nor had it saved an experience for this context before. The results began to appear from the second experience through until the tenth experience demonstrating the application and software code were operating correctly.

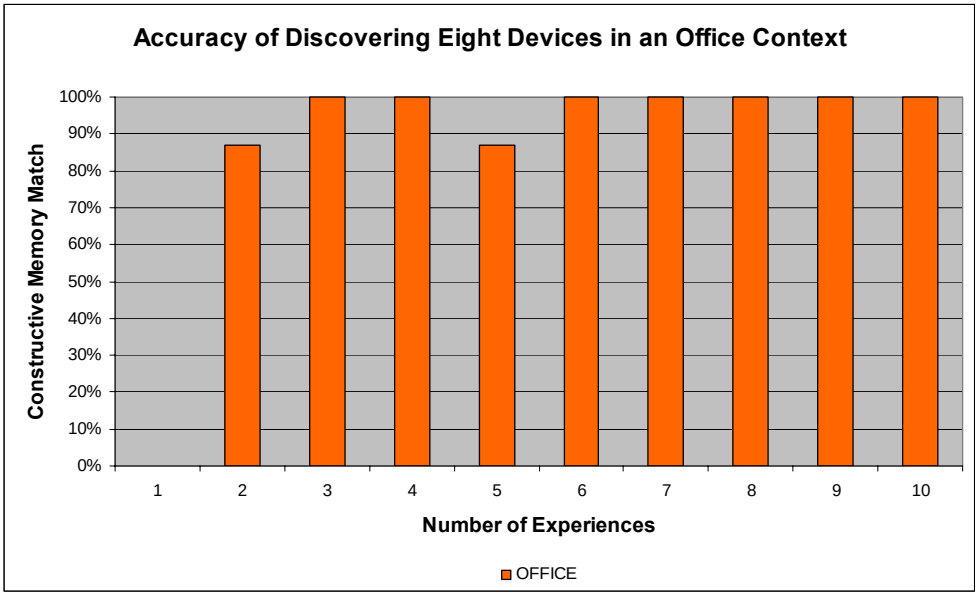

Figure 5.1. The accuracy of constructive memory-matches of devices in an Office context.

The constructive memory-match of the second Office context experience was 87%, which means only seven of the eight devices were discovered in this current experience. The device not discovered was the Printer Multifunction. It can only be assumed that the Bluetooth radio connection dropped out of range or there was some radio interference in the period during which this experiment was being conducted. Experiences three and four provided constructive memory-matches of 100% in discovering and identifying all eight devices.

The fifth Office context experience was only able to discover seven of the eight devices as well, providing a constructive memory-match of 87%. The device not discovered was the DELL Desktop device. The reason for this was that the non-discovered device, in this case a desktop computer, went into hibernation mode, powering down the Bluetooth radio connection. When conducting the experiment for this specific experience, fifty minutes of time had elapsed. There had been no physical activity or interaction – that is, movement of the mouse or keystrokes on a keyboard – occurring on the desktop, and based on our computer settings forced the device to hibernate.

For experiences six through ten, the mouse was moved on the DELL Desktop device to keep it active and the Bluetooth radio connection alive. There was a constructive memory-match of 100% for the remaining experiences (six through ten). Having five consecutive successful discoveries of eight devices and 100% matches for each experience suggests the application was performing at a level that validates the implementation and overall aim of this research in discovering devices and using experience to determine a context.

Discounting the first experience in this experiment, due to it being a training experience to map the eight devices to an Office context, the accuracy of discovery in encountering nine experiences is 97%. The results of this experiment have demonstrated that this implementation can correctly and accurately sense devices without any extensive reasoning over many application iterations for a known context.

5.2 Experiment Two: Discovering Nine Devices in a Context

5.2.1 Aim and Motivation

The total number of devices the application had been discovering is between six and eight. This experiment was motivated by a desire to demonstrate the application is able to extend its ability to discover and reason about a ninth device in a context. The aim of this experiment was to illustrate the application is able to inductively learn and reason from previous experiences, although never previously discovering more than eight devices in a context at one time. This also relates to the modified architecture requirement of scalability.

