Augmenting Autobiographical Memory: An Approach Based on Cognitive Psychology

A thesis submitted in fulfilment
of the requirements for the degree
of
Doctor of Philosophy

in
Computer Science

at
The University of Waikato

by
Andrea Schweer

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Department of Computer Science
Hamilton, New Zealand
2011

Dedicated to the memory
of Liesel Schweer née Maas

Abstract

This thesis investigates how an interactive software system can support a person in remembering their past experiences and information related to these experiences. It proposes design recommendations for augmented autobiographical memory systems derived from Cognitive Psychology research into human memory – a perspective missing from prior work.

Based on these recommendations, a conceptual design of an augmented autobiographical memory system is developed that aims to support users in retrieving cues and factual information related to experiences as well as in reconstructing those experiences. The retrieval aspects of this design are operationalised in an interactive software system called the Digital Parrot. Three important factors in the design and implementation are the context of an experience, semantic information about items in the system and associations between items.

Two user studies evaluated the design and implementation of the Digital Parrot. The first study focused on the system's usability. It showed that the participants could use the Digital Parrot to accurately answer questions about an example memory data set and revealed a number of usability issues in the Digital Parrot's user interface. The second study embodied a novel approach to evaluating systems of this type and tested how an improved version of the Digital Parrot supported the participants in remembering experiences after an extended time period of two years. The study found that the Digital Parrot allowed the participants to answer questions about their own past experiences more completely and more correctly than unaided memory and that it allowed them to answer questions for which the participants' established strategies to counteract memory failures were likely to be unsuccessful.

In the studies, associations between items were the most helpful fac-

tor for accessing memory-related information. The inclusion of semantic information was found to be promising especially in combination with textual search. Context was used to access information by the participants in both studies less often than expected, which suggests the need for further research.

Identifying how to appropriately augment autobiographical memory is an important goal given the increasing volume of information to which users are exposed. This thesis contributes to achievement of this goal by stating the problem in Cognitive Psychology terms and by making design recommendations for augmented autobiographical memory systems. The recommendations are confirmed by the design and implementation of such a system and by empirical evaluations using an evaluation method appropriate for the field.

Acknowledgements

First of all I thank my two main supervisors, Annika Hinze and Steve Jones. Together and individually, you provided a mix of inspiration, enthusiasm, constructive criticism, encouragement and reassurance. I am grateful for this opportunity to grow both as a person and as a researcher. Many thanks also to Ian Witten, the third member of my supervisory panel.

I also thank everyone who discussed my research with me, in the Computer Science department and at conferences. In particular I'd like to thank Dave Nichols, Angela Martin and Judy Bowen. Past and present members of the ISDB group throughout the years provided constructive criticism and opportunities for discussion; thank you all.

Thanks to Stefan Mutter, Judy Bowen and Andreas Löf for proofreading parts of this thesis.

Thanks go to everyone working behind the scenes who ensured that I could conduct and write up the research presented in this thesis without major hang-ups, in particular Clint Dilks and the secretaries.

I am very grateful for the friendship with, and support from, current and former fellow PhD students in the Computer Science department. Special mention to Andreas Löf, Matt Jervis, Simon Laing, David Milne, Kathryn Hempstalk and Doris Jung. And an extra-special mention to Stefan Mutter, without whom this thesis wouldn't be what it is – nor would I be who I am now.

Other friends, in the department and further away, deserve my gratitude, for things as different as feeding me (special mention: Tim and Jacqui Elphick), distracting me (everyone with whom I went climbing or played recorders, boardgames or video games), asking me how my thesis was going and *not* asking me how my thesis was going. My thanks go to you all. Particular thanks go to Katrin Credner for supporting me from the other end of the world, sending the occasional care parcel and generally putting things into perspective. Thanks also go to the various supportive flatmates I've had while working on my PhD, particularly Petra Malik and Briar Paul. My family, especially my father Wolfgang Schweer, my mother Ulrike Engelsing and my aunt Bärbel Schweer, supported me in many small and big ways and also have my gratitude.

Finally, I cannot thank Andreas Löf enough for all his support throughout the years and in particular for coping ("well, but not too well") with my all-but-absence in the weeks leading up to the submission of this thesis. I owe you – and I'm sure you'll have plenty of opportunities to collect when it's your turn to write up!

I gratefully acknowledge the financial support of the Computer Science department at the University of Waikato. Further financial support came from Google Australia/NZ in the form of the Anita Borg Memorial Scholarship in 2008.

My sincere thanks to all participants in the user studies I conducted for my PhD research.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

We are capable of remembering astonishing amounts of information at an astonishing level of detail. However, at the same time our memory is fallible, often in seemingly unpredictable ways. Mnemonic strategies and tools to counteract memory failures have been known for millenia (see Aristotle's *On Memory* from the 3rd century BCE, Bloch 2007), often involving verbalisations and visualisations of the material to be remembered. Today, tasks involving verbalisations and visualisations are often performed with the aid of interactive software systems.

This thesis investigates ways to support a person in remembering past experiences and facts related to those experiences using interactive software systems.

Specifically, this thesis proposes recommendations for such systems, which are derived from an analysis of Cognitive Psychology research. It proposes the conceptual design of a system that can be used to record and remember experiences and introduces a selective implementation, the Digital Parrot. To evaluate the design and its implementation, this thesis uses a two-phase method tailored to systems designed for personal memories. Two user studies, including one in which participants attempted to remember personal experiences after two years, showed that the Digital Parrot meets its design goals.

1.1 Research agenda

This section outlines the research agenda for the work presented in this thesis. It first defines the main objective and hypothesis and then refines the objective into a number of research questions. The contributions made by this thesis in relation to each of the research questions are summarised briefly.

1.1.1 Objective

The main objective of this thesis is

> *to develop an interactive software system that helps people remember past experiences and related information.*

The central hypothesis underlying this work is that

> *it is beneficial to use Cognitive Psychology as a basis for such an approach.*

1.1.2 Research questions

The main objective of this thesis, helping people remember past experiences and related information with the use of an interactive software system, raises a number of questions:

1. What does it mean to help someone remember?
2. How can an interactive software system help someone remember?
3. How can we determine whether an interactive software system helps someone remember?
4. Does the system introduced in this thesis help people remember?

Each question is developed in more detail below.

What does it mean to help someone remember?

The seemingly simple objective of “helping someone remember” is actually quite complex and this research question has several subquestions:

- How do experiences turn into memories? What is memory, and what types of memories can be distinguished?
- What does it actually mean to remember something?

- What strategies do people use when they consciously try to bring something back to mind that they are sure is “there” somewhere but not remembered?
- What does “helping” mean in this context – what can go wrong during remembering, and what constitutes “better” remembering?

This thesis focuses on a particular type of memory: a person’s memory for their past experiences. The sub-questions above need to be answered both for memory in general and also for this type of memory.

The research presented in this thesis examined research results from Cognitive Psychology to answer this research question and its subquestions. This served two functions. Firstly, it clarified the terms and concepts used throughout the rest of the thesis. Secondly, it led to six recommendations for methods that seek to fulfil the thesis objective.

Three of the recommendations describe choices which must be made in the design of an interactive software system that aims to meet the thesis objective. The remaining three recommendations advocate the use of three factors that were identified as important in remembering experiences. These three factors are

- the context of an experience;
- semantic information about information items in the system; and
- associations between memories.

The working hypothesis for the remainder of this thesis is that inclusion of these three factors in an interactive software system is beneficial to support a person’s memory for their past experiences.

How can an interactive software system help someone remember?

This question can be answered by describing existing interactive software systems whose goal is to help their users remember. It can also be answered by proposing a novel solution. The research presented in this thesis does both.

The solution space was narrowed down through an analysis of existing systems that aim to support people in managing personal memories and other personal information. Strengths as well as shortcomings were discovered in the treatment of the three factors (context, associations and semantic information) and in the effectiveness of existing systems.

Following the second way to answer the question, a new system was developed that combines strengths of existing approaches while avoiding their weaknesses. The new system is described on two levels: as a conceptual design and as an operationalisation in the form of a selective implementation, the Digital Parrot.

The conceptual design incorporates the three factors identified in relation to the first research question and addresses the entire memory process. It focuses on support for retrieving facts and cues related to semi-structured and potentially important experiences and for reconstructing such experiences. These characteristics make experiences easier to capture automatically and allow the assumption that users may spend at least some time annotating captured information. Examples for this type of experiences are visits to academic conferences, travel and scientific fieldwork.

The Digital Parrot also incorporates all three factors and aims to support the remembering phase of the memory process, the area in which the biggest gap in existing research was identified. The planned end-user studies of the approach determined which components of the conceptual design were selected for implementation. The retrieval and data-storage components were implemented but not those components related to data capture and input – data input in the studies was performed by the researcher rather than by the study participants.

How can we determine whether an interactive software system helps someone remember?

Any system put forward as an answer to the previous question needs to be evaluated to determine whether it actually meets the objective. A number of standard evaluation methods exist in the areas of Human-Computer Interaction and Information Retrieval. However, evaluations of systems dealing with personal information in general and personal memories in particular pose challenges that set them apart from evaluations of other systems.

The research presented in this thesis reviewed the challenges involved in evaluating systems dealing with personal memories and other personal information. A new approach to evaluating systems designed for personal memories was developed to address these challenges. A major challenge in evaluating any system for personal information centres around the data used in evaluations. Artificially generated data collections lead to non-naturalistic interaction with the system and thus to less meaningful evaluation results. On the other hand, truly personal data collections are difficult to obtain and make it hard to compare results across participants. Evaluating systems designed for personal memories comes with the additional challenge that long timespans may be required between an experience and studies for remembering, to allow study participants to forget all or parts of the experience.

The new evaluation method allows to study the effectiveness of a system for personal memories. Comparability of results across participants is achieved using a task-based design in which questions are personalised for each participant but fall into shared categories. The new evaluation method allows for more meaningful evaluations of systems for personal memories than methods traditionally used for such systems.

Does the approach introduced in this thesis help people remember?

Once an evaluation method has been identified, it can be used to determine the effectiveness of a proposed solution. Answering this question determines whether the solution put forward actually fulfils the objective.

Two end-user studies of the Digital Parrot were conducted and are described in this thesis. The first study used standard methods for usability testing. The second study evaluated the Digital Parrot's effectiveness using the method introduced in this thesis; it involved a gap between experience and remembering of approximately two years. The memory data collections used in the studies describe experiences at academic conferences as an exemplar for the type of experiences targeted by the system. For reasons of scope, both studies focus exclusively on the remembering phase of the memory process. This focus also helps minimise the impact of the study on memory formation.

The evaluations show that the Digital Parrot successfully supports its users in answering questions about someone else's experiences (first evaluation) as well as their own experiences (second evaluation). Regarding the three factors identified above, the results of the evaluations suggest that information about the context of an experience may not be as useful in supporting memory as expected based on the survey of Cognitive Psychology research and especially based on the treatment of context information in existing software systems. The results suggest that associations between memory items are very useful in supporting remembering. Semantic information about items in the system is also useful; evaluation results indicate that the Digital Parrot's user interface component for this type of navigation leaves room for improvement and suggest that it may benefit from a stronger integration with textual search.

1.2 Structure of this thesis

This section explains how the research questions and the contributions made in this thesis map to the structure of this thesis.

Chapter 2 answers the first research question, about the meaning of "helping someone remember". Relevant research results from Cognitive Psychology are summarised to establish terms and concepts that are central to the remainder of this thesis: different types of information to be remembered, different kinds of remembering, different ways to remember as well as common memory failures. The chapter concludes with a list of recommendations.

Chapter 3 contributes to answering the second research question, how an interactive software system can help someone remember. It analyses the strengths and shortcomings of existing software systems that aim to support memories; the two main areas that are considered are Capture, Archival and Retrieval of Personal Experiences (CARPE) and Personal Information Management (PIM). The concepts and recommendations identified in Chapter 2 form the basis for this analysis.

Chapters 4 and 5 complete the answer to the second research question. Together they show two aspects of a new approach to support a person's memory for past experiences. Chapter 4 proposes the conceptual design of a system to help people remember. The conceptual design takes into account the recommendations and the results of the analysis of existing approaches. Chapter 5 introduces the Digital Parrot, an implementation of those aspects of the conceptual design that relate to remembering.

Chapter 6 addresses the third research question, about methods to evaluate systems designed for personal memories. It does this by reviewing challenges around evaluating such systems and existing strategies to overcome these challenges.

Chapters 7 and 8 describe end-user studies of the Digital Parrot that answer the fourth research question, about the effectiveness of the approach introduced in Chapters 4 and 5. The results of these studies support the central hypothesis of this thesis. The first study uses a traditional evaluation method, while the second study uses a new evaluation method tailored to systems designed for personal memories.

Finally, Chapter 9 employs a wider perspective by summarising the work presented in this thesis, including its contributions and answers to the research questions. It discusses the implications of the findings of the work presented in this thesis for augmenting autobiographical memory, describes its limitations and points out opportunities for future work.

Chapter 2

Augmenting memory: a Cognitive Psychology perspective

This chapter addresses the first research question: What does it mean to help someone remember? To answer this question, this chapter examines theories and models from Cognitive Psychology research that are relevant to the work presented in this thesis.

Figure 2.1 shows a process view of a person's interaction with his or her memory. Experiences are encoded into memories. Later, cues can cause recollections to be remembered, i. e. recalled from memory. Terms related to remembering are not used consistently in the literature reviewed in this chapter. In this thesis, the following main terms are generally used when not quoting terms used by others:

Remembering The act of bringing back something from memory to consciousness.

Recalling Used interchangeably with "remembering", mostly in phrases such as "in the recall process" or "during recall". A slight prefer-

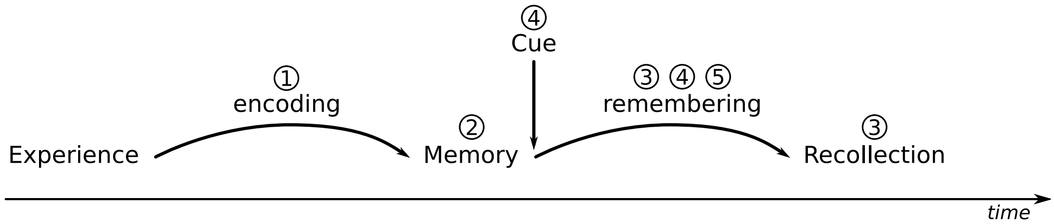

Figure 2.1. Memory lifecycle

ence is given to “recalling” over “remembering” to indicate situations in which a person consciously attempts to remember something.

Retrieving Deliberately recalling information without apparent effort; implies that the information to be retrieved is relatively small and unaltered by the retrieval process. Can refer to retrieving information from somewhere external to a person’s memory, for example from a computer system. “Remembering an experience” and “recalling an experience” stand for the retrieval of a memory item, the representation of an experience or fact.

Recollection The product of the recall process; that which is remembered.

The numbers in Figure 2.1 refer to the subquestions of the first research question (see Section 1.1.2):

(Q1) How do experiences turn into memories? What is memory, and what types of memories can be distinguished?

(Q2) What does it mean to remember something?

(Q3) What strategies do people use when they consciously try to bring something back to mind that they are sure is “there” somewhere but not remembered?

(Q4) What does “helping” mean in this context – what can go wrong during remembering, and what constitutes “better” remembering?

These questions apply to memory in general and also to the specific type of memory central to this thesis: memory for past experiences and for information related to these experiences.

This chapter is structured as follows. Sections 2.1 through 2.3 summarise Cognitive Psychology research to answer the four questions (Q1–Q4). The differences between general memory on one hand and memory for past experiences and related information on the other hand are addressed throughout these sections. Section 2.4 discusses the implications of the research described in this chapter on two levels: for the research presented in this thesis and for the wider research community in general. The chapter concludes with a summary in Section 2.5.

Note that this chapter does not provide a comprehensive introduction to the Cognitive Psychologists' view of human memory. Rather, it is an overview of those theories and models from Cognitive Psychology that pertain to the research presented in this thesis. A more general introduction can be found in (Eysenck and Keane, 2005, ch. 6).

2.1 Formation and types of memories

This section reviews how an experience turns into a memory, what memory is and what types of memory can be distinguished (Q1). These questions are important to consider because the answers clarify what kinds of information augmented memory systems need to work with. This section first gives an overview of current memory models and then takes a more in-depth look at four sub-types of memory that are particularly relevant for this thesis.

Cognitive psychologists distinguish different kinds of memory (Eysenck and Keane, 2005, part 2). The main distinction is according to the timespan for which information remains in memory. From shortest to longest lifetime, these are: sensory stores, short-term memory and long-term memory.

Data perceived through the senses arrive in the *sensory stores*. Data in these stores have a lifetime measured in milliseconds. Sensory stores are limited in capacity and newly incoming data replace older data.

Short-term memory can hold a limited amount of information for up to several minutes. After this timespan, information in short-term memory will decay unless it is rehearsed. Rehearsal, for example by repetition, of information in short-term memory can lead to it being encoded to long-term memory. Information can enter short-term memory from the sensory stores and from long-term memory. Information from both sources will only make it into short-term memory if it is given attention.

In recent years, cognitive psychologists have extended the rather simple model of short-term memory to that of *working memory*. Working memory consists of separate components for visuo-spatial, phonological and episodic information and a central executive. In the working memory model, the central executive coordinates the other components. It plays a role similar to that of attention in the classical model of short-term memory but also takes into account the self-concept of the person.

Information in *long-term memory* has a virtually unlimited lifetime. The type of memory with which this thesis is concerned, memory for a person's experiences and related information, is part of long-term memory. Hence, the next section takes a more in-depth look at long-term memory.

2.1.1 Long-term memory

Cognitive psychologists distinguish different kinds of long-term memory, according to the type of information that is stored. *Procedural memory* is knowledge on how to perform certain tasks (such as how to ride a bicycle). In contrast, *declarative memory* consists of factual knowledge. This thesis is concerned with declarative memory; hence, declarative memory is described in more detail below.

Figure 2.2 shows a taxonomy of the types of memory that have been described so far, with the types that are of importance for this thesis highlighted in boldface.

There are several criteria by which declarative memory can be further subdivided. In early memory models (see Eysenck and Keane, 2005, p. 233), declarative memory was in turn subdivided into *semantic memory* for general facts on one hand (such as what a bicycle is) and *episodic memory* for facts about individual events on the other hand (such as recollections of when one learnt to ride a bicycle). In more recent models, this distinction has become blurred. Semantic memory is now seen as derived from episodic memory through generalisation and abstraction (Cohen, 1996, p. 146).

Brewer (1986, p. 26ff) proposes an alternative, more elaborate model of long-term memory that classifies types of memories according to three

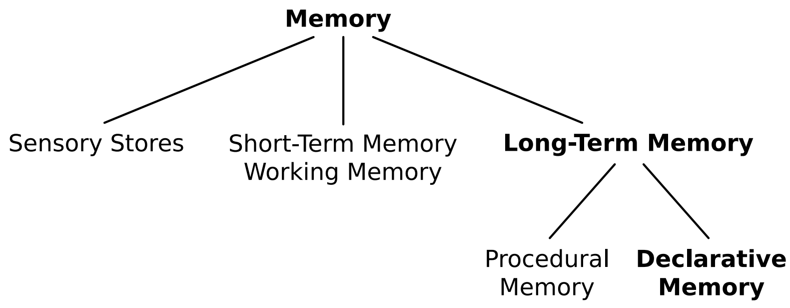

Figure 2.2. Types of memory

dimensions. These dimensions, as shown in Figure 2.3, are:

Acquisition condition: single instance versus repeated exposure. According to Brewer, experiences that occur repeatedly tend to blend into a generic form of memory, while single experiences stand for themselves.

Content: self versus depersonalised. Depersonalised memories are further subdivided into those with visuo-spatial (objects, places), visuo-temporal (events, actions) and semantic content.

Form of representation: imaginal versus nonimaginal. Brewer does not explain which kind of experiences will lead to which form of representation.

Examples for some of the resulting categories are given in Table 2.1.

The following two sections go into even more detail and describe two areas of long-term memory that are particularly important for this thesis. The first is autobiographical memory – memory for past experiences. The second area sits on the boundary between memory and knowledge and consists of concepts, scripts and schemata.

2.1.2 Autobiographical memory

Autobiographical memory is commonly defined as the memory for the events of one's life (Conway and Rubin, 1993, p. 103). Brewer (1986, p. 33f) considers all those memories to be autobiographical which have a connection to the self – the topmost slice of the cube shown in Figure 2.3. Similarly, a revised definition of episodic memory is based on "autonoetic

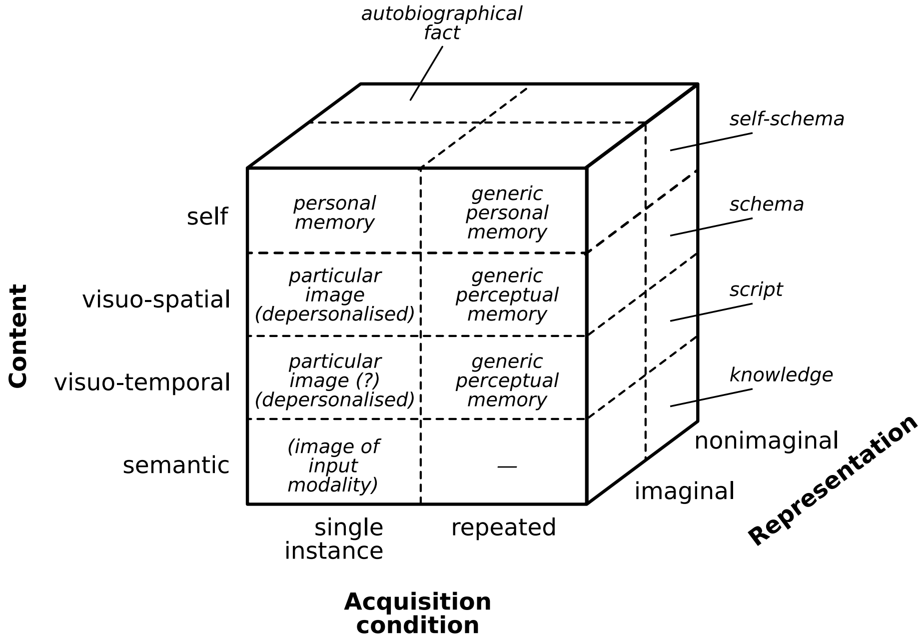

Figure 2.3. Brewer's classification of long-term memory

1. "A very clear picture in my mind of sitting with A. on that pebble beach on the way to Cathedral Cove, looking out over the water and talking about scuba diving"
   Personal Memory. Single instance; self; imaginal.

2. "I was there three times, with three different people"
   Autobiographical fact. Single instance; self; nonimaginal.

3. "Memories of sitting in the passenger seat of a car; it's a very windy road and there's a steep hill going up to the right; to the left, it's just a few metres down to the sea; waves breaking onto a pebble beach"
   Generic personal memory. Repeated; self; imaginal.

4. "There still are feral kiwi birds on Coromandel Peninsula"
   Semantic memory. Single instance; semantic; nonimaginal.

5. "A picture in my mind of the outline of Coromandel Peninsula as it would appear on a map"
   Generic perceptual memory. Repeated; visuo-spatial; imaginal.

Table 2.1. Memories retrieved with the cue "Coromandel Peninsula" and their category, acquisition condition, content and mode of representation

awareness" – a subjective feeling of the person that she or he is remembering a past experience (see Eysenck and Keane, 2005, p. 233f). Following either classification, autobiographical memory is mostly episodic but also contains some semantic aspects.

Autobiographical memory differs from other long-term memory in its functions and its structure.

Function

Autobiographical memory is strongly connected to the self. Consequently, it has a number of functions that affect the self. Autobiographical memory provides the person with a sense of self; it enables the person to predict the future based on their past experiences; it enables the person to connect with others by sharing memories and thus communicating a certain self-image (Conway and Pleydell-Pearce, 2000).

Structure

Theories about the structure of autobiographical memory vary (Cohen, 1996, p. 152; Conway and Pleydell-Pearce, 2000), but they all have in common that they see information in autobiographical memory as organised hierarchically into lifetime periods, events and event-specific knowledge. This is illustrated in Figure 2.4.

Lifetime periods are longer timespans. These periods can follow several different themes and periods can overlap. Examples are work: when one had the student job as a tutor, one's first job after university; and relationships: before one met A, during the time when B was one's best friend.

Events are short periods of time that are perceived as distinct from one another: the visit to conference C, the holiday in Barcelona. Events are located in one or more lifetime periods and can be nested within each other. For example, each conference day is a sub-event of the whole conference visit.

Event-specific knowledge is fine-grained information associated with a certain event: impressions of the room in which the conference's first

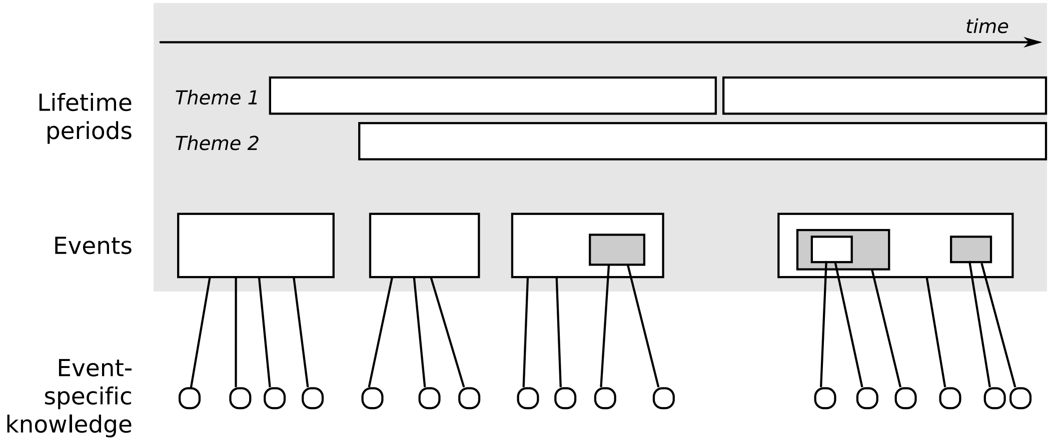

Figure 2.4. The structure of autobiographical memory. Based on Figure 1 in (Conway and Pleydell-Pearce, 2000, p. 265).

keynote was held, the name of the hotel in which one stayed in Barcelona. Event-specific knowledge is independent of time.

2.1.3 Concepts, scripts and schemata

Semantic memory holds generic and factual information. Generally, this information has been derived from past experience via generalisation or abstraction (Cohen, 1996, p. 146).

It is important to look at semantic memory because memory for past experiences is not purely autobiographical. Generic knowledge, in the form of so-called *schemata*, plays a big role in structuring and storing new information as well as in retrieving memories (Cohen, 1996, p. 77f).

Schema theory was introduced in the early 1930s. There are different types of schemata, depending on the type of entity (Cohen, 1996, p.76f).

Concepts: generic information about objects (Eysenck and Keane, 2005, p. 293f). In Brewer's model (see Section 2.1.1), concepts are a subset of the knowledge category.

Scripts: generic information about events (Cohen, 1996, p. 138f). This is a subset of the "scripts" category in Brewer's model.

Action schemata: generic information about actions and sequences of actions (Cohen, 1996, p. 137); they are a generic form of procedural memory. Brewer's model does not distinguish action schemata from scripts.

Of these, concepts and scripts are described further because memory for actions is not relevant for the objective of this thesis.

The schema-plus-tags model (Cohen, 1996, p. 141) describes how memories about individual events are stored in memory: Each event refers to the appropriate scripts plus additional information (tags) that is not represented in the script. The schema-plus-tags model is now thought to be oversimplified and has been extended (Cohen, 1996, p. 144ff), but for our purposes the basic model as described above is sufficiently expressive.

Cognitive psychologists assume that these abstraction and generalisation processes are a countermeasure to limitations in mental capacity (Barclay, 1986, p. 89). Their use goes beyond memory processes: schematic knowledge also plays a crucial role in planning and performing actions as well as in understanding other people's behaviour (Reiser et al., 1986, p. 102).

As mentioned, concepts and scripts also play a role during remembering. This role is described in detail in Section 2.2.2.

2.1.4 Memory for planned actions

This section briefly describes a type of memory that on first glance seems closely related to the objective of this thesis: memory for actions to be performed in the future. An examination of Cognitive Psychology research shows that in fact this type of memory operates quite differently from memory for past experiences and for facts related with such experiences.

Memory for actions to be performed in the future is called *prospective memory*, in contrast to *retrospective memory* for past experiences (Eysenck and Keane, 2005, p. 288ff). Prospective memory can either be time-based or event-based. Time-based prospective memory concerns itself with actions to be performed at a specific point in time (for example, "it's my friend's birthday on Saturday, I must remember to call them"). Likewise, event-based prospective memory concerns itself with actions to be performed when specific conditions are met (for example, "I must remember to buy milk when I pass the supermarket on the way home from work tonight").

Prospective memory may be perceived as long-term memory because it operates over longer timespans than are normally associated with other types of memory. The working memory model, however, includes an explanation of how this type of memory works as a part of short-term memory, especially for event-based prospective memory (Eysenck and Keane, 2005, p. 290). In this model, planned actions are stored as a combination of trigger and action. The central executive dedicates a portion of its processing power to a background loop that checks the current situation for the presence of the stored trigger.

2.2 Remembering

The previous section gave an overview of how memories are created and their structure and functions. This section considers another important aspect of memory: remembering. It addresses the third and fourth question covered by this chapter: What does it mean to remember something (Q2)? What strategies do people use when they consciously try to bring something back to mind that they are sure is "there" somewhere but not remembered (Q3)? The answers to these questions clarify what processes need to be supported in augmented memory systems and provide a basis for the interaction design of such systems.

The remainder of this section first describes recollections as the product of the remembering process. It then describes the process of remembering.

2.2.1 Recollections

There are two views in Cognitive Psychology about the nature of recollections; these views differ fundamentally. Both are relevant to the research presented in this thesis because they lead to different requirements for interacting with information in augmented memory systems.

The two main theories about the recall process are the copy theory and the reconstruction theory. Proponents of the *copy theory* believe that remembering is retrieval of existing, fixed memories; this would make the recollection a copy of the original experience. In contrast, proponents of the *reconstruction theory* believe that remembering is a reconstruction of likely experiences that are consistent with smaller existing, fixed memories as well as with the person's schematic knowledge and possibly other factors such as the person's self-image and current goals.

Brewer (1986, p. 40ff) gives a short overview of both theories and proposes a more moderate partial reconstructive view according to which some memories are reconstructed while some memories are retrieved (see also Cohen, 1996, p. 162). In accordance with his theory about the structure of long-term memory (see Section 2.1.1), he argues that memories for single events – the left half of the cube shown in Figure 2.3 – are more likely to be copies of the original experience, while repeated events – the right half of the cube in Figure 2.3 – become generalised and are more likely to be reconstructed during recall.

Whether recalled directly or reconstructed, recollections of past experiences can be of two different types (see Gardiner and Richardson-Klavehn, 2000). One, generally simply called *remembering*, is accompanied by a strong subjective feeling of at least partially experiencing the original event again. This type of remembering is referred to as *re-living* in the remainder of this thesis to avoid confusion. The other, generally called *knowing*, is accompanied by a subjective feeling that the recollection is correct but not by a feeling of re-living the experience.

2.2.2 Cues and process

Figure 2.1 on page 9 shows the remembering process as being triggered by a cue. However, a person does not necessarily need to be aware of the cue in order to remember. In fact, the person is not necessarily aware of the remembering process itself. With or without an obvious cue, a recollection may suddenly appear in a person's mind without apparent effort (something "just springing to mind"). This type of recall is called *spontaneous recall*. In contrast, *generative recall* is a conscious process in which a person actively tries to recall some memory, more or less deliberately choosing different approaches if necessary.

Conway et al. (2001) describe the so-called “retrieval mode”, a mode into which the brain goes while remembering. Remembering, here called retrieval, occurs in a feedback loop where cues trigger the activation of associated information, which in turn may be used as cues in a next retrieval step. This is repeated until some verification criteria are satisfied. This process may occur consciously or unconsciously.

The remainder of this section describes three components of this retrieval mode in more detail: cues, strategies and verification criteria.

Cues

An important type of cue is context. The *encoding specificity principle* states that recall of a particular memory item is better the more overlap there is between the information available at encoding time and the information available at recall time. This context encompasses both external state (such as location, other people nearby, weather conditions) and internal state (such as mood). This is summarised by Brown and Craik (2000, p. 98f).

One of the implications of the encoding specificity principle is that recognition is generally easier than recall: During recognition tasks, a lot more of the context overlaps with the context available at encoding time.

Strategies

Barsalou (1988, p. 215f) describes how people navigate between topics and clusters of topics during recall. Specific topics can serve as pivot points to switch between two groups of memories. Typical pivot points for autobiographical memory are activities, locations, time and participants.

Reiser et al. (1986) and Cohen (1996) give examples for strategies for reconstructing experiences. Reiser et al. (1986, p. 101) sees remembering as “a process of *reunderstanding* the experience”. He lists a number of typical questions that a person might ask herself/himself while trying to recall an autobiographical memory: “Why would I have been doing that? What might I have been doing when that occurred? What would have led to such an event? What would have happened following such an event?” In more abstract terms, he lists the types of reasoning in autobiographical memory search as strategies based on goals, activity settings, enabling conditions and consequences of events.

Cohen (1996, p. 180) extends this list with inferences based on meta-knowledge (“It can’t be true because I would have known it if it were true”), set inferences (for example subset inheritance of characteristics: “All conferences have keynote talks, so I’m sure this conference had keynote talks too”) as well as spatial and temporal inferences. She points out that inferences can be negative as well as positive.

Verification criteria

One criterion to determine whether a remembered given event has actually been experienced (and not just imagined) is the presence of sufficient event-specific knowledge associated with the event (Conway and Pleydell-Pearce, 2000). Examples for such event-specific knowledge are remembered images and smells.

The working memory model was described in Section 2.1. Conway and Pleydell-Pearce (2000) extend the central executive part of this model to the *working self* and explain its role in the retrieval of autobiographical memories. According to Conway and Pleydell-Pearce, the working self keeps track of the person’s current goals and plans. In retrieval mode, the working self suppresses the activation of information that is in conflict with these goals and plans. Likewise, the working self influences the verification criteria of the feedback loop.

2.3 Memory failures

The two previous sections described how experiences turn into memories and how memories turn into recollections. This section describes typical categories of failures during the processes involved. It addresses the fifth research question covered by this chapter: What can go wrong during remembering, and what constitutes “better” remembering (Q4)? The section first describes general memory failure and then failures of autobiographical memory. The answer to this question clarifies leverage points for augmented memory systems; it shows in which ways and during which stages of the memory lifecycle memory can fail.

2.3.1 General memory

Schacter (1999) defines seven categories of failures to which memory in general is susceptible. He offers an explanation for the existence of these memory failures: They could be side-effects of adaptive processes that make perception and memory more efficient. The following describes each of these categories. The descriptions are all based on (Schacter, 1999).

Absent-mindedness

One reason for not being able to remember an experience or fact is that the original information was not completely encoded to memory. One reason for incomplete encoding is distraction and thus division of attention. A reason for such distraction can be that the task that is carried out is performed frequently and thus automatically.

Another reason for incomplete encoding is shallow, superficial processing. According to the levels-of-processing theory, shallow processing of items leads to a decreased ability to remember them later. An example for encouraging deep processing in studies that require participants to learn lists of words is to ask the participants to answer semantic questions about the words that are presented. Shallow processing, on the other hand, is encouraged when participants are asked about non-semantic characteristics of the words, for example whether they are presented in uppercase or lowercase.

A third reason for incomplete encoding is that people generalise their observations. This can lead to “change blindness” – not noticing changes in an observed scene. Using the schema-plus-tags model introduced in Section 2.1.3 as the basis for explanation, change blindness occurs when changes are consistent with the schema and the changed information was not explicitly encoded as tags.

Transience

Even memory items that are encoded deeply can decay over time. This process is called *transience*. Transience occurs both in short-term and in long-term memory. Typically, the rate of decay is high when the memory item is fresh and then slows down. Transience can be counteracted by rehearsal, i. e. repeated exposure to the information as described in Section 2.1. However, there is neurophysiological evidence that the initial encoding of a memory item has an influence on the likelihood for this memory item to be forgotten, which suggests that rehearsal may not always be effective.

There are two types of transience: one where the actual memory item is lost from memory, and one where the memory item is still in memory but not accessible any longer. Some cognitive psychologists believe that all transience is of the second type, but there is some evidence for transience of the first kind.

Blocking

Blocking is the temporary inability to access a memory item, usually accompanied by the strong conviction that this memory item is actually present in memory.

An example is the tip-of-the-tongue state, where one feels that a particular word is “almost” there. It comes in an “ugly sisters” variant in which one can recall a similar word but not the word which one wishes to recall. Both apply, for example, to recall of names. Other examples for blocking are temporary inabilities to recall a certain word or fact as experienced by students during an exam, or by actors on stage.

Misattribution

Misattribution occurs when a memory item is remembered but placed in a wrong context. Three types of misattribution are distinguished.

The first is characterised by correctly recalling a memory item but attributing it to an incorrect source, with a strong subjective feeling of remembering the memory item.

The second type is when this occur without the subjective feeling of remembering the item. This can lead, for example, to instances of unintended plagiarism when one perceives an idea as one's own original thought when in fact it has been heard elsewhere before.

The third type, also called confabulation, is the feeling of remembering something that in fact never occurred.

Suggestibility

Suggestibility leads to the feeling that an event or a detail of an event is being remembered, even though it actually never happened, due to suggestions by other people. This can concern episodic memory (for example, childhood or adult experiences) and semantic memory.

Bias

Bias causes distorted encoding or recall of memory items due to preexisting knowledge and beliefs. As described in Section 2.1.3 and in 2.2.2, semantic information such as schemata, scripts and concepts have a strong influence on encoding experiences to memory and on remembering. Bias occurs when such generic knowledge leads to *distorted* encoding or recall of memory items, for example leading one to "remember" schema-typical details of a particular event when this detail in fact deviated from the schema in this particular instance.

Bias can also be caused by one's mood or feelings, at the time of encoding or during remembering. There is some evidence that consistency plays a role here; people remember their own past opinions and attitudes as more similar to their current opinions and attitudes, especially when they believe that their opinions and attitudes did not change over time.

Persistence

The six categories described so far all concern memory failures in which a memory is not remembered, or remembered incorrectly. Persistence is the reverse type of memory failure: the inability to forget memory items that one would prefer to forget. Examples are fears, phobias and memories of disturbing or traumatic events. One variant of persistence is the overpresence of negative autobiographical memories compared to positive autobiographical memories, especially in people with a negative self-image.

2.3.2 Autobiographical memory

In addition to the generic memory failures described above, there are a number of memory failures that are specific to autobiographical memory.

Brewer (1986, p. 35) presents evidence that people typically strongly believe their personal memories (i. e. those with imaginal representation, see Section 2.1.1) to be true. This increases the effects of misattribution but decreases suggestibility for this type of recollection. Recollections of autobiographical facts (i. e. those with nonimaginal representation) are not accompanied to this extent by a strong feeling of remembering. Thus, it is easier to convince a person that their recollection of an autobiographical fact is inaccurate. Because of these effects, Brewer cautions against using the degree to which a person is convinced about the accuracy of their recollections in determining their actual accuracy.

One of the reasons for bias, describe above, includes one's mood or feelings as well as a desire for consistency. This particularly applies to autobiographical memory. For the same reason, pleasant events are remembered better than unpleasant ones. According to Linton (1986, p. 60), this explains why free recall tends to produce more positive than negative memories.