5.2.2 Experiment Procedure and Conditions

The setup of this experiment involved using pre-existing device-friendly names that have been discovered for each context. Each of the contexts was tested individually, using a variety of friendly names for the nine devices that the application was trying to discover. Prior to running this experiment, each context had five saved experiences in the application. In each of these five previous experiences, the total number of devices discovered ranged from a minimum of six to a maximum of eight. Across these five experiences per context, there were a total of between ten to thirteen device-friendly names discovered. In some instances, the device-friendly names did overlap into other contexts, such as friendly names for mobile phones (Nokia N80) and laptop computers (IBM Laptop).

Before beginning this experiment, each device was assigned a friendly name that the particular context had previously seen. This involved paying attention to the context in which the experiment was being conducted, and assigning the correct friendly names to the devices in the specific context. Duplicating friendly names assigned to devices or adding new friendly names to a context was not desirable, as this would introduce conflicting results and did not support the motivation or aim of this experiment. The aim of this experiment was focusing on discovering a ninth *known* device in a context in which this device had not been previously seen. Adding an *unknown* device to a context not seen previously seen is the focus of another experiment, in Section 5.3. Table 5.1 lists the nine device-friendly names assigned to each context.

	Meeting Room	Lecture Theatre	Office	Library	Computing Studio	Home
1	Media Remote	Sony Vaio	IBM Desktop	Bookshelves Main Catalogue	Desktops	BT LAN Access Point
2	Plasma Display	Projector	Web Camera	Bookshelves Reference Catalogue	Projector	Nokia N95
3	Projector	Student Mobile Phones	Desk	Bookshelves Archive Catalogue	Cameras	Entertainment Unit
4	MacBook	BLIP Nodes	Filing Cabinet	BLIP Nodes	Scanners	Bookshelf
5	IBM Laptop	Presenter Device	Bookshelf	Sony Ericsson W600i	BLIP Nodes	BT Dongle Receiver
6	Audio Speakers	Audio Speakers	Sony Ericsson W600i	Photocopiers	Audio Mixers	Playstation 3
7	Meeting Table	Lectern	Printer Multifunction	Desktops	Max/MSP sensors	IBM Laptop
8	Polycom Conference	Plasma Display	Tablet PC	Slide & Microfilm Scanner	Plotter A0/A1	Printer Multifunction
9	DELL Notebook	A/V Controller	Palm PDA	Palm PDA	Printer A3/A4	Port Replicator

Table 5.1. The friendly names assigned to nine devices per context.

5.2.3 Results and Analysis

Figure 5.2 illustrates a graph for the constructive memory-matches of each context during its sixth experience in discovering nine devices. Provided below is a performance analysis of each context, and how the constructive memory-matches were obtained for the particular context are explained. In all the sixth experiences for the specific context, a small amount of common devices were discovered from other contexts, with the constructive memory-match percentages ranging from 0% to 33%. These constructive memory-matches for a predicted context were weak compared to the context selected for the sixth experience.