When people are asked to recall memories from their lives, there is a distinctive pattern to the distribution of the number of memories recalled per decade of the person's life (Conway and Rubin, 1993, p. 114ff): Almost no memories are recalled from very early childhood; there is a steep incline in the number of memories recalled for the late teens and early twenties; from then on, there is a decline in the number of memories that can be approximated by a power law; this is followed by a very steep increase in the number of memories recalled (as much as 50% of all recalled memories) from the decade prior to the study. This pattern is called the reminiscence bump. It has been observed for autobiographical memories in people over 40 years of age, and for significant memories in younger people as well.

To summarise, this means that there are two main issues concerning recall of autobiographical memories that need to be taken into account. Firstly, the degree to which a person is convinced about the accuracy of an autobiographical recollection (often called veridicality in Cognitive Psychology research) cannot be taken as a measure for the actual accuracy of the recollection. Secondly, there are biases towards certain types of memories: at least in free recall, people tend to recall more positive than negative memories and more memories from certain times of their lives.

2.4 Discussion

This section discusses the implication of the research summarised in this chapter, both for this thesis and beyond. It reviews the most important terms and concepts and re-states the hypothesis of this thesis using these concepts. It briefly outlines design guidelines proposed by others for interactive systems that aim to help people remember. It then makes recommendations for such that are derived from the Cognitive Psychology research introduced above; they incorporate and extend the design guidelines proposed by others.

2.4.1 Definition of terms

The objective of this thesis was stated in Section 1.1.1 as to create an interactive software system that *helps people remember past experiences and related information*. Based on the terms and concepts from Cognitive Psychology introduced in this chapter, the components of the thesis objective can now be phrased more formally and with more detail:

Past experiences: Autobiographical memory is the memory for past experiences. It is part of declarative memory which in turn is part of long-term memory. It is a type of retrospective memory. Autobiographical has a specific structure in which events belong to lifetime periods.

Related information: The type of autobiographical information that best describes this part of the thesis objective is event-specific knowledge, specific individual pieces of information that may or may not have a link to the self and that may or may not be remembered in sensory form.

Remember: Remembering can take one of three forms: re-living the original experience, knowing about the experience or reconstructing a likely version of the experience. Contextual cues, semantic information and associations play a big role during recall of memories. Contextual cues make it easier to recall experiences made in a similar context. Semantic information allows inferences based on generalised experiences. Memory items can be connected through associations, for example when they share characteristics. Remembering can be spontaneous or conscious and deliberate.

Help: Most memory failures lead to experiences that cannot be remembered or that are remembered incorrectly. The points of failure can lie either at the time of encoding or at the time of recall. Especially with autobiographical memory, some of the memory failures actually serve important functions.

A Computer Science approach to meet the thesis objective involves the creation of an interactive software system. Such interactive software systems that help individuals to remember past experiences and related information are called *augmented autobiographical memory systems* in this thesis.

2.4.2 Related design approaches

The strategy pursued in this thesis, using Psychology as a foundation on which to base Computer Science approaches for helping people remember, is finding increasing support in the community. Van den Hoven and Eggen (2007), Elsweiler et al. (2007) and Sellen and Whittaker (2010) all propose guidelines for the design of augmented autobiographical memory systems and similar systems. Their guidelines are based on surveys of Cognitive Psychology research that are similar to the survey in this chapter but have partially different results.

Van den Hoven and Eggen (2007) examine autobiographical memory. They take a narrower view than this thesis; in the terms used in this thesis, their recommendations focus mostly on past experiences and only to a lesser degree on related information. The focus of their research is on specific functions of autobiographical memory, namely to provide a person with a sense of self and on connecting with others by sharing experiences (see Section 2.1.2). Within this context, they specifically focus on the role of physical artefacts. This focus is reflected strongly in their recommendations. However, van den Hoven and Eggen (2009) could not confirm the superiority of tangible cues over other types of cues. Very brief textual cues led to the biggest amount of recalled memories in their study. Their recommendations suggests that their approach is to provide opportunities for semi-spontaneous recall and to reinforce memories by repetition. This is only one option in the spectrum of counteracting memory failures described in Section 2.3 and summarised in the previous section.

Van den Hoven and Eggen make no explicit distinction between the types of remembering. However, their descriptions make it clear that they mainly aim to support the system's user in re-living experiences. They do acknowledge that autobiographical memories can be reconstructed and that multiple reconstructions of the same memory can vary. For the function of autobiographical memory that their approach seeks to support, this is seen as beneficial. Similarly, their recommendations address the flexibility that is required in the system's interpretation of information related to memories.

Elsweiler et al. (2007) focus on the area of Personal Information Management (PIM; discussed in more detail in Section 3.3) and on the information items commonly found in PIM systems: photographs, e-mails, electronic calendar and contact data as well as other electronic documents. All this information is depersonalised; Elsweiler et al. advocate allowing the user of a PIM system to access information in the system based on memories they may have of the information and of the user's interaction with the information. Their focus is especially on memory lapses and the lessons that can be learned for PIM tool development from studying everyday memory lapses.

Their design principles for PIM tools suggest to include access to information items based on the context of the information items as well as on the context in which the user interacted with the items. Further, they suggest to allow "retrieval journeys" involving small steps based on partial recollections. They promote to show cues to the user during retrieval sessions that facilitate the retrieval of further information objects.

Sellen and Whittaker (2010) analyse research in the area of Continuous Archival and Retrieval of Personal Experiences (CARPE; discussed in more detail in Section 3.2). Their analysis is conducted on the background of Cognitive Psychology research. Their design guidelines show opportunities for a stronger incorporation of Cognitive Psychology research into systems in the CARPE area. Sellen and Whittaker advocate tailoring such a system to the memory failures that the system aims to counteract; to focus on retrieval cues rather than aim to capture copies of experiences; to tailor the system to the types of remembering that the system aims to support; and to build on the strengths of human memory by aiming to supporting rather than to substitute it.

2.4.3 Recommendations

The guidelines reviewed in the previous section form a starting point for designing autobiographical augmented memory systems. However, few of the guidelines can directly be translated into recommendations of factors that should be present in such a system; not all characteristics of memory for experiences and related information reviewed in this chapter are covered by the guidelines.

This section makes recommendations for augmented autobiographical memory systems based on the research in Cognitive Psychology summarised in this chapter that extend the guidelines summarised in the previous section.

Re-live, retrieve or reconstruct. Opinions differ among Cognitive Psychology researchers on whether recollections are copies of the original experience or whether they are reconstructed during the remembering process (Section 2.2.1). An in-between type of recollection is knowledge about the original experience that is not accompanied by sensory recollections. Memories that are reconstructed at recall time will not necessarily be reconstructed in the same manner every time they are recalled (see also Sections 2.2.2 and 2.3.1). This may or may not be desirable (see Section 2.1.2).

R1 When constructing an augmented autobiographical memory system, the decision has to be made whether the system should

1. provide cues that enable its user to re-live experiences, or
2. act as an objective repository of facts which the user can retrieve to verify their recollections, or
3. support the user in reconstructing an experience according to the user's current goals.

Van den Hoven and Eggen (2007) and Sellen and Whittaker (2010) similarly advocate basing the system design on the type or types of remembering to be supported. Reconstruction of experiences is mentioned by van den Hoven and Eggen and implicitly also by Sellen and Whittaker but not assigned much importance.

Support, not supplant. Clearly, an augmented autobiographical memory system cannot store the vivid, subjective memories involved in re-living.

It may be able to store pictures, sounds or even smells and other modalities, but it cannot hold the subjective connection with the user (the self-link in Brewer's classification, see Section 2.1.1). The reason is that the link between an individual's autobiographical memories and the individual's sense of self is intrinsic to the individual. It can, however, aim to store information that can then act as cues for the user's during the remembering process, regardless of the type of remembering that is being aimed for.

However, the distinction is less clear between cues that can be stored in such a system on one hand and the actual information to be remembered on the other for autobiographical facts and especially for depersonalised information items that are related to autobiographical events (see Sections 2.1.1 and 2.1.2).

R2 Augmented autobiographical memory systems must support, not supplant natural memory. Consequently, they must aim to store cues rather than memories or copies of experiences. Knowledge and facts related to experiences, as well as depersonalised information items, can be stored and serve as cues.

Sellen and Whittaker (2010) include a similar design guideline.

Experience-time or recall-time support. Memory failures can occur both when an experience is made and during recall of a memory (see Section 2.3). Both phases provide a range of opportunities for an augmented autobiographical memory system to support natural memory:

1. While an experience is being made, such a system can encourage its user to pay more attention to the experience or to engage with it more deeply. Information items related to an experience that do not have a self-link can be recorded by such a system in the case that the system's user does not encode this particular item to memory at all. Such a system can counteract the effect of absent-mindedness.
2. While the system's user is attempting to remember an experience or related information, the system can show cues to the user to trigger one of the types of remembering listed in recommendation R1. It

can provide the user with options to navigate among and view the memory items stored in the system. Such a system can counteract the effects of transience, blocking, misattribution, suggestibility and bias.

Even when the system's user is not actively attempting to remember an experience or related information, the system can show stored memory items to the user. This allows the user to engage with the memory items, which helps to fixate the user's natural memories and consequently counteracts the effect of transience.

R3 When constructing an augmented autobiographical memory system, the decision has to be made whether the system should support natural memory during the experience, i. e. at encoding time, or when remembering, i. e. at recall time. Another option for the system is to expose its user to their memories even when the user is not explicitly trying to remember.

Sellen and Whittaker (2010) and to a lesser degree also van den Hoven and Eggen (2007) similarly advocate choosing the types of memory failures to counteract with an augmented autobiographical memory system.

Memories use context. Recall of a particular memory item is better the more overlap there is between the information available at encoding time and the information available at recall time (see Section 2.2.2). This information is called the context of an experience. Context encompasses both external state (such as location, other people nearby, weather conditions) and internal state (such as mood). Certain types of context are also used consciously during the recall process (see Section 2.2.2): The context of experiences, such as time, location and other people present, are used to pivot between groups of recollections. The presence of event-specific knowledge can serve as a verification criterion in retrieval mode; aspects of an experience's context can turn into event-specific knowledge. The time and place of experiences can be used in inferences made while reconstructing experiences.

R4 Augmented autobiographical memory systems should make use of context in storing and retrieving information.

Context could help in an augmented autobiographical memory system in two ways. The system could allow traversal of information items in the system along contextual dimensions, such as finding all other information items that share context with a given information item or finding information items by contextual aspects. The system could also allow access to an information item's context.

The design guidelines by Elsweiler et al. (2007) and by Sellen and Whittaker (2010) advocate using context; types of context mentioned are time and place by the former and time, place, people and events by the latter. The only type of context promoted by van den Hoven and Eggen (2007) are tangible artefacts related to the experience.

Memories rely on semantic information. When similar information is encountered repeatedly, a generic memory of this information is formed (see Section 2.1.3). Sections 2.2.2 and 2.3.1 described how such generalisations, collectively called semantic information in this thesis, are used in structuring information and in reconstructing experiences. An experience may be stored in memory by reference to its type (i. e. an underlying generalisation) with the addition of information that deviates from or elaborates on the generalisation (schema-plus-tags model). Semantic information is used mainly in reconstructing experiences, for example to support set inferences. The type of an activity or event can also serve as pivot points in retrieval mode, similar to the way context can serve as pivot point.

R5 Augmented autobiographical memory systems should make use of semantic information in storing and retrieving information.

Semantic information could help in an augmented memory system in several ways. If the experiences to be represented in the system conform to types to at least some degree, the schema-plus-tags model could serve as a guide for how to store information in the system, at least on a conceptual level. The system could allow access to information items with the same type as a given information items and to information items of a given type. Finally, the system could allow access to information items based on typical characteristics of the represented information types and of the relationships between them.

None of the design guidelines reviewed in Section 2.4.2 advocate the use of semantic information. This may be because semantic information plays a bigger role in the reconstruction of experiences compared to other types of remembering and reconstructing is not emphasised in any of the guidelines.

Remembering follows associations. When remembering occurs as a conscious process of searching for a memory, there are a number of strategies that people typically employ (see Section 2.2.2). Most involve a series of small steps, navigating along a chain of associations. This may include, but is not limited to, backtracking and pivoting. While recommendation R4 refers to the context of an experience, these associations can be seen as the “context” of a memory item.

R6 Augmented autobiographical memory systems should provide means for navigation in small steps along associations and for retrieval journeys.

An augmented memory system could incorporate associations between experiences and facts by allowing arbitrary connections between information items in the system. The system could enable traversal of information items in the system along chains of connections.

The recommendation to incorporate associations in an augmented memory system is similar to a guideline proposed by Elsweiler et al. (2007).

2.5 Summary

This chapter presented an analysis of research in Cognitive Psychology that answers the first research question. The analysis was conducted for two reasons. The first reason was to clarify the objective of this thesis, “helping individuals remember past experiences and related information". The second reason was to derive requirements for an approach to achieve the objective.

The objective of this thesis was clarified by examining what central terms used in stating the original objective mean in terms of Cognitive Psychology. Section 2.4.1 provides definitions for "past experiences", "related information", "remembering" and "helping". A name was given to interactive software systems that fulfil the objective: augmented autobiographical memory systems.

The requirements were formulated as a set of design recommendations for augmented autobiographical memory systems. Section 2.4.2 briefly reviews design guidelines and principles proposed by others in related areas. Section 2.4.3 then makes six recommendations for such systems based on the analysis of related work in this chapter. The first three recommendations are more global in perspective and describe options for the design of augmented memory systems; they outline choices that have to be made when developing such a system. The remaining three recommendations paint a very high-level view of components that augmented autobiographical memory systems should have: support for context, support for semantic information and support for associations.

Together, these two parts of the chapter allow the working hypothesis of this thesis to be phrased as:

> *An interactive software system that combines the context of experiences, semantic information and associations is a suitable means to support individuals in reconstructing autobiographical memories and in retrieving cues and facts related to such memories.*

The phrasing of the hypothesis reflects a decision made for the research presented in this thesis with regards to the types of remembering that are addressed. The research presented in this thesis focuses on the second and the third type of remembering in recommendation R1: retrieving/knowing and reconstructing.

The next chapter analyses how these recommendations are realised in existing software systems for augmenting autobiographical memory and in related areas.

Chapter 3

Augmenting memory: a Computer Science perspective

This chapter addresses the second research question: How can an interactive software system help someone remember? The chapter contributes to answering this question by reviewing how others in the field of Computer Science have addressed this and similar problems.

This chapter is structured as follows. Section 3.1 explains the focus of this chapter, the scope of the analysis and the criteria used. Sections 3.2 and 3.3 then apply these criteria to approaches in two areas. Section 3.4 discusses the implications of the research summarised in this chapter for this thesis and for augmenting autobiographical memory in general. The chapter concludes with a summary in Section 3.5.

3.1 Focus

This chapter analyses Computer Science approaches that are related to the objective of this thesis. This section gives more details about the scope of the analysis, i. e. which areas are covered, and about the criteria used.

3.1.1 Scope

The approaches in this chapter fall into two main categories. The first category comprises approaches that share the objective of this thesis in supporting an individual in remembering past experiences and related information. Most of these approaches capture a person's everyday experiences and make them available for later retrieval. Generally, these approaches make use of mobile devices and employ various degrees of context-awareness. These approaches are analysed because they show how experiences, memory for experiences and recall of experiences as well as the three factors – context, semantic information and associations – are treated in interactive software systems.

The second category comprises approaches that are concerned with a specific form of autobiographical memory: an individual's memory for their interaction with personal information available in digital form. These approaches are analysed because they give further examples of the integration of the three factors in interactive software systems.

3.1.2 Criteria

The recommendations for augmented autobiographical memory systems made in Section 2.4.3 are used as criteria for the analysis. These are:

Remembering (R1)
: The type of remembering that is addressed by the system: Re-living, knowing (i. e. retrieval of information from the system) or reconstructing.

Information (R2)
: The type of information that the system aims to store: Cues or a copy of the experience.

Phase (R3)
: The phase of the remembering process that the system is aiming to support: Experiencing (i. e. encoding of an experience to memory), remembering or both.

Context (R4)
: Whether the system allows its user to use the experience's context when remembering an experience with the system's help.

Semantic Inf. (R5)
: Whether the system allows its user to use semantic information about experiences when remembering an experience with the system's help.

Associations (R6) Whether the system allows its user to use associations between memory items when remembering an experience with the system's help.

As explained in Section 2.5, the first three recommendations consist of groups of choices that must be made in the design of an augmented memory system. The analysis describes what choices were made, implicitly or explicitly, for the analysed systems. The remaining three recommendations, R4 through R6, describe factors that should be included in augmented autobiographical memory systems based on the Psychology research examined in the previous chapter. The analysis describes whether or not the factor was used in the analysed system.

3.2 Continuous archival and retrieval of personal experiences

This section gives an overview of approaches that capture a person's experiences for later retrieval. One name of this field is *Continuous Archival and Retrieval of Personal Experiences* (CARPE). As summarised by Gemmell and Sundaram (2006), a typical strategy in this field is to continuously capture certain types of data. The types of data that are captured vary among different systems. Commonly captured types are audio and video streams, but there are systems that capture data such as the user's location, proximity of other people or the user's interaction with software systems (e. g. web browsing history, files opened or modified). This captured data is assumed to equal or represent the user's experiences.

Continuous capture leads to large amounts of data, which makes retrieval difficult. Systems in this field typically cross-index the different types of data they capture. In addition, machine-learning techniques are employed to automatically label segments of data and identify particularly interesting moments.

A very early vision of an interactive system that augments its user's memory was described by Bush in 1945. He described how technology could support researchers in managing their documents, notes and other information. Many approaches described in this chapter draw their inspiration, explicitly or implicitly, from elements of Bush's vision (interpretations of which vary, see Veith, 2006).

One component of the "Memex", Bush's proposed system, was a storage device built into a researcher's desk that would hold all documents encountered by the researcher. The user of this system would be able to easily add new documents and retrieve those he or she has already seen. In addition to that, the system would allow for the creation of connections between documents. Another component of Bush's vision was a wearable device, wirelessly connected to the main system. This device would record photographs, voice comments and timestamps while the researcher is working in the field or in the laboratory.

Bush's vision was fueled by technological advances made in communication and office technologies at the beginning of the 20th century. While it ostensibly draws on characteristics of human memory, these are presented in anecdotal form only. Technological advances, rather than a solid foundation in Psychology research, similarly were a driving factor when the CARPE area emerged from the area of Wearable Computing in the early 1990s (Weiser, 1991; Norman, 1992).

3.2.1 Systems

This section describes a selection of CARPE systems using the criteria stated in Section 3.1.2. A summary is given in Table 3.1. Most surveyed CARPE systems aim to support retrieval of information at recall time. Their goal is generally to capture copies of experiences rather than cues, ignoring fundamental issues with this approach (see recommendation R2). The context of an experience is used by almost all systems in this area, while semantic information and associations are supported by only a few systems, and in most cases only partially.

Note that this section is not a complete survey of the CARPE field. A more comprehensive survey was conducted by Truong and Hayes (2009); it focused on applications in the workplace, in educational settings and for personal uses as well as on common capture techniques. This section also excludes approaches that specifically aim to help people re-live

	R1			R2		R3					
System	re-live	retrieve	reconstr.	cue	copy	exp.	in-betw.	recall	**R4**	**R5**	**R6**
Forget-me-not		•		[•]	[•]	[•]		•	✓	[✓]	
Jimminy		•		•	•		•	[•]	✓		[✓]
Conf. Assistant		•		•	•	•		•	✓		[✓]
Familiar/inSense		•			•			[•]	[✓]	[✓]	
eyeBlog		•			•			•	✓		
Life logs		•			•			•	✓		
iRemember		•		•	•			•	✓		
MyLifeBits		•			•		[•]	•	✓	[✓]	[✓]
ButterflyNet		•		•	[•]	[•]		•	✓		✓
Affective Diary	[•]		[•]					[•]	[✓]		

Table 3.1. Comparison of CARPE systems. The criteria are explained in Section 3.1.2. R1: targeted type of remembering, R2: nature of stored information, R3: targeted phase of memory process, R4: use of context, R5: use of semantic information, R6: use of associations. Symbols used: • and ✓ – supported, [•] and [✓] – partially supported, blank cell – not supported.

their own past. Some of these approaches are entirely software-based (Peesapati et al., 2010) while others combine software systems with physical artefacts (van den Hoven, 2004; Petrelli and Whittaker, 2010; Petrelli et al., 2009). These approaches are not further described here because support for re-living is not part of the objective of this thesis.

Studies into the usefulness of approaches in this field are summarised in the next section.

Forget-me-not

Forget-me-not (Lamming and Flynn, 1994; Lamming et al., 1994) was an early wearable memory aid for workplace-related information. A prototype implementation ran on PDA-style devices. Forget-me-not continuously captures data about its user's context and makes it available for querying using a graphical command language. Types of data captured in the prototype include the user's location; encounters with other people (who would wear another Forget-me-not device so that the devices can exchange identifying information); activities at the user's workstation (mainly e-mail sent and received, but also files opened/modified and programs started); files shared and printed; and phone calls received and made by the user. The conceptual description lists further types that apparently were never implemented (Lamming and Flynn, 1994).

Captured data is associated with a timestamp, with events occurring within a short time frame being grouped together into "episodes" (Lamming and Flynn, 1994). The wearable device shows these episodes in temporal order, using graphical icons to represent different types of activity. The list of episodes can be filtered by specifying sets of icons that must occur in an episode for it to be shown. Filtering allows for quite complex interactions, for example: "[Mike] is trying to locate a document he passed to Marge during a meeting involving Grouch, Peter and Professor [. . .] Mike also remembers that the meeting was held in his office" (Lamming and Flynn, 1994, p. 7; icons omitted from quote).

The Forget-me-not project explicitly refers to autobiographical and episodic memory; diary studies were conducted to investigate types of memory failures and the requirements analysis was linked to findings in Cognitive Psychology (Lamming et al., 1994).

Jimminy, the Wearable Remembrance Agent

Jimminy, also called the *Wearable Remembrance Agent*, was developed to explore application areas for wearable devices (Rhodes, 1997, 2003). It is the wearable version of a desktop-based application, the Remembrance Agent (Rhodes and Starner, 1996). Memory theory research is briefly referenced as the motivation for choosing this application area (Rhodes, 1997).

Jimminy runs on a wearable device that includes a one-hand chorded keyboard and a heads-up display, making it much closer to a true wearable system than Forget-me-not. It runs in the background in a word processor; the user's interaction with Jimminy is entirely text-based. When the user types words, similar documents from the user's personal archive are suggested. Additionally, suggestions can be triggered when there is a change in the user's context or when a scheduled event is drawing near.

Jimminy's personal archive comprises typed notes, e-mails and other electronic documents written or read by the user. In addition, Jimminy uses its user's current location and other people present as the basis for its suggestions. All notes that are created within the system are annotated with the user's physical location at the time they created the note.

Jimminy is intended as a general-purpose memory aid, without a specific application area in mind. However, the technology used to automatically determine the user's location and people nearby depends on specialised hardware in the environment. These features can only be used where this hardware is present; otherwise, these types of context have to be typed in by the user.

The system does not allow the user to query the database; instead, implicit retrieval is used throughout where suggestions are shown that are similar to the current text within the word processor and the user's current context.

Conference Assistant

The *Conference* Assistant was developed as an example for a context-aware application on a mobile or wearable device (Dey et al., 1999). It is intended to be used while and after the user attends an academic conference. More specifically, it can be used "to help users decide which activities to attend, to provide awareness of the activities of colleagues, to enhance interactions between users and the environment, to assist users in taking notes on presentations and to aid in the retrieval of conference information after the conference concludes" (Dey et al., 1999, Sect. 2.1).

The Conference Assistant receives some information about the conference from the conference organisers, such as the conference schedule and details of other conference attendees. It is aware of the user's location and the current time as well as the location of (previously specified) colleagues attending the same conference and these colleagues' interest in the presentation they are currently attending, if applicable. If the user is attending a presentation, the system also has access to details of the presentation being attended such as the presenter's details and the currently displayed web page or presentation slide. The user can attach notes to individual web pages or slides. After the conference, the system also has access to audio/video recordings of presentations where available.

The system's user interface shows the conference schedule and a timeline. Both are annotated with events and further information that were captured during the conference. The user can view their notes and other information as well as replay audio or video recordings, if available.

The Familiar and inSense

The *Familiar* was developed to semi-automatically construct a partial diary of its user's activities based on sensor data (Clarkson et al., 2001). The Familiar's user would wear a video camera, a microphone and movement sensors that continuously collect data. Machine learning algorithms are used to extract a number of higher-level features from the sensor data; these features are based on the recognition of faces, speech and gestures. Data and extracted information are clustered hierarchically into "more complicated scenes such as shopping for groceries, being at home, and going to work" (Clarkson et al., 2001, Sect. "Preliminary Results").

The Familiar never made it past the prototype stage and does not appear to have had a user interface to retrieve any of the data and extracted information. It was planned to construct a diary of the user's activities that the user would then be able to annotate further.

A later project, *inSense*, combines the capturing stage with the recognition and mining stages to detect situations that should be captured (Blum et al., 2006). This project captures audio data as well as photographs or video. Capture is triggered when a point of interest was identified through changes in location, posture, activity (a higher-level concept derived from the previous two) and through presence and type of speech.

The Familiar and inSense projects are included in this survey because they are examples of the sub-field of CARPE that focuses on extracting higher-level information from data collected by wearable sensors. This field, sometimes called "Reality Mining", is not directly concerned with helping the user to remember, but results from research in this field have an impact on data collection by wearable memory aids.

eyeBlog

The *eyeBlog* system posts video captured by a wearable device to a weblog (Dickie et al., 2004). Video capture is triggered manually when the user presses a button or automatically. The video feed is monitored for eye contact with a subject in the video and for the occurrence of special glyphs that can be attached to arbitrary objects. Both situations will cause video to be recorded.

Recorded video is posted to a weblog for annotation by the user. The user can then use the annotations or the weblog's facilities for temporal navigation to retrieve captured data.

Life Logs

The *Life Logs* project uses a wearable device to record a video feed and context data from various sources that can then be used to find specific parts of the video (Aizawa et al., 2004a,b). Context data can come from various sensors such as those for the user's location (both raw GPS data and human-readable addresses resolved via a geocoding service), brainwave activity and motion. Other types of context data used in the project are the weather at the user's location, news headlines on the day, web pages visited and e-mails sent and received by the user, as well as documents viewed by the user on their desktop computer. Face-detection algorithms are also run on the video data to detect presence of other people. The user can further annotate scenes in the video using keywords.

All types of context data can be searched or browsed to find a specific part of the recorded video. Aizawa et al. (2004b) acknowledge that retrieval poses the biggest challenge for their system. Their approach is to limit the recording of video to scenes that are likely to be of later interest. They experimented with introducing context-based triggers for video recording such as sampling when the user changes their speed or the direction in which they are moving, or when a face or conversation is detected.

Key ideas from the Life Log project were later incorporated into a “ubiquitous home environment” where video and context were not only captured via wearable devices but also via devices placed in the environment (de Silva et al., 2006).

iRemember

iRemember is an audio-based memory aid (Vemuri et al., 2004; Vemuri and Bender, 2004; Vemuri et al., 2006). It runs on a mobile device and records audio at the user’s request. Audio is then transcribed to text using speech recognition technology. The user can search within the transcribed audio and also navigate it using a timeline. The speech recognition technology used by iRemember is not 100% accurate; text is shown in different shades in iRemember’s user interface, with the shade of each word corresponding to the system’s confidence that the word has been recognised correctly. The search employed by iRemember is text-based but also includes a component that matches words if they sound similar to the query term.

iRemember was developed to overcome memory problems, in particular transience and blocking (see Section 2.3.1). Its goal is not necessarily to contain any information that the user might be looking for, but to enable the user to find triggers that will then lead to a recollection of the required information. The iRemember system has been evaluated, with encouraging results for this approach.

MyLifeBits

MyLifeBits is a system that integrates capture of data using a wearable device with data collected on the user’s computer(s) and makes it available for later retrieval (Gemmell et al., 2002, 2006). It is likely the longest-running and best-publicised system in the CARPE area. Originally called CyberAll (Bell, 2001), it set out to explore possibilities and challenges that arise from collecting as much digital information as possible about a person's life.

Capture of information in MyLifeBits is done via wearable devices such as SenseCam (Hodges et al., 2006). Examples for captured data used by MyLifeBits are automatically taken photographs and video recordings and GPS data. Various data from the user's desktop computer is also included in the collection. Examples are document files, e-mails and instant messaging logs. Additional information that is collected includes incoming and outgoing phone calls as well as radio and television programmes watched by the user. A relational database is used for storage. Documents can be annotated manually with textual and voice comments. The user can create links between items, such as connecting a photograph with the place it was taken and with the people who appear in it, that supplement automatically captured data.

MyLifeBits allows for several types of retrieval. An early, manual "story generation mode" was replaced by querying based on location and time (Gemmell et al., 2005). Results of a query can be visualised with timelines or as a list. A screensaver that displays random photographs encourages serendipitous encounter of information; the screensaver allows the user to annotate the material shown for retrieval in the future.

ButterflyNet

The *ButterflyNet* project explores the integration of paper notebooks with electronic data, both captured automatically and collected deliberately, for biology fieldwork (Yeh et al., 2006). It allows biologists to combine information from paper notebooks with photographs, GPS data and data from environmental sensors as well as with physical specimens. ButterflyNet creates cross-associations of information from different sources automatically using timestamps. The user can also directly connect information by drawing certain placeholder gestures into their paper notebooks with a digitising pen and by using 2-dimensional bar codes in photographs and on envelopes holding physical specimen. The user can flip through a digitised version of their scanned paper notebooks in which each placeholder gesture is replaced by a visualisation of the information associated with this placeholder. ButterflyNet also provides a time-based view of the collected information.

The Affective Diary

The *Affective Diary* is similar to most other systems presented in this section in that it captures information from wearable sensors and the user's interaction with their mobile phone (Ståhl et al., 2009). However, the Affective Diary project takes a unique approach in that its goal is to provide its user with the means to analyse and reflect on their experiences, rather than providing an objective account. It intentionally visualises data in an ambiguous way to allow for and stimulate the user's interpretation.

Sensor data from a wearable device, namely movement and physiological arousal level, is combined with information from the user's mobile phone – text messages received and sent by the user, photographs taken and presence of Bluetooth-enabled devices in the vicinity.

A visualisation of the captured data can be viewed on a tablet PC. The user's activity is presented as an anthropomorphic figure along a timeline, with the figure's posture and colour at a given point corresponding to the amount of movement and the arousal level of the user. Cell phone activity is overlaid on this timeline and the user can add further annotations. The posture and colour of the figure can also be changed by the user.

3.2.2 User studies

Most of the systems described above explicitly aim to supplement their user's memory. The approaches taken by the systems vary greatly. This section examines which evidence there is that these systems meet their aim and also whether there is evidence that favours certain approaches over others.

Evaluation of systems in the CARPE area is in most cases restricted to tests of functionality and usability, answering questions such as "can the system do what it promises to do" and "how easy is the system to use".

Few evaluations focus on general desirability of such a system by end-users or on effectiveness; exceptions are the evaluations of iRemember by Vemuri et al. (2006) and of ButterflyNet by Yeh et al. (2006).

Kalnikaitė and Whittaker (2007, 2008) as well as Sellen et al. (2007) conducted studies with a more general focus to investigate under what circumstances and to what extent CARPE systems actually help their user's autobiographical memory. The results of their studies challenge basic assumptions made by CARPE researchers. The issues discovered lie both with the desirability of such systems – whether people actually like and would use CARPE systems – and with the effectiveness of such systems – whether these systems actually help people remember.

Kalnikaitė and Whittaker (2007) compared unaided memory with pen-and-paper notes, a dictaphone and a system that correlates written notes with audio recordings. They found that in some cases, people prefer to use their own memory rather than a CARPE-style system, even if they rate their own memory as less accurate than the system. This occurs when they expect that using their own memory will be faster than using a memory prosthesis or when they feel that their own memory is sufficiently accurate. They found complex metacognitive processes at work, where people were good at knowing whether they knew something; this corresponds to the recall strategies using meta-knowledge described in Section 2.2.2. Their findings suggest that the strong focus on capture technology in the CARPE area, to the detriment of research into visualisation and retrieval, is particularly unfortunate.

The same researchers then explored the usefulness of hand-written notes for triggering recollection (Kalnikaitė and Whittaker, 2008). They found that taking notes generally increases recollection even if the notes are never referred to again. Based on these findings, they challenge two assumptions prevalent in the CARPE area: that automatic capture is superior to manual capture and that note-taking and other explicit interaction with a memory aid system must necessarily distract from the experience in all cases, leading to worse encoding and later recall of the experience.

Sellen et al. (2007) are among the few researchers in this area to address the fundamental difference between re-living (called “remembering” in their work) and knowing about one’s past – the same distinction made in this thesis between these types of remembering, even though reconstruction of experiences is still not considered in their work. They argue that CARPE systems confuse these two different issues and claim to capture experiences while in fact the systems merely capture data that may then act as cues to trigger true remembering. Based on the findings of their study, they speculate that passively captured data (in their case, photographs) does aid in remembering events but that this effect is still subject to forgetting over time.

In a further study, Kalnikaitė et al. (2010) investigated the usefulness of photographs and location information for triggering recollection. Location data was visualised as a path on a map. Their findings suggest that images are helpful to trigger re-living, while location data mostly assists with metacognitive and inferential reconstruction processes.

3.3 Personal Information Management

Personal Information Management (PIM) is the area of Computer Science that is concerned with “both the practice and the study of the activities people perform to acquire, organize, maintain, retrieve, use and control the distribution of information items [. . .] for every day use to complete tasks [. . .] and to fulfill a person’s various roles” (Jones and Teevan, 2007a, p. 3). This section focuses on a selected part of PIM research; a more comprehensive overview was published by Jones and Teevan (2007b).

The previous section described approaches and systems that aim to help their user to recall their own experiences. These experiences are mostly genuine real-world experiences; some systems also consider interactions with digital artefacts on a computer. This section describes approaches and systems that mostly deal with the organisation of digital artefacts on a computer by individuals.

	R1			R2		R3					
System	re-live	retrieve	reconstr.	cue	copy	exp.	in-betw.	recall	**R4**	**R5**	**R6**
Stuff I've Seen		•			•			•	[✓]		
Phlat		•			•	[•]		•		✓	
LifeStreams		•			•			•	✓		✓
PCT		•			•	[•]		•	✓	[✓]	
eVITAe		•			•	[•]		•	✓		
TimeSpace		•			•			•	✓		[✓]
PhotoMemory		•			•			•	✓	[✓]	✓
Jourknow		•			•			•	✓	[✓]	
Haystack		•			•			•		✓	✓
SemanticLIFE		•			•	[•]		•	✓	✓	
Semex		•			•			•		✓	✓
iMemex		•			•	[•]		•		[✓]	✓
Gnowsis		•			•	[•]		•		✓	✓

Table 3.2. Comparison of PIM systems. The criteria are explained in Section 3.1.2. R1: targeted type of remembering, R2: nature of stored information, R3: targeted phase of memory process, R4: use of context, R5: use of semantic information, R6: use of associations. Symbols used: • and ✓ – supported, [•] and [✓] – partially supported, blank cell – not supported.

3.3.1 Systems

This section describes a selection of PIM systems using the criteria stated in Section 3.1.2. A summary is given in Table 3.2. The systems are grouped by their main interaction factor – re-finding, context and semantic information. Most surveyed PIM systems aim to support retrieval of information at recall time, though some systems also include partial support at experiencing time. All systems seek to store copies of the original information; in contrast to CARPE systems, this does not conflict with recommendation R2 because information in PIM systems does not directly represent experiences.

The systems described in this section are a selection of PIM systems which are related to the research described in this thesis. Three main strands of systems are considered. The first comprises systems that address re-finding of information already known to the user at some earlier point. The second comprises systems that take into account the user's context in some form. The third comprises systems that include semantic information about items in the system.

Some PIM systems specifically aim to support prospective memory (as introduced in Section 2.1.4) by providing time-based or event-based reminding functionality. Since this thesis focuses on supporting retrospective memory, this type of PIM system is not included in this survey. A survey of PIM systems conducted by Jones (2007b) employs a more general view.

Stuff I've Seen

Stuff I've Seen is a search engine for previously accessed information such as document files, e-mails, calendar items and web browsing history (Dumais et al., 2003). The system was created as a research tool to explore possibilities for personal search engines (as compared to general Internet search engines).

Stuff I've Seen allows the user to conduct textual searches and to narrow down the result sets using a range of other criteria, such as the type of a document or its age. Search results are shown in a rich list view that includes a preview of the item (pictures for graphics files and presentation slides, text snippets for textual documents). The results list can be ordered by rank or by date.

A timeline view for information indexed by Stuff I've Seen is also available (Ringel et al., 2003; Cutrell et al., 2006a). Called Memory Landmarks, it shows timestamps as well as "landmarks", semantically labelled events. Landmarks are generated from the user's personal events, including calendar appointments and photographs, as well as from public holidays and important world events.

Phlat

Phlat is a system for searching and browsing personal information such as document files, e-mails and multimedia files (Cutrell et al., 2006b). It builds on experiences gained with the Stuff I've Seen system described above (Cutrell, 2006). Phlat indexes content on the user's personal computer and makes it available for retrieval. The user can locate information items using combination of keyword search, typed search and faceted browsing. The user can search by date, path, person, information type and tag. Previously executed queries are stored and can be re-executed on request.

One design consideration in the development of Phlat was to allow iterative refinement of both keyword and metadata-based queries. The goal was to blur the boundaries between searching and browsing. Further blurring the boundaries between different modes of interaction, Phlat allows the user not only to query by tag but also to apply tags to their personal collection (Cutrell, 2006).

LifeStreams

LifeStreams is a timeline-based system for managing electronic documents (Freeman and Fertig, 1995). Temporal organisation of documents was chosen as "an alternative for the desktop metaphor" (Fertig et al., 1996b) based on early research into people's ways of organising (Fertig et al., 1996a). A prototype implementation supported text files, e-mails, calendar items and a few specialised document types such as timesheets and stock reports, but the concept extends to any electronic document. All types of documents are visualised as a stack, with older documents being drawn further in the background. Documents currently being edited are shown in a special place to the side of the stack. A further organisation tool are so-called *streams*, which are (potentially nested) partial views of the collection. Documents are added to streams either automatically by some criterion or manually.

Documents can be associated with a future time and date, for example to set a reminder. The user can change the viewing time to the past or the future to see an older state of their document collection or to see items scheduled for the future.

Personal Chronicling Tools

The Personal Chronicling Tools (PCT) project applies ideas from the fields of life-logging and personal information management to enterprise settings (Kim et al., 2004). In its conceptual form, the PCT capture the user's interaction with computers at work and their other work-related activities, as well as business events. The prototype that was implemented focused on the user's interaction with their work computer (document files, e-mails and web browsing history).

Part of the PCT is a dedicated button that is added to the window border of each application. The user can click this button to designate an "interesting moment" (Kim et al., 2004, p. 61). This causes the PCT to capture as much information as possible about the current state of this application window, including such information as the computer's IP address and a timestamp. Additionally, the user can manually annotate the snapshot further, for example by adding a free-text comment.

Retrieval methods available with the PCT are keyword search and filters based on certain criteria, for example the type of event, the event's age and its visibility (private, a specific part of the enterprise or the whole enterprise). All free-text comments are scanned for keywords which are then cross-linked to WordNet entries. This enables the PCT to provide search for synonyms of query terms.

The initial version of PCT ran on top of the Lotus Notes applications suite and its database. A later version, eChronicles, is more strongly based on events and uses a custom storage mechanism (Kim et al., 2005).

eVITAe

eVITAe is a tool to manage personal information in various media, such as photographs and videos but also document files (Singh et al., 2004). An extended version also includes support for voice annotations (Pinzon and Singh, 2005).

eVITAe was developed following the recommendations for experiential computing as introduced by Jain (2003). Consequently, its user interface consists of several different views on the data and interaction by the user with one of the views causes the other views to be adjusted as well. The views are a timeline, a map and a details view of the particular item being examined. Metadata, such as temporal and spatial data, can be extracted from various file types or added manually by the user.

Experiential computing advocates keeping the user interface for performing queries identical to the user interface for showing results, and to reveal details of a user interface item on request. eVITAe achieves this by letting the user find information according to time and location by zooming in on relevant parts of the timeline or map. The extended version also supports textual search that can be performed on audio comments that have been automatically transcribed to text.

TimeSpace

TimeSpace organises the user's document files along a timeline (Krishnan and Jones, 2005). Files are also grouped by activity, where the user defines which files belong to a particular activity. A file can belong to more than one group. TimeSpace also allows the user to arrange their files spatially and maintains this visualisation when showing documents along a timeline.

PhotoMemory

PhotoMemory is a tool to manage a user's personal collection of photographs (Elsweiler et al., 2005). It is a proof-of-concept prototype to explore issues around PIM systems that allow retrieval using multiple types of context. PhotoMemory narrowly focuses on photographs because these are linked quite strongly to autobiographical memory, allow for richer interaction than text and at the same time can be classified automatically (Elsweiler et al., 2005, p. 3).

In PhotoMemory, photographs can be annotated with free-text descriptions and concepts. Images can also be grouped semantically. Annotation can happen at any time, for example when images are added to the system or at retrieval time. During retrieval, PhotoMemory allows the user to perform keyword search and also to specify several different types of context: by time; by location of the image on the screen (where the system has been designed such that the location of a given image on the screen stays constant during retrieval); and by semantic group.

Jourknow

The Jourknow system addresses management of "personal information scraps" – short personal notes and reminders that typically are not integrated into PIM systems (Bernstein et al., 2008a,b). Based on the assumption that traditional PIM tools are too heavyweight to manage this type of information, Jourknow aims to simplify entry of information scraps as much as possible. It uses a text-based language for data entry. Jourknow consists of a desktop program and a mobile phone application, JourMini. Both clients synchronise stored information scraps.

Information scraps entered into Jourknow are automatically annotated with some contextual information, such as a timestamp, the user's current location, running applications (in the case of the desktop client) and music playing on the user's computer. The user can specify additional context, including the type of information item (e. g. todo list item or reminder) and semantic tags (Van Kleek and Shrobe, 2007). These annotations are stored into an RDF-based data model. All types of context information are available for faceted browsing at retrieval time.

Haystack

Haystack is a metadata-based personal information repository (Karger et al., 2005). It visualises and makes available for retrieval classical PIM information such as the user's document files, e-mails and calendar items. A strong emphasis is placed on connections between pieces of information.

Haystack's user interface consists of recursively rendered components that are arranged based on the structure of information shown. The user can drag and drop components and invoke context menus on components. The user can also use textual search to locate information in Haystack. Several different types of main view are supported, for example a tabular view (e. g. for viewing the user's e-mail inbox) and a calendar view. With views and actions being highly customisable, Haystack's user interface is quite complex, a fact acknowledged by its creators. Interface definitions can be saved and shared with other users; it was anticipated that power-users would heavily customise the interface of their own Haystack installation and then share their customisations with other users.

Haystack's underlying data model uses RDF as a realisation of the original conceptual model (Adar et al., 1999; Huynh et al., 2002). Haystack makes some use of type information but does not enforce strict adherence a some schema. For example, the type of a piece of information determines which component will be used to render it – an "e-mail" component might contain components for the sender, the title (subject), the time at which the e-mail was sent and the body. However, a tabular view with columns for the sender, title and time can also hold an address book entry. In this case, the sender and time cells of this entry would be empty, with the contact's name being shown in the title column.

SemanticLIFE

SemanticLIFE is a personal information management platform that uses semantic annotations for organising the user's content (Ahmed et al., 2004; Mustofa and Tjoa, 2006; Latif and Tjoa, 2006). Typical PIM information (such as document files, e-mails, calendar items and browsing history) is automatically collected from various sources. Some information is annotated with semantic information automatically, but most annotation is performed by the user. SemanticLIFE uses RDF as its underlying data model, combined with ontological information. The user can then make use of the semantic annotations while retrieving information.

One particular focus of SemanticLIFE is the organisation of the user's collection of photographs (Latif et al., 2006). Items in the collection can be annotated with various types of context. Some context is automatically determined, such as the photograph's creation time and date, and some is manually specified, such as tags, people visible in the photograph and the occasion at which the photograph was taken. The user can designate certain photographs to be "landmarks"; these are then shown more prominently to help the user find photographs that were taken close (in time) to a landmark photograph. A location-based view is also available, in which photographs are placed on a map according to their location annotation.

Semex

Semex is a database for personal information items and their connections (Dong et al., 2004; Cai et al., 2005). It extracts information from the user's document files, e-mails, address books and other personal information repositories. Information from external sources can also be integrated into Semex, either temporarily or permanently.

All information within Semex can be queried by the user. Semex offers keyword search as well as typed search that performs keyword search on information items belonging to specified classes only. Results are visualised in a tree structure. The tree structure also allows the user to query along connection chains of length 2.

Semex uses a triplet-based data model and an ontology to keep track of the classes and association types. A particularly strong emphasis is placed on automatic discovery of entities and relationships and on reconciliation of multiple references to the same entity (Dong and Halevy, 2005).

iMemex

iMemex is a personal dataspace that lets the user find and annotate personal information (Dittrich et al., 2005; Blunschi et al., 2007; Dittrich et al., 2009). iMemex sits between the file system and the applications on the user's computer. It provides keyword search over the user's document files, e-mails and other electronic information items.

The user can annotate their information using *iTrails*, a language for specifying associations between sets and classes of information items (Dittrich et al., 2008). This allows the user to specify, for example, that queries executed on a given folder in the file system should also take into account e-mails in a given folder in the user's e-mail account. iTrails provide further query rewriting features such as cross-language search and user-defined shortcuts.

It is part of iMemex's design philosophy that the user can define annotations in a "pay-as-you-go" manner – at query time rather than up-front. These annotations are then added to iMemex and are used in future queries. This allows the information within iMemex to be transformed gradually into a semi-structured form, making it possible for iMemex to provide more tailored search result than standard desktop-search systems.

Gnowsis

Sauermann and Heim (2010) report on an evaluation of their Gnowsis tool. Gnowsis is a semantic desktop system that indexes information on the user's computer and allows the user to annotate files, e-mails, address book entries and web bookmarks with semantic concepts (Sauermann et al., 2006; Sauermann and Heim, 2010). It also contains a semantic wiki as well as an ontology-based instance browser and editor. Instances can be linked using relations.

3.3.2 User studies

The research area of PIM emerged in the 1980s as part of efforts to make office work more efficient (Jones and Teevan, 2007a). Even though some of the research in this area – just like CARPE research – is inspired by Bush's Memex (see Section 3.2), PIM research has traditionally been more orientated towards people and their behaviour rather than towards technology as its driving factor. Consequently, there are many end-user evaluations of PIM systems reported in literature. This section reviews such evaluations in three areas: *(a)* PIM, search and re-finding, *(b)* PIM and context and *(c)* semantic PIM.

PIM, search and re-finding

One approach to PIM is to "Search Everything" (Russell and Lawrence, 2007). Systems in this category index the user's personal information and let the user perform searches on this information. Keyword (textual) search, similar to that provided by web search engines, is usually supplemented by search options based on metadata such as the information item's type and creation date. Of the systems described above, Stuff I've Seen and Phlat fall into this category. Similar systems are part of, or available for, most desktop environments in modern operating systems; examples are Apple Spotlight, Google Desktop Search and Beagle[1].

It has been noted that user's search behaviour in their own personal information space differs from that in more general information spaces such as the WWW (Cutrell et al., 2006a). One major point of difference is the knowledge that a user has about the information, and its structure, in their own information space. Another is that users performing searches over their personal information typically wish to *find again* information that they have encountered previously, which means that both recall and recognition play a role in the re-finding process (Capra and Pérez-Quiñones, 2005). This is a connection to the type of information with which this thesis is concerned.

Teevan et al. (2004) studied people's search behaviour both over their personal information and on the Web and both for new and previously encountered information. They observed two distinct strategies that they call *orienteering* versus *teleporting*. Orienteering uses a sequence of small steps to navigate to a goal, using contextual clues at each step to determine the direction of the next. Teleporting behaviour jumps directly to a goal. An example for orienteering behaviour is to locate a phone number by opening an e-mail program, navigating to a specific folder, using keyword search to find all e-mails from a given sender and then opening from the search results an e-mail containing the phone number. An example for teleporting behaviour is to locate a phone number using a query such as "*name* phone number" with a desktop or web search engine.

[1] http://beagle-project.org

Teevan et al. (2004) found that their participants used orienteering rather than teleporting much more often than expected. Their explanation is two-fold. They speculate that, on one hand, teleporting (i. e. search engine) technology at the time of their study simply did not work well enough to make teleporting behaviour successful. On the other hand, they speculate that orienteering reduces cognitive load on the user by making available contextual clues at each step. They believe that even if the "perfect search engine" existed, some people would still prefer the orienteering approach regardless of the quality of search tools. Gwizdka (2006) takes a similar position, showing that the personal information space itself provides context that can help in retrieving information items from it.

Similar observations were made by Capra and Pérez-Quiñones (2003), who studied re-finding behaviour of web-based information. They found heavy use of *waypoints* in re-finding tasks. Waypoints are identifying characteristics of steps along the way in an iterative search process, such as the URL or title of a web page. They speculate that users remember more about their initial search process than about the actual path they took, and that the additional contextual hints available at each waypoint allows users to recall the next step they took when they initially found the information.

Bergman et al. (2008) studied whether navigational preferences of users had changed since the studies described above, given improvements in search technology. Their study examined strategies for locating files in the user's filespace. They found no significant changes in search behaviour, with the majority of participants still preferring to use folder-based, orienteering style re-finding behaviour over pure search. Similar to Capra and Pérez-Quiñones (2003), Bergman et al. (2008) note that memory of the filing process might help during retrieval; they speculate that the reason may be that the user can then rely on procedural in addition to declarative memory.

Jones (2007a) cites several studies into users' behaviour around organising and structuring their personal information space. These studies found that the structure of an information space can reflect the user's understanding of the area to which the information in the space is related. Structuring, organising and re-organising information items helps users with goals such as planning tasks and actions as well as with understanding a topic.

To summarise, the studies reviewed in this section show that the act of organising and structuring personal information items, when they are first encountered or subsequently, have benefits for the user beyond merely leading to a personal information collection that is more or better organised. This suggests that search by itself cannot replace the richness of other observed information organisation behaviour.

PIM and context

Barreau (1995) studied which types of context people use to acquire, classify and retrieve electronic documents in their personal information space at the workplace. Our discussion here focuses on her findings related to retrieval of information. To find a document, participants most often used the location of a document, its name or title (with the name often chose to reflect the topic or function of the document or the name of a person associated with it) or its date. Tools used to retrieve documents included directory listings sorted by name or date, browsing through their filespace, using the application that had been used to create the document (which shows files of this particular type stored in the application's default location) and, in one instance, a hard copy catalogue. Barreau (1995) explicitly notes that search was extremely rarely employed by the study participants.

Ringel et al. (2003) evaluated the Memory Landmarks add-on to the Stuff I've Seen system described above to determine the value of adding landmarks to a timeline visualisation of search results for e-mail messages. Landmarks are significant events derived from public holidays, news headlines, the user's photographs and appointments in the user's calendar. They compared retrieval times of their timeline+landmark visualisation to a pure timeline visualisation (i. e. with dates only). Their participants completed the set tasks significantly faster with the visualisation that included landmarks and reported that they liked the timeline+landmark visualisation of search results.

An evaluation of PhotoMemory showed that the availability of multiple types of context is useful. Elsweiler et al. (2005) observed that participants tended to stick to one type of context during a retrieval "burst" (part of a retrieval session) until either they were convinced that they could not complete a given task using this type of context or they saw a photograph that reminded them of a different type of context to try. This is consistent with findings in cognitive psychology about retrieval strategies, compare our summary in Section 2.2.2 of observations by Barsalou (1988).

Elsweiler et al. (2008, 2009) studied which types of context people remember about e-mail messages they are trying to re-find in their personal store and which types of recollection help in re-finding e-mail messages. They found that their study participants remembered a wide range of attributes about an e-mail they were looking for, including who sent it, why it was sent and what the e-mail was about. These three attributes were remembered in almost every case. Additional attributes that participants remembered were temporal information, information about other recipients of the e-mail and presence of attachment, images or links within the message.

Surprisingly, Elsweiler et al. (2009) found that retrieval speed did not necessarily improve with the number of attributes remembered about an e-mail. Remembering multiple attributes, and in particular the topic of an e-mail, seemed to slow down participants in re-finding the e-mail. Elsweiler et al. speculate, firstly, that people find it difficult to decide which recollection to choose as a starting point to re-find the e-mail; and secondly, that semantic information is less easily translated into retrieval strategies than other types of attributes. This would lead people to take longer to use their recollection of an e-mail's topic for retrieval. Temporal information, when remembered, had a particularly positive effect on retrieval speed. Recollection of the sender of an e-mail did not appear to have a noticeable effect on retrieval speed, though this may be an artefact introduced by the study itself.

One of the recommendations that Elsweiler et al. (2009) make based on their findings is improved support for semantic categorisation and retrieval (for example using faceted browsing) in future PIM tools.

Semantic PIM

The predominant approach in the sub-area of PIM that focuses on semantic information about items has been dubbed “Structure Everything” (Catarci et al., 2007) and “Unify Everything” (Karger, 2007). These names refer to two main research topics in the area: the derivation and addition of structure to originally unstructured data and the challenge to reconcile multiple references to potentially the same object.

Of the analysed systems, Haystack, SemanticLIFE, Semex, iMemex and Gnowsis fall into this category. End-user studies of these systems that focus on their usefulness and effectiveness have not been reported, with the exception of Gnowsis.

Gnowsis was evaluated in two studies (Sauermann and Heim, 2010). The first study involved eight participants and was conducted over two months. The second study was conducted as a longitudinal study over two years with two participants. Sauermann and Heim reported moderate uptake of the system: the participants used the wiki and made use of semantic annotation of files and other information items. Ontology-based retrieval was also used by the participants and rated as more important than textual search (although participants in an earlier, shorter study reported frequent use of gnowsis’ desktop search). The users made some slight customisations to the generic PIM ontology distributed with the system. The main factors that prevented further customisation were lack of expressiveness in the language used for the ontology (specifically, classes could not be linked with documents) and missing functionality in the user interface. Similarly, missing or faulty functionality of the user interface meant that the participants did not use some of Gnowsis’ features and completed tasks outside of Gnowsis that Sauermann and Heim (2010) expected could be completed using it.

Limitations in the evaluation method make it unclear what the implication of these studies are for the effectiveness of Gnowsis and of this sub-area of PIM in general (Sauermann and Heim, 2010).

3.4 Discussion

This section discusses the implication of the approaches analysed in this chapter for the research presented in this thesis. The section is organised along the recommendations for augmented autobiographical memory systems (Section 2.4.3). For each recommendation, relevant approaches in the surveyed research are reviewed along with insights into the effectiveness of the approaches where such insights are available.

3.4.1 Type of remembering

The survey in this chapter showed that most of the systems that were analysed do not distinguish between the memory types described in recommendation R1: re-living, knowing and reconstructing. A few systems explicitly aim to help with re-living; these were not described in detail because the focus of this thesis lies on knowing and reconstructing.

All systems that were analysed, with the exception of the Affective Diary, allow the user to perform straightforward retrieval of information using browsing by a variety of criteria, textual search and customised query techniques. In PIM systems, the actual information objects stored in the system are digital artefacts and superficially unrelated to experiences. Consequently, all of these systems can only aim to support retrieval.

Of the three types of remembering mentioned, reconstructive remembering is the only type that is not explicitly addressed by any system. This may be because the system cannot know which information about an experience is directly available to the user's natural memory at any given time and which information would need to be reconstructed. A comparison between photographs and visualised location information found that location information is more likely than photographs to aid with reconstructive processes. One explanation for this effect may be that the location visualisations did not contain as much information as the photographs, making reconstruction more likely to be necessary in the case of location visualisations. Another explanation may be that photographs are closer to an imaginal representation and thus closer to the representation of the experience in memory.

A very early system used inferences and other reconstructive strategies in the system itself to build up a knowledge repository about a person's experiences (Kolodner, 1983; Barsalou, 1988, p. 203ff). The goal of that system, however, was to mimic and model human memory processes rather than to support them.

3.4.2 Type of information

The lack of distinction between types of remembering observed in the previous section also leads to a lack of distinction between cues and experiences (R2). Most CARPE systems assume that the multimedia records that are stored are somehow identical with the experience. A wide range of different representations of stored data can still be observed, such as video streams, audio streams, photographs and textual representations.

PIM systems typically are able to store their main information objects directly because they mainly deal with information that is already in digital form. Again, a wide range of such information objects is found in PIM systems, such as e-mails, image files, calendar items, addressbook entries and various types of electronic documents.

3.4.3 Phase

The majority of systems surveyed aim to support their users during the remembering phase (R3). A few systems, more so in the area of PIM than in the area of CARPE, also support their users when the user first encounters a digital information item or makes an experience.

While most CARPE approaches claim that interaction with the system at this stage distracts from the experience, Section 3.2.2 reviewed user studies which found that support at the encountering stage in fact leads to improved remembering. Conversely, automatic capture of content without intervention by the user, as is done by most CARPE approaches, does in itself not improve remembering performance. Allowing the user to annotate or organise information in the system has several positive effects. It facilitates later retrieval because it leads to richer data in the system. It also helps the user to understand the information or encourages the user to engage with the information. Very few systems explicitly aim to improve remembering by repeatedly exposing their users to information in the system.

3.4.4 Context

The context of an experience (R4) is used in almost all systems that were analysed. Contextual information is often visualised in the surveyed systems and used to allow access to other information items. The types of contextual information used in the systems range from low-level sensor data, timestamps and geospatial location to high-level information about the user's activity.

The thesis objective includes support for remembering information related to experiences. This kind of information includes information about the context of an experience, leading to a dual role of contextual information both as "context" (i. e. less important data attached to the actual information items) and as first-class information items.

Contextual information is generally accepted to be useful, both for CARPE and for PIM systems. Automatic capture of context data is technically feasible and does not necessarily suffer from the same issues as automatic capture of content does.

3.4.5 Semantic information

Semantic information (R5) is used by approximately half of all surveyed systems; it is used much more frequently in PIM systems than in CARPE systems. The use of semantic information by reconstructive processes in unaided memory suggests that semantic information in augmented autobiographical memory systems may be able to play a similar role. However, almost all existing solutions treat semantic information merely as metadata. Given the rich inferential processes of memory, this suggests that the design space for semantic information in augmented autobiographical memory systems has not yet been exhaustively explored. The suggestion is supported by studies that showed decreased performance in retrieving personal information based on partially remembered content where such content was semantic in nature. This decreased performance is thought to be caused by difficulties for the user to translate such partially remembered semantic information into interaction strategies with the system.

The effectiveness of semantic information in augmenting autobiographical memory and managing personal information has not been studied extensively.

3.4.6 Associations

Associations (R6) can be interpreted as the context of a memory item, as opposed to the “context of an experience” as in recommendation R4 (see Section 3.4.4). However, these associations are rarely made explicit in the surveyed systems. Instead, the use of associations is mostly limited to the ability to pivot on an information item between types of context. A few systems show how associations can be taken advantage of to formulate queries along chains of connections, sometimes also in combination with semantic information.

Remembering follows chains of associations. If the context of an information item is visualised, steps in such retrieval journeys can be transformed from retrieval tasks to recognition tasks, which are generally seen as easier. The lack of support for this type of interaction in current systems, and consequently the lack of evaluations of its effectiveness, suggest that it is valuable to explore this factor further.

3.5 Summary

This chapter analysed Computer Science approaches related to the objective of this thesis. It showed that most systems which aim to support memory for experiences continuously and indiscriminately capture context and content data of various types; this is supposed to equal the capture and externalisation of experiences. Most of these systems incorporate context as a method to access stored information, although in general the retrieval side of systems in this area has not received much attention. Systems that aim to support the management of personal information that is already in digital form show how the three factors can be exploited to allow retrieval of such information.

A review of studies into the effectiveness of these approaches showed that no single approach achieves the goal of this thesis and that no single approach meets all recommendations made in the previous chapter. The analysis revealed elements onto which new approaches can build to fulfil the thesis objective; examples are automatic capture of the context of an experience and interaction techniques for semantic information. However, the analysis also showed that the roles of reconstruction, semantic information and associations are particularly poorly explored in existing systems.

The following chapters introduce a new approach to augmenting autobiographical memory that addresses the gaps in the surveyed approaches with regards to the recommendations and that builds on the discovered strengths of existing approaches.

Chapter 4

Conceptual design of an augmented memory system

This chapter introduces the conceptual design of an augmented memory system. The design is based on the results of the previous two chapters: on the recommendations for augmented autobiographical memory systems derived from the survey of Cognitive Psychology research and on the strengths and shortcomings of existing systems for personal memories and other personal information derived from the analysis of related research in Computer Science.

The chapter is structured as follows. Section 4.1 gives a high-level description of the focus employed in the design of the proposed augmented autobiographical memory system. Section 4.2 introduces the conceptual architecture and Section 4.3 states the requirements for systems implementing the architecture. Section 4.4 gives examples of how end-users can record and remember experiences using such an implementation of the design and Section 4.5 explores some user interface design ideas in more detail through design sketches. Section 4.6 discusses how the conceptual design relates to insights from the preceding chapters. The chapter concludes with a summary in Section 4.7.

Early versions of parts of this chapter have been previously published elsewhere (Schweer and Hinze, 2007a,b,c).

4.1 Focus

The aim of this thesis, as introduced in Section 1.1.1, is to support people in remembering their past experiences and information related to these experiences. Using terms and concepts from Cognitive Psychology, the central hypothesis of this thesis was stated in Section 2.5 as follows:

> *An interactive software system that combines the context of experiences, semantic information and associations is a suitable means to support individuals in reconstructing autobiographical memories and in retrieving cues and facts related to such memories.*

This chapter puts forward a novel solution that combines automated and semi-automated capture of information with manual annotations by the system's user. It also combines the use of context, semantic information and associations to allow an individual to retrieve and reconstruct their past experiences and related factual information (see recommendations R1 and R4–R6 in Section 2.4.3). It seeks to store cues that can trigger memories of experiences in the user (R2) and to support its user during all phases of the memory process (R3).

4.1.1 Targeted situations

To overcome problems associated with automatic capture of context and content, the design focuses on experiences that are semi-structured and that are potentially important to the user of such a system.

Semi-structured experiences have at least some degree of inherent structure. For example, academic conferences typically consist of sessions that in turn contain presentations; travel often involves visits to a number of sights. This characteristic makes it easier to automatically capture information related to the experiences because the structure can be used as a guide.

Potentially important experiences are those that the person is likely to wish to remember at a later point in time. This characteristic allows the assumption to be made that the system's user will invest time and effort into structuring and annotating the captured information. Structuring and annotating captured information allows the information in the system to be more meaningful and also has its own benefits, as seen in the analysis of CARPE systems.

Most people have developed strategies to remember, and remind themselves of, absolutely critical information; however, information that is only *potentially* important may remain unrecorded. This means that support for this type of experience is insufficient in existing software systems.

Examples for situations with these characteristics are visits to academic conferences and trade conventions, travel and scientific fieldwork. Often, a person will experience several similar events in these categories – for example, a scientist may conduct many fieldwork trips a year, all with a similar structure. The psychology literature reviewed in Chapter 2 suggests that repeated, similar events may be stored in a generic form and that recall of such events may consequently involve a high degree of reconstruction.

To illustrate the type of situation targeted by the research presented in this thesis, the following describes typical events and example information needs for two types of situations: attending an academic conference and travelling.

Attending an academic conference

Typical types of activities while attending an academic conference include:

- travelling to and from the conference city;
- commuting between the hotel and the conference venue;
- checking in at the conference reception;
- meeting other conference attendees – some the first time, some who are already known;
- attending presentations, demonstrations, keynotes, panel discussions, poster sessions and other events scheduled in the conference programme – alone or as part of a group;

- talking to other conference attendees about a wide range of topics, for example discussions about professional and social events at the conference, research ideas, other professional topics, plans for the evening, travel advice for the conference city and its surroundings, or other personal topics;
- taking part in excursions or social events with other conference attendees; and
- exploring the conference city, alone or with other conference attendees.

Examples for information that someone may wish to remember about attending an academic conference:

- At which place/time/event did I meet this person?
- Which topics did I discuss with this person (last time we met/at a given event/. . .)?
- With whom did I speak (about a given topic/at a given event)?
- Whom did I tell about this place/person/conversation?
- To whom did a given colleague introduce me at a given event?
- At the conference lunch on Thursday of this given conference, there was someone sitting at my table, two seats to my right. What was her/his name?

Travelling

Typical types of activities while travelling include:

- getting from one place to another;
- visiting sights such as a museum, a waterfall or a temple;
- participating in more or less structured activities, such as a sea kayaking trip or a guided city tour;
- staying at hotels, campgrounds or other places of accommodation;
- talking to other travellers – in particular, exchanging opinions and recommendations for sights, activities, places of accommodation, means of transportation etc; and
- taking pictures/movies.

Examples for information that someone may wish to remember about their travels:

- When was I at a given place?
- With which organiser/company did I undertake a given activity at a given place and time?
- When/where did I take a given picture?
- Which pictures did I take at a given place?
- What are the contact details of a person I met at a given place/time?
- What recommendations did I get from a given person or about a given activity/place?
- Which path did I take between two given locations?

4.1.2 Links to related work

Chapter 3 analysed existing interactive software systems related to augmenting memory. The aim of the analysis was to determine how the three core factors – context, semantic information and associations – have been used in systems designed for a person's memories and for other personal information.

One outcome of the analysis was that context of experiences as well as records of experiences in various representations can be captured automatically, to form an archive of such records that can then be referred back to. However, the analysis also showed that there are conceptual limitations to such approaches and that automatic capture by itself is not actually as effective in practice as is often assumed in the design of these systems. The research presented in this thesis targets situations with characteristics that can make a combination of automatic capture with manual annotations more feasible.

Another outcome of the analysis was that there are several ways in which context, semantic information and associations can be used to facilitate access to personal information. Most of these can be adapted for use with information about personal memories.

4.2 Conceptual architecture

This section introduces the conceptual architecture of an augmented memory system that builds on the results of the previous two chapters. Figure 4.1 shows the roles and actions involved in interacting with the system. The system includes a capture component, a general-purpose user interface and a storage component. Interaction with the system is shown separated into three phases: experiencing, revising and remembering. The following sections describe each phase in turn.

4.2.1 Experiencing

The experiencing phase takes some aspects from approaches in the area of CARPE as analysed in Section 3.2. As outlined in Section 3.4.4, automatic capture of context and content data related to an experience allows the user to focus on the experience. The system introduced here can trigger capture, for example at regular intervals or when a potentially interesting moment has been identified. Examples for the types of context and content data that could be captured by such a system are the user's location, timestamps and the presence of others as well as higher-level events such as conversations. Captured data can also include media files such as photographs, video and audio recordings.

One of the insights gained in Section 3.2.2, however, is that automatic capture by itself is not necessarily beneficial for augmenting autobiographical memory. For that reason, the conceptual architecture shown in Figure 4.1 allows the user to manually trigger capture by marking a moment as interesting. The effect is that as much data as possible is captured about the moment and stored for annotation and integration into other information in the system by the user during the revising phase.

The focus on semi-structured experiences makes it more likely that the system will be able to identify and capture information about high-level events. This can either be from data about the structure that is available from other sources (e. g. conference program, the user's calendar) or with the use of machine learning techniques such as those employed by some CARPE systems. Other generic context data such as time, location

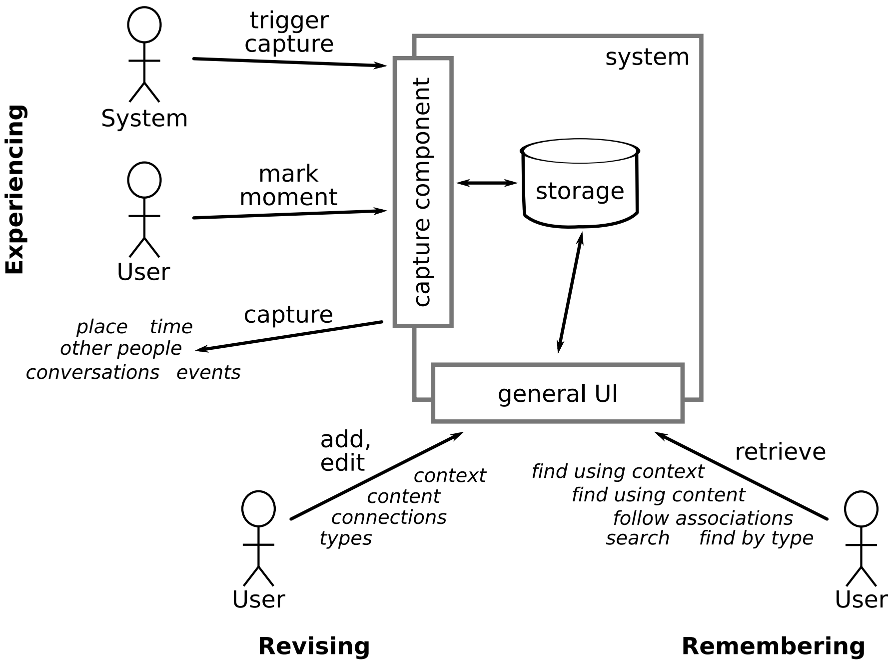

Figure 4.1. Roles and actions

and presence of others can be captured with the techniques that CARPE systems use.

4.2.2 Revising

The user of the augmented memory system can edit and add to the information stored in the system. As indicated in Figure 4.1, changes can be made to the context and content data in the system, as well as to connections between items and to the types assigned to an item.

The revising phase can coincide with the experiencing phase, if the user of such an augmented memory system chooses to enter information about an experience immediately. Otherwise, the user can defer this phase to a later time, for example by marking a moment as described in the previous section. Information in the system can also be edited at any later time.

The focus on at least potentially important experiences allows the assumption that the user will spend time to annotate and edit at least some of the captured data. Editing the data after capture serves two functions.

Firstly, it corrects any mistakes in the automatically captured data. Secondly, it adds information beyond that which is possible to capture automatically.

Interacting with the data may additionally improve the user's memory by encouraging deeper processing and decreasing the likelihood of memory failures related to shallow processing (as introduced in Section 2.3.1); compare findings reported in Section 3.2.2 that handwritten notes taken at a meeting improve recollection of the meeting even if they are never referred to again.

4.2.3 Remembering

The system's user interface for retrieval allows the user to find information using fully or partially remembered content and context as well as by following connections between items in the system. This follows recommendations R4 and R6 made in Section 2.4.3. Certain kinds of context allow for specialised user interface components; text search caters to some kinds of context and also to textual content. Semantic information about items in the system and about associations between items can also be used to find information in the system (R5).

The remembering phase can take place at any time after the original experience. In line with the objective of this thesis, "remembering" includes the retrieval of facts from the system and the reconstruction of facts and experiences based on information in the system. The retrieval side of the system incorporates aspects of the PIM systems surveyed in Section 3.3.

4.3 Requirements

This section summarises the requirements for a system that follows the high-level approach and the conceptual architecture. The requirements are grouped by system component: overall system, user interface, capture component and storage component. They are stated on a fairly abstract level and are intended to be read in conjunction with the interaction examples in Section 4.4 and the user interface design sketches in

Section 4.5.

4.3.1 Overall system

The overall system needs to

(S1) provide a capture component, a user interface and a storage component, as shown in Figure 4.1;
(S2) provide an infrastructure for interaction between the storage component on one side and capture and user interface components on the other side; and
(S3) have a component that triggers automatic capture at certain conditions (the "System" actor at the top left in the figure).

These requirements are a formalisation of the system components named in the conceptual architecture (Section 4.2).

4.3.2 User interface

The user interface needs to

(U1) visualise information items, representing experiences and related facts, as well as connections between information items, representing associations between memory items;
(U2) allow traversal of connections between information items;
(U3) allow retrieval of information items using full or partial context of the experience represented by the item;
(U4) allow retrieval of information items using semantic information;
(U5) allow retrieval of information items using text search for full or partial item content; and
(U6) allow addition, modification and deletion of context and content, of semantic information for items in the system as well as of connections between items.

The first four requirements are derived from recommendations R4–R6 (Section 2.4.3). The fifth requirement adds typical functionality of computer systems and the sixth requirement formalises manual information input into the system.

4.3.3 Capture component

The capture component needs to

(C1) automatically acquire context and content data; and
(C2) allow the user to trigger capture of context and content data, for example in situations that the user wishes to record in the system but where fully manual entry of information is impractical.

These requirements formalise aspects taken over from related work as summarised in Chapter 3.

4.3.4 Data model and storage

The storage component needs to be able to fulfil all requests from the capture and the user interface component. In particular, the storage component and the underlying data model need to

(D1) accommodate information items and connections between items, i. e. a graph structure consisting of nodes and edges;
(D2) allow information items and connections to be typed, where each item or connection can have more than one type;
(D3) accommodate information items that represent content and information items that represent context;
(D4) allow customisation of types available for items and connections to match the vocabulary of the system's user; and
(D5) accommodate information items in a variety of representations, textual and non-textual.

The first four requirements are needed to support the user interface requirements. The last requirement reflects information representations available in systems reviewed in Chapter 3.

4.4 Interaction examples

To illustrate the conceptual design, this section gives examples for user interaction with an augmented autobiographical memory system that follows this design. It describes two fictitious users and four scenarios: one experiencing/revising scenario and one remembering scenario for each user. These fictitious users are similar to personas (Cooper, 2004, ch. 9) in that they are user archetypes; they and the scenarios are informed by the literature reviewed in Chapter 2; they are not based on ethnographic data.

4.4.1 Sarah the Earth Sciences lecturer

Sarah is a senior lecturer in Earth Sciences. As a researcher in Earth Sciences, Sarah frequently goes on field trips and travels to academic conferences. Travel is one aspect of her job that she enjoys immensely. Apart from seeing new places, she gets to meet fellow Earth Sciences researchers from all over the world at these occasions. She likes talking to people, particularly about her research field, and she also knows that she will benefit from a good network throughout her career. She just wishes she was better at keeping track of the people she speaks to and at staying in touch with them afterwards.

Sarah teaches undergraduate Earth Sciences papers and also supervises a few research students (Honours and Masters level and as of last year also a PhD student). Sarah really likes teaching but wishes it would take up less time. She enjoys introducing students to doing research in Earth Sciences and tries to show them what it means to be an academic. For example, whenever she is at a conference and sees a presentation that might be of interest to one of her research students, she makes sure to point them to it, to give them a sense of connectedness to the greater research community. This can get tricky to remember though; Sarah usually tries to speak to as many people at conferences as possible and when she gets back home, some of the people and conversations just blend together.

Sarah is not exactly an early adopter when it comes to technology, but she is a confident computer user and does enjoy her gadgets. She owns an iPhone and is very happy with it – partially because it comes with a GPS unit. Being an Earth Scientist, Sarah is very geographically minded. Places mean a lot to her, and often she remembers where she met someone or heard something but not whom she met or what the conversation was about.

4.4.2 Sarah stores a conversation at a field trip

Sarah is on a field trip with some of her research students and with Martin, a visiting Earth Sciences lecturer from Germany. While they are busy collecting rock samples, Sarah and Martin chat about their research and academic life in general.

Martin points out one particular rock formation and tells Sarah that Christine, one of his German colleagues, is doing research about these formations. He also mentions some surprising facts that Christine discovered about them. Sarah is always looking for research topics that would make good student projects. She thinks that these formations may eventually be connected to an open Master's project that she's been thinking about for a while but that she hasn't really thought through yet. If she ever finds a student who is interested in this type of project, then Christine's publications may make a good starting point. She does not really have time right now to write down all the information, so she tells the system to "mark this moment" (C2). She takes a picture of the rock formation with her iPhone and asks Martin to say Christine's full name so she can record it (U6, D3, D5). The system automatically determines that both the picture and the audio recording belong to the captured moment (C1).

Later, in the car back to the camp where they are staying for the night, Sarah has a bit of spare time. She decides to get out her iPhone and tidy up some memories she stored during the day. She goes to the list of today's captured moments. They are sorted by time and she goes through them chronologically (U3). When she gets to the conversation with Martin about the rock formation and his colleague student Christine, she tells the system that this moment actually was a conversation (U6, D2). The system had noticed already at the time that Martin was nearby (C1). When Sarah changes the "marked moment" to a conversation, the system suggests to specify Martin as a partner in the conversation. Sarah accepts the suggestion.

The picture and the audio recording are already associated with the conversation; Sarah adds the type *rock formation* and the appropriate scientific term as text in case she will want to be able to find this using type or text search sometime in the future (D4, U6). The system has given her a rough transcription of the recorded name; however, it did not do a too good job at the German name so Sarah has to correct the spelling. She specifies that Christine is a *person* and also a *topic* in the conversation with Martin (D2). She adds a *researcher at* connection to Christine; the system gives her a list of all universities and other research institutes in the system to help her in filling in Christine's affiliation (U2, D1). Martin works at the same university, so its name is in the system already and Sarah can just pick it from the list.

4.4.3 Sarah tries to remember a conversation

Sarah is trying to remember a conversation she had with Aroha, one of her colleagues in the Earth Science department. They spoke about applying for promotion and about something else career-related – she cannot remember exactly what it was, but she thinks it might have been advice for job interviews. One of her research students is about to finish his Master's degree and asked her if she can recommend anything to him about this topic. She does remember that the conversation happened while she and Aroha were taking a walk around the university campus, and she assumes it was towards the end of last year because their last applications for promotion had been due then.

Sarah starts the memory program on her office computer. It automatically synchronises with her iPhone, so both devices have the same data, but she prefers to use the desktop version if possible because the screen is bigger and it is easier to type than on the iPhone.

Sarah tells the program to show only things that took place on the university campus (U3). There is still so much shown that she cannot spot the conversation, so she tells the program to show only things that happened towards the end of last year, while keeping the location limited to the university campus. She had hoped that she would be able to spot the conversation, but still no luck. She searches for Aroha's name (U5). When she selects the "Aroha" item, only two connected conversations are shown and she quickly determines which is the right one (U1, U2). The other topic in this conversation was salary negotiation tips for women; Sarah reads through her notes and is happy that she could use the program to remember, even though it turned out that this will not help her student.

4.4.4 Eric the Maths PhD student

Eric is a PhD student in mathematics. His undergraduate degree is a double major in maths and linguistics. He went to France for eight months as an exchange student in the third year of his undergraduate degree.

Eric sometimes attends academic conferences, but as a PhD student, funding is always hard to come by; most of the conferences he has attended were local. When he does get to go, he always tries to make the most of the experience, both by going to as many talks as possible and by talking to everyone he meets. He hopes this way he will get some ideas of what to do after his PhD, and of course to make some useful connections to others in his research area.

Even though he eventually settled for mathematics, Eric is a bit of a language geek. He used to compete in Scrabble but has had to take a break now while he is concentrating on his PhD. Between all the maths and languages, his brain is wired to spot patterns everywhere.