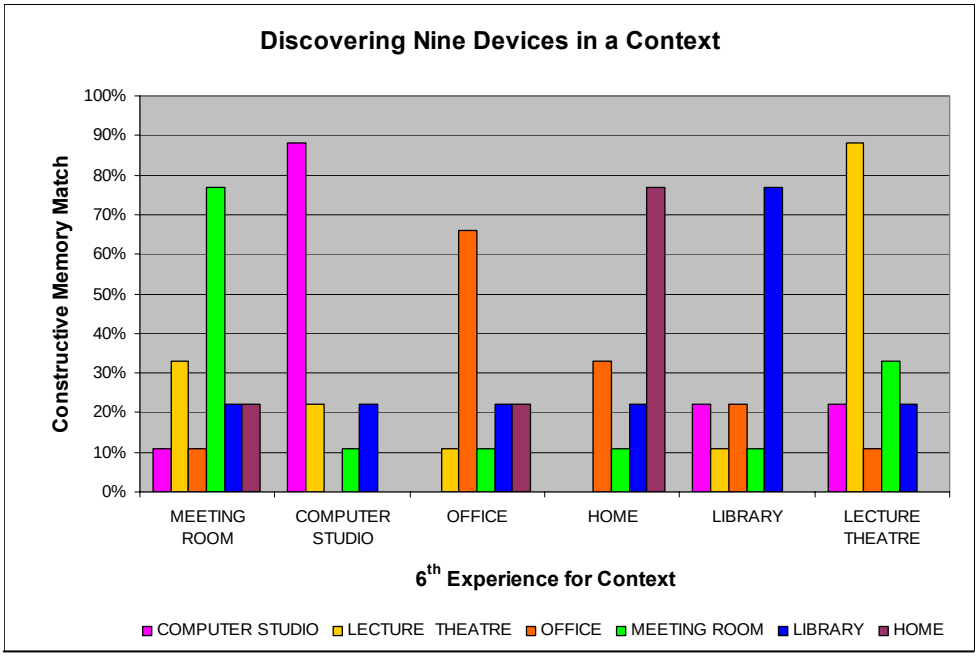

Figure 5.2. Constructive memory-matches for discovering nine devices in a context.

The performance of the Computing Studio context was one of two contexts with the highest constructive memory-match of 88%. In the sixth experience for the Computing Studio context, the application matched eight of the nine devices discovered and the best match was with the first experience the application stored and saved as a Computer Studio context. The other highest performance with a constructive memory-match of 88% was the Lecture Theatre context. The sixth experience for the Lecture Theatre context also matched eight of the nine devices discovered. However, there were two previously saved experiences successful matches, but there is a difference of one device-friendly name between the two experiences. The first match occurred with the first experience the application stored and saved as a Lecture Theatre context. The second constructive memory-match occurred in the fifth experience the application saved as a Lecture Theatre context. The difference between the two experiences is a single device, meaning that seven common devices were previously saved in the first and the fifth experience in the Lecture Theatre context.

The next best performance for memory construction of a context in this experiment were the Meeting Room, Library and Home contexts, with constructive memory-matches of 77%. All matched a total of seven devices previously stored and saved as an experience for their respective contexts out of a possible nine discovered. Firstly, the sixth experience saved as a Meeting Room context was able to best match with the third experience saved for this particular context. Secondly, the first experience saved by the application as a Library context was the best constructive memory-match for its sixth experience. Lastly, the best constructive memory-match

for the sixth experience saved as a Home context was the second experience stored and saved by the application for this context.

The performance of the Office context had the lowest constructive memory-match of 66% when comparing it to other context classes. This context was only able to match six of the nine devices discovered in its sixth experience. The best match for the discovered devices occurred from the fifth experience the application saved as an Office context.

The results of experiment two indicate the application was still able to learn and reason about more than the total amount of devices it has seen in previous experiences saved (which was between six and eight devices) without any additional system training. All of the contexts performed particularly well, with five of the six contexts above 75% in their respective matches for devices it had previously seen. Overall, these results demonstrate that the constructive memory component helps identify and recall a context specific to the devices it discovers in an experience.

5.3 Experiment Three: Adding a New Device to a Context

5.3.1 Aim and Motivation

So far in the experiments conducted, the application has stored and saved a total of 36 experiences; six experiences per context. These experiences have common device-friendly names in each context, but the total devices discovered varied from six to nine per experience, which in turn varied the constructive memory-match percentages. This experiment was motivated to demonstrate how the application reacts to discovering an *unknown* device in a context it had never previously seen before and in some cases, the application had never seen at all. The aim of this experiment was to prove that by adding a completely new unknown device to a context, it does not disturb or upset how this application reacts to reason and construct an appropriate context based on other devices discovered in the current application interaction.

5.3.2 Experiment Procedure and Conditions

To set up this experiment, a list of device-friendly names previously discovered for each context across all experiences was compiled. The reason for this was to avoid conflict with the aim of discovering a new unknown device for a given context. After this list was compiled, the experiment began to assign friendly names to devices and additionally assigned an unknown friendly name to a device. When setting this experiment up, the potential unknown devices that might be discovered, based on *real world scenarios,* of a specific context were kept in mind. Real

world scenarios refer to a device carried by a user into a context as well as a static device or object one might discover in a particular context. Table 5.2 lists the new unknown device-friendly names assigned to each context. In the Home context, the device-friendly name Motorola RAZR V3xx, was the only device the entire application had never seen before at all. Other device-friendly names assigned had been discovered in other contexts, but not the context they were assigned for this experiment.

Meeting Room	Lecture Theatre	Home	Office	Library
BT Dongle Receiver	Tablet PC	Motorola RAZR V3xx	BT Dongle Receiver	Apple iPhone
	IBM Laptop			MacBook

Table 5.2. New devices added to context classes.

5.3.3 Results and Analysis

Figure 5.3 illustrates the results of adding a new device to a context. The average constructive memory-match percentage for this experiment was approximately 70%. This percentage is a strong result for discovering a new unknown device in a context it has previously seen before. A few of the results for this experiment will now be discussed. The Meeting Room context had a constructive memory-match of 66% when the new unknown device was discovered for the tenth experience of this context. The best constructive memory-match for this current discovery occurred in the sixth experience previously stored and saved for this context. The current experience was able to match six of the nine devices from memory. Although the memory-match percentage in this experience fell below the overall experiment average, it can still be regarded as a good result.

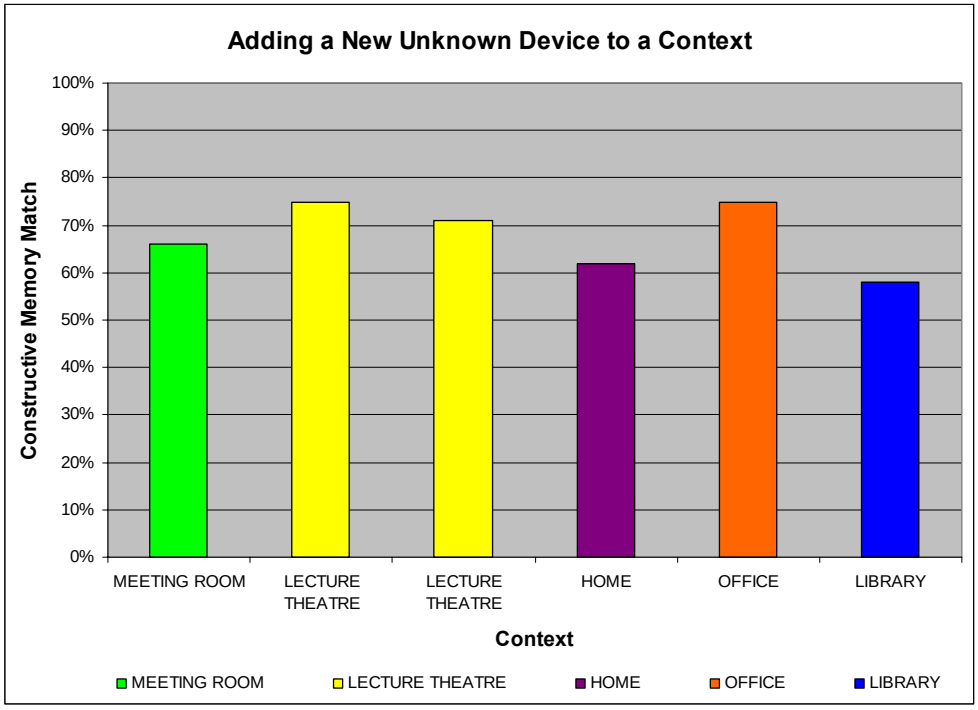

Figure 5.3. Constructive memory-matches for adding a new device to a context.