Eric is a confident enough computer user. He owns a Nokia N900 Internet tablet which he uses for all kinds of things, including jotting down ideas for his research.

4.4.5 Eric takes notes at a conference talk

Eric is at a conference. He is attending a talk about a topic that is relevant for his own PhD research. The speaker mentions a paper that Eric thinks he might want to read at some point, so Eric picks up his N900 and opens the memory application.

The application has figured out from his calendar that he is at the conference and is now showing an item for each talk in the session he is attending (U3, C1). Eric adds a new item of type *publication* and fills in the title of the paper (U6, D2). He connects the newly created item with the item for the current talk, specifying that the connection is of type *mentioned in* (U6, D1, D2). He also looks up the paper online so that he can add items for the paper authors and connect them with the paper as well.

4.4.6 Eric tries to remember a book recommendation

Eric is in a bookstore. He is trying to recall a book that someone recommended to him a while ago. He wants to know which book it was so he can look for it in the store. All he can remember is that this happened at a conference during the poster session. However, he has been to poster sessions at three different conferences and he cannot remember which one it was.

Eric opens the memory program on his N900. He tells the program to show all *conversations* that are connected to *poster sessions* (U4, U1). There is still quite a lot of information shown, so he tells the program to show all *books* connected to conversations connected to poster sessions. Unfortunately, the system now shows only one book and he knows that this is not the one he is after.

Eric goes back to looking at conversations connected to poster sessions and now looks for all *people* connected with these conversations. He spots a name that he thinks is the person who recommended the book to him. He tells the system to show the topics of all conversations connected to this person; now he sees that he spoke to this person during the poster session but also later at the same conference, and he got the book recommendation in the second conversation rather than during the poster session. Shaking his head a bit about the strange ways in which his mind works sometimes, but happy that his memory program helped him find what he was after, he goes off to see whether the store has this book on the shelf.

4.5 Design sketches

This section presents ideas for the realisation of a system that follows the conceptual design described in this chapter. Design sketches are shown first for information visualisation in the system and then for interaction with the system during the experiencing, revising and remembering phases.

4.5.1 Visualisation

The left part of Figure 4.2 is a photograph of memories from a conference visit, represented through business cards and sticky notes that were arranged on a whiteboard and connected with drawn lines. Business cards

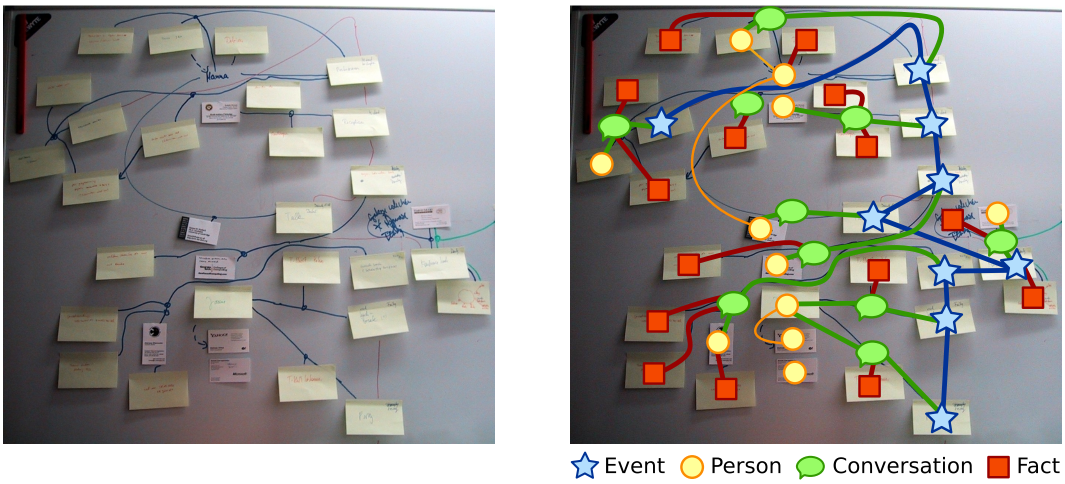

Figure 4.2. Representation of memories from a conference visit. Connections and type information are overlaid on the right.

represent people. Sticky notes stand in most cases for events or additional information, though some stand for people for whom no business card was collected during the conference visit. Events are connected with a line that indicates the temporal order of these events. Events, people and pieces of additional information are connected with lines that visualise associations, for example that a person was present at an event. These annotations are shown overlaid onto the photograph in the right

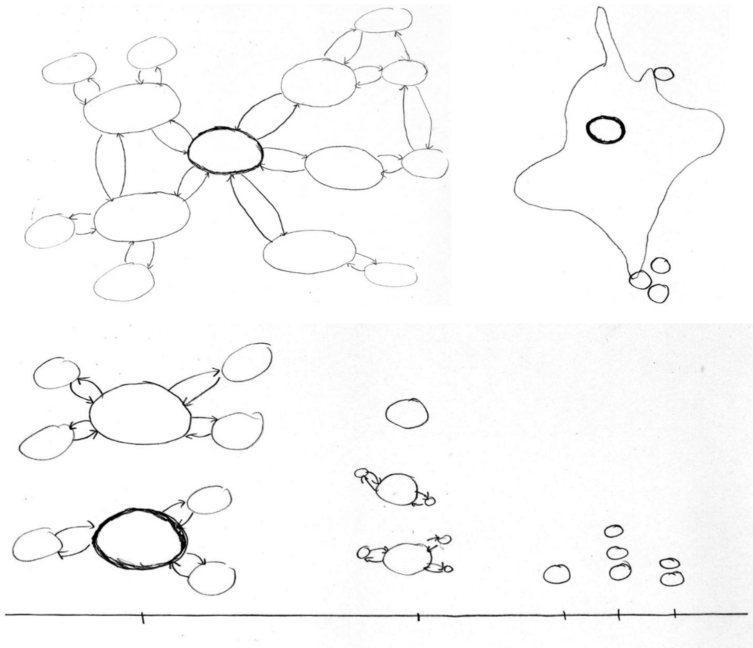

Figure 4.3. Design sketch: information visualisation as a graph (top left), on a map (top right) and along a timeline (bottom)

part of the figure. Additionally, overlaid symbols indicate the type of entity represented by a sticky note or business card.

The top left portion of Figure 4.3 shows a sketch for an information visualisation that roughly follows the structure shown in Figure 4.2. The resulting structure is similar to a mind map, with the exception that it is a graph rather than a tree and that relationships between nodes are not hierarchical. Information items are visualised as nodes in the graph and connections are visualised as edges between nodes.

The other parts of Figure 4.3 shows subgraphs overlaid on a map and on a timeline.

Creating these design sketches revealed that overlaying information on a map or timeline poses the problem of where to draw information that does not belong to a particular location or time. Two options were considered for such items. One option was to designate certain parts of the visualisation area for such items. The second option was to show such items only when they are adjacent to an item that does belong to

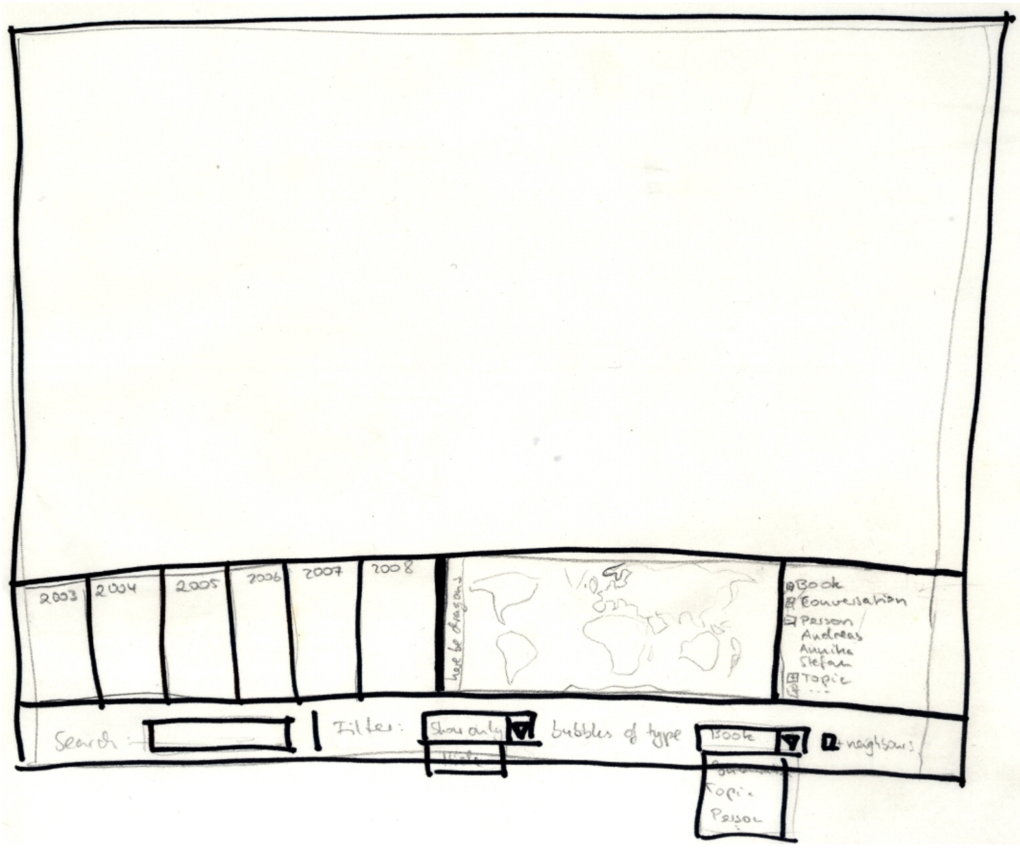

Figure 4.4. Design sketch: main view and controls

a particular location or time, and then close to that item. The drawback of the first option is that it can cause adjacent items to be shown quite far apart. The drawback of the second option is that this could cause items to be shown more than once, if they are adjacent to more than one item with a location or time. These and other potential solutions for this problem were not explored further in the design sketches.

Figure 4.4 shows a design sketch for the controls present in a software system that allows for the types of interaction described earlier in this chapter. Underneath the main visualisation area, the sketch shows a timeline, a map, a tree of information items by their type, a search box and further options for filtering by type (top left to bottom right). The placement and design of these controls were not finalised at this stage. The versions shown in the design sketch should merely be taken as place-

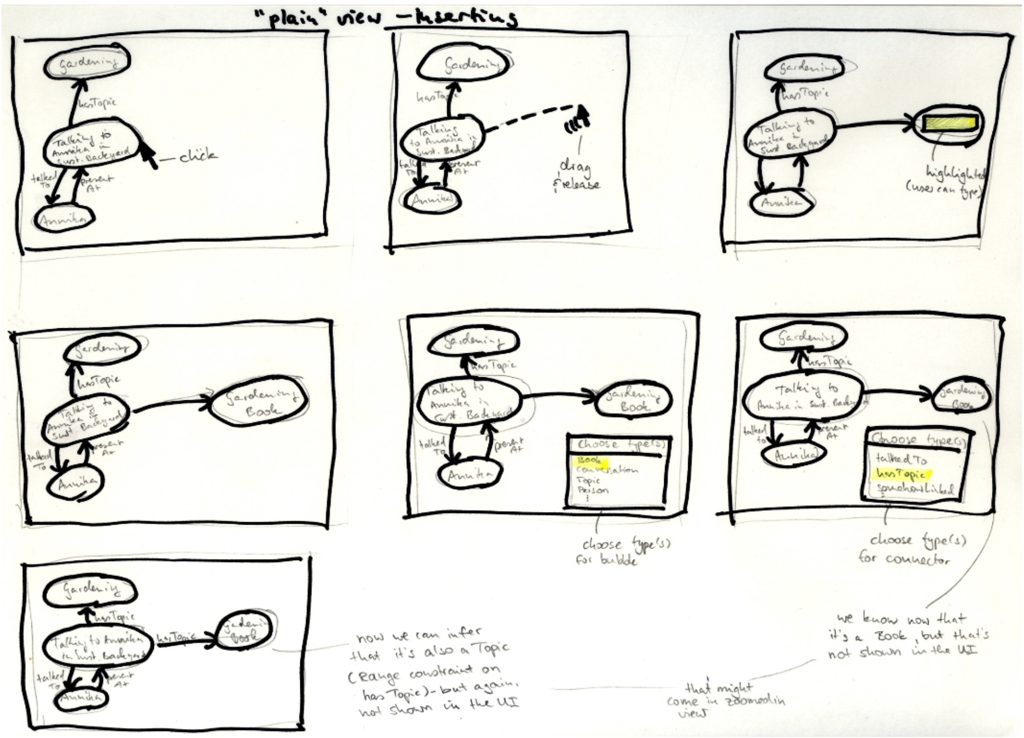

Figure 4.5. Design sketch: inserting information

holders that indicate, for example, that it must be possible to filter items by the location of the represented experience.

Initially, the main screen shows either all data in the system or a selection of relevant items. The system could use information about the user's current context, such as the user's current location or the presence of other people, to determine which items may be relevant. It could then show items connected with that location or with these people. Alternatively, the system could show the items with which the user last interacted when at this location or around these people.

4.5.2 Experiencing and revising

Figure 4.5 shows a sequence of sketches that illustrate how information items could be inserted into the system by attaching them to information items already in the system. In the sketch, the user first clicks into an information item and then drags the mouse out of it. A new information item “bubble” is created and the user types in a name. The user chooses a type for the information item (“book”) and a type for the connection (“has topic”). The choice of types suggested for the connection is based on the type of the initially existing information item. The sketch also shows how type inference can be possible based on the subject and object types of a connection predicate: the newly added item now is also of type “topic” because “has topic” connections link a “conversation” item to a “topic” item. The sketch indicates that the type of an information item or a connection is not necessarily shown in the user interface.

This sketch also shows that information items and connections are typed. Initial assumptions about the nature of these types were:

- an information item can have more than one type;
- a connection type can know what item types to expect for its subject and object items; and
- types can form “is a” hierarchies of sub- and supertypes.

Figure 4.6 shows a sequence of sketches that illustrate how information that was automatically captured by the system can be used during manual input of information. The user clicks on the empty background to indicate that they wish to create a new information item. A new “bubble” is created; the user first types in the name and then chooses a type (“conversation”). The system automatically suggests a connection of type “talked to” to an already existing information item of type “person”. This suggestion is based on a number of facts:

- the presence of the person represented by this item was detected for the time that the “conversation” item was created,
- the system knows by some mechanism that the detection of presence typically describes real-world objects that are represented in the system with the type “person”,
- “talked to” is a connection type that links items of type “conversation” to items of type “person”.

Similar automatic linking would occur for other types of context such as time and location.

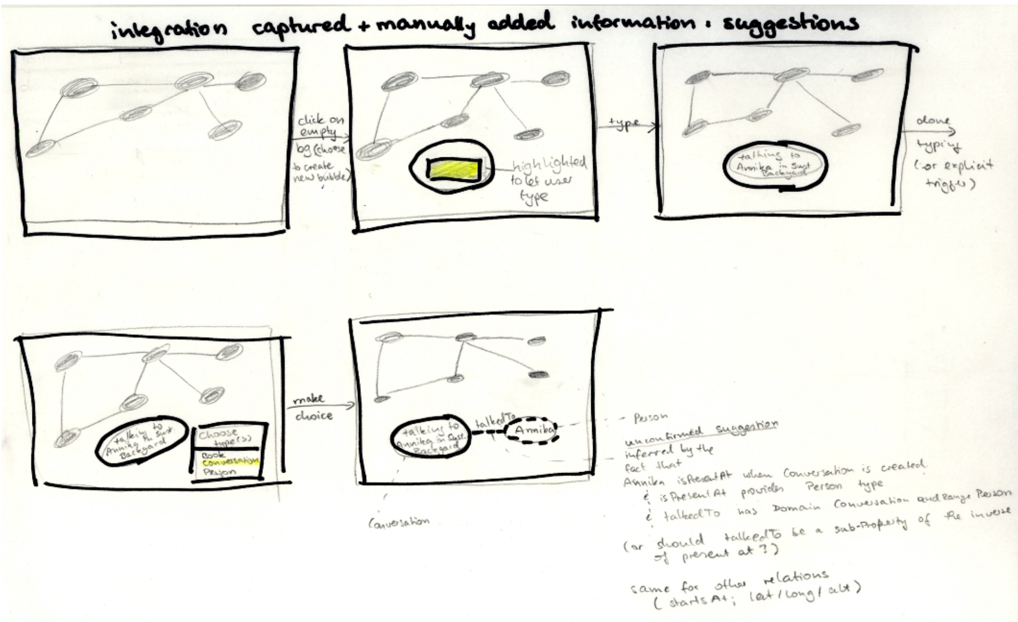

Figure 4.6. Design sketch: integrating captured and manually added information

The design sketches in Figure 4.7 illustrate editing a marked moment. When the user chooses to mark a moment, the system captures as much context as possible. All marked moments are available for retrieval. When the user edits a particular marked moment, the captured context as well as the user's current context are used as the basis for suggestions of new information items and connections.

4.5.3 Remembering

The design sketches in Figure 4.8 illustrate how information connected to an experience can be found again using the time and the location of the experience. From the initial view, the user switches to the view in which information is overlaid on a map. Only those information items are shown that belong to a particular location, shown on the currently visible portion of the map; potentially, information items adjacent to these are also shown. The user zooms in on a particular location to see only those items from to a reasonably confined area; adjacent items without location

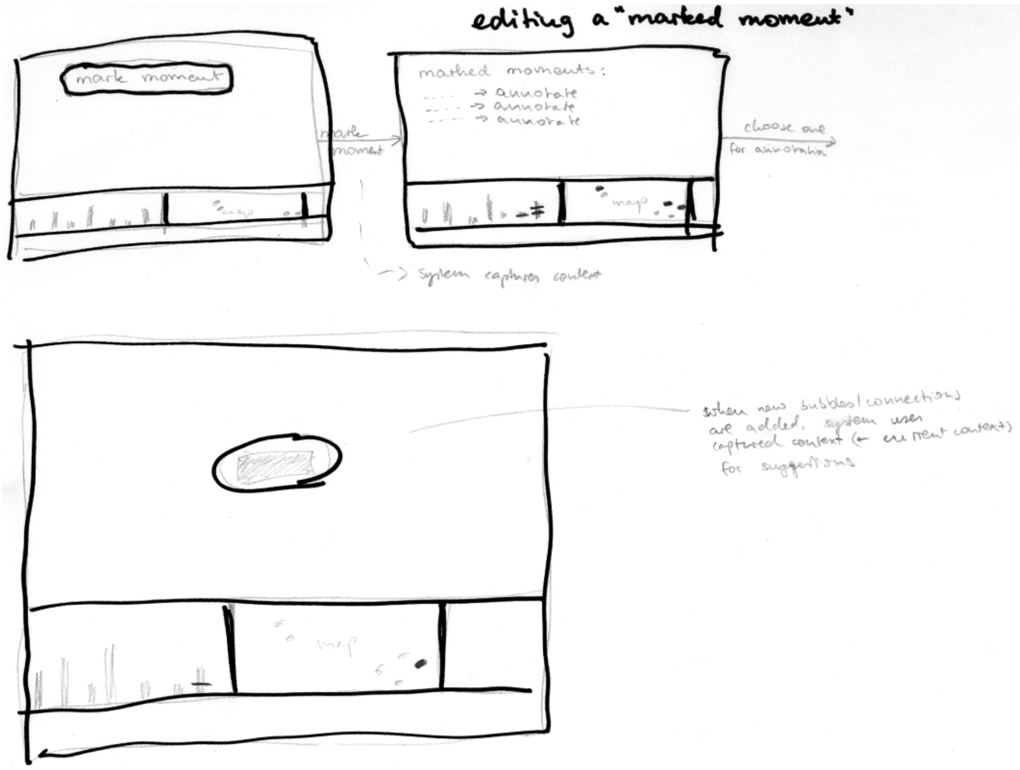

Figure 4.7. Design sketch: editing a marked moment

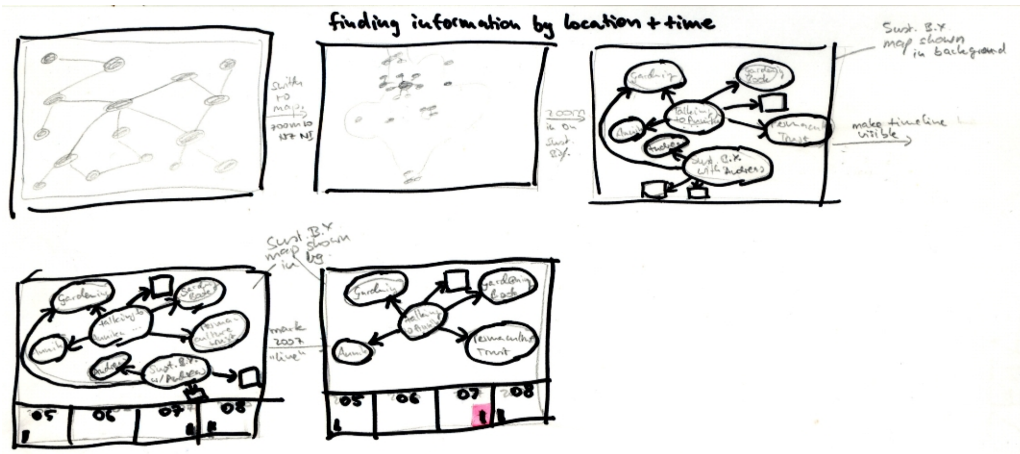

Figure 4.8. Design sketch: finding information by location and time

information are also shown.

The user then makes the timeline visible in addition to still being in the map overlay view. The timeline indicates how many items belong each to time period (e. g. month). The sketch allows for the possibility that the type of time can be changed, listing such time types as the time the item was added, the time(s) the item was edited and the time(s) the item was explicitly viewed in the system. The user chooses a particular timespan by selecting the line indicating items at this location associated with timespans towards the end of the year 2007. The system now shows only those items that have a location within the area shown on the map *and* a time within the selected timespan. Adjacent items without a time and without a location are also still shown.

The design sketches in Figure 4.9 illustrate how to find information with the help of semantic information. The user chooses to "show only bubbles of type conversation and neighbours". This causes all other information items to be hidden. Three choices can be made when using the type-based filter: whether the selected type specifies items to be hidden or items to be shown; which type to use; and whether adjacent information items are supposed to be affected by the filter.

The lines circling a selection of information items in the sketches are a result of interaction with the sketches and are unrelated to the scenario described here.

4.5.4 Paper prototype

Figure 4.10 shows a paper protoype based on the design sketches described in the previous sections. The paper prototype consists of a laminated piece of paper with a printed main window similar to Figure 4.4, sticky paper "bubbles" for information items and laminated, sticky maps at several scales (in this case showing the world, New Zealand's North Island, Hamilton, and the University of Waikato's Hamilton campus). The main window is approximately 16 cm wide by 11 cm tall. All components are contained in a small cardboard folder that also stores the unused maps and any spare bubbles. The size was chosen to represent a mobile device while allowing for the lower resolution of the paper prototype

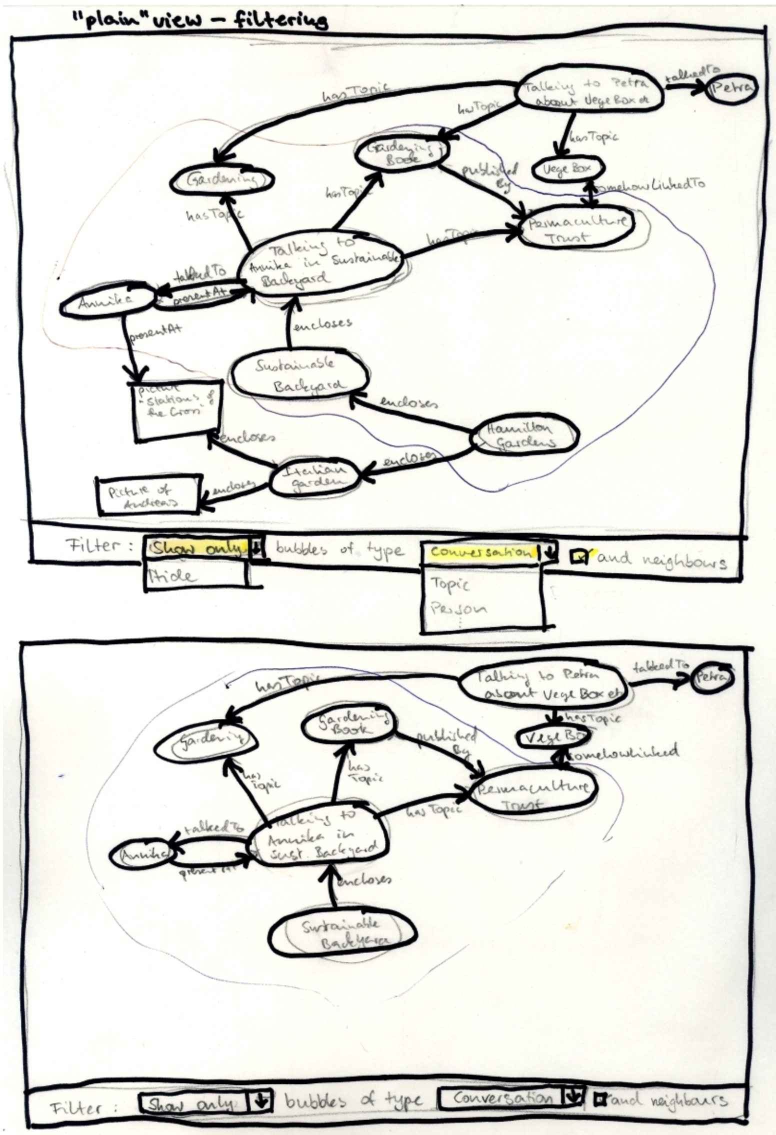

Figure 4.9. Design sketch: finding information by semantic type

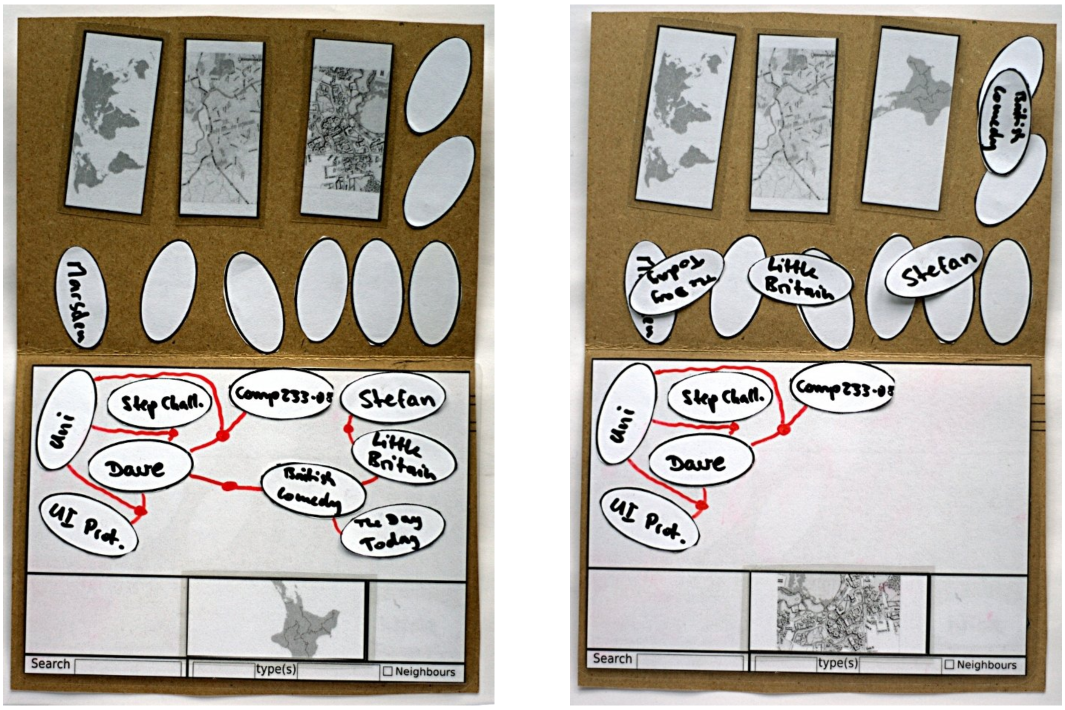

Figure 4.10. Paper prototype: all nodes (left) and nodes at a specific location (right)

medium.

Bubbles can be written on, stuck to the main window, moved within the main window or replaced to the cardboard folder. The main window and the maps are laminated to allow annotations and connections to be drawn with OHP or whiteboard markers and erased.

Walkthoughs of scenarios similar to those described in Section 4.4 quickly showed that a paper prototype was not suitable to fully explore and evaluate the effectiveness of the conceptual design. Even quite small scenarios, such as conversations with two different people about a handful of different topics (see Figure 4.10), require so many information items that they are barely manageable within reasonable timespans. On the other hand, significantly more information items are needed to explain all aspects of the system's functionality and to convey the necessity and usefulness of the various filter methods.

4.6 Discussion

This section discusses the conceptual design with regards to the six recommendations made in Section 2.4.3.

4.6.1 Type of remembering

As explained in Section 4.1, the decision was made to focus the research presented in this thesis on support for remembering experiences via knowing/retrieving and via reconstructing. Consequently, only these types of remembering are addressed by the conceptual system design presented in this chapter. The third type of remembering an experience, re-living, is considered to be fundamentally different and more complex than the other two types; it is not aimed for in the conceptual design but not specifically discouraged either.

4.6.2 Type of information

The conceptual design described in this chapter acknowledges the distinction between factual information and cues, which can be stored in a system, and memories, which cannot be externalised. A system following this design allows its user to retrieve factual information that may be what the user is trying to remember. Information items retrieved with the system can also act as cues that cause the user to remember experiences represented by, or related to, these information items. Information items can be in a variety of formats to approximate imaginal and nonimaginal representation of information in memory.

4.6.3 Phase

The conceptual design contains components to support remembering during all three phases identified in recommendation R3. Semi-automatic capture combined with manual annotations encourage deep processing of experiences, which in itself may cause the system's user to better remember these experiences later. The system helps in remembering experiences by allowing the user to retrieve factual information and cues.

The system also aids its user's memory at all other times by allowing the user to continuously engage with the information items in the system.

4.6.4 Context

The context of an experience is used in several ways in the conceptual design. During the experiencing and revision phases, automatically captured context is used to enrich information items and, together with semantic information, to suggest the creation of connections between information items. During the remembering phase, the information items shown in the system can be filtered to those with a particular context.

Two types of context, the time and the location of an experience, are explicitly visualised in the design sketches. These visualisations are based on similar visualisations of time and location in the CARPE and PIM systems analysed in Chapter 3. Other types of context, for example people present at an experience, can still be represented by integrating them into the main view and connecting them to related information items.

4.6.5 Semantic information

Similar to contextual information, semantic information is used at all three phases of interaction with the system. During the experiencing and revision phases, semantic information is used together with the semi-structured nature of the targeted events to make captured information more meaningful and to help link it with existing information items. During the remembering phase, semantic information can be used together with connections to find chains of information items of a given type or sequence of types.

Of the three factors identified in recommendations R4–4.6.6, semantic information was the least used in the PIM systems and particularly in the CARPE systems analysed in Chapter 3. This made it necessary to experiment with new user interfaces for this type of information. The design sketches show a relatively simple use of semantic information during the remembering phase and a more advanced use of this information during content capture and input.

4.6.6 Associations

While experiences and facts are represented in the conceptual design as information items, associations between experiences or between experiences and facts are represented as connections between the information items. More formally, this structure is a graph with the information items as nodes and the connections as edges. In the design sketches, this graph structure is visualised explicitly. The neighbours of an information item in this structure form the context of the item, which can be used during the remembering phase. Visualising the links between information items may also encourage spontaneous recall when the user interacts with the system because it makes related items easily accessible.

It was noted in Chapter 3 that associations between memory items, and consequently connections between arbitrary information items, are rarely made explicit in the analysed CARPE and PIM systems. The design sketches show a direct visualisation of the underlying graph structure, which is not commonly found in user interfaces targeted at general audiences but is similar to visualisations of the hierarchical tree structure of mind maps.

4.7 Summary

This chapter contributes to answering the second research question identified in Section 1.1.2, how an interactive software system can help someone remember, by proposing the conceptual design of such a system. The system's architecture is grounded in results from Cognitive Psychology, using context, semantics and associations. It incorporates aspects of existing Computer Science approaches to augmenting memory, combining automatic with manual capture of context and content and with retrieval options based on context, semantic information and associations.

The description of the conceptual architecture is supplemented with a list of requirements for systems implementing the architecture, with examples for user interaction with the system and with design sketches. The conceptual design is described on a relatively high level; different implementations are conceivable that all fulfil the requirements and allow for interaction similar to that given in the usage examples. Room for variation is left in particular regarding the realisation of the capture component, regarding the types of context used in the system and regarding the visualisation of items and their connections.

Even though design sketches and a paper prototype were used to explore the conceptual design, the amount of information items necessary even for simple scenarios made this impractical for a full evaluation. None of the software systems reviewed fulfil all requirements for the targeted type of situations. Consequently, a selection of the conceptual design was implemented in a software system to allow end-user evaluations to be conducted. The following chapter introduces the Digital Parrot, an implementation of the retrieval aspects of the conceptual design. It uses temporal and spatial context of experiences and includes a straightforward, graph-based visualisation of memory items and their associations.

Chapter 5

The Digital Parrot: a selective implementation

The previous chapter describes the conceptual design of a novel technique for augmenting autobiographical memory. This chapter describes the Digital Parrot, a selective implementation of this technique that focuses on the retrieval aspects. The selection of aspects of the conceptual design to be implemented was guided by the components necessary for the evaluations described later in this thesis (Chapters 7 and 8).

This chapter is structured as follows. Section 5.1 gives more details about the focus taken in the implementation. Sections 5.2 and 5.3 explain how the Digital Parrot meets the requirements for the overall system and the user interface outlined in Section 4.3. Section 5.4 links the Digital Parrot to the scenarios introduced in the previous chapter; it illustrates the steps described in the scenarios with screenshots from the Digital Parrot. Sections 5.5 and 5.6 explain how the Digital Parrot meets the remaining requirements, those related to the capture component and the data model. Section 5.7 briefly describes the implementation environment. The chapter closes with a discussion of the Digital Parrot in relation to the insights from preceding chapters in Section 5.8 and a summary in Section 5.9.

An early version of parts of this chapter was previously published elsewhere (Schweer et al., 2009).

5.1 Focus

The focus of the research described in this thesis is on remembering. Thus, the Digital Parrot is a partial implementation of the conceptual design described in the previous section, with the spotlight on the "remembering" phase and the system components required for this phase. The following summarises the decisions made with regards to the first three recommendations in Section 2.4.3.

Just like in the conceptual design, the types of remembering (R1) to be supported with the Digital Parrot are retrieval of facts and reconstruction of experiences. Re-living of experiences is not discouraged but not particularly encouraged either.

The types of information to be stored in the system (R2) are, as in the conceptual design, information items that represent facts and experiences related to facts. These information items can be retrieved either for their own sake or to act as cues for the user's memory. The Digital Parrot currently supports only textual representation of information items, although support for other types (such as images, audio recordings, video) could easily be added.

The Digital Parrot focuses on support during the remembering phase (R3). The main reason for this is that the hypothesis underlying the work presented in this thesis, and consequently the end-user studies, focus on the remembering phase. Evaluating systems designed for personal information, such as autobiographical memory, poses several challenges; these challenges are reviewed in Chapter 6. The focus on the remembering phase helps avoid confounds that may arise from addressing both the experiencing and the remembering stages.

5.1.1 Links to conceptual design

Figure 5.1 repeats the visualisation of roles and actions associated with the conceptual design that was given in Figure 4.1. Roles and actions that were implemented in the current version of the Digital Parrot are shown in black while roles and actions that were not implemented are shown in a lighter colour.

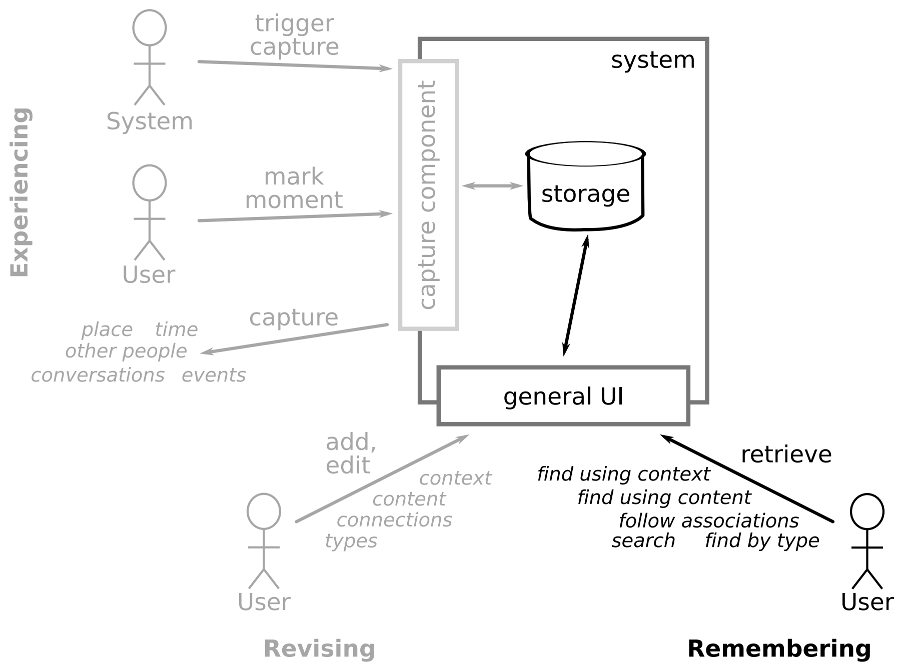

Figure 5.1. Implemented parts of the conceptual design

Due to the focus on remembering, requirements S1 and S2 apply only in part and requirement S3 as well as the capture component requirements (C1 and C2) do not apply at all. Requirement U6 does not apply because it belongs to the revising phase. Additionally, the current version of the Digital Parrot has been developed to support textual data only, thus fulfilling requirement D5 only in part.

Context of an experience in the Digital Parrot is primarily the experience's geospatial and temporal context. These two types of context are catered for directly in the user interface. However, the boundaries between content and context items in the Digital Parrot are fuzzy. This means that other types of context can be included in the Digital Parrot as content items, without special treatment in the user interface.

The implementation followed the design sketches described in Section 4.5 as far as was practicable. The following sections describe the Digital Parrot; deviations from the design sketches are noted where applicable.

5.2 Overall system

This section explains how the Digital Parrot fulfils the requirements S1 and S2. The third requirement for the overall system in Section 4.3 does not apply to the focus chosen for the current version.

The Digital Parrot consists of a user interface and in-memory storage. Interaction between these components loosely follow the Model-View-Controller paradigm. Interface components register themselves with appropriate instances of the data model to receive events on updates; they can also make changes to instances of the data model.

The current version of the Digital Parrot was developed for desktop computers. The mobile/wearable component mentioned in the design sketches and scenarios is used mostly during the experiencing and revising phases. Since the Digital Parrot focuses on the remembering phase, it was decided to implement it for desktop computers only. A version for mobile devices is conceivable with an adapted user interface that offers the same functionality and includes data synchronisation mechanisms between devices.

5.3 User interface

This section explains how the Digital Parrot fulfils the requirements related to the general user interface that were stated in Section 4.3.2.

The user interface of the Digital Parrot consists of a main view and four different navigator tools. Figure 5.2 shows an example of the Digital Parrot's user interface. The main view is described in-depth in Section 5.3.1. In the same window as the main view are buttons that allow the user to activate and deactivate the navigators. The navigators influence the information shown in the main view by highlighting some items and hiding others. There are two contextual navigators, the timeline and the map navigator, as well as a type navigator and textual search. Each navigator, with one exception, has its own window and all can be shown and hidden individually. The navigators are described in-depth in Sections 5.3.3 through 5.3.5.