New unknown devices were added to the eighth and tenth experiences for the Lecture theatre context. Firstly, in the eighth experience, a Tablet PC was added as the new device that had never been sensed before in this context and discovered by the application. Eight devices were discovered in this current experience and the constructive memory-match was 75%. The best match for this current discovery occurred in the first experience with six common devices discovered among eight saved. There was also a match of six devices in the sixth experience; a total of nine devices were saved in this experience. However, the higher constructive memory percentage of devices discovered in the current experience to experiences previously saved was selected as the best match; therefore, in this case occurring in the first experience. Secondly, in the tenth experience, an IBM Laptop was added as the new unknown device that had never been sensed before in this context. In this current experience, only six devices were discovered by the application and the constructive memory-match was 71%. The best match for this current experience occurred both in the seventh and ninth experience previously stored and saved as the Lecture Theatre context. Both experiences successfully matched five of the seven devices discovered in the current experience. Two other experiences also matched the five devices discovered in the current experience. However, the constructive memory percentages were lower then 71%. The total devices matched in the second and eighth experience were five of eight, giving a constructive memory percentage match of 62%.

In the Home context, a device was added which had never been discovered in this context before and additionally had never been seen by the application at all. The device was a mobile phone and its friendly name is Motorola RAZR V3xx. The application discovered a total of eight devices, including this new device for this context in the tenth experience. The constructive memory percentage calculated by the application was 62%. The best match for this current discovery occurred in the second experience saved for this context, matching five of the eight devices discovered. Although the application had never seen the Motorola RAZR V3xx device, the system was still able to reason and construct a context from memory, based on discovering devices it had seen from previous experiences.

Overall, when new devices were added to a context, the experiment showed consistency in being able to construct from memory an appropriate context for all devices discovered in that particular experience. The results of experiment three justify using constructive memory as a method in the discovery and sensing process, to learn and recall contexts without any modifications to the application implementation.

5.4 Experiment Four: Consistency of Context over 10 Experiences

5.4.1 Aim and Motivation

Over a period of time, the application will begin to reason about devices discovered and provide the user with a predicted context. This experiment was motivated to investigate the consistency by the application in constructing and predicting a context based on the total number of devices discovered and the experiences encountered. The aim of this experiment was to demonstrate how constructive memory is useful in context-aware systems in terms of the training required to map devices discovered to a context, based on what the application had seen in previous experiences.

5.4.2 Experiment Procedure and Conditions

This experiment continued to use devices discovered and saved from experiments one, two and three, extending the amount of experiences in a context to ten. Therefore, based on how many experiences a context already had stored and saved by the application, based on previous experiments, each context was required to have a total of ten experiences. When storing and saving a new experience for a context, the application was able to discover between six and nine devices. Over the course of the ten experiences, each context would have a total of between twelve to fifteen device-friendly names. Some of these devices were stored and saved in more then one context, due to the portable nature of the device entering particular contexts, such as

mobile phones or laptop computers. When the application stores and saves each context with additional experiences, care was taken in assigning the correct friendly names from the total number of devices discovered and saved already otherwise the application would encounter inaccuracies in the constructive memory-match percentages. Additionally, when assigning the names to the devices, duplicating device-friendly names may cause application inaccuracies as well. However, when designing and developing this application, this problem was raised and a solution implemented. The application is able to handle an experience when two devices with the same friendly name are discovered. For example, if two devices in a Meeting Room context are discovered and both device-friendly names are *Palm PDA*, the application is able to recognise this discovery as two individual devices. The constructive memory application implementation treats the devices individually, matching only a single device at a time during previous experiences when it has seen the device *Palm PDA*.

5.4.3 Results and Analysis

Figure 5.4 illustrates the consistency demonstrated by the application in constructing a context across ten experiences. The first experience was considered a training experience in mapping the devices saved for the first time to a context class. Therefore the constructive memory-match for the first experience is always 0%. From the third experience onwards, the application began to learn and construct from memory a context based on the devices discovered and what the system had seen in previous experiences. In some context classes, there were increases in the learning rate as the experiences with that particular context increased. Some of the contexts are analysed further where varying rates of consistency occurred.

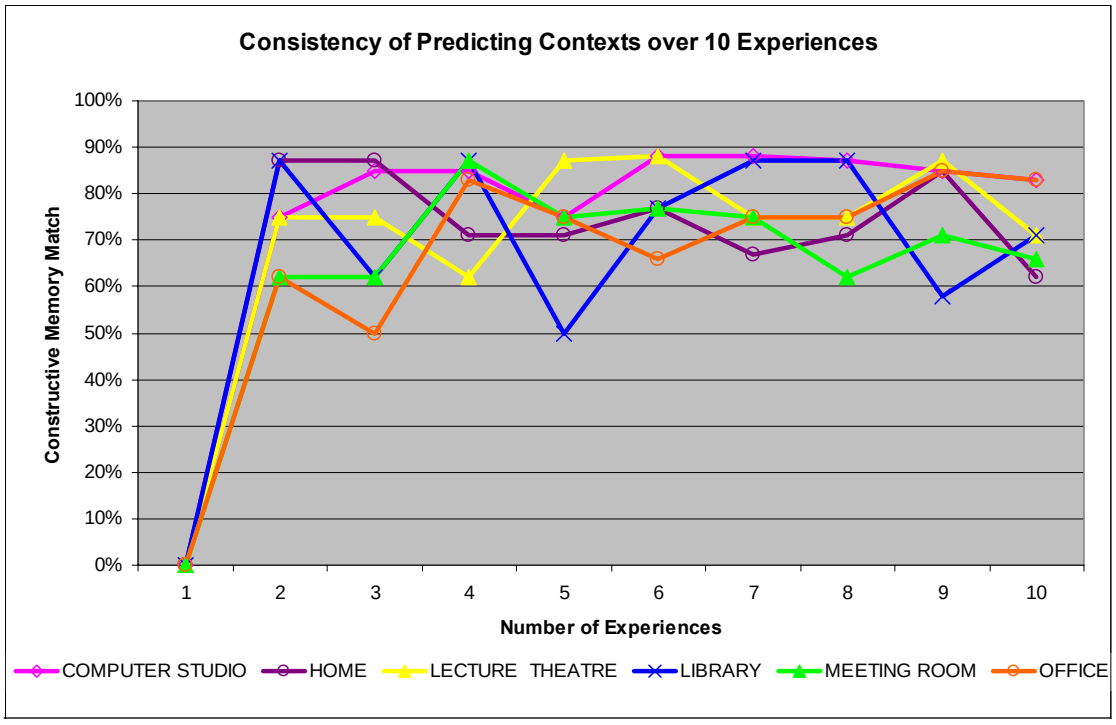

Figure 5.4. Consistency of predicting a context over 10 experiences.

The consistency rate for a Computer Studio context across ten experiences was very strong. Discounting the first experience as the training experience, Figure 5.5 illustrates that there were seven experiences above 80%. The application was able to accurately predict the Computer Studio context as the most appropriate context for the devices it discovered for the experiences with which the application interacted with. Referring to Figure 5.5, there was a dip in the constructive memory-match of only 75% for the fifth experience. In this particular experience, a total of eight devices were discovered. However, when the application constructed the memory-match for this experience, it found previous experiences with some of these saved devices stored, but only the first experience had a maximum match of six devices. Therefore, the reason why the match was low, relative to the other higher rating experiences, was because two new devices, never stored and saved by this context before, were discovered in this fifth experience.