Navigation tools in the Digital Parrot affect the information shown in

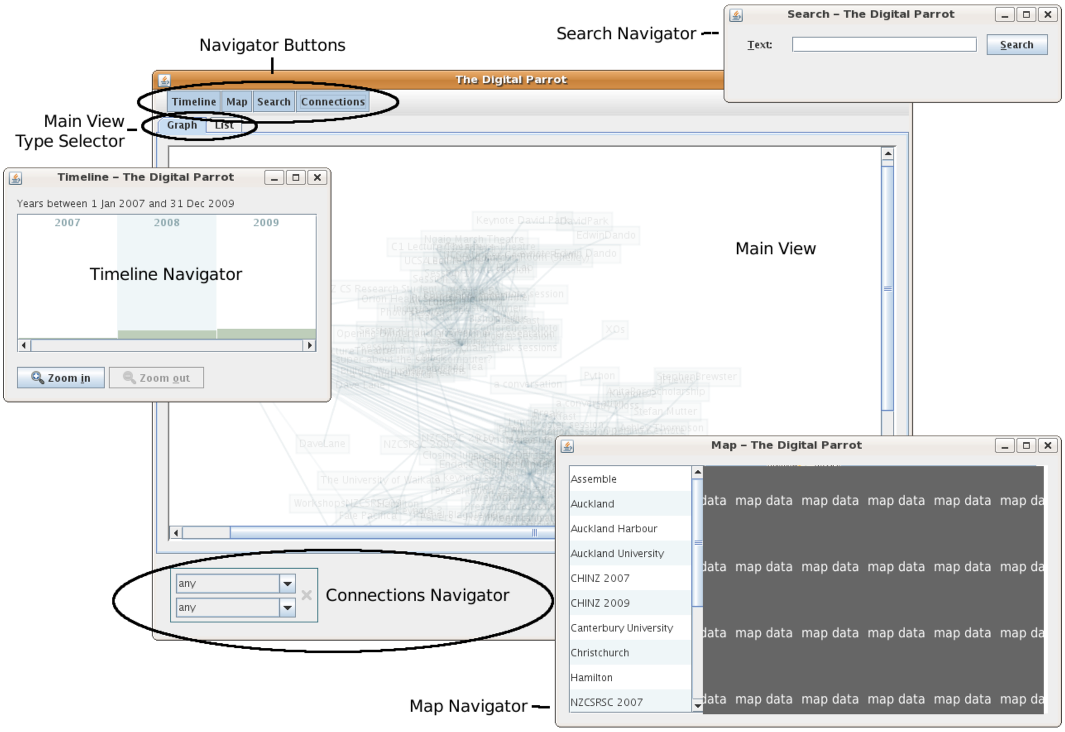

Figure 5.2. The Digital Parrot: Main view and all navigators. Map data removed from screenshot for copyright reasons.

the main view in two ways: through *restricting* or through *highlighting*. By restricting the view, navigation tools request that some information items be hidden from the main view. Restrictions are disjunctive; all information items are hidden from the main view that are hidden by at least one active navigation tool. Highlighted information items are shown more prominently in the user interface. Highlights are sequential; only those information items are highlighted that were requested to be so in the most recent interaction with a highlighting navigator. When the user deactivates a previously active navigation component, all its restricting and highlighting requests are removed from the main view. They are restored when the navigation tool is re-activated.

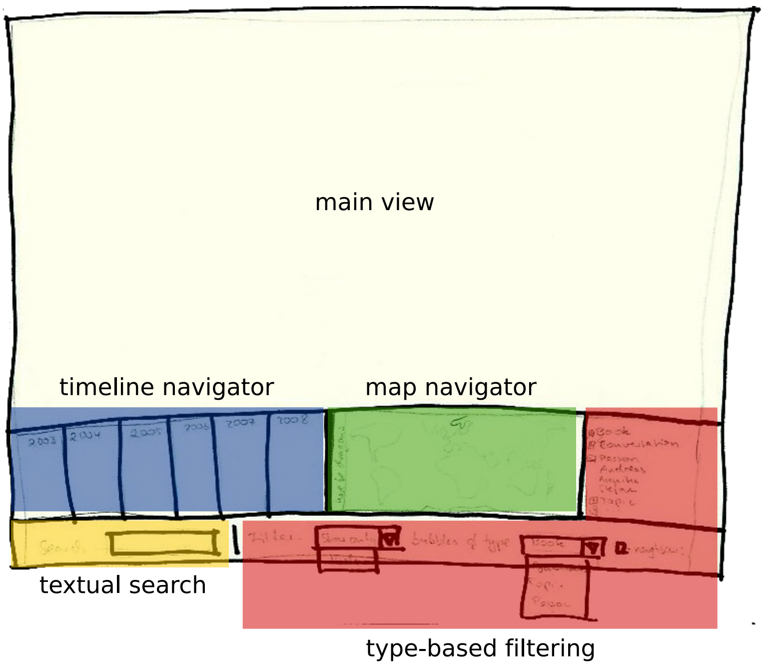

Figure 5.3. Early design sketch of the user interface

Differences between design sketches and implementation

The design sketches in Section 4.5 combine all navigators and the main view into a single window, as shown in Figure 5.3. This integrated user interface was based on principles of experiential computing (Jain, 2003). In particular, it was planned for interaction with the data through one component to be reflected immediately in all other components; for example, selecting an item in the main view that has associated geospatial information would cause the map to adjust to show the item's location.

This approach was abandoned during the move from initial design to implementation. The main reasons were issues on the conceptual level, centered around the dual role of some components as both display and filter. It was decided not to focus on resolving these issues but rather to separate the navigator windows from the main window. At the same time, the navigators were changed from being both display and filter to being filters only. The move to separate windows was intended to simplify the user's mental model of the navigators' function.

5.3.1 Main view

The main view visualises memories stored in the system, fulfilling requirement U1. Information in the system is represented as a set of statements, where each statement consists of a subject, a predicate and an object. Subject and objects are information items that represent memory items; predicates form connections between information items that represent associations.

5.3.2 Main view types

The memory structure consisting of information items and connections (requirement U1) directly corresponds to the standard definition of a graph consisting of nodes and edges. It was decided early on in the design of the Digital Parrot to visualise this graph structure directly in the user interface.

The hypothesis was that this visualisation would be beneficial for the users of the Digital Parrot because connections between information items are immediately apparent. Users without a background in Computer Science and related areas may not be familiar with data visualisations as graph networks; however, the graph view of the Digital Parrot was assumed to be similar enough to mind maps that it can be learned by all users.

Another assumption was that users of the Digital Parrot might like to be able to arrange memory items spatially. This assumption is similar to that followed in the design of the PhotoMemory system in which photographs in a collection remain at the same position on screen even when filters are applied to the collection (Elsweiler et al., 2005). This assumption also influenced the decision to make node positions persistent across runs of the Digital Parrot.

A list of statements view was developed as an alternative visualisation of the underlying structure. It was assumed that some users might prefer the list view, which is ultimately based on text, to the graph view.

The graph view is shown by default; the user can switch between main views using tabs. Restrictions, highlights and selections are synchro-

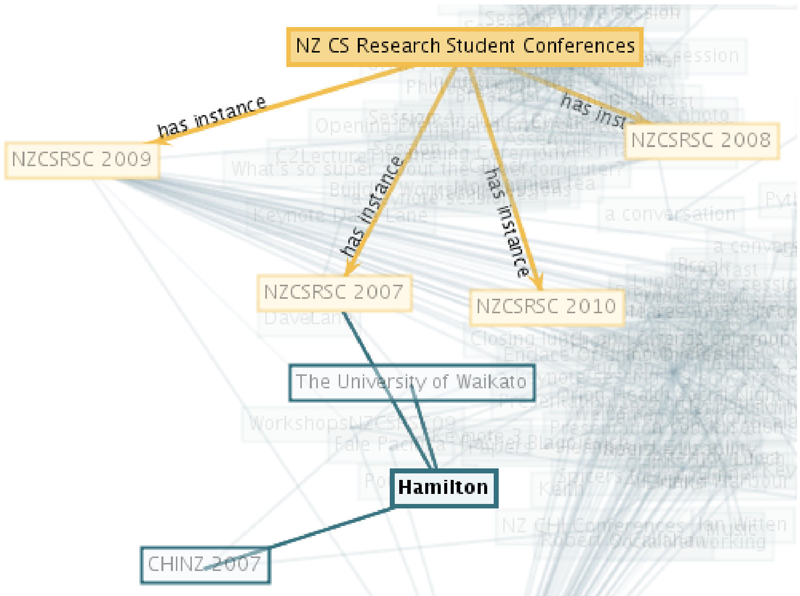

Figure 5.4. Detail of the graph view. The node “Hamilton” is highlighted, the node “NZ CS Research Student Conferences” is selected.

nised between main views.

Graph view

The graph view shows the information in the Digital Parrot as a directed graph consisting of nodes (subjects/objects) and edges (predicates). All nodes are labelled. Edges are initially shown undirected and unlabelled. Figure 5.4 shows a close-up of the graph view.

When the Digital Parrot is started, all memory items in the system are shown in the graph, drawn semi-transparently.

When the user selects a node by clicking on it with the mouse, this node becomes opaque with an orange fill colour and an orange outline – see “NZ CS Research Student Conferences” in Figure 5.4. All nodes adjacent to the node, and all connecting edges, are drawn in a different shade of orange (e. g. “NZCSRSC 2009” in the figure). All these nodes are drawn on top of all other nodes. All incident edges of a selected node are labelled with the type of connection and are shown with arrows indicating the direction of the label (“has instance” in the figure).

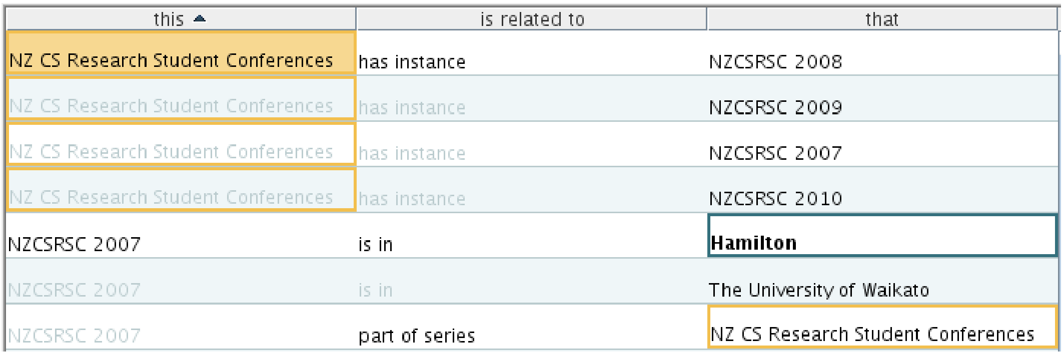

this	is related to	that
NZ CS Research Student Conferences	has instance	NZCSRSC 2008
NZ CS Research Student Conferences	has instance	NZCSRSC 2009
NZ CS Research Student Conferences	has instance	NZCSRSC 2007
NZ CS Research Student Conferences	has instance	NZCSRSC 2010
NZCSRSC 2007	is in	**Hamilton**
NZCSRSC 2007	is in	The University of Waikato
NZCSRSC 2007	part of series	NZ CS Research Student Conferences

Figure 5.5. Detail of the list view. The item "Hamilton" is highlighted, the topmost subject "NZ CS Research Student Conferences" is selected.

Nodes that are highlighted by a navigator are shown with less transparency than normal nodes, are outlined in dark blue and the label uses a boldface font – see "Hamilton" in Figure 5.4. All nodes adjacent to the node, and all connecting edges, are outlined in a different shade of blue (e. g. "CHINZ 2007" in the figure). Highlighted nodes are drawn on top of normal nodes but underneath selected nodes. Selection takes precedence over highlights – i. e. when the user selects a highlighted node, it is drawn as a selected node instead.

The user can switch between two view modes. The default mode allows the user to move nodes by clicking on and dragging them; clicking and dragging in the other mode allows the user to pan the graph.

List of statements view

The list view shows the information in the Digital Parrot as a list of statements, one statement per row. Each statement is separated into its three components (subject, predicate and object). The background colour of rows alternates. Figure 5.5 shows a close-up of the list view.

When the Digital Parrot is started, all information items in the system are shown in the statement list.

When the user selects an item by clicking on it with the mouse, it is shown in the same way as selected nodes in the graph view – see the "NZ CS Research Student Conferences" subject in the topmost row in

Figure 5.5. Most items occur more than once in the statement list. All other occurrences of a selected item, whether as a subject or as an object, are highlighted with an orange border – see all other occurrences of “NZ CS Research Student Conferences” in the figure. This facilitates finding other occurrences of the selected item.

All occurrences of items that are highlighted by a navigator are shown in the same way as highlighted nodes in the graph view – see the “Hamilton” item in Figure 5.5. Just like in the graph view, selection takes precedence over highlights.

The list of statements is initially sorted alphabetically by subject in ascending order. This can be changed to sorting by predicate or by object, each in ascending or descending order, by clicking on the column headers.

5.3.3 Contextual navigation

The Digital Parrot supports contextual navigation for geospatial and temporal context, fulfilling requirement U3. Navigation based on geospatial context is provided by the map navigator. Navigation based on temporal context is provided by the timeline navigator. This section describes both contextual navigators.

Map navigator

The map navigator affects all information items for which geospatial data is available. The Digital Parrot attempts to infer geospatial data when it is not directly available, for example from enclosing information items with geospatial information. Figure 5.6 shows a close-up of the map navigator’s user interface.

The map navigator’s user interface contains a list of places (on the left in Figure 5.6) and a map (on the right in the figure). The map shows markers for all information items whose geospatial information places them on the currently visible part of the map. The user can interact with the map by panning and zooming. The list of places is sorted first by size of place and then alphabetically. The current scale of the map determines

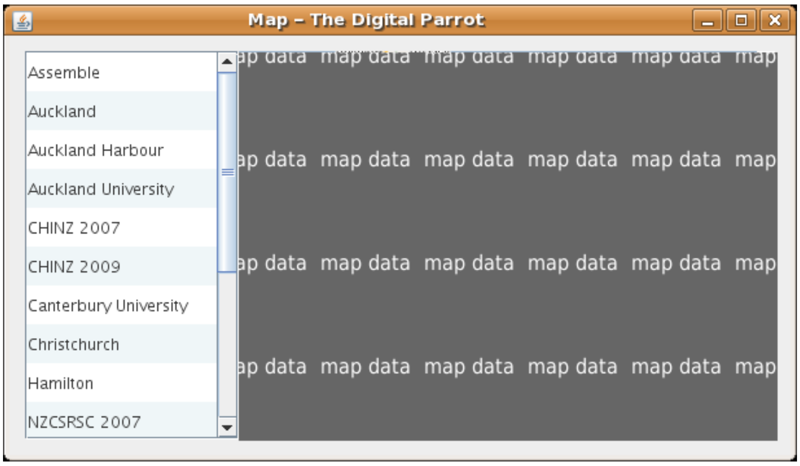

Figure 5.6. Detail of the map navigator. Map data removed from screenshot for copyright reasons.

the size range of places shown. When the user clicks on a place in the list, the map is centered on this place and scaled to fit the place's size.

The map navigator can both restrict and highlight. While the map navigator is active, only those information items are visible in the main view that have no associated geospatial data and those whose geospatial data places them within the boundaries of the map at the current zoom level. Information items annotated with geospatial data are shown on the map. Selecting information items on the map highlights them in the main view.

In the main view, the user can request an information item with geospatial data to be shown on the map. This is done via the item's context menu (which can be brought up in the usual fashion for the operating system – e. g. right-click or command-click).

Timeline Navigator

The timeline navigator affects all information items for which temporal data is available. The Digital Parrot attempts to infer temporal data when it is not directly available, for example from spanning information items with temporal information. Figure 5.7 shows close-ups of the timeline navigator's user interface.

The major part of the timeline navigator's user interface is taken up

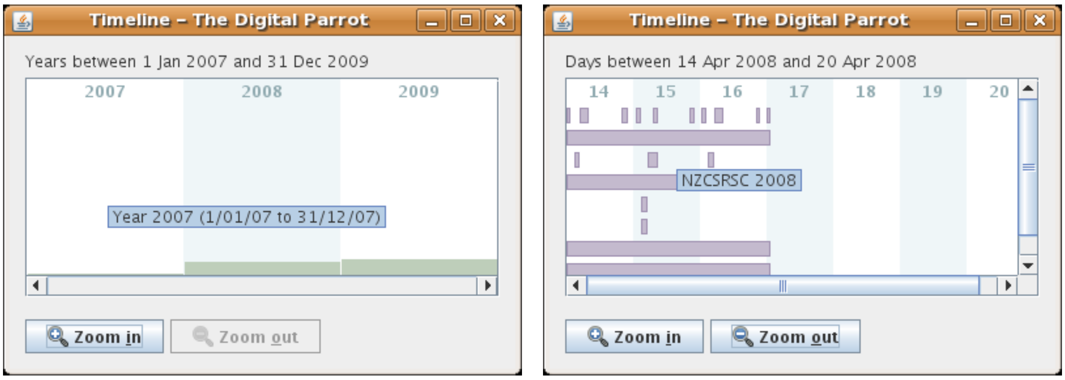

Figure 5.7. Detail of the timeline navigator: Initial view (left) and view zoomed in to show days (right). The day view shows some memory items.

by the actual timeline. The timeline is divided into slices, where each slice corresponds to a year, a month, a week or a day, depending on the length of the currently shown interval. The slice background alternates. Information items that are long enough compared to the size of a slice are shown as purple rectangles (see right side of Figure 5.7). Each slice also contains a histogram-type area at the bottom, drawn in dark green, that indicates the number of information items within the interval represented by that slice.

The user can adjust the interval shown in the timeline by zooming in and out, causing the timeline to show a shorter or longer interval, respectively. Zooming out is done via the "Zoom out" button. Zooming in can be done in three ways: via the "Zoom in" button, which zooms in by a fixed percentage; by double-clicking a slice, which zooms in to fit the slice; and by selecting a consecutive range of slices with the mouse and double-clicking the selection, which zooms in to fit the selected range.

The timeline navigator can both restrict and highlight. While the navigator is active, only those information items are visible in the main view that have no associated temporal data and those whose temporal data places them within the interval currently shown by the timeline. The user can select parts of the currently shown interval; information items within the interval are then highlighted in the main view.

5.3.4 Type navigation

The connections navigator of the Digital Parrot fulfils requirement U4. It allows users to build connected chains of information items and their types, restricting the main view to the information items on the chain. Chains let users focus on a narrow portion of their information items. The connection navigator's user interface is displayed at the bottom of the main window. The user can still show and hide the connections navigator with the connections navigator button. Figure 5.8 shows close-ups of the connection navigator's user interface.

To answer a question such as "Which book was recommended to me when I spoke to someone about hypertext at a conference in Auckland?", the user could start a chain with the information item "Auckland". The user could then add to the chain the type *Conference* and then the type *Conversation* to restrict the main view to conversations that are connected to conferences that are connected to Auckland.

The chain is visualised in the user interface, allowing for easy backtracking by removing the most recently added link and switching between types and instances. The user can add a "blank" link to the end of the chain, requesting that all items directly connected to the end of the chain be shown in the main view as well. All items on the chain are highlighted, with the exception of those items that match a "blank" link at the end. The start of the chain can be changed to an item within the chain, discarding the part of the old chain up to the link chosen as the new starting point.

Differences between design sketches and implementation

The design sketches for type-based filtering described in Section 4.5.3 let the user find instances of one type at a time. The type-based filtering in Figure 5.3 allows the user to restrict the main view either to instances of one type or to instances of all types except one (bottom right), in the first case optionally also showing items directly connected to these instances. It also shows a tree structure of the type hierarchy and all instances of a type (higher up on the right). Moving from design to implementa-

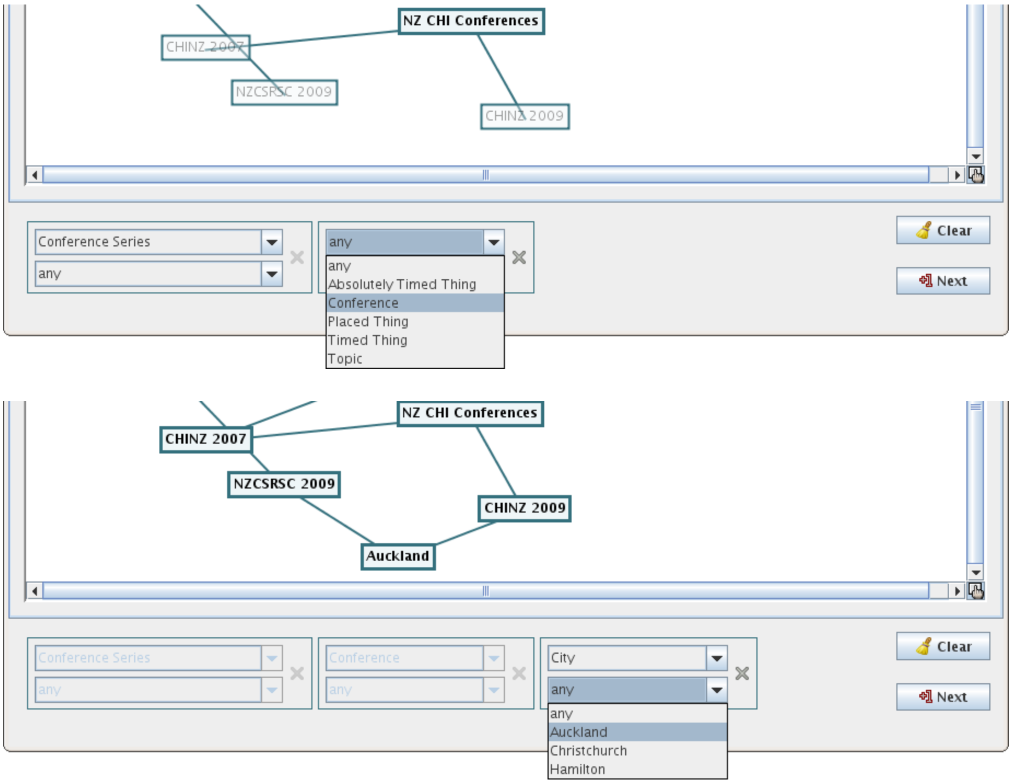

Figure 5.8. Detail of the connections navigator: Showing any node connected to a conference series instance (top) and selecting a city instance connected to a conference instance connected to a conference series instance (bottom).

tion, it was decided to give the user more fine-grained type-based control over the information shown in the main view. User testing revealed that the initial implementation was very hard to use and major changes were made as a result. More details are given in the next chapter.

The connections manager in its current form essentially provides a means to formulate subgraph queries over the data in the Digital Parrot. It extends the Feldspar query tool introduced by Chau et al. (2008a,b). Like Feldspar, the connections manager allows the user to build up chains of types to arrive at information items.

There are two major differences between Feldspar and the Digital Parrot's connections manager. Firstly, Feldspar is written specifically as a frontend for Google Desktop and supports only those datatypes that are

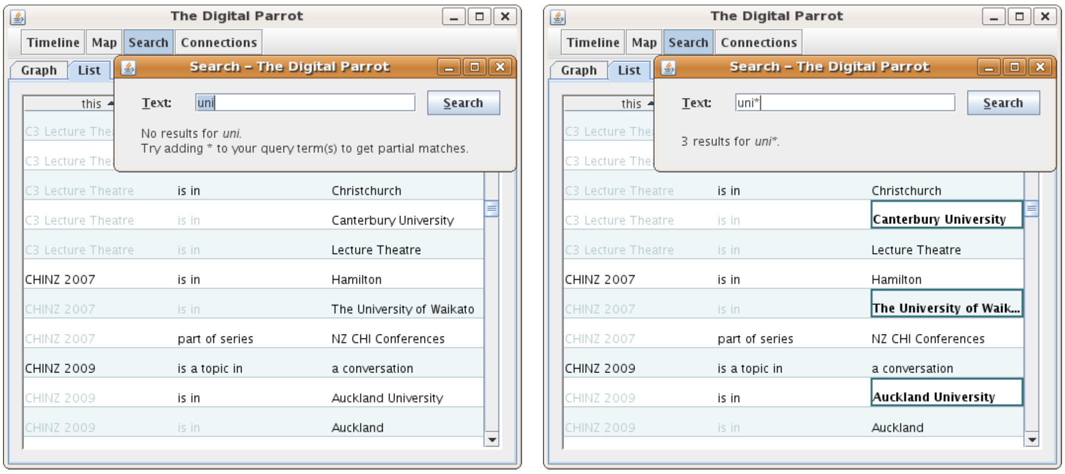

Figure 5.9. Detail of the search navigator: No results (left) and results for wildcard search (right).

available through the Google Desktop Search API. In contrast, the Digital Parrot and its connections manager support arbitrary types of information items. Secondly, Feldspar combines the display of information items with the query tool, while the Digital Parrot contains a more sophisticated view for the information itself. This allows the Digital Parrot's user to see more easily what intermediate results there are.

5.3.5 Textual search

The text search component of the Digital Parrot fulfils requirement U5. Figure 5.9 shows close-ups of the search navigator's user interface. The window of the text search component consists of a text box for the query terms and a button to initiate a search. Information items that match the query are highlighted. The text search window shows the number of search results. If no matches are found, the window additionally shows a brief help text for the query syntax.

5.4 Scenario walkthroughs

The previous section introduced the Digital Parrot's user interface. This section shows how to use these elements to perform the actions described in the two remembering-phase scenarios in Section 4.4. The screenshots were produced using an artificial data set created specifically for these walkthroughs.

5.4.1 Sarah uses the Digital Parrot

This section revisits the scenario in Section 4.4.3, in which Sarah tries to remember the topic of a conversation. For each step that Sarah makes in the scenario, this section shows screenshots of the appropriate parts of the Digital Parrot's user interface and explains how to perform this step using the Digital Parrot. Deviations from the scenario are noted where applicable.

Step 1: "Sarah starts the memory program on her office computer. It automatically synchronises with her iPhone [. . .]."

Figure 5.10 shows the Digital Parrot's user interface on startup. This walkthrough uses the graph view as the main view. As described in Section 5.3.1, the graph view initially shows all information items in the system, but draws them semi-transparently to reduce visual clutter.

Step 2: "Sarah tells the program to show only things that took place on the university campus."

This is achieved with the Digital Parrot using the map navigator. The "uni campus" item is selected from the list of placenames to adjust the map to the area of the university campus. Figure 5.11 shows the Digital Parrot with the map navigator zoomed in on the university campus and the corresponding restricted main view. The main view now shows all information items except those whose geospatial context information places them outside of the area visible in the map navigator (note that information items for example in the top left of the main view are now missing from the main view).

The scenario describes that "there is still so much shown" after this action. In the graph view, information items are shown semi-transparently by default; it is still possible to see that the main view in Figure 5.11 contains almost as many information items as the main view in Figure 5.10. It should also be noted that the data file used to produce the screenshots

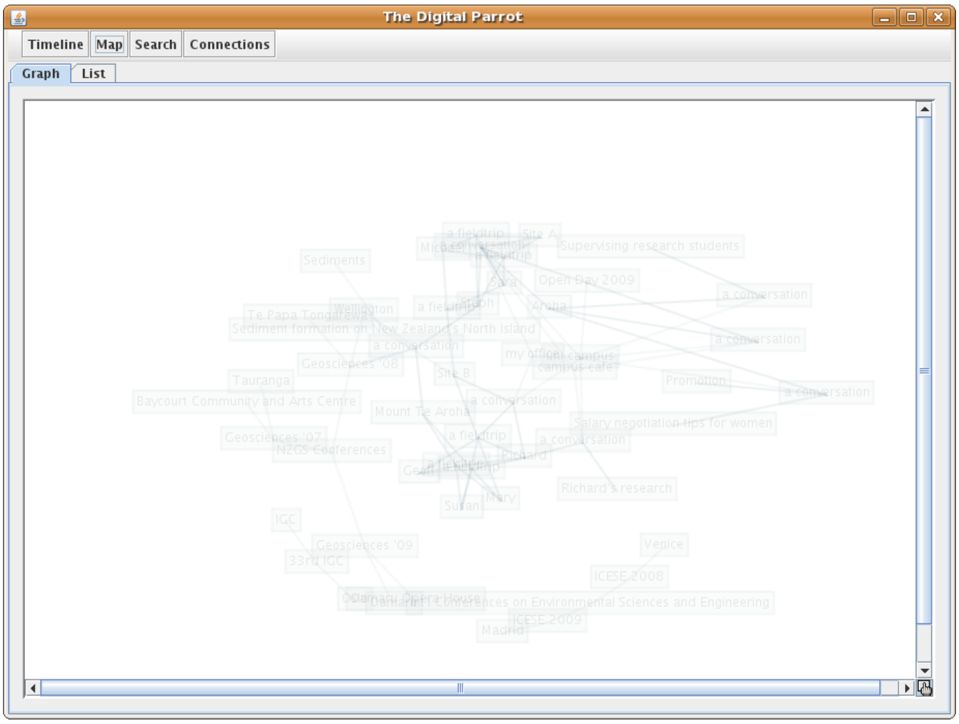

Figure 5.10. The Digital Parrot's user interface on startup: Main window with graph main view showing all information items in the system. This corresponds to Step 1 of the scenario.

contained only enough information items to show the Digital Parrot's usage rather than the number of information items that would be in the system after several years of use.

Step 3: "[...] [S]he tells the program to show only things that happened towards the end of last year, while keeping the location limited to the university campus."

This is achieved in the Digital Parrot using the timeline navigator in combination with the map navigator as in the previous step. In the timeline navigator, the timeline is first adjusted to the previous year and then the appropriate part of the previous year is selected. Figure 5.12 shows the restricted main view with the map and the timeline navigator. The main view now additionally hides all information items that have associated temporal context information which places the item outside of the interval visible in the timeline (i. e. the previous year). The main view

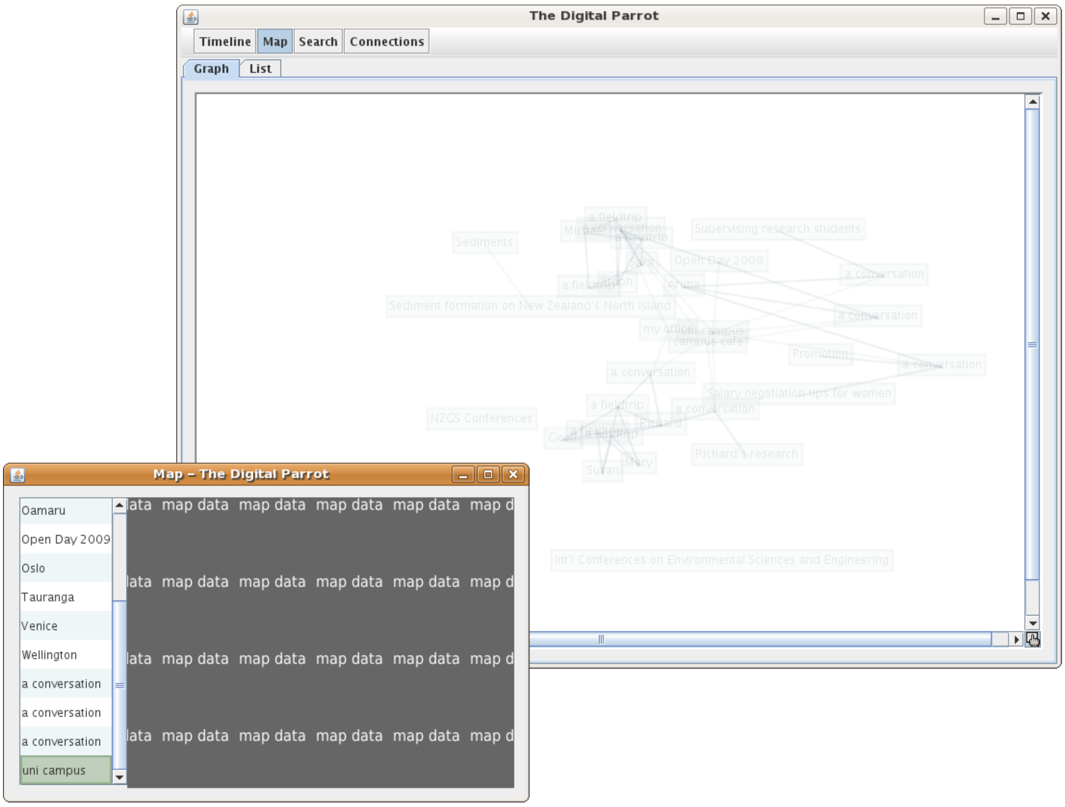

Figure 5.11. Using the map navigator to hide memory items outside the university campus, as in Step 2 of the scenario. Map data removed from screenshot for copyright reasons.

also highlights all information items that have associated temporal context which places the item inside the interval selected in the timeline.

Step 4: "[. . .] She searches for Aroha's name."

This is done in the Digital Parrot using text search. Figure 5.13 shows the Digital Parrot's user interface after a search for "Aroha" has been performed. The restrictions by time and location are still in place. The main view now highlights all items that match the search.

Step 5: "When she selects the "Aroha" item, only two connected conversations are shown and she quickly determines which is the right one. The other topic in this conversation was salary negotiation tips for women [. . .]."

This is done in the Digital Parrot by clicking on the "Aroha" informa-

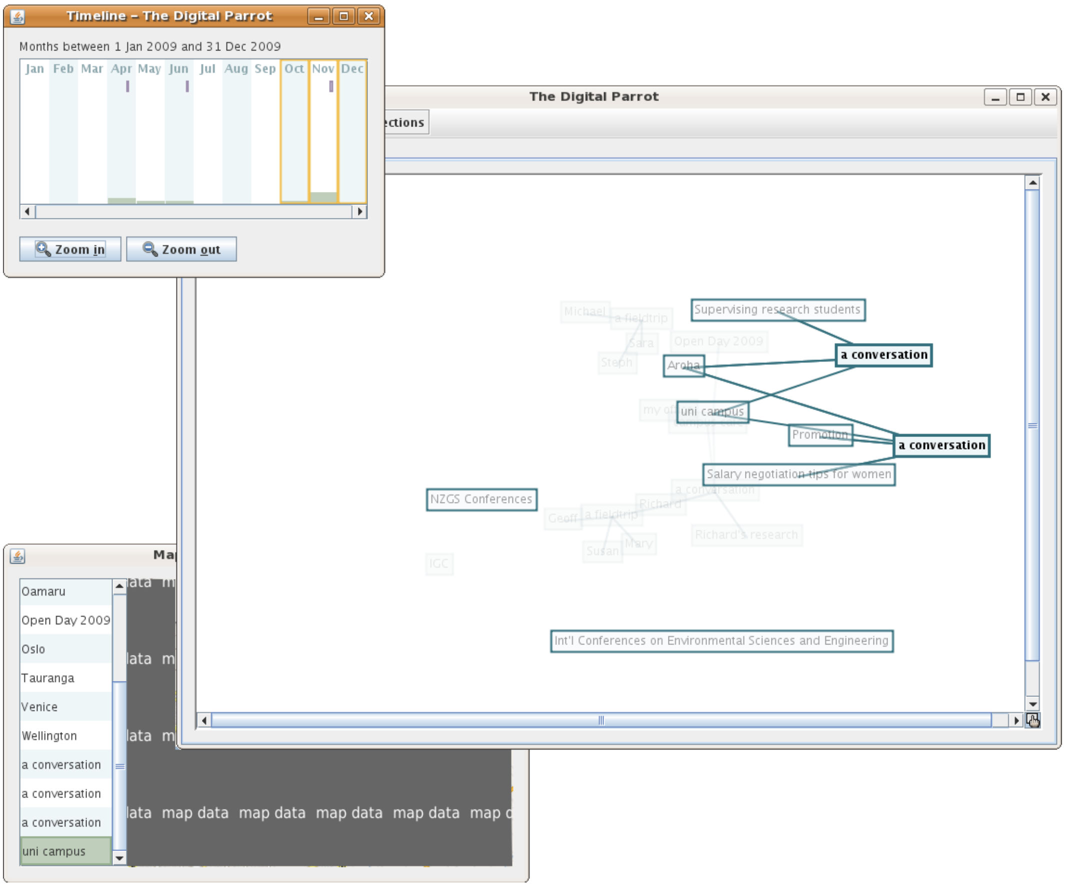

Figure 5.12. Selecting the second half of the previous year in the timeline. The main view is now restricted by location and time and some items are highlighted by time, as in Step 3 of the scenario. Map data removed from screenshot for copyright reasons.

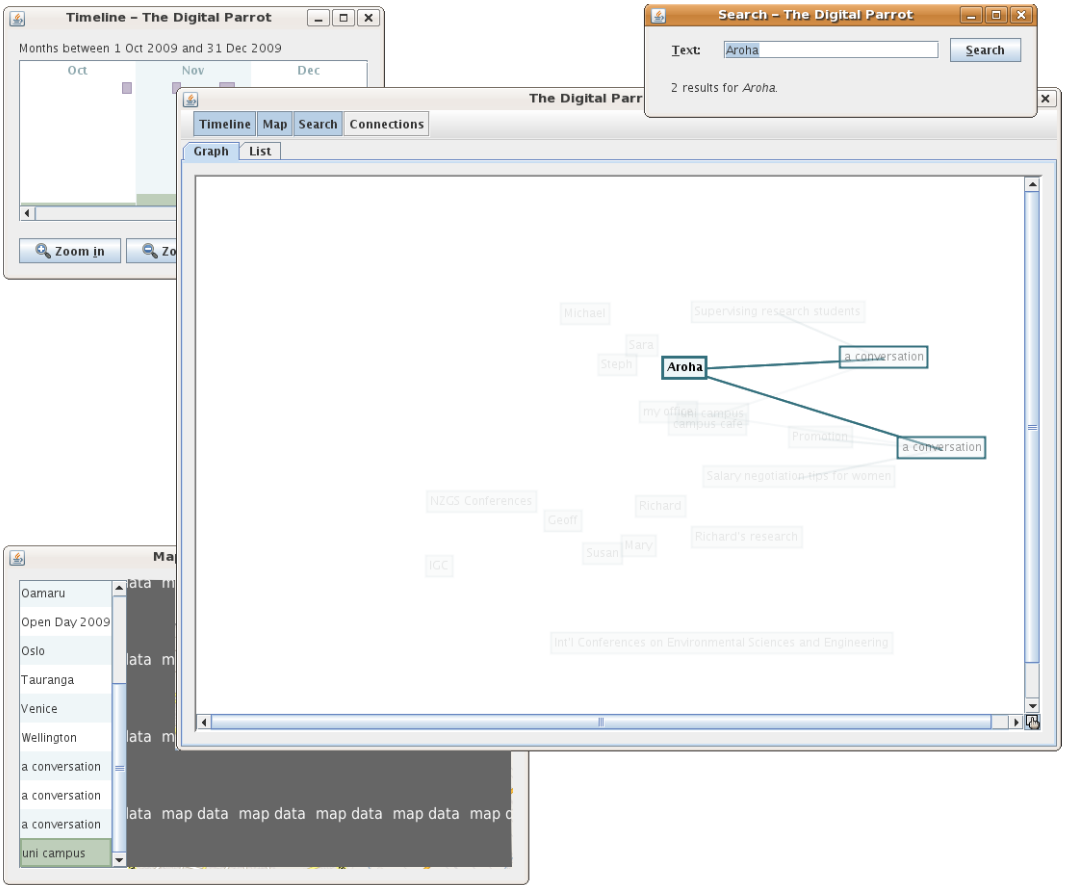

Figure 5.13. Using text search to find “Aroha”, as in Step 4 of the scenario. Search results are highlighted in the main view. The main view is still restricted by time and place as in the previous step; the timeline has now been fitted to the interval selected in the previous step to restrict the main view to the end of the previous year. Map data removed from screenshot for copyright reasons.