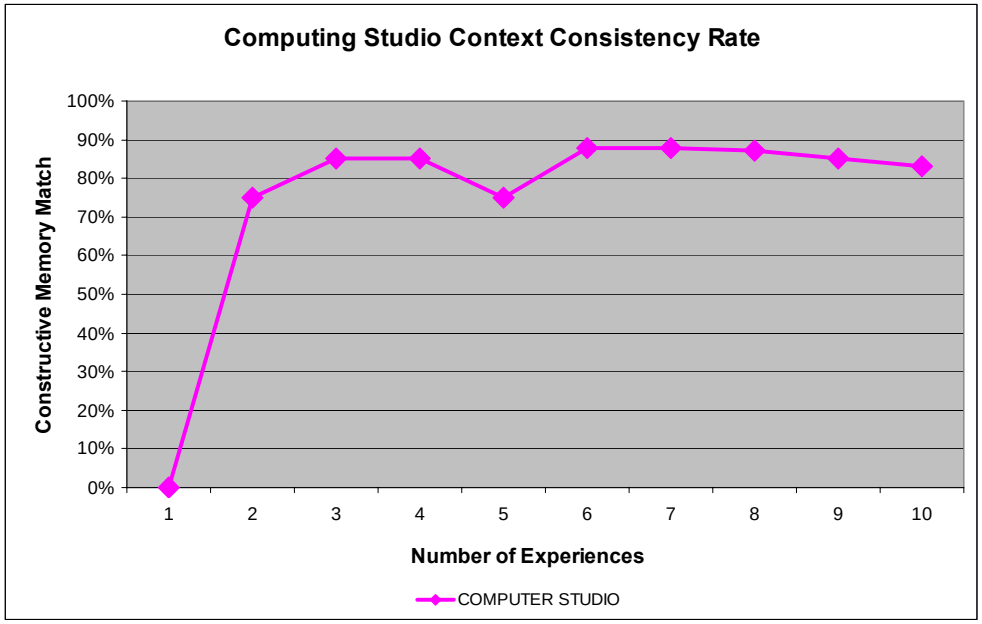

Figure 5.5. Consistency of predicting a Computer Studio context over 10 experiences.

The consistency rate for the Office context across ten experiences was strong. Focusing on the graph in Figure 5.6, the learning rate became progressively more accurate as the number experiences for this context increased. In this context, six experiences rated 75% or higher in constructing from memory that an Office context was the most appropriate for the devices discovered across ten experiences. In the third experience, the constructive memory percentage was the lowest across all experiences, with only a 50% match. The reasoning behind this low match is that only six devices were discovered, and two of these six devices were new to the Office context. The application computed this discovery and concluded after only seeing two experiences prior to this specific encounter, that the best match was the first experience, matching four devices. It is logical to assume that because four devices were matched with the current experience, the constructive memory percentage should be 66%. However, due to the way in which the constructive memory component is implemented, the match occurred with only previously saved experiences. In this instance, both experiences one and two had matches but experience one had the maximum percentage in matching four of the six devices discovered in our current experience. The second Office context experience stored had only three of the six devices saved. Looking at Figure 5.6, the last five experiences show a steady increase in the memory-matches, as there was a new experience. This trend might have continued if more application experiences for an Office context had been conducted.

Figure 5.6. Consistency of predicting an Office context over 10 experiences.

The consistency rate of the Library context is best described as *inconsistent*, as illustrated in Figure 5.7. Overlooking the first experience of the Library context, as it is considered a training experience, the constructive memory percentages varied dramatically. The highest memory-matches experienced by the Library context occurred for experiences two, four, seven and eight, with a memory-match of 87%. The second experience stored for this context discovered a total of eight devices, but it could only compare the devices it had sensed in its current experience to the first experience. Therefore, in this instance seven devices in the current experience were found to have been in common with the first experience. Analysing the eighth experience, a total of eight devices were discovered. Due to previous interactions the application encountered when a Library context was selected as the predicted context, it had the advantage of comparing the current experience with seven prior experiences. In this instance, the devices discovered in the current experience matched with all seven previous experiences, but the maximum match occurred with the fourth experience saved as a Library context. Alternatively, analysing the lowest constructive memory-match, experiences three, five and nine all had matches that were approximately 60% or less. Looking at the ninth experience in more detail, the constructive memory percentage was 58%. In this experience a total of six devices were discovered, of which two devices had never been sensed and stored as a Library context before. Therefore, there were only four devices with which the application was able to reason and use constructive memory to match to other experiences in this specific context. Given there were only four common devices and there were eight prior experiences for comparison, there would be many experiences that matched with these common devices. Nevertheless, the best match was with

the fifth experience saved as a Library context. The memory system computed that four out of a total of seven devices were common for this current experience. Overall, the Library context was only able to provide constructive memory-matches of above 75% for five of the ten experiences. When comparing this rating with other context consistency rates, this performance is poor and inconclusive. It appears that the more contexts there are in which more than one device appears in common, the lower is the accuracy of the constructive memory-matching algorithm.