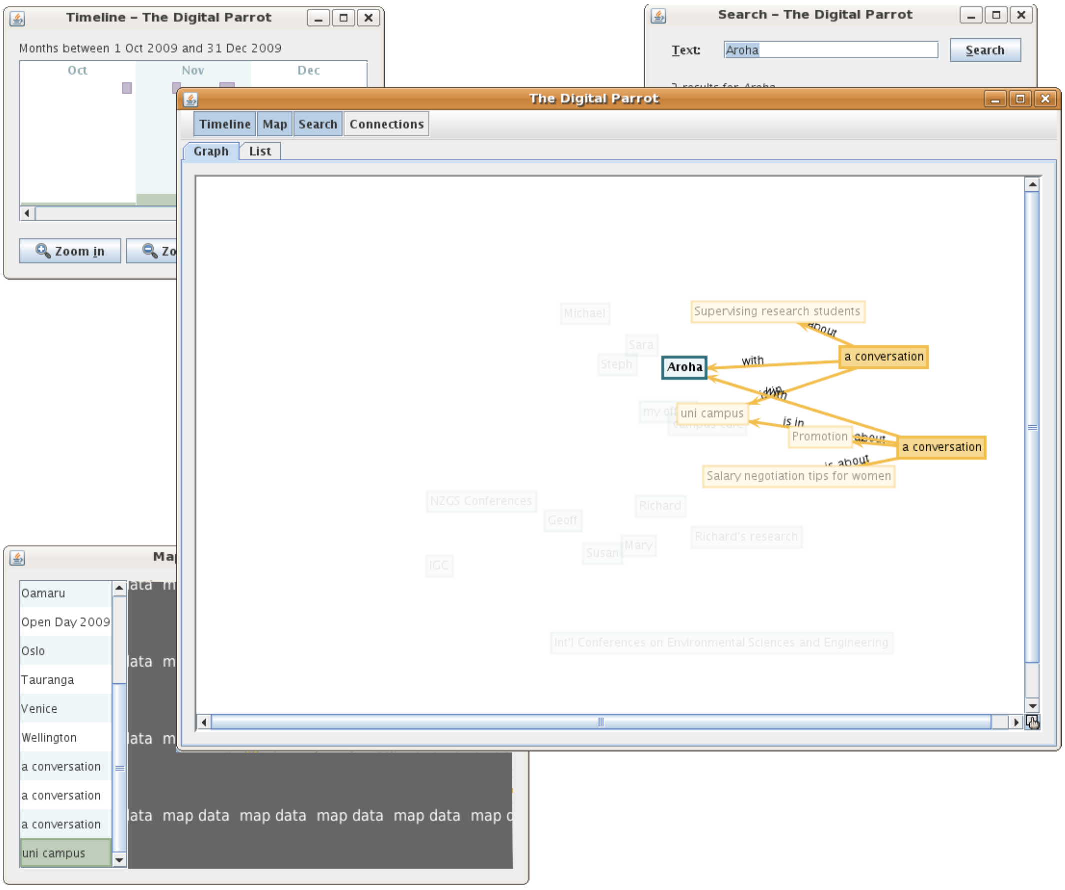

Figure 5.14. Selecting the conversation items to see their topics, as in Step 5 of the scenario. Restrictions and highlights as in previous step. Map data removed from screenshot for copyright reasons.

tion item, selecting the item. Items connected to the selected item are drawn more prominently than normal items. To better see the topics of the conversation, in the Digital Parrot the two conversations should be selected to view all connected items, including the conversation topics. This is shown in Figure 5.14.

5.4.2 Eric uses the Digital Parrot

This section goes through the scenario in Section 4.4.6, in which Eric tries to remember a book recommendation. For each step that Eric makes in the scenario, this section shows screenshots of the appropriate parts of the Digital Parrot's user interface and explains how to perform this step using the Digital Parrot. Deviations from the scenario are noted where applicable.

Step 1: "Eric opens the memory program on his N900."

The Digital Parrot's user interface on start-up is shown in Figure 5.15. This walkthrough uses the statement list as the main view, for two reasons: firstly, because the statement view of the current, desktop-only version of the Digital Parrot can be used directly on a mobile device such as a Nokia N900 Internet tablet; secondly, to give more examples of the list view. It also fits in with Eric's characterisation as a "language geek". As explained in Section 5.3.1, the list view initially shows all information items in the system.

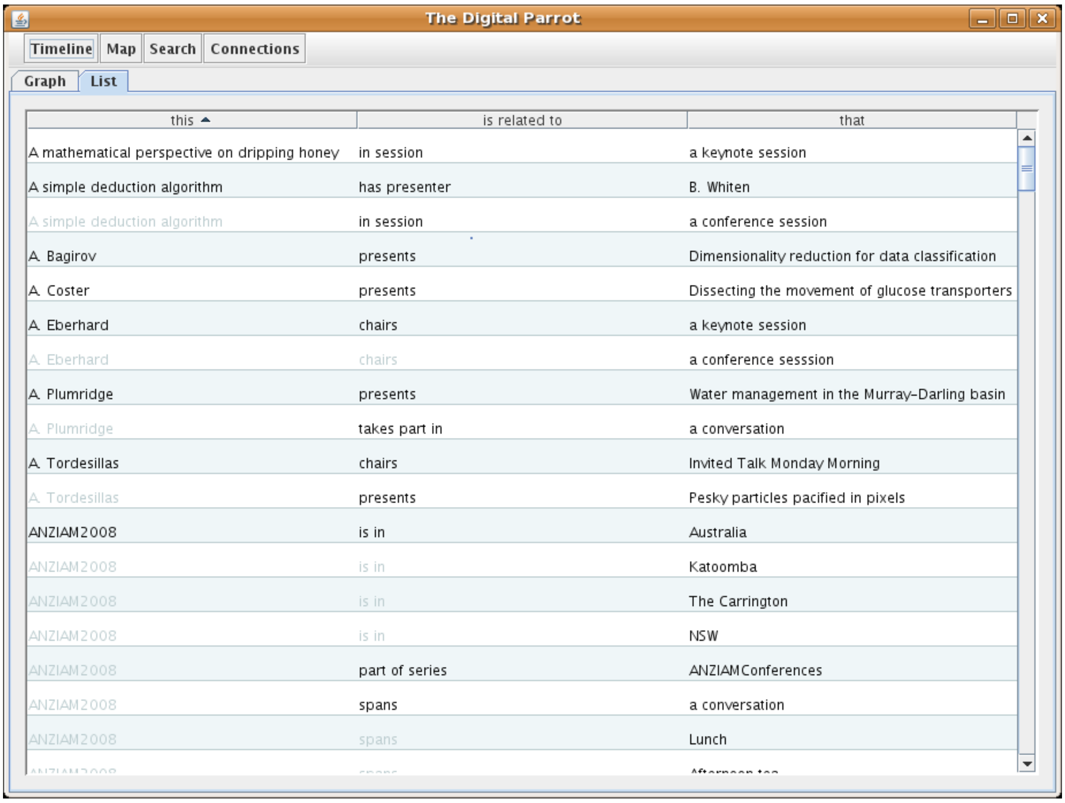

Figure 5.15. The Digital Parrot's user interface on startup, showing all information items in the system. This corresponds to Step 1 of the scenario.

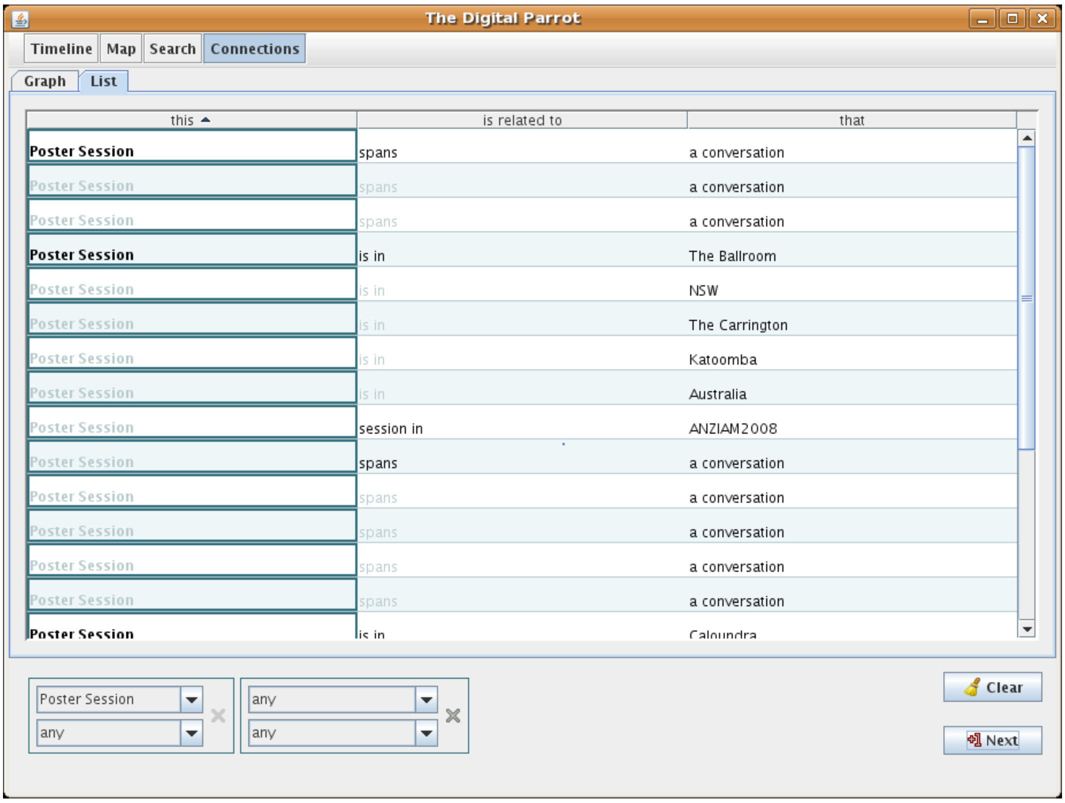

Figure 5.16. Restricting the main view to statements whose subject is of type "Poster Session", as at the beginning of Step 2 of the scenario.

Step 2: "He tells the program to show all *conversations* that are connected to *poster sessions*. [...] [H]e tells the program to show all *books* connected to conversations connected to poster sessions."

This is accomplished in the Digital Parrot using the connections navigator. The chain is constructed by first setting the type of the first link to "Poster Session". An additional empty link needs to be added to the chain for any statements to be shown. This is necessary in the list view because this view always shows complete statements (which consist of *two* class instances), unlike the graph view which can show individual nodes. The empty link can be appended by clicking the connections navigator's "Next" button. Figure 5.16 shows the Digital Parrot's user interface at this stage; the main view is restricted to the list of statements whose subject is of type "Poster Session".

To extend the statements shown to those about "books connected to conversations connected to poster sessions", two more changes to the chain are needed. First, the rightmost chain link is set to the type "Conversation". This restricts the main view to all statements whose subject is of type "Poster Session" and whose object is of type "Conversation". The result of this action is shown on the left side of Figure 5.17; the main view now shows nine conversations that are connected to three different poster sessions. Then, a third chain link is added using the "Next" button and its type is set to "Book". This removes some statements from the main view and adds others:

- statements whose object is not connected to a book are removed;
- statements whose object is connected to a book are retained; and
- all statements are added
  - whose subject is an object in one of the statements retained from the previous chain and
  - whose object is of type "Book".

The result is shown on the right side of Figure 5.17; the main view now contains one conversation that is connected to a poster session and also connected to a book ("Stochastics in Networks").

In the scenario, "[t]here is still quite a lot of information shown" when Eric requests the system to show all conversations connected to poster sessions. As can be seen on the left of Figure 5.17, the list view of the Digital Parrot does not actually show much information directly. However, there are nine different conversations shown and Eric would have to go through each of them individually to find the correct one if he did not want to extend the chain.

Step 3: "[. . .] Eric goes back to looking at conversations connected to poster sessions and now looks for all *people* connected with these conversations."

This is done in the Digital Parrot by changing the type of the rightmost chain link from "Book" to "Person". Figure 5.18 shows the Digital Parrot's user interface after this change. The main view now shows statements related to nine people, each of whom is connected to a different

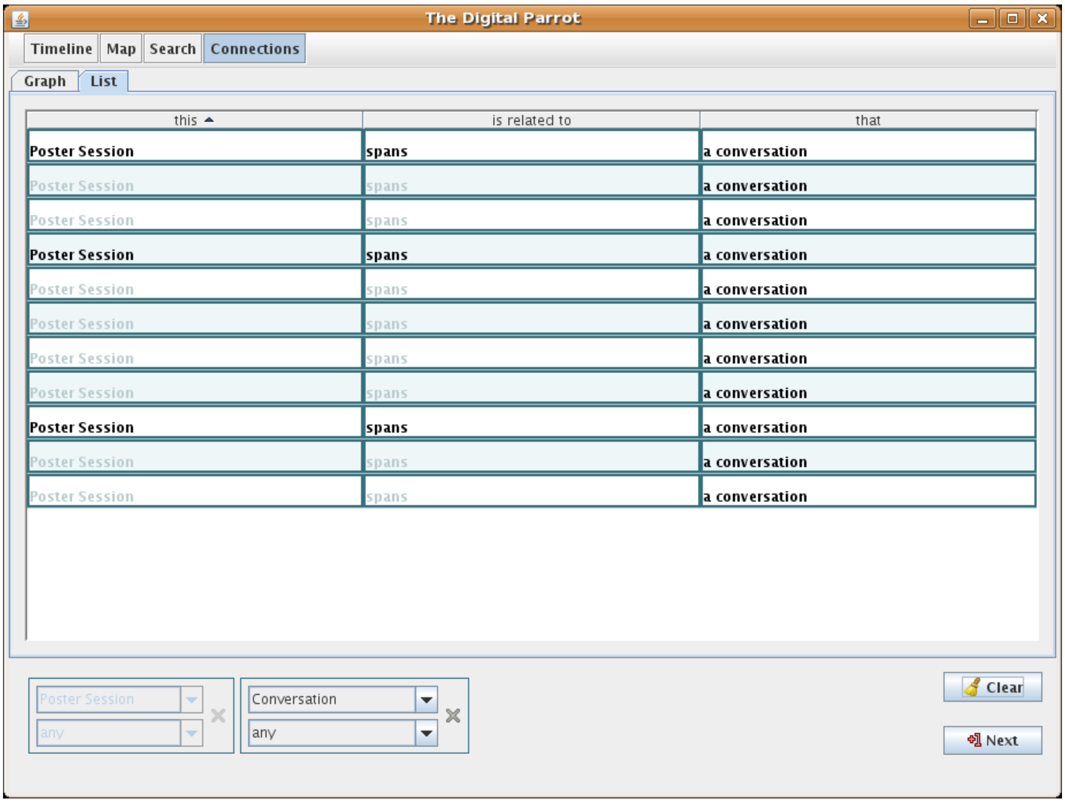

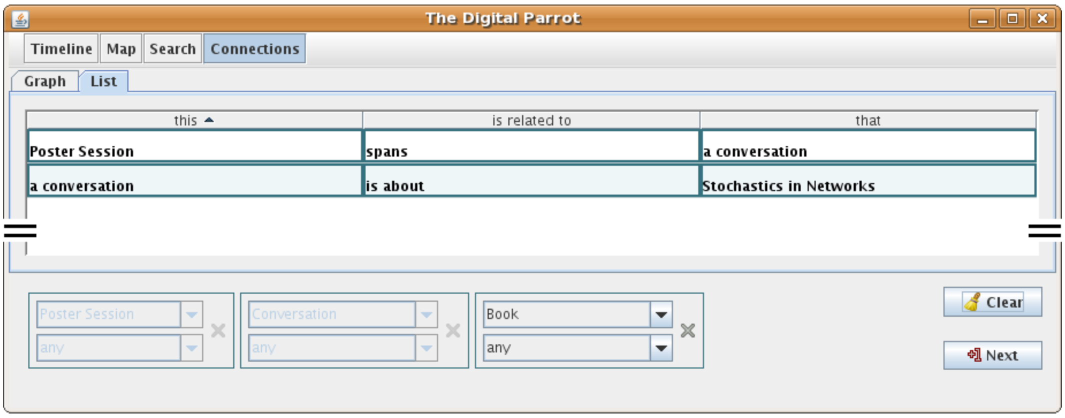

Figure 5.17. Restricting the main view further by appending to the chain the type "Conversation" (left) and then additionally the type "Book" (right), as in the second part of Step 2 of the scenario.

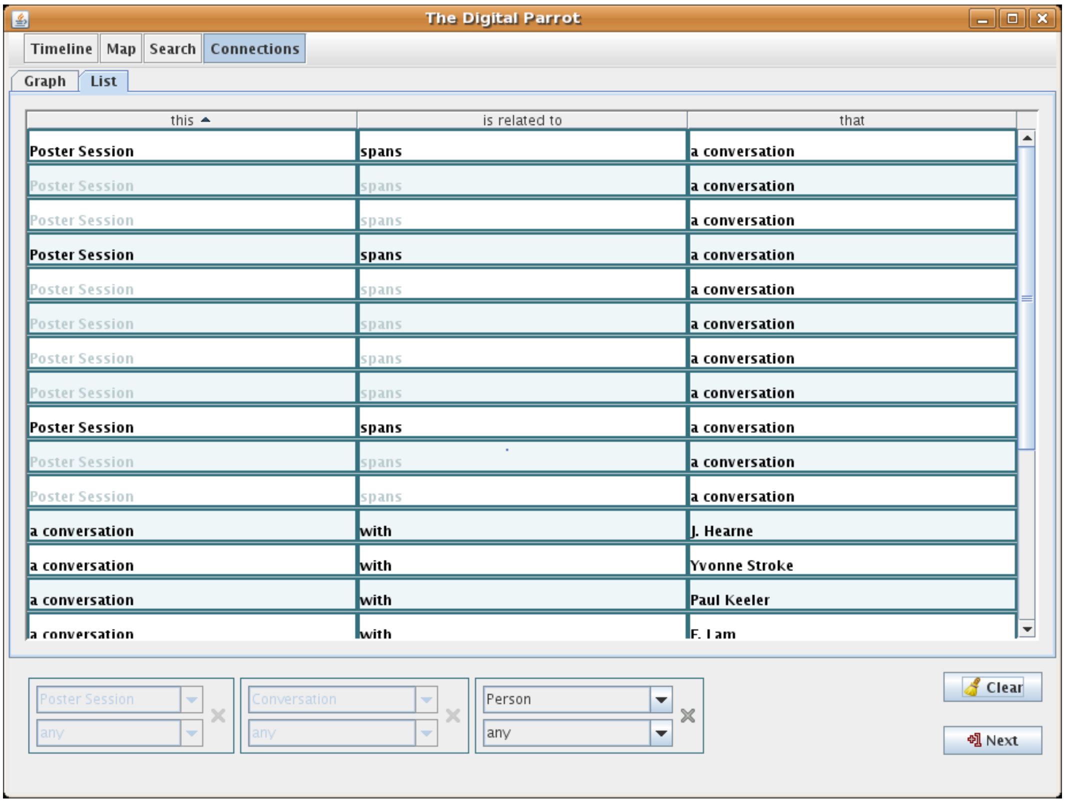

Figure 5.18. The Digital Parrot after "Person" has been substituted for "Book" in the chain, see Step 3 in the scenario.

conversation. Each conversation is in turn connected to one of three different poster sessions.

Step 4: "[...] He tells the system to show the topics of all conversations connected to this person [...]"

This can be done in the Digital Parrot by restarting the chain at the item corresponding to the person and then adding "Conversation" and "Topic" to the new chain. First, the chain is restarted at the item corresponding to the person (here, "Yvonne Stroke"). This can be done in two ways:

1. by clearing the chain using the "Clear" button, then setting the remaining chain link's type to "Person", choosing the person from the list of instances and appending an empty chain link; or
2. by double-clicking the person's item repeatedly – this first sets the instance of the original chain's "Person" link to the person, then

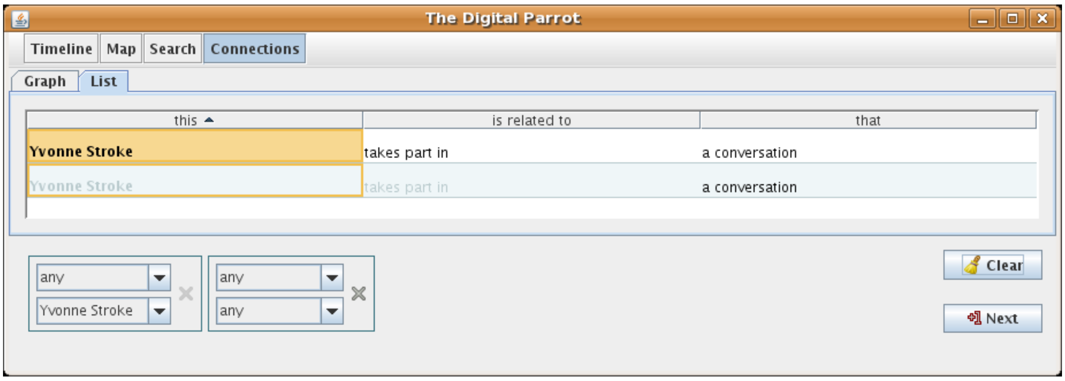

Figure 5.19. Restarting the chain at the person "Yvonne Stroke", as at the beginning of Step 4 of the scenario.

appends an empty chain link and finally re-starts the chain at the person's item.

Following either way, the main view then shows all statements whose subject is "Yvonne Stroke", as shown in Figure 5.19. After that, the type of the rightmost chain link is set to "Conversation"; this does not change the statements shown in the main view because the objects of both statements shown already are of type "Conversation". Then, an empty link is appended to the chain using the "Next" button to show all statements starting at one of the two conversations. The result is shown on the left side of Figure 5.20. Finally, the type of this link is set to "Topic" to hide all of these statements whose object is not a conversation topic.

The result is shown on the right side of Figure 5.20. The book that Eric tried to remember is selected ("Industrial Mathematics"). The Digital Parrot's type model allows an item to have more than one type (in this case, "Book" and "Topic"); Eric could have chosen "Book" as the type for the final chain item instead of "Topic".

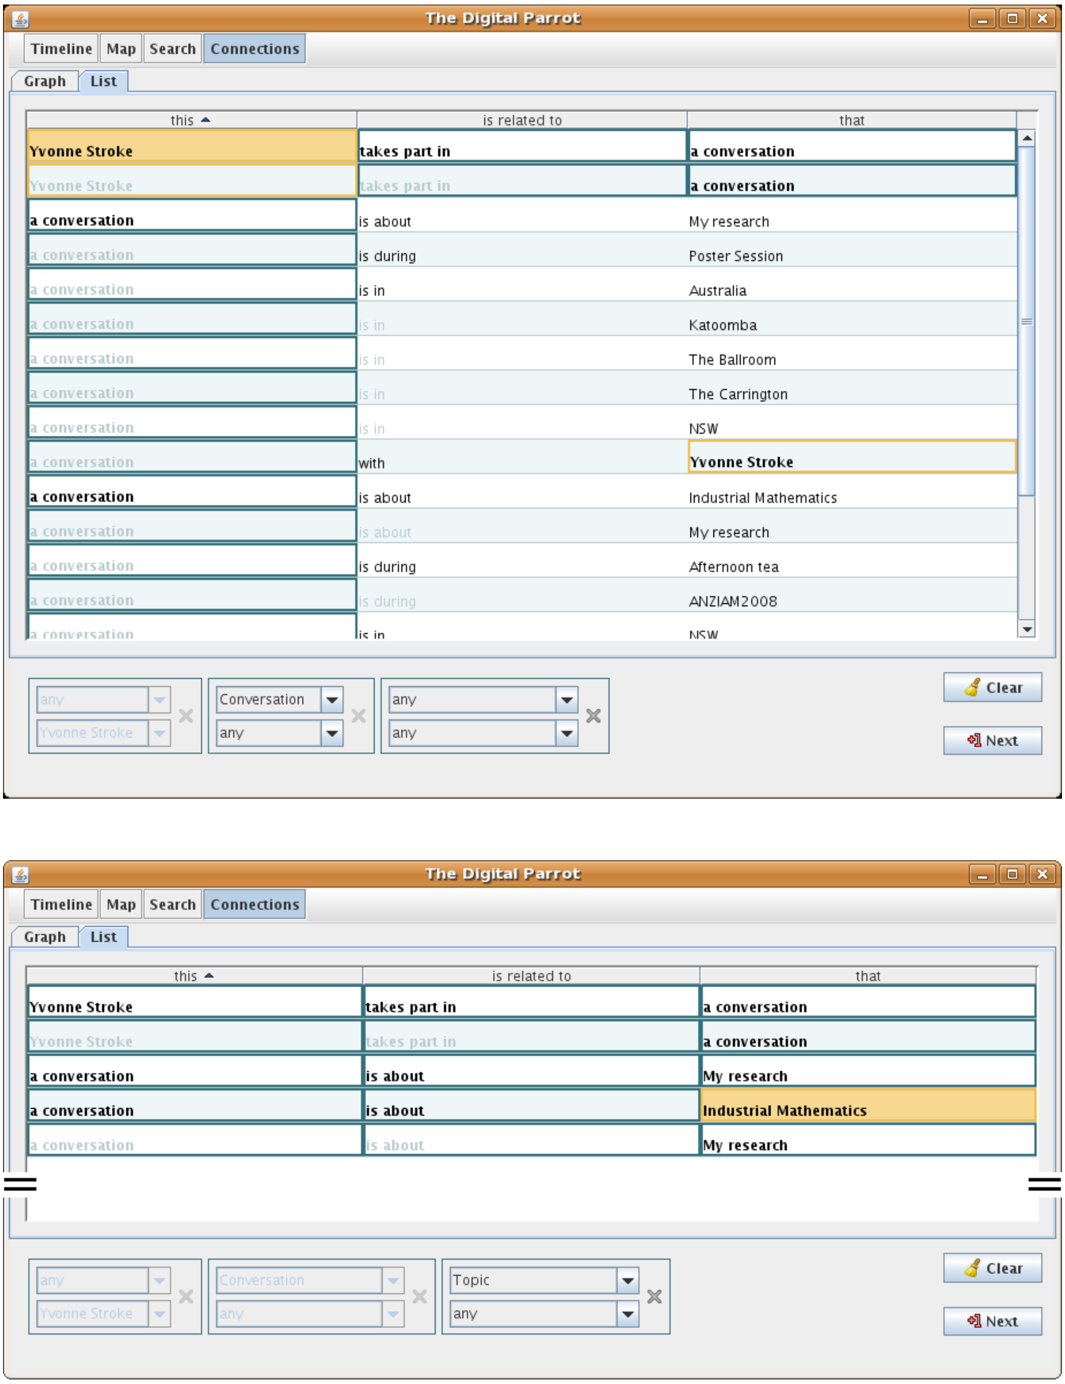

Figure 5.20. Viewing conversations (left) and topics of conversations (right) connected to a person, as at the end of Step 4 of the scenario. The final result is selected on the right.

5.5 Substitute for capture component

As explained in Section 5.1, the current version of the Digital Parrot focuses on the remembering phase. It does not include the capturing component that is part of the conceptual design, nor does it seek to address requirement U6 that is related to adding information to the system and modifying information in the system.

Instead of using the capture component, the Digital Parrot receives the information to be shown from a file. This file is read by the Digital Parrot once on start-up, when it is parsed to populate the data model.

5.6 Data model

This section describes how the Digital Parrot fulfils the requirements related to the storage component and data model, requirements D1 to D4.

5.6.1 Typed, graph-based model

To accommodate the graph structure consisting of information items and connections between information items (requirement D1), the data model of the Digital Parrot is based on the Resource Description Framework (RDF)[1]. Information in RDF format is represented in triples that directly map to the subject–predicate–object statements used in the Digital Parrot's user interface (see Section 5.3.1). Since the object of one statement can be the subject of other statements, a set of statements forms a (potentially disconnected) graph, with subjects/objects as nodes and predicates as edges.

Figure 5.21 shows the corresponding graph and RDF statements for the English language statement "CHINZ 2009 is located in Auckland". This and all following RDF and OWL examples are given using N3 notation (Berners-Lee and Connolly, 2008).

To allow memory items and connections to be typed (requirement D2), the data model uses ontologies defined in the Web Ontology Language (OWL; Bechhofer et al., 2004). The Digital Parrot uses the OWL DL sub-

[1] http://www.w3.org/RDF/

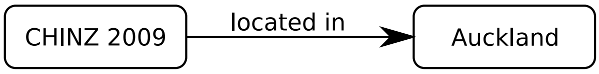


```
:CHINZ2009 timeplace:locatedIn :Auckland .
```

Figure 5.21. Graph and RDF representations of the statement "CHINZ 2009 is located in Auckland"

set of OWL version 1.

OWL ontologies provide types for subjects/objects and predicates by defining classes and properties. The subject of an RDF triple is an instance of one or more OWL classes. The predicate of an RDF triple is an instance of an OWL property. The object of a triple is either an instance of one or more OWL classes (where the predicate is an object property) or a literal (where the predicate is a datatype property). A literal can be typed (e. g. date, number or string) or untyped. OWL properties can have a defined domain (subject classes) and range (object classes). Classes and properties form inheritance hierarchies, since a class or property can extend other classes or properties.

Not all type information in OWL needs to be declared explicitly. Since OWL semantics are built on description logic, further information can be inferred from explicitly declared information. Figure 5.22 shows declared type information for the statement shown in Figure 5.21 as well as further inferred type information.

The Digital Parrot's code explicitly references only two ontologies: the Time and Place ontology, for contextual information, and the Digital Parrot ontology, for annotations of types and properties to be shown in the user interface. Both ontologies are described in more detail below (Section 5.6.2 and 5.6.3). They are given in full in Appendix A.1. All other ontology data is customisable by the user; this is described in Section 5.6.4.

5.6.2 Context ontology

A custom ontology, the Time and Place ontology, was developed to describe temporal and geospatial context of memory items in the Digital Parrot (requirement D3). This section describes the content of the on-

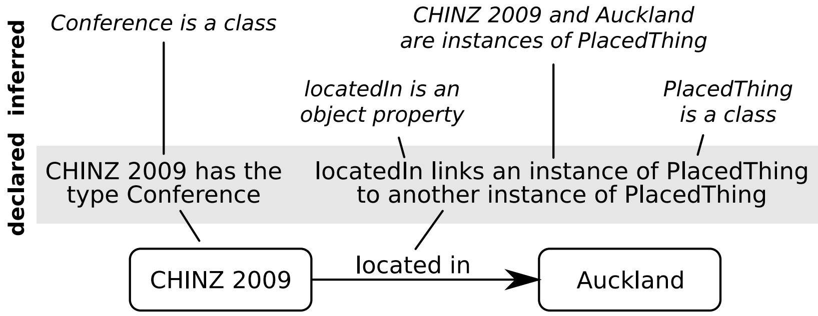


```
:CHINZ2009 a conf:Conference .

timeplace:locatedIn rdfs:domain timeplace:PlacedThing ;
  rdfs:range  timeplace:PlacedThing .
```

Figure 5.22. Declared and inferred OWL information for the statement "CHINZ 2009 is located in Auckland"

tology and how it is used in the Digital Parrot. Figure 5.23 shows a visualisation of the classes, properties and individuals in this ontology. Appendix A.1.1 gives the full listing.

Classes and properties

The Time and Place ontology contains one base class each for items with geospatial and temporal information: PlacedThing and TimedThing. Context information for instances of these two classes can be vague. Both base classes have extensions for items which can be anchored in space or time unambiguously. These extensions are AbsolutelyPlacedThing and AbsolutelyTimedThing. An item can be anchored in space unambiguously when it has coordinates for both latitude (property lat) and longitude (property long). Similarly, an item that can be anchored in time unambiguously when it has both a starting time (property startsAt) and an ending time (property endsAt). Both coordinate properties take string literals as values. Both time properties take XML Schema datetime literals as values.

The ontology contains some additional object properties. Most of these describe relationships between pairs of PlacedThing instances or between pairs of TimedThing instances. The property encloses and its inverse, lo-

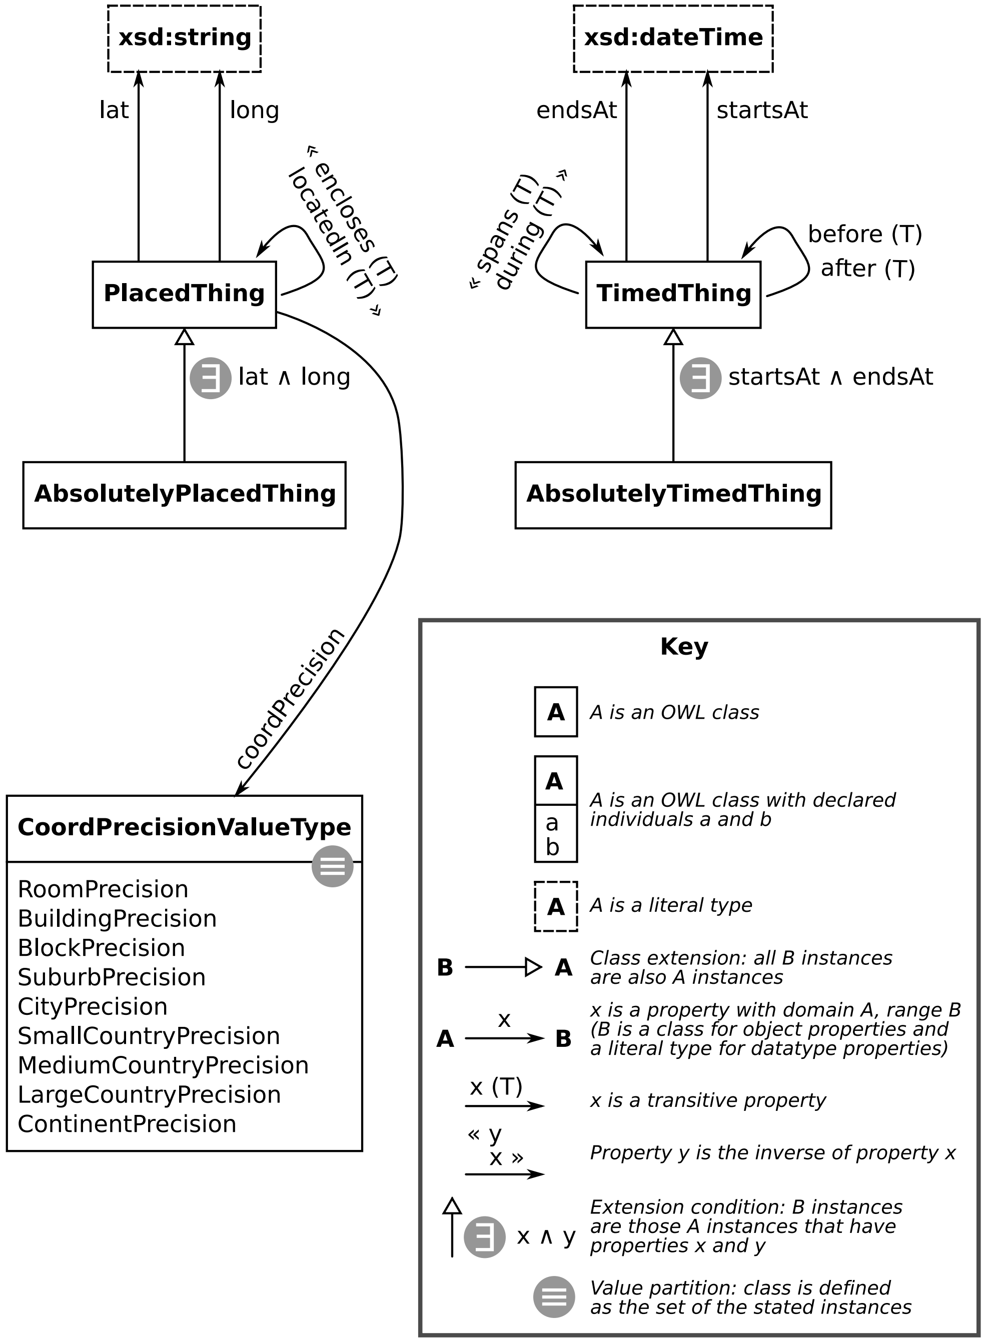

Figure 5.23. Visualisation of the Time and Place Ontology

```
:CHINZ2009 timeplace:locatedIn :AucklandUni ;
  timeplace:startsAt "2009-07-06T10:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-07T18:00:00+13:00"^^xsd:dateTime .

:AucklandUni timeplace:locatedIn :Auckland .
  timeplace:lat "-36.85167"^^xsd:string ;
  timeplace:long "174.769574"^^xsd:string ;
  timeplace:coordPrecision timeplace:BlockPrecision .
```

Figure 5.24. Context information in RDF

catedIn, relate PlacedThing instances to each other. Likewise, the property spans and its inverse, during, related TimedThing instances to each other, as do the two properties before and after. All of these properties are transitive.

An additional object property, coordPrecision, describes the approximate precision of the coordinates of a PlacedThing instance. Its range is CoordPrecisionValueType, which is a partition value type and defined as the set of its nine instances as shown in Figure 5.23.

Figure 5.24 shows context information for CHINZ 2009, the conference used for the examples in the previous section.

Use in the Digital Parrot

The Digital Parrot uses this ontology to determine which items are affected by the map navigator and the timeline navigator (Section 5.3.3). Both navigators show only those items that are of a suitable type. Restricting and highlighting by both navigators also works only on these items.

The map navigator works on all instances of AbsolutelyPlacedThing, both declared and inferred. Additionally, the map navigator works on instances of PlacedThing that are not instances of AbsolutelyPlacedThing (i. e., they do not have values both for latitude and longitude) if they have a locatedIn relationship with an instance of AbsolutelyPlacedThing. In this case, the values for latitude, longitude and coordinate precision of the enclosing item are used if they are not specified for the item itself.

Similarly, the timeline navigator works on all instances of Absolutely-TimedThing, whether declared or inferred. It also works on all instances of TimedThing that are not instances of AbsolutelyTimedThing (i. e., they do not have values both for starting and ending time) if they have a during relationship with an instance of AbsolutelyTimedThing. In this case, the values for starting and ending time of the spanning item are used if they are not specified for the item itself.

Design of the context ontology

The Digital Parrot uses a custom-developed ontology to describe the geospatial and temporal context of information items. There is an existing ontology for geospatial information, the WGS84 Geo Positioning ontology[2], which could have been used for geospatial context. However, this ontology is expressed in RDF Schema rather than OWL and cannot easily be used in conjunction with OWL ontologies.

Parts of the Digital Parrot's context ontology that describe geospatial context replicate vocabulary of the WGS84 Geo Positioning ontology in OWL terms: latitude and longitude properties are almost identical in both ontologies, and the Digital Parrot's PlacedThing and AbsolutelyPlacedThing are similar to SpatialThing and Point in the WGS84 Geo Positioning ontology.

The treatment of both temporal and geospatial context in the current version of the Digital Parrot is designed for the current focus of the implementation (see Section 5.1). Future extensions of the Digital Parrot to include a true capture component may need to include modifications to the Digital Parrot's treatment of context.

5.6.3 The Digital Parrot ontology

The second ontology that the Digital Parrot's code explicitly refers to is the Digital Parrot ontology (full listing in Appendix A.1.2). Figure 5.25 visualises the content of the ontology.

This very small ontology is used to determine which additional classes

[2] http://www.w3.org/2003/01/geo/wgs84_pos

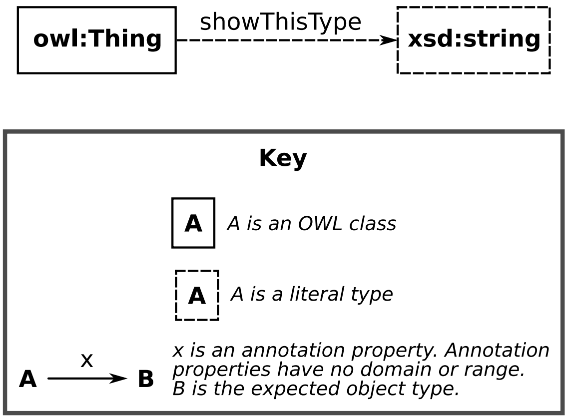

Figure 5.25. The Digital Parrot ontology

```
conf:Conference parrot:showThisType "primary"^^xsd:string .

timeplace:startsAt parrot:showThisType "secondary"^^xsd:string .
timeplace:endsAt parrot:showThisType "secondary"^^xsd:string .
```

Figure 5.26. The showThisType annotation

and properties should be considered by the Digital Parrot. It contains only the annotation property showThisType. The Digital Parrot expects this property's value to be one of the string literals "primary" or "secondary". Figure 5.26 shows the property's usage.

When a class or property is annotated with a showThisType value of "primary", it is included in the "type" drop-box of each chain element in the connections manager (Section 5.3.4). The instances of the class or property are directly shown in the Digital Parrot's user interface. In contrast, when a class or property is annotated with a showThisType value of "secondary", the class is not included in the connections manager and its instances are shown in a less prominent manner (currently, in a tooltip).

Instances of classes or properties that are not annotated with one of these two showThisType are not shown in the Digital Parrot's user interface at all. Where an item or predicate is instance of more than one class or property, the highest associated showThisType value is used.

5.6.4 Custom ontologies and user data

The Digital Parrot allows the use of additional ontologies to describe the memory data in the system. Custom classes and properties are shown by the Digital Parrot according to their annotation as described in the previous section.

Two further custom ontologies were created, one for the domain of academic conferences and one for more general interaction between people. The Conference ontology defines classes and properties such as Conference, ConferenceSeries, PosterSession, hasChair and sessionIn. The Interaction ontology defines classes and properties such as Conversation, Topic, hasConversationPartner and topicIn. Both ontologies refer to an ontology created by the Friend-of-a-Friend project[3] to describe people.

Figure 5.27 shows statements that use custom classes and properties. Appendix A.2 gives the full RDF file used for producing the screenshots in Section 5.3. Appendix A.3 gives the Conference ontology and the Interaction ontology.

5.6.5 In-memory storage

The Digital Parrot uses no dedicated storage component; all RDF and OWL data is read once from the text file and then held in memory. For the amounts of data used in the evaluations of the Digital Parrot that were conducted in the scope of this thesis, in-memory storage was sufficient in terms of both memory requirements and speed.

Future extensions of the Digital Parrot may need to deal with larger amounts of data or with changes to the data during the runtime of the program. There are a number of dedicated triple stores available that can be used.

5.7 Implementation environment

The Digital Parrot was written in Java. Java Swing is used for the user interface and a number of external Java libraries are used for specialised

[3] http://www.foaf-project.org/; OWL version at http://www.mindswap.org/2003/owl/foaf

```
:Session2NZCSRSC08 a conf:PresentationSession ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C3LectureTheatre ;
  timeplace:startsAt "2008-04-15T14:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T15:50:00+13:00"^^xsd:dateTime .

:TalkingToAnnAndSeanAtBreakfast a interact:Conversation ;
  timeplace:during :BreakfastTuesdayNZCSRSC09 ;
  interact:hasConversationPartner :Ann, :Sean ;
  interact:hasTopic :AnnsResearch, :Web2Dot0 .
```

Figure 5.27. Data described with custom classes and properties

parts of the program. Most of these libraries are available under open-source licenses. The parts and libraries are:

RDF and ontology data: Jena Semantic Web Framework for Java[4]. Some additional reasoning is provided by Pellet: the Open Source OWL Reasoner[5].

Graph view: Java Universal Network/Graph Framework, JUNG 2.0[6], for the visualisation itself and also for some underlying graph algorithms.

Timeline navigator: Joda Time API[7] for internal representation of time-related objects and for time-related calculations.

Map navigator: WebRenderer Swing Edition[8] to embed a web browser showing Google Maps[9]. The backend uses a QuadTree implementation from OpenMap[10], which was used instead of WebRenderer and Google Maps in an early version of the Digital Parrot.

Text search: Apache Lucene for Java[11] as the indexer and search engine.

4 http://jena.sourceforge.net/
5 http://clarkparsia.com/pellet/
6 http://jung.sourceforge.net/
7 http://joda-time.sourceforge.net/
8 http://www.webrenderer.com/products/swing/
9 http://maps.google.com/
10 http://www.openmap.org/
11 http://lucene.apache.org/java/docs/

Java was chosen mainly because of the availability of external libraries, especially those for working with Semantic Web data and for graph visualisation. Java's cross-platform availability was also a factor. All tests of the Digital Parrot were conducted using Linux. However, the only platform-specific code is that used by WebRenderer, which is available for all three major operating systems.

5.8 Discussion

Section 5.1 explained the decisions in focus that were made in implementing the Digital Parrot with regards to the first three recommendations made in Section 2.4.3. This section discusses the Digital Parrot with regards to the remaining three recommendations.

5.8.1 Context

As in the conceptual design, the main types of context represented in the Digital Parrot are temporal and geospatial context – the time and the location of an experience. Temporal context is expressed both using timestamps and in a semantic way; it is visualised on the timeline (for items that, directly or indirectly, have temporal context in the form of timestamps) and using "content" objects (using the time-related classes and properties in the context ontology). Similarly, spatial context is expressed both using GPS latitude/longitude coordinates and in a semantic way. Spatial context is visualised in three ways: using markers on the map and in the list of items with location information shown next to the map (for items that, directly or indirectly, have temporal context in the form of GPS coordinates), as well as using "content" objects (using the location-related classes and properties in the context ontology).

Information items can be traversed using their context. The geospatial context of an information item, beyond that expressed through other information items, can be accessed from the main visualisation by requesting to show a specific information item on the map. The temporal context of an information item, beyond that expressed through other information items, can be accessed from the main visualisation via tooltips on information items.

The overlay of the main visualisation onto a timeline or map that was included in the design sketches (see Section 4.5.1) was not included in the Digital Parrot because of conceptual issues around the representation of items without temporal/spatial information in these visualisations. More explicit ways to access an information item's context were considered but not implemented due to conceptual issues.

5.8.2 Semantic information

Well-known mechanisms from the Semantic Web are used to express both the graph structure of information items and connections and semantic information. Semantic information is included in the Digital Parrot as OWL classes (types of information items) and properties (types of connections). OWL and RDF were chosen over more traditional relational or object-oriented data models because they allow set-based inferences similar to those occurring during reconstructive remembering and because their flexibility is better suited to the semi-structured nature of the information to be represented.

Access to information items via their types is provided in the connections navigator. The user interface of the connections navigator goes far beyond the simple type-based interaction included in the design sketches. It is similar to other research prototype user interfaces found in the literature.

5.8.3 Associations

The graph structure of information items and connections is expressed via RDF triples. The graph view follows the design sketches and makes this structure explicit. The graph view makes it easy to traverse the information in the system by following connections between items. A list of statements view is provided to allow for comparison with a more text-based information representation.

5.9 Summary

This chapter contributes further to answering the second research question identified in Section 1.1.2, how an interactive software system can help someone remember.

It introduced a an implementation of the conceptual design described in the previous chapter, the Digital Parrot. The Digital Parrot focuses on the remembering phase of the interaction described in Section 4.1. Thus, it allows its user to retrieve facts and cues for memories with interaction methods grounded in results from Cognitive Psychology as surveyed in Chapter 2. It incorporates aspects of existing Computer Science approaches to augmenting memory as analysed in Chapter 3.

Chapter 6

Evaluation methods for augmented memory systems

The previous two chapters described the conceptual design and a selective implementation of an augmented autobiographical memory system. The goal of this system, following the thesis objective, is to support its users in remembering past experiences and information related to these experiences. To determine whether the system reaches this goal, the system needs to be evaluated.

The third research question in Section 1.1.2 asks about methods that can be used to evaluate augmented autobiographical memory systems. This chapter contributes to answering this question in two ways: it reviews challenges associated with such evaluations and it describes strategies used to overcome these challenges in evaluations of similar systems.

The chapter considers evaluations of augmented autobiographical memory systems as well as systems that deal more broadly with personal information in general (i. e. PIM systems, see Section 3.3). This is because most of the challenges faced in evaluating systems for personal information also apply to evaluating augmented autobiographical memory systems. Evaluations of augmented autobiographical memory systems can build on strategies developed for evaluating PIM systems. A much wider range of evaluation methods has been published for PIM systems than for augmented autobiographical memory systems and other systems designed for personal memories (such as those in the area of CARPE, see Section 3.2). However, evaluations of augmented autobiographical memory systems need to overcome additional challenges that are specific to such systems.

This chapter is structured as follows. Section 6.1 reviews challenges involved in evaluating PIM systems and systems designed for personal memories. Section 6.2 describes evaluation methods and strategies for both types of systems that have been developed to overcome these challenges. Section 6.3 discusses the implications for the research presented in this chapter for this thesis. The chapter concludes with a summary in Section 6.4.

6.1 Challenges

This section describes challenges involved in evaluating PIM systems and systems designed for personal memories.

6.1.1 PIM systems

Two main reasons make PIM systems difficult to evaluate (Kelly, 2006; Kelly and Teevan, 2007). Firstly, PIM systems deal with personal information, which differs greatly between individuals. People also typically have a lot of implicit knowledge about the information used in these systems (Cutrell et al., 2006b). Secondly, people are very used to their established ways of dealing with this information. This makes it difficult to genuinely evaluate a new approach.

Kelly and Teevan (2007) advocate a combination of naturalistic, longitudinal, case study and laboratory approaches for evaluating PIM systems. They identify challenges in all types of evaluation related to participants, collections, tasks, baselines and measures. A similar discussion is presented by Elsweiler and Ruthven (2007). The challenges involved in evaluating PIM systems can be summarised as follows:

Participants for PIM system evaluations are generally hard to find because of the time and effort required to participate, particularly for evaluation methods that are not short-term and not lab-based.

Data collections used in PIM system evaluations pose two main challenges. One one hand, artificially generated collections may not lead to realistic results because they are not linked to the participant and thus the participant lacks additional information about the data. On the other hand, using natural collections (for example, the user's own collection of photographs or e-mails) makes it harder to compare results across participants and also raises privacy concerns.

Tasks in PIM system evaluations need to be generic enough to allow for participants' idiosyncratic information management behaviour, while at the same time specific enough to allow for comparison across participants and/or techniques. If the data collection differs between participants, then the tasks need to either be flexible enough to be used across collections or also differ between participants. At one end of the spectrum are self-identified tasks and settings in which a system is deployed and the participants' integration of the system into their routines is observed. At the other end of the spectrum are precisely defined sets of tasks, possibly identical across participants, that are typically used in laboratory-based studies. Another challenge is to create tasks that are realistic enough for the participants to make evaluations meaningful.

Measures in PIM system evaluations need to be chosen with the PIM use context in mind. Standard measures from the field of Information Retrieval, namely precision and recall, are difficult to apply in this context because they rely on the availability of clear relevance judgements and on the knowledge of all potentially retrievable relevant information.

Standard usability measures, such as effectiveness, efficiency and satisfaction, do not pose many constraints on the evaluation design. However, they are not sufficient to fully evaluate PIM systems because they do not take into account whether they support the user in their needs with regards to managing their personal information. Other measures assess behavioural changes caused by PIM systems, such as whether and why people adopt a new system into their PIM behaviour. Yet other measures assess subjective and affective aspects.

Kelly and Teevan (2007) advocate using a mixture of qualitative and quantitative measures in PIM system evaluation.

6.1.2 Systems for personal memories

The challenges in finding participants for PIM evaluations apply equally to evaluations of augmented autobiographical memory systems. Most augmented autobiographical memory systems aim to support the user's natural memory over long periods of time (i. e., years rather than months or even shorter periods). If evaluations of such systems are to be naturalistic and longitudinal, they thus too need to be conducted over long periods of time. Participants for such studies need to be available throughout the study period. Another challenge in finding participants for memory-related evaluations is the high societal value that is placed on having a "good memory". Participants may be reluctant to sign up for such evaluations because they fear that their memory performance will be rated and exposed as not good enough.

Data collections are even more challenging for evaluations of augmented memory systems than for PIM systems. One important distinction between PIM systems and augmented autobiographical memory systems is that the information actually stored in augmented autobiographical memory system is only a part of the whole; in most cases, it is only a cue that can trigger remembering of the experience and the memory item itself is not stored in the system. This makes it even more important to use natural data collections in evaluating the effectiveness of augmented memory systems.

PIM system evaluations often have the option to acquire participants' already established natural collections that were created with other systems than the one to be evaluated. However, typically the cues stored in an augmented memory system are not already present in digital form for participants who do not already use the system. To create natural data collections, either the system needs to be used over a long period (preferably several years), or data needs to be obtained by some other means and then transferred into the system prior to the evaluation. The first approach is rarely possible in a research context, particularly because the system needs to be robust enough in that case to be used "in the wild" (Kelly and Teevan, 2007).

When it comes to task construction, evaluations of augmented autobiographical memory systems face the same challenges as evaluations of PIM systems. Additionally, there are a few challenges specific to augmented autobiographical memory systems. Evaluations of these systems often force the participants to remember their past, usually by requiring the participant to use the system to answer specific questions about previously captured experiences. This means that all tasks used in these evaluations are artificial to some extent because they assume that the participant does not remember the requested information and currently wishes to remember this particular experience.

Measures for evaluations of augmented autobiographical memory systems need to be chosen following the same criteria as for measures for PIM system evaluations. While Kelly and Teevan (2007) still consider traditional Information Retrieval measures, precision and recall, as applicable to PIM systems in some cases, these are hard to transfer to evaluations of augmented autobiographical memory systems. Autobiographical memory is strongly linked to the self; its functions related to the preservation of identity and self-image mean that unaided remembering of the same experience at different points in time does not necessarily lead to identical recollections. An augmented autobiographical memory system may wish to support or counteract these differences between recollections, and measures must be chosen for evaluations of the system to reflect the system's goal. Other possible measures for evaluating augmented autobiographical memory systems are

- the correctness of recollections made using the system, either as judged by the participant or against some external verification criterion;
- whether use of the system leads to larger quantities of recollections; and
- whether use of the system leads to more detailed recollections.

Time and time-related aspects have a much stronger impact on evaluations of augmented autobiographical memory systems than on evaluations of other systems. Most memory failures occur over time and are essentially unpredictable. Generally, time needs to pass for an experience to be forgotten. Evaluations that use natural data collections need to allow sufficient time between experience and remembering. One implication is that task order is extremely important; finding the answer to one task may make it easier to recall or find the answer to other tasks related to the same experience. For the same reason, tasks cannot be repeated within a short timeframe, which rules out straightforward within-subject experiment designs that require the participant to perform the same task using different systems.

6.2 Evaluation methods and strategies

This section reviews methods and strategies that have been used to address the challenges described in the previous section. Methods and strategies are described first for PIM systems and then for systems designed for personal memories.

6.2.1 PIM systems

To overcome challenges related to finding study participants, some studies of PIM systems used the principal researcher as the only participant (e. g. Rhodes, 2003). However, results from such an evaluation are generalisable only to a limited extent (Sellen and Whittaker, 2010). Another option is to lower the boundaries to participation, for example by conducting the evaluation over a shorter period (e. g. Sellen et al., 2007; Kalnikaitė et al., 2010).

To address challenges related to data collections and tasks, Elsweiler and Ruthven (2007) as well as Elsweiler et al. (2008) introduced a laboratory-based approach that combines the advantages of natural collections and personalised tasks with the comparability normally only afforded by using the same set of tasks for all participants. Interested specifically in tasks related to web and e-mail re-finding, they first conducted a diary study to capture tasks in this domain as they occurred naturally for their participants. From these tasks and supplemented by further in-depth investigation of their participants' practices for web and e-mail re-finding, they derived three general categories of tasks. They then constructed artificial tasks for each category, with the motivation for conducting a specific task taken from motivation for similar tasks recorded in the diaries.

The generated tasks were tailored specifically to each participant's data collection (messages in the university e-mail accounts of undergraduate and postgraduate students as well as academic staff). However, some tasks referred to e-mails that were sent to entire sub-groups of participants (such as seminar announcements sent to all staff or job vacancies sent to all students) and could thus be re-used for other participants in the same sub-group.

6.2.2 Systems for personal memories

This section describes two studies that show how challenges related to participants and data collections can be addressed in evaluations of approaches related to personal memories.

The evaluation of iRemember, an audio-based memory aid (see Section 3.2.1), by Vemuri et al. (2006) shows how to capture information about experiences without an augmented memory system that can later be transcribed and used to evaluate such a system. Vemuri et al. made audio recordings of conversations between the principal researcher and a small number of participants in the study over several years. This allowed them to create a natural data collection over a long timespan without the challenges associated with deploying a research prototype "in the wild". Since their research focused on retrieval technology, this phase of their study was conducted without the system under test. These recordings were then fed into the system and used for a second evaluation phase where participants were asked to use iRemember to answer questions about these conversations.

The study described by van den Hoven and Eggen (2009) that investigated the role of tangible artefacts in triggering re-living shows another method for creating experiences that can then used in a remembering phase. Their study explored which types of cues are most beneficial in evoking recollections, independent of any software system. The study used a two-phase approach in which participants were exposed in a semi-controlled way to experiences that were thought to create memories: participants visited a theme park and took part in a set of activities that were accompanied by additional stimuli (audio, scent). Participants were then later asked to recall activities they experienced at the theme park; presence and absence of stimuli related to the experience (audio, video, photograph, scent or created artefact) were used as controls in the experimental set-up of the second phase.

6.3 Discussion

The strategies reviewed in the previous section have to make various types of trade-offs to address the challenges described in Section 6.1. These trade-offs broadly fall into three categories: participants, data collections and tasks. Each category is described in more detail below. Some of these categories are related to one another. For example, the tasks in a study need to be tailored to the data collection: In most cases it is pointless to ask a participant to remember an experience using an augmented autobiographical memory system if no information related to this experience is stored in the system.

Participants. A typical trade-off made with regards to participant selection is to limit the number of participants involved in the study. It is not uncommon for evaluations especially in the area of CARPE to have very few participants, sometimes even only one. Another trade-off is often made with regards to the representativeness of the study participants, especially in longitudinal studies. Participants are often chosen from groups that have some close connection with the researchers conducting the study; this makes it easier to ensure that participants do not drop out where long time periods are involved.

Data collections. There are not many trade-offs that can be made with regards to the data collection used in an evaluation that still allow the evaluation to be meaningful. One trade-off is to collect data by some other means than with the system that is to be tested; this allows data to be collected before system development is complete and thus extends the timespan between experience and remembering. Another trade-off is to generate the experiences that are then later to be remembered. This may lead to slightly artificial experiences but on the other hand allows to test several participants' memories of the same or very similar experiences.

Tasks. Trade-offs around task selection involve the same interdependency between comparability and personalisation as in evaluations of PIM systems. Task-based evaluations (Elsweiler and Ruthven, 2007) offer a balance between the two factors.

6.4 Summary

This chapter reviewed existing work related to two aspects of the third research question: What challenges are involved in evaluating whether and how well an augmented autobiographical memory system supports its user in remembering past experiences and related information? What strategies have been developed to overcome these challenges? It described how evaluations of augmented autobiographical memory systems and of similar systems make trade-offs, typically related to one or more of three factors: participants, data collections and tasks.

Based on the findings of this chapter, the decision was made to conduct two user studies of the Digital Parrot. The first study evaluated the usability of an early version of the Digital Parrot and its effectiveness in helping to answer memory-related questions. In this study, ten participants answered the same set of questions about an artificial data collection based on another person's memories. This evaluation made its main trade-offs with regards to the data collection and the tasks.

The second study evaluated the effectiveness of an improved version of the Digital Parrot. In this study, four participants answered questions about their own experiences two years after the experiences had been made and recorded for use in the study. Strengths of the study were the use of mostly realistic data and the timespan of two years, a much longer timespan than is typical for evaluations of augmented autobiographical memories systems. It made its main trade-offs on the number and background of participants as well as on the tasks, which were standardised into categories but still personalised for each participant.

These two studies of the Digital Parrot are described in the following two chapters.

Chapter 7

Evaluating the usability of the Digital Parrot

This chapter contributes to answering the fourth research question, about the effectiveness of the approach to augmenting autobiographical memory introduced in this thesis. It describes an end-user study which was conducted to evaluate the Digital Parrot's usability. The study investigated how the Digital Parrot aided users in answering questions about a data set that was based on another person's memories of experiences and related information.

This chapter is structured as follows: Section 7.1 explains the focus employed in designing the study. Section 7.2 introduces the experimental design. Section 7.3 states the quantitative findings of the study, while Section 7.4 describes further observations and participants' comments. Sections 7.5 and 7.6 relate the findings to the two study goals. The chapter concludes with a summary in Section 7.7.

7.1 Focus

This section describes the focus employed in designing the study. It outlines the goals of the study and the trade-offs made to overcome the challenges involved in evaluating augmented memory systems.

7.1.1 Goals of the study

The study had two main goals. The main goal was to find usability issues with the Digital Parrot's user interface so that these could be resolved before conducting a second user study (described in the next chapter). The second goal was to determine how people use the Digital Parrot to answer questions about memories.

7.1.2 Trade-offs

The main trade-offs made in designing this study are related to the data collection and to the tasks. It was decided to use an artificial memory collection for the study. Likewise, the tasks are identical across participants to allow for cross-participant comparisons. Consequently, no actual remembering was expected to occur during the study.

This has the advantages that the number of participants in the study could be reasonable for a usability study and the pre-study effort per participant could be kept low. Careful selection of tasks still ensured that usability issues in the Digital Parrot could be detected with the study.

7.2 Experimental Design

This section gives details about the experimental design of the study: the method, participants and procedure of the study, the data collection and tasks, the computing environment and the types of data captured in the study.

7.2.1 Method

The study was a computer-based laboratory study in which the participants used the Digital Parrot to complete four tasks. To cut down on the effort required by the participant to learn to use the Digital Parrot, each participant used only one of the two main view options (graph view and list view, see Section 5.3.1) in a between-subjects design. The participants in this study were randomly assigned to one of the options, with an equal number of participants per option.

Since the study focused on gaining insights of the participants' use and perception of the Digital Parrot, the time for each session was not constrained.

7.2.2 Participants

The study had ten participants. All participants were members of the Computer Science Department at the University of Waikato; six were PhD students and four were members of academic staff. Two participants were female, eight were male. The ages ranged from 24 to 53 years (median 38 years, IQR[1] 15 years).

Participants were recruited via e-mails sent to departmental mailing lists and via personal contacts. Participants were not paid or otherwise rewarded for taking part in the study.

7.2.3 Procedure

The study was conducted in the researcher's office, using the researcher's desktop computer and account. However, care was taken to remove personal items from the desk to ensure that participants would not feel that they were intruding on the researcher's personal space.

After the researcher had obtained the participant's consent, the participant was provided with a workbook. A copy of the workbook and of all other material provided to the participants is shown in Appendix B. The workbook gave a quick introduction to the purpose of the study and a brief overview of the Digital Parrot's features. Once the participant had read the introductory page of the workbook, the researcher started the Digital Parrot and briefly demonstrated the main view and the four navigation options. It was explained to the participant that the Digital Parrot contained a part of the researcher's memories related to several academic conferences.

[1] The Interquartile Range (IQR) is a statistical measure of dispersion, i. e. of the variation around the central tendency. It is calculated as the difference between the third and the first quartile and thus describes the length of the interval within which the central 50% of the data fall. A larger IQR indicates a higher variation in the data. It is robust against outliers and can be applied regardless of the distribution of the underlying data.

The participant was then asked to use the Digital Parrot to perform the four tasks stated in the workbook. After the tasks were completed, the participant was asked to fill in a questionnaire about their experience with the Digital Parrot and to answer some background questions. The session was concluded with a brief discussion; the researcher typically asked some questions to follow up on observations made while the participant was working on the tasks, and the participant was invited to share any comments they might have had about the Digital Parrot and the study.

7.2.4 Memory data

The memory data provided to the participants was the same data that was used to produce the screenshots of the Digital Parrot in Section 5.3; an anonymised full listing is given in Appendix A.2.1. The memory data describes some of the researcher's memories, based on notes taken while attending five academic conferences in three years (2007–2009). All conferences took place in New Zealand: two conferences took place in Auckland, two in Hamilton and one in Christchurch. The five conferences belonged to two conference series; one series was represented with two conferences and the other series with three.

The memory data contains basic information about the two conference series and six conferences, including the time and the location for five of the conference. Temporal information is given as timestamps. Location information is given in most cases both as a semantic label ("C2 Lecture Theatre", "Auckland University") and as a latitude/longitude pair.

For two conferences, and to a lesser degree also a third, more detailed information about sessions and keynote presentations is included that has been taken from the conference programs. The data contains a total of 26 different sessions and seven keynote presentations.

The memory data contains information about ten conversation with nine different conversation partners; some of the conversations had more than one conversation partner and some people took part in more than one conversation. Seven additional people are mentioned in the data; these are keynote speakers at the conferences and not connected to any conversations. The memory data lists nine conversation topics; between none and two for each conversation. Again, some conversation topics are shared between several conversations. All conversations are linked to conference events based on time (such as "before the conference opening" or "during breakfast").

When the data file was loaded by the Digital Parrot, 144 distinct nodes were shown, forming 731 statements. This means that the graph view showed 144 nodes connected through 731 edges and that the list view showed 731 rows.

The choice to use memory data that described experiences of a person other than the participant was made to allow a larger number of participants to be recruited for the study. The creation of a naturalistic data collection using each participant's own memories is costly (see Section 6.1.2). Furthermore, a natural data collection is most beneficial in combination with a long time interval between an experience and its recall by the participant in the study. In the design of this study, it was decided to use another person's memories rather than using naturalistic data and a shorter time interval. The participants were recruited from a background that ensured that they were familiar with the domain of the memory data; some participants even shared some of the memories in the data collection.

7.2.5 Tasks

The tasks were the same for all participants. They were chosen based on expectations of typical tasks for the Digital Parrot. Furthermore, they were chosen such that they covered a wide range of strategies with regards to the types of information asked for and the types of information given in the task description.

All four tasks were phrased as questions about the researcher's experiences as recorded in the Digital Parrot. The four tasks were:

(T1) "To whom did I talk about scuba diving? Write their name(s) into the space below."

(T2) "Which conferences did I attend in Auckland? Write the conference name(s) into the space below."

(T3) "At which conference(s) did I speak to someone about Python during the poster session? Write the conference name(s) into the space below."

(T4) "In which place was the New Zealand HCI conference in 2007? Write the place name into the space below."

7.2.6 Computing environment

The user interface of the version of the Digital Parrot used in this study differed from that described in Section 5.3. The version described in Section 5.3 already includes changes made to the Digital Parrot based on the results of this study. Details of these changes are given in Section 7.5.2. Even though some of these changes are quite noticeable in the appearance of the Digital Parrot, none of the changes fundamentally affected the way the Digital Parrot works. In the following, the name "trail navigator" refers to the initial version of the navigator using semantic information. In the revised version of the Digital Parrot and in the description in Section 5.3.4, this navigator is called "connections navigator".

7.2.7 Data collected

Each participant was asked to think aloud while using the Digital Parrot. The researcher took notes throughout each session.

As part of the study, the participants were asked to rate the Digital Parrot on the System Usability Scale, a generic usability evaluation tool introduced by Brooke (1996). The wording of the questions in the study followed the minor modifications described by Bangor et al. (2009, p. 115), except that the phrase "the system" in each question were replaced with the phrase "the Digital Parrot". Table 7.1 shows the questions used. Each question was answered by giving a rating on a five-point Likert-like scale. The full questionnaire is shown in Appendix B as part of the participant workbook.

1. I think that I would like to use the Digital Parrot frequently.
2. I found the Digital Parrot unnecessarily complex.
3. I thought the Digital Parrot was easy to use.
4. I think that I would need the support of a technical person to be able to use the Digital Parrot.
5. I found the various functions in the Digital Parrot were well integrated.
6. I thought there was too much inconsistency in the Digital Parrot.
7. I imagine that most people would learn to use the Digital Parrot very quickly.
8. I found the Digital Parrot very awkward to use.
9. I felt very confident using the Digital Parrot.
10. I needed to learn a lot of things before I could get going with the Digital Parrot.

Table 7.1. System Usability Scale: Questions

7.2.8 Pre-study expert review

Prior to the study, a usability review of the Digital Parrot was conducted with two experts on usability in context-aware systems. Both experts were familiar with the conceptual design of the Digital Parrot. The goal of the expert review was to gain a first indication of serious usability issues with the Digital Parrot.

The expert review was conducted as a cognitive walkthrough of typical usage scenarios for the Digital Parrot. The memory data used was the same as for the usability study. The review revealed a small number of genuine faults in the Digital Parrot's behaviour. These were repaired before the study. Additionally, the expert reviewers requested some changes and additions to the Digital Parrot's functionality. The main issues noted by the reviewers were:

- The graph view is "too cluttered".
- The map and timeline give no indication which locations or time-spans contain data.
- The trail navigator is cumbersome to use, mainly because nodes can be added and removed only individually.
- The navigator windows get lost behind the main window.

In response to these issues, the following changes were made to the

Digital Parrot prior to the usability study:

Graph view: Nodes are now initially shown greyed out; the colour scheme was changed to more clearly distinguish a selected nodes from its neighbours; edge labels are now hidden; pairs of antiparallel directed edges are now collapsed into a single edge.

Map navigator: A histogram was added to indicate areas of interest; cities are now shown less obtrusively (in the study: not at all due to deadlocking issues with the underlying library); the colours for water and land were reversed.

Timeline navigator: A histogram was added to indicate areas of interest; the label for “week” slices was changed to be less confusing.

Trail navigator: The trail can now be replaced with a given node; clearing the trail is now possible with a single operation.

Other: The navigator windows now stay on top of the main window; some additional statements were added to make the memory data more balanced.

7.3 Quantitative findings

This section reports on the participants’ ratings of the Digital Parrot on the System Usability Scale (SUS) and on the participants’ accuracy in performing the tasks. Since the SUS is a usability measure, the participants’ ratings on this scale give an indication of the usability of the Digital Parrot and are thus related to the first goal of the study as described in Section 7.1.1. The participants’ accuracy in performing the tasks gives an indication of the effectiveness of the Digital Parrot in allowing the participants to answer questions about experiences and memories; it is thus related to the second goal of the study.

7.3.1 System Usability Scale

The median SUS score of the Digital Parrot is 65 out of 100 (min = 30, max = 92.5, IQR = 35), below the cut-off point for an acceptable SUS score (which is 70). The overall score of 65 corresponds to a rating between “ok” and “good” on the adjective scale introduced by Bangor et al. (2009).

The median SUS score in the graph condition alone is 80 (min = 42.5, max = 92.5, IQR = 40), which indicates an acceptable user experience and corresponds to a rating between “good” and “excellent” on the adjective scale. The median SUS score in the list condition is 57.5 (min = 30, max = 77.5, IQR = 42.5). The difference of the median SUS score between conditions is not statistically significant (Mann-Whitney[2] $U = 6$, $n_1 = n_2 = 5$, $p > 0.1$ one-tailed).

The average normalised score for Question 4 is slightly higher than the average across all questions: more than half of the participants did not think that they would need the help of a technical person to use the Digital Parrot. This high number may be explained by the Computer Science background of the participants. Question 5, related to the degree to which the various functions of the Digital Parrot were integrated, received on average a score that was lower than the average score across all question.

Figure 7.1 shows boxplots of the normalised SUS scores for each question, for both conditions combined.

7.3.2 Accuracy

Table 7.2 gives an overview of the correctness of answers for each task and for both main view conditions.

Four of the ten participants gave complete and correct answers in all four tasks and another three participants completely and correctly answered three tasks. Two participant completely and correctly answered two tasks, while the remaining participant did not give a complete and correct answer for any of the four tasks. This participant gave answers to Task 3 and Task 4 that did not even match the type of the required an-

[2] The Mann-Whitney test is a test for statistically significant difference between two independent groups of samples. Its function is similar to that of Student’s t test, but the Mann-Whitney test is non-parametric and can be used for ordinal scales (such as Likert scales).

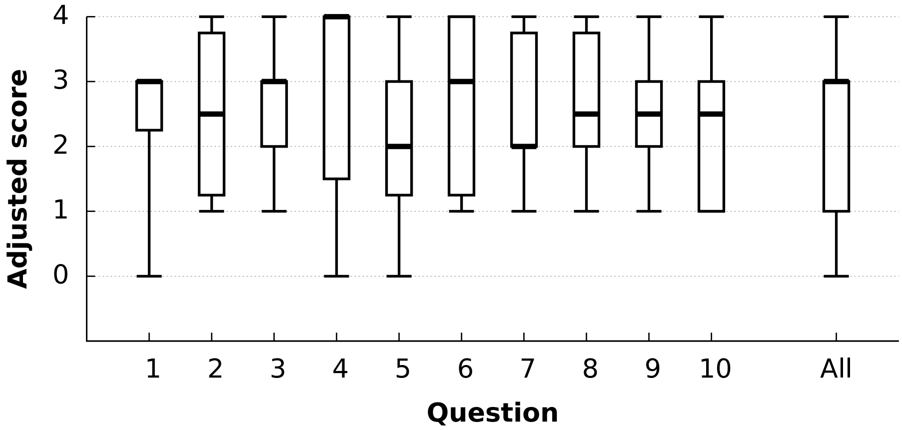

Figure 7.1. Boxplot of the SUS scores for each question. Scores have been adjusted such that a greater score correspond to a better experience in using the Digital Parrot regardless of the direction of the question.

	Task 1				Task 2				Task 3				Task 4			
Cond.	ok	extra	part	wrong	ok	extra	part	wrong	ok	extra	part	wrong	ok	extra	part	wrong
List	4		1		2	3	2		2		1	1	4			1
Graph	5				4	1			3		2		5			
Both	9		1		6	4	2		5		3	1	9			1

Table 7.2. Correctness of answers, by task and by condition. Columns: “ok” – a complete and correct answer; “extra” – an answer that had some correct information but also some extraneous information; “part” – an answer that missed some of the required information or (T3) is the result of a failure to distinguish between two items; “wrong” – an answer that is incorrect. An answer can fall into both the “extra” and “part” categories.

swer; the participant's answers to these tasks suggested that the participant failed to follow one further step which could have led to the correct answer (they gave the name of the conversation partner in Task 3 and the name of the conference in Task 4).

Task 1 and Task 4 appeared to be easiest to answer, with nine participants completely and correctly answering both. The remaining participant gave a partial answer in Task 1 and, as described above, an incorrect answer in Task 4.

Task 2 and Task 3 appeared harder to answer; they were answered completely and correctly by about half the participants each. In Task 2, two participants named only one of the two correct answers but gave additional items that were not in fact conferences. Two further participants named both correct answers and gave additional items. In Task 3, three participants chose the wrong answer from the two possibilities. One participant gave an incorrect answer, as described above. One participant, in the graph condition, gave up on this task; the participant located an item using search but then gave up when they could not determine how to use the trail from there.

There appeared to be no significant connection between the accuracy of a participant in performing the tasks on one hand and the participant's subjective experience of using the Digital Parrot on the other hand. Specifically, there was no significant connection between the number of correctly answered tasks and

- the overall SUS score given by a participant (Spearman[3] rank-order correlation coefficient $r_s \cong 0.50$, $n = 10$, $p \cong 0.068$ one-tailed);
- the ease of use question in the SUS questionnaire (Question 3 in Table 7.1; $r_s \cong 0.38$, $n = 10$, $p > 0.14$ one-tailed);
- the confidence question in the SUS questionnaire (Question 9 in Table 7.1; $r_s \cong 0.24$, $n = 10$, $p > 0.2$ one-tailed).

[3]Spearman's rank-order correlation coefficient indicates the degree of statistical correlation between two samples. It serves the same function as Pearson's product-moment correlation coefficient but can be used for ordinal data.

7.4 Observations

This section reports on observations made during the study. It first describes the strategies that the participants used to solve the four tasks. It then states overarching observations.

The quantitative findings reported in the previous section indicate that there were some usability issues with the Digital Parrot's user interface as used in the study but that it nevertheless allowed the participants to answer questions about another person's memories. The observations in this section are organised first by task and then by component. They describe how the participants went about answering the tasks using the Digital Parrot. These qualitative findings provide insights both into the nature of the usability issues and into the roles of the Digital Parrot's components. Consequently, they are related to both study goals.

7.4.1 Task-related observations

This section describes observations related to the four study tasks.

Task 1: Conversation partner by topic

The first task required the participants to find the names of all partners in one or more conversations, given the conversation topic:

> To whom did I talk about scuba diving? Write their name(s) into the space below.

The data provided to the participants contained one conversation about this topic. The conversation had two conversation partners.

The expectation was that the participant would locate "scuba diving" using textual search, and would then use either the trail or the main view's statement structure to complete this task.

These expectations were mostly met. Seven participants answered this question using search and the trail. One participant (in the graph condition) spotted "scuba diving" straight away and answered the question using the trail, without searching. Two participants did not use the trail. Both participants located "scuba diving" using textual search. One of these (in the graph condition) then answered the task by following the statement structure; the other one (in the list condition) used the same strategy but missed one of the two names that are the correct answer. This same participant had previously tried to use the map “to find things in the water that would relate to diving” before they realised that it was a conversation *about* scuba diving that was being asked for.

Task 2: Conferences by location

This task required the participants to find all conferences at a given location:

> Which conferences did I attend in Auckland? Write the conference name(s) into the space below.

The data provided to the participants contained two conferences that had occurred in Auckland.

The expectation was that the participants would answer this question following one of these strategies:

- Use the map to restrict the main view to items in Auckland and then use the trail to narrow the view down to items of type “conference”.
- Locate “Auckland” via text search and then use the trail to find connected items of type “conference”.

This expectation was mostly met, but overall the participants seemed to find the map navigator less useful than expected. Four participants used the map navigator in answering this question, while seven participants used text search (with three participants using both). Seven participants used the trail navigator. However, only one participant answered the question using the map+trail combination; all other users of the map resorted to text search when the map did not give them the desired result. Four participants directly opted for the search+trail approach.

One participant, in the list condition, answered this question using only the trail navigator. This participant added “Auckland” (which was directly visible at the start of this task) to the trail and then scrolled through all related statements, noting down everything that "looked like" a conference (including one item that was in fact a building).

One participant, in the graph condition, answered this question using no navigators at all, simply by looking at the information shown in the main view; however, this participant appeared to randomly pick out items that to them looked like a conference. Two of the three items selected as the answer by this participant actually were conference sessions. This is particularly interesting because this participant attended one of the conferences in question themselves.

Task 3: Enclosing event of a conversation

This task required the participants to determine all conferences with a conversation about a given topic during a given type of session:

> At which conference(s) did I speak to someone about Python during the poster session? Write the conference name(s) into the space below.

The data provided to the participants contained two conversations involving this topic; however, both took place at different conferences, with only one occuring during a poster session.

The expectation was that the participants would use text search to locate "Python", and then use the trail to find conversations. The expectation was that the trail navigator would be used to determine which of the two possible conversations was the correct one.