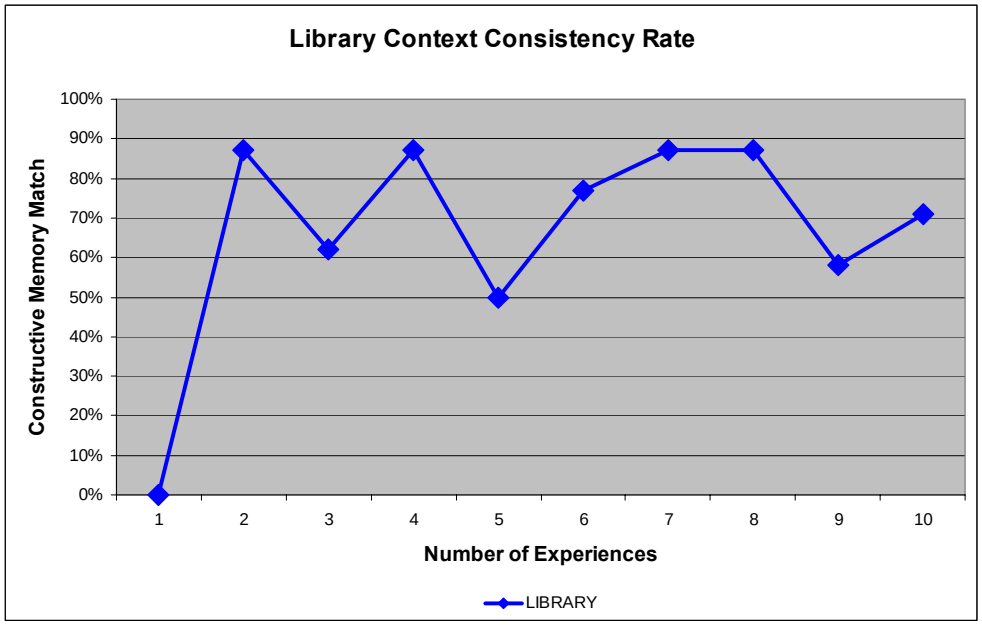

Figure 5.7. Consistency of predicting a Library context over 10 experiences.

5.5 Evaluation

This chapter has presented a series of experiments to demonstrate the application and implementation of the research approach. Each experiment had a different aim and presented different results. For a more detailed description of each experiment, it is recommended that each experiment's analysis and discussion be read for further insight. This evaluation is more a summary of the underlying concepts that have been concluded from all of the experiments.

Experimental results show that the application is capable of discovering devices and constructing contexts from previous application experiences through its use of the constructive memory model. Another interesting finding is that the application is able to develop an adaptive context repository, through constructing a memory during which the application coordinates the systems experience and context in a situated manner. The system exhibits adaptive behaviour.

The characteristics of the application's adaptive behaviour are related to a number of factors:

- the strength of the constructive memory-matching algorithm;
- the experience structures developed for a context; and
- the degree of application interactions.

It is worth noting that the application's performance in discovering devices for a specific context is not necessarily accompanied by an increase in the volume of its experience. Rather, it is associated with the strength and structural features of the experience, as well as the degree of application interactions involved in constructing memory to determine a context.

Experience structures, in terms of sensory, perceptual and conceptual experience, are hierarchical data structures on which the constructive memory and applications experience are modelled. The degree of interactions refers to the number of experiences in a specific context and the application interaction that takes place. These interactions, which encompass reasoning and learning, constitute a constructive memory process. The experiments have demonstrated that by using a constructive memory approach in context awareness, the system can draw on previous experiences to induce and construct potential contexts. What is constructed is the relation between past experiences and details in the current experience. This has the potential to support context awareness and therefore to enhance the efficacy of training such applications.