This expectation was met, with all but one participant following this strategy. The exception was one participant who answered this question using textual search alone; this participant (in the list condition) searched for "Python", then scrolled through the list to locate the search result and found the correct conversation by chance. However, this participant made no attempt to double-check their answer and they could just as easily have found the incorrect conversation instead.

Task 4: Location of event by type

The final task required the participants to find the place in which the conference belonging to a given series was held in a given year:

> In which place was the New Zealand HCI conference in 2007? Write the place name into the space below.

The data provided to the participants contained only one conference matching these criteria. To encourage the participants to use means other than textual search to solve this task, the item representing the series of New Zealand HCI conferences was named “NZ CHI conferences”.

The expectation was that the participants would use the timeline to narrow down the statements shown to events in 2007, then use text search to find the series of New Zealand HCI conferences, then either use the trail to find its 2007 instance or find the conference just by looking at the (by then very few) statements shown.

The timeline was used less extensively than expected, with only four participants using this strategy. One of these answered the question using the timeline and the trail, while two attempted to use the timeline first but then changed their strategy, in one case to use just the trail and in the other to use search and then the trail. The fourth attempted to answer the question using text search first; however, the view was still restricted to items in Auckland from previous tasks and this attempt failed. The participant eventually found “CHINZ 2007” and “NZCSRSC 2007” in the timeline navigator and chose their answer from these options.

A fifth participant commented that they knew the answer to this question from their own memory, but would probably have used the timeline if that had not been the case. This participant still selected the correct answer in the main view (in the graph condition) before writing it into the workbook.

Two participants (both in the graph condition) answered this question directly after spotting the “CHINZ 2007” item and its related statements in the main view. A third participant first searched for “CHI” but then spotted “CHINZ 2007” (even though it was not highlighted, since the search only considers whole-word matches) and answered the question directly.

Out of the remaining two participants, one participant answered the question using text search only and the other combined search with the trail.

7.4.2 Component-related observations

This section reports on the researcher's observations and participants' comments made throughout the study. Most are related to individual components, while some are related to the navigators in general.

Graph view

Initial view The version of the Digital Parrot used in this study initially showed all nodes in their standard, unselected state: as grey boxes with a black outline and black text. Highlighted nodes were outlined in blue and used a boldface font, while selected nodes were outlined in orange. The z-order of nodes was essentially random and not related to the state of each node (unselected, highlighted, selected). This meant that selected or highlighted nodes could be fully or partially obscured by unselected nodes.

Several participants in the graph condition commented on the main view, describing it as "cluttered" and "confusing". They felt that too many items were visible initially and that the overlaps between items made it hard to see "what's there".

One participant suggested changing the graph layout to a radial layout, in which the currently selected node is always at the center of the graph. The participant felt that this would make it easier to understand what is related to the current node.

Relationships A few participants tried to verify their answers, which typically included trying to double-check the nature of relationships between items in the graph. They were disappointed to see that this was not possible; when asked whether edge labels would help, they agreed.

However, they said they would want edge labels to be shown on demand rather than all the time, for example only for edges adjacent to the currently selected item.

Spatial position of nodes One participant, in the graph condition, expressed concern after they had moved a node on the screen: "now I moved your data". They also asked whether there was an undo function for rearranging the nodes. Questions revealed that this seemed to be a concern about changing the study conditions for other participants (which was not actually the case – the spatial arrangement of items was re-set between sessions and identical for all participants in the graph condition) rather than a concern about the experience for the participant themselves. Other participants moved nodes in the graph view without voicing any concerns.

List view

Statement structure In the version of the Digital Parrot used for this study, a cell in the list view was blank if it contained the same text as the cell directly above it. This was the case even when the cell directly above it belonged to a different statement (i. e. in the case of predicate cells: is associated with a different subject; in the case of object cells: is associated with a different predicate and/or a different subject). Additionally, every other row had a darker background, regardless of the underlying statement structure. Users seemed to not understand the "empty" cells at all and to find it very hard to see which cells were associated with which subject.

Statements in the list view were shown in the order subject, predicate, object from left to right, following the normal reading direction in English. One participant in particular was observed to use the right column in the list almost exclusively. When asked why, they explained that there was "more to click" there. This referred to the number of cells in the left column were blanked out due to repetition, which was higher in the left column because statements were sorted by that column.

Comparison with graph view Most of the participants who were in the list condition were shown the graph view during the discussion at the end of their session. All said that they would prefer the graph view over the list view, although this was based on a brief look at the graph view rather than actual use of it.

Text search and the list In the version of the Digital Parrot used for this study, text search in the list condition highlighted all matching cells by showing their text (if any) in boldface. These matches could be off the screen; in this case, users got the impression that the text search navigator was not in fact working. When instructed to scroll through the list to find highlighted items, the boldface text was overlooked by almost all participants because it did not stand out strongly enough.

All navigators

Several observations were made that related to all navigators.

Usage of navigation components Usage of the Digital Parrot's interface components varied across participants. Overall, the contextual navigators were not used as often as expected. Table 7.3 shows how many participants in each condition used each component, for each of the tasks as well as overall. Figure 7.2 shows a boxplot for the number of tasks in which each component was used.

All ten participants used the search function for at least one of the tasks; the search was used on average for 3.5 tasks (min = 1, max = 4, IQR = 1). All but one participant used the trail navigator for at least one task and all but one participant (in the list condition) used the statement structure as visualised in their respective main view. The trail navigator was used on average for 3 tasks (min = 0, max = 4, IQR = 0.5) and the statement structure on average of 2 tasks (min = 0, max = 2, IQR = 1).

Only three participants used both the map navigator and the timeline navigator for at least one of the tasks each. One further participant used the map for one task (T2) and commented that they would have used the

Task	Condition	Timeline	Map	Search	Trail	Structure
T1	Graph			4	4	5
	List		1	5	4	1
T2	Graph		1	3	4	2
	List		3	4	3	3
T3	Graph			5	5	2
	List			5	4	1
T4	Graph	1		1	1	3
	List	3		4	3	1
All	Graph	1	1	5	5	5
	List	3	3	5	4	4
Total (Participants)		4	4	10	9	9
Total (Tasks)		4	5	31	28	19

Table 7.3. Component usage by task and condition, in number of participants

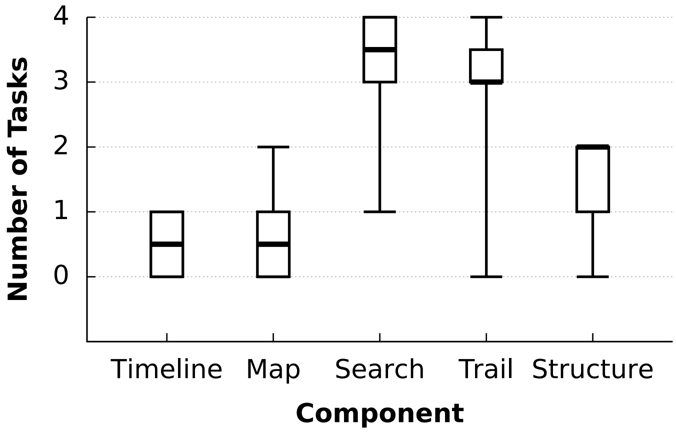

Figure 7.2. Component usage in average number of tasks, per participant

timeline for another task but knew the answer without using the Digital Parrot. One further participant used the timeline for one task (T4). This means that half of all participants never used any contextual navigation at all. Usage of contextual navigators was slightly higher in the list condition than in the graph condition, with all three participants who used both types of contextual navigation being in the list condition. However, the difference between participants using contextual navigators in the list condition compared to the graph condition is not statistically significant (Fisher Exact Probability test[4], $p > 0.5$ one-tailed).

In line with the expectations, the timeline was used only for the task that specified a time (T4). The map was used almost exclusively for the task that specified a place (T2), with the exception of one participant who attempted to use it for another task (T1). The timeline and map were used on average for 0.5 tasks (timeline: min = 0, max = 1, IQR = 0.5; map: min = 0, max = 2, IQR = 0.5).

Use of the search was generally high, with the last task being an exception. Use of the statement structure was slightly lower for the third task but was the sole means to complete the fourth task for three of the participants.

Navigator windows Almost all participants were confused by the fact that each navigator is activated when its window is shown and deactivated when its window is hidden. The version of the Digital Parrot used for this study showed navigator windows centered on the main view when they were first opened and always forced navigator windows to be drawn in front of the main window. This meant that navigator windows frequently obscured information in the main view and that the effect of navigators on the main view often was not immediately obvious, especially when the participant's attention was focused on the navigator window.

Several participants did not deactivate the map or timeline navigator after having explored it, moving them to the side of the screen instead.

[4]The Fisher Exact Probability test is a test for statistical association between two variables with two subcategories each. It serves the same function as the χ^2 test but can be used for sample sizes that are too small for the χ^2 test.

Typically this led to unexpected behaviour in subsequent tasks because the participant did not realise that the main view was still restricted by choices made in the navigator.

Timeline

Interaction One point of confusion with the timeline was its interaction with the main view, as shown by a comment from one participant: “How do I get from timeline to main view?” Participants seemed to expect interaction between the timeline and the main view that was different from the selecting and highlighting that in fact occurred. In particular there seemed to be an expectation that clicking on an item in the timeline should have a more direct effect.

Personalisation One participant who did not use the timeline stated that the timeline might have helped if the timespans in the study had been longer. When the concept of personalised timelines was introduced to this participant, they said that they liked this idea and gave examples of personalised timespans that they might consider useful: “during the time in NZ”, “when I lived in *city*”, “when I was in a relationship with *partner*”.

Map

Interaction Participants showed confusion about the purpose of the map that is very similar to that about the purpose of the timeline: “how do I get from the map to the data?”, “I don’t really understand the purpose of the map”.

Map style The map in the version of the Digital Parrot used in this study was very simple, with only landmass boundaries being shown. One participant explained that they had difficulties deciding in which direction to pan the map because they did not remember exactly where on the map some of the places actually were, particularly with the absence of placenames and other major landmarks (rivers, motorways).

Another participant found it laborious having to interact with the map to switch between places. When asked whether a list of placenames would help, they agreed but stated that this should be an addition to the map rather than a replacement of the map.

Markers and labels The version of the Digital Parrot used in the study showed markers and marker labels only when certain thresholds were reached with regards to the current zoom level of the map and the approximate size of the location represented by the marker. Markers and labels for locations too small or too big for the current zoom level where not shown. Instead, histogram-like circles were shown in those cells of an invisible grid overlaid onto the map that would contain markers when zoomed closer, with the approximate number of markers indicated by the opacity of the circle's fill colour.

Several participants expressed confusion about the histogram circles. Participants were also confused that marker labels were missing at the initial zoom level: "oh, there is writing on the map" (after zooming in). They were also confused that clicking on a marker did not appear to have any effect, and one participant asked why they couldn't move the markers on the map.

Search

Preference of search over other navigators One participant explicitly stated that they were using search to solve the first task "'cause I'm used to Google".

Syntax The text search in the version of the Digital Parrot used for the study searches full words only (e. g. searching for "CHI" would not find "CHINZ"). Some participants asked what type of search would be performed; all seemed to expect that it would "do what Google does" in terms of query syntax.

Search and types Some participants tried to search for types (e. g. "conference") and were disappointed that nothing suitable was found. One participant suggested to allow for typed search (i. e. searching for a particular string only in instances of certain types).

Trail

The trail was generally seen as the most difficult component of the Digital Parrot. One participant remarked that they considered using the trail for one of the tasks but didn't because "it's complicated and I don't want to look stupid". One participant (in the graph condition) also commented that they did not believe that participants with a background other than Computer Science would be able to understand and use the trail navigator, mostly because "they don't understand what graphs and trees are". One participant commented that they would have liked more opportunities to practice using the trail navigator.

Effort Several participants had problems remembering how to use the trail navigator. The version of the Digital Parrot used in the study required the user to bring up an item's context menu, typically via a right click with the mouse, to start the trail at this item or to add the item to the trail. A few participants had to be reminded to bring up the context menu via right click. Most participants appeared to feel that using the trail was laborious; one reason was the required right click and another reason was having to navigate nested menus when adding a type to the trail.

Some participants commented that they did not normally use context menus because the operating system on their computers, Mac OS, does not generally make much use of them.

Visibility of types Several participants criticised that they had to start the trail at a specific item. One participant explained that this was cumbersome when they already knew which type they wanted to add: "I need to click on something that looks like a conference to be able to add the

Conference type to the trail". Another participant explained that easy access to the types would make it easy to find all instances of a specific type: "It would be nice if there was a conference root [shared ancestor of all conference instances in the graph], 'cause I want to be sure I grabbed all the conferences".

One participant suggested having a list of all types "somewhere to the side of the user interface" that, when clicked, could highlight all items of this particular type.

Trail mechanics One participant appeared to have difficulties distinguishing what it meant to add a type to the trail and what it meant to add an item, i. e. an instance of a type, to the trail. This may be related to the observations around the visibility of types – in the version of the Digital Parrot used in the study, both actions were performed from the context menu of an item.

Some confusion was also observed around the meaning of the order of items on the trail and around how to retrace one's steps. One participant explained that to them, the word "trail" suggested that it was related to what one had previously seen. Another participant suggested a history of recently used trails that would make it easier to undo changes to the trail.

General

Double-checking Some participants double-checked most or all of their answers, while some did not do so at all. One of the participants who did not attempt to double-check their answers said that they shared some of the memories in the data file. They explained that they used their own memories to evaluate the plausibility of the answers they had found using the Digital Parrot.

Cues vs memories Several participants stated their confusion about the type of information stored in the system – "what's in there". They spent some time trying to get to "the actual memories" before realising that there was "nothing to open". They explained that they had expected to be able to access something that was "more like actual memories" and were confused at first that they could not distinguish between "things you have done and things you know about", experiences and other information: "Am I to infer that because it's there you attended it? I'll assume that it's 'was at' not 'have knowledge of'. But that could mean that you *gave* the keynote speech – or that you *went to it*?" The participants who were confused about this generally had not attended the conferences mentioned in the data.

Desirability At least one participant commented when filling in the SUS questionnaire that they would need to know more about how the memories "get in there" before they could decide whether they would like to use the Digital Parrot frequently. One participant expressed their enthusiasm about the graph visualisation and explained that a similar visualisation would help them in one of their hobbies, which involves having to make sense of highly connected information.

7.5 Usability of the Digital Parrot

This section relates the findings reported in Sections 7.3 and 7.4 to the first goal of this study as introduced in Section 7.1.1: to detect any major usability issues in the version of the Digital Parrot used. It first describes the usability issues and then the changes made to the Digital Parrot's user interface to remedy these issues.

7.5.1 Usability issues

This section identifies the main usability issues discovered in this study and describes possible resolutions for each issue. The issues are grouped by interface component. Several of these issues had already come up during the pre-study expert review (see Section 7.2.8); the study made it clear that the changes made after the review did not fully resolve these issues.

Graph view. The SUS scores and participants' comments suggest that the graph view is reasonably easy to use. One main issue with the graph view was that the way nodes were rendered made the interface appear too cluttered; the changes made after the expert review apparently had not gone far enough. Another issue was that it was impossible to determine the nature of the relationships between nodes. Both issues could be resolved by making changes to the way the graph is drawn. Additionally, the number of nodes shown on program start-up could be restricted by some criterion.

List view. The SUS scores and participants' comments suggest that the list view had more severe usability issues than the graph view. The main issue with the list view was that the statement structure did not become clear. This issue could be addressed with minor changes to the way the list is drawn.

Another issue with the list view was that it made it difficult to determine how many, if any, matching items were found when a textual search was conducted. This issue could be resolved by allowing the user to navigate between search results, scrolling the list to the location of the next result. Another, simpler solution would be at least to indicate whether the search led to any results and if so, how many.

Navigators. From the observations and participants' comments, it is obvious that participants found it confusing that navigators and main view are all in separate windows. As described in Section 5.3, early designs of the Digital Parrot's user interface integrated the main view and the navigators into a single window. This approach was later discarded because of unresolved questions around the interaction mechanism. This suggests that resolving the issue completely may lie outside of the scope of this work. To ameliorate the situation, changes could be made to window placement and window options for the navigator windows.

Map navigator. The main presentation issues with the map were related to the map style, particularly the absence of detail that would allow the participant to orient themselves, and to the markers on the map. The pre-study expert review had given an indication of usability issues with the map and the changes after the review quite obviously had not resolved these issues. They could be addressed by changing the underlying map to one that is more familiar to typical users, such as Google maps.

The behavioural issue with the map, that participants were unsure about its purpose, equally applies to the timeline. It is most likely related to the general issue with navigators described above. If this issue cannot be resolved in time for later studies, the usage of navigators to restrict and highlight information in the main view should be explained more thoroughly to participants in those studies.

Search navigator. There were a few occasions when the search function did not behave according to the participant's expectations. This could be addressed by modifying the default search options.

In both main view conditions, text search highlights all items matching the query. However, often these matches are not visible. This could be because the search window is obscuring the main view, because of overlapping items (in the graph view) or because all highlighted items are currently off the screen (in the list view). As described above, this could be addressed either by modifying the main views or by adding information about the search results to the search navigator.

Trail navigator. The trail navigator appeared to be perceived as particularly difficult to use by the participants. The new features added after the pre-study expert review were used by the participants but essential functionality was still lacking. In part the issues with the trail navigator certainly arose because none of the participants had ever used similar interaction techniques. However, one issue stands out that is related more to the specific interface that was implemented rather than to the interaction technique in general. Requiring the use of context menus to start a trail or add to an existing one does not appear to have been a good choice. It contributes to confusion between instances and their types, it requires the user to first find an instance of a type before the type can be added to the trail, and it makes it very hard to determine which types are known by the system. It also requires a high degree of physical effort (right clicks, navigation of nested menus).

Addressing the issues around the trail navigator most likely would require a complete redesign of its user interface. Interaction with the trail should be possible without having to interact with individual items in the main view. Types and instances should be separated more clearly.

7.5.2 Changes made to the Digital Parrot

This section describes the changes made to the Digital Parrot as a consequence of the usability issues describe above.

Graph view

Changes were made to the graph view to reduce the impression of clutter and to make it easier to determine the nature of relationships between nodes. The three most important changes were:

- The z-order of the nodes is now determined such that selected nodes and their neighbours are drawn on top of highlighted nodes and their neighbours (rather than having an essentially randomised z-order);
- Nodes that are not selected nor highlighted, nor a neighbour of a selected or highlighted node, are drawn semi-transparently to make them less prominent.
- Edge labels are shown for selected edges and edges incident to a selected node.

Figure 7.3 shows screenshots of the graph view before and after the changes had been made.

List view

Changes were made to the list view to improve the way the statement structure is shown. The main change was: instead of blanking out cells

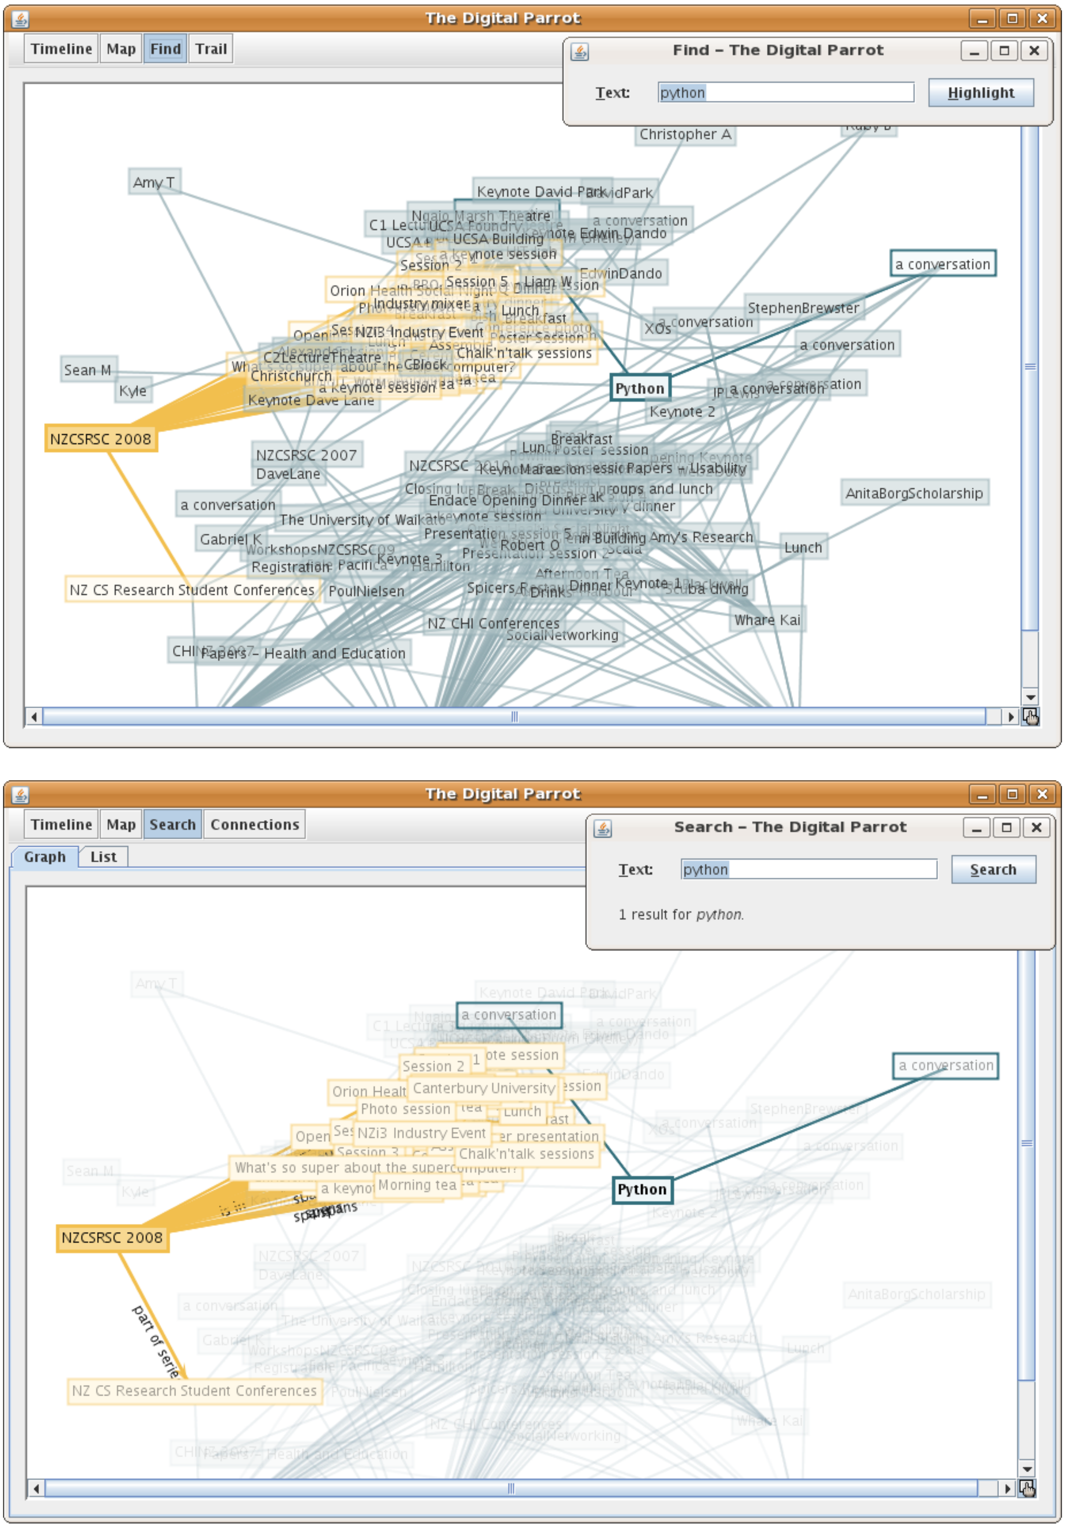

Figure 7.3. Graph view, before (top) and after (bottom) changes

that would repeat the value of the cell directly above, these cells now contain the text but in a colour with little contrast to the background. Alternating the background colour between groups of statements, rather than between rows, was too difficult to implement in the time available. Figure 7.4 shows screenshots of the list view before and after the changes had been made.

Navigators

The window options and default placement of the navigator windows were changed to reduce obscuring of data in the main window by the navigators. Navigator windows are now shown to the sides of the screen, only partially overlapping the main window.

Integrating of the navigators and the main view into a single window was not feasible in the time available. As indicated in Section 5.3, such a change would require further conceptual work and likely a fundamentally different approach to the navigators, going far beyond changes to the user interface alone.

Map navigator

Changes were made to the map navigator to address the confusion that participants felt with the map, particularly in relation to markers, and to improve navigation between items on the map. The two main changes were:

- The underlying map used is now Google Maps, displayed using an embeddable browser component provided by WebRenderer[5]. The earlier version of the map navigator used OpenMap[6].

- A list of placenames was added to the map navigator. A double-click on a placename adjusts the position and zoom level of the map to display the selected placename.

[5] http://www.webrenderer.com/products/swing/product/
[6] http://www.openmap.org/

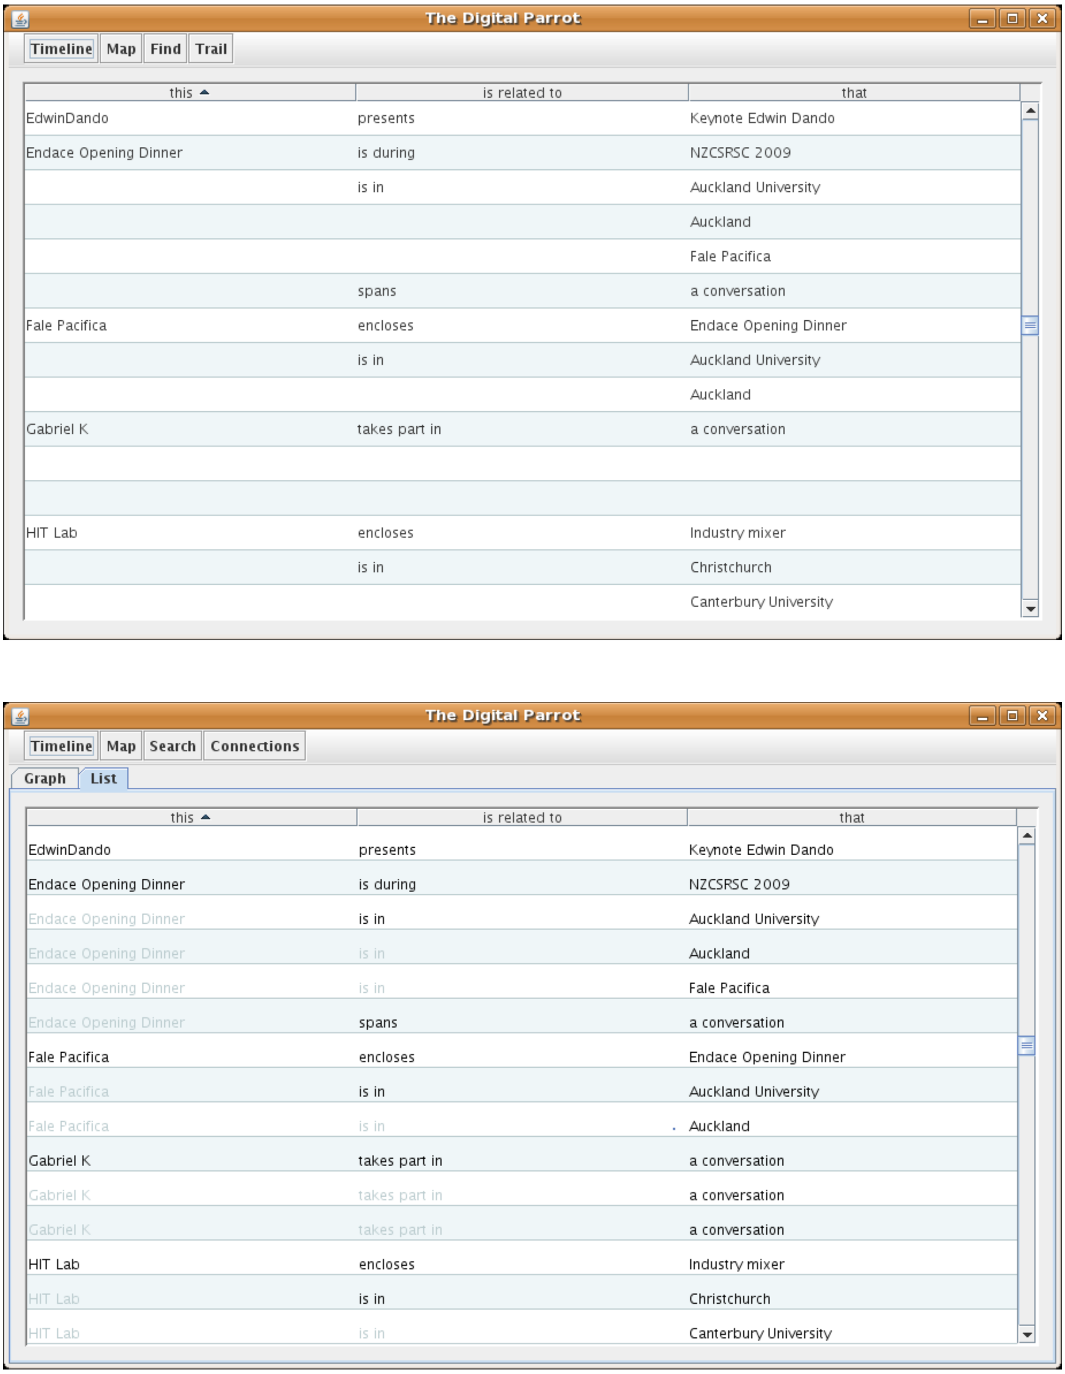

Figure 7.4. List view, before (top) and after (bottom) changes

Figure 7.5 shows screenshots of the map navigator before and after the changes had been made.

Search navigator

Changes were made to the search navigator to address usability issues regarding the query syntax and regarding the display of search results. The three main changes were:

- The search navigator now indicates the number of search results for the current query. Adding functionality to scroll the list so that the next search result is in view was too difficult to implement in the time available.
- When a query leads to no result, the search navigator suggests to add the wildcard character "$*$" to query terms if subword matches are required. It was decided against modifying the query terms to add wildcards programmatically because of possible performance issues with the query engine. It was also decided against modifying the query programmatically to achieve a conjunctive query mode rather than the disjunctive query mode that is the default mode of the query engine. Instead, participants of future studies would be informed about the query syntax used by the system.
- Items that match the current query are now outlined in blue in addition to using a boldface font. This makes it easier to see search results, particularly when scrolling through the list.

Figure 7.6 shows screenshots of the search navigator before and after the changes had been made. The list view was chosen for these screenshots; consequently, the changes made to the rendering of search results are also shown.

Trail navigator

The user interface for retrieval using semantic information, the trail navigator in the original version of the Digital Parrot, was redesigned com-

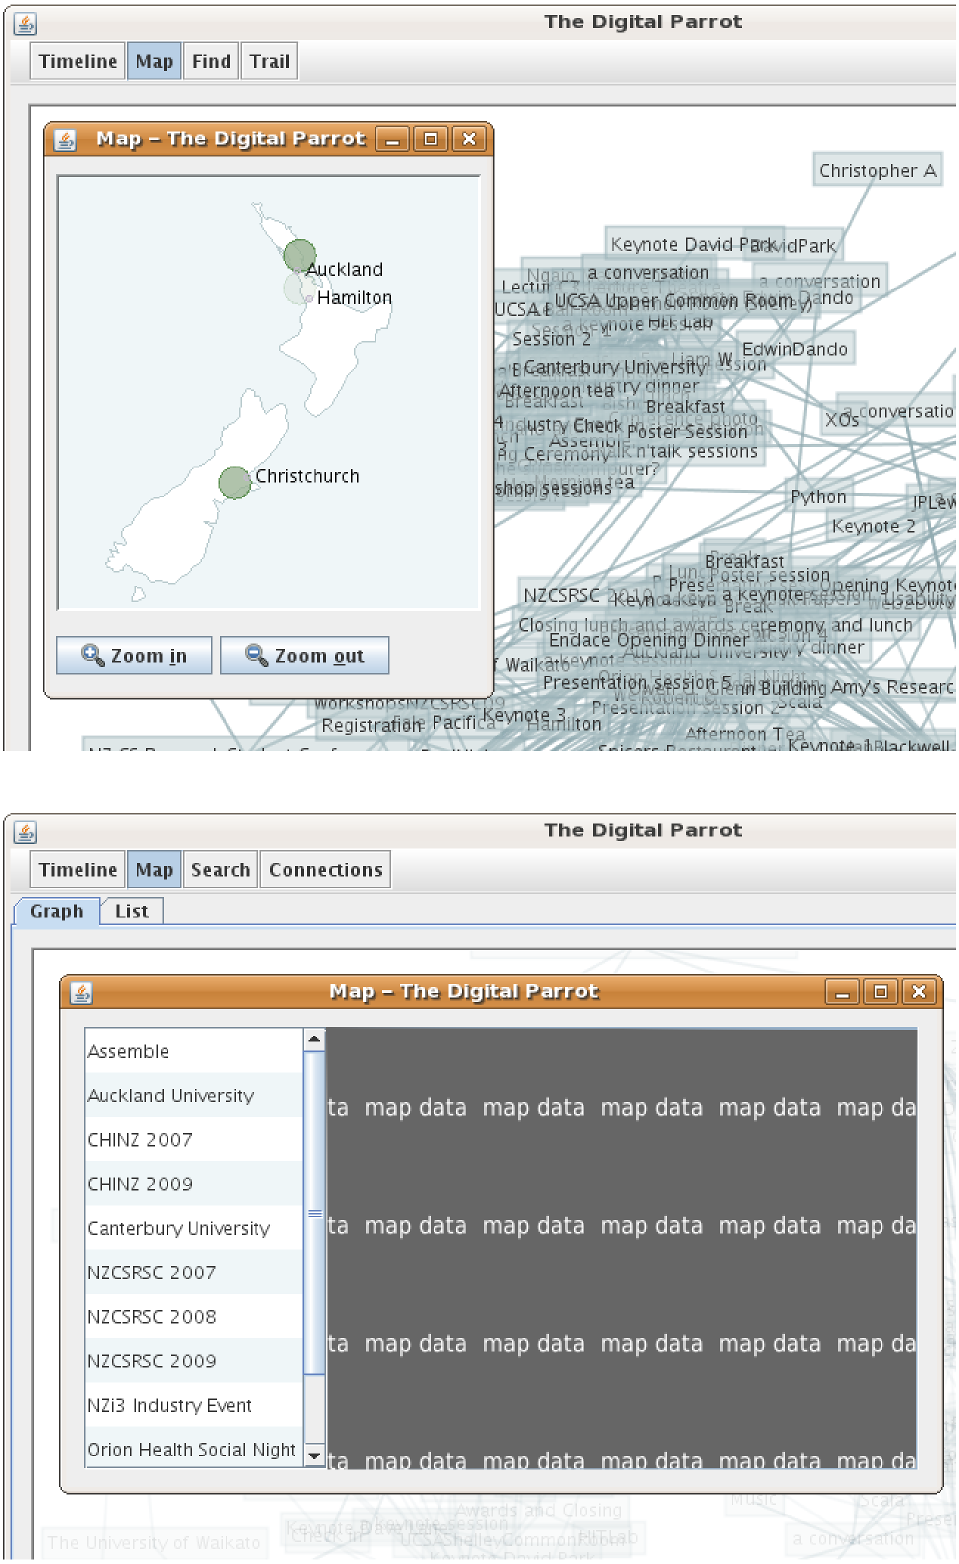

Figure 7.5. Map navigator, before (top) and after (bottom) changes. Map data removed from screenshot for copyright reasons.

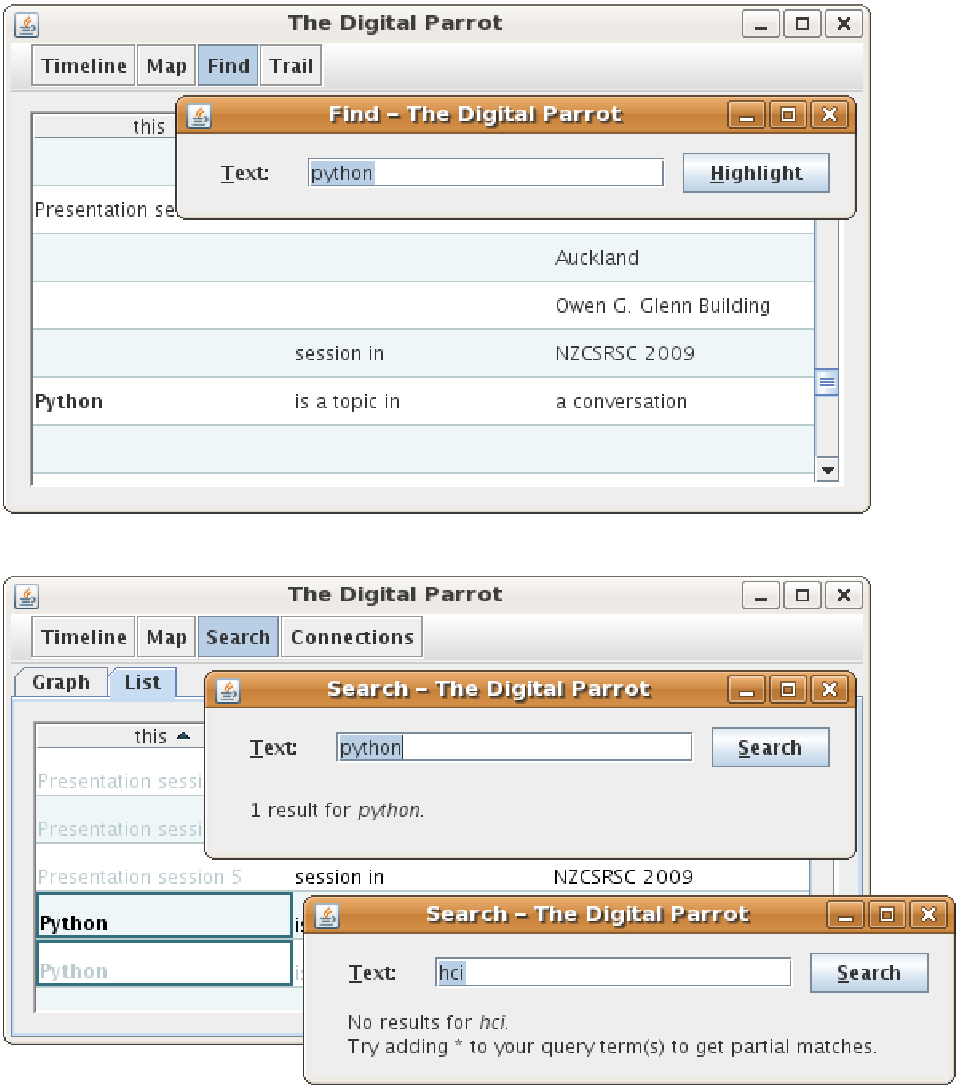

Figure 7.6. Search navigator, before (top) and after (bottom) changes. The "after" image also shows a second screenshot of the search navigator, for a search that has no results.

pletely. It is now called the connections navigator and shows the behaviour described in Section 5.3.4. Major changes included

- the separation of types and instances in the connection navigator's user interface;
- the move from interaction based on context menus on items in the main view to interaction with drop-down boxes separate from the main view, which also allows the connections navigator to be used without first having to locate a start item;
- the ability to configure whether only items on the trail/chain are to be shown, or items on the trail/chain and their direct neighbours; and
- easier ways to backtrack in the chain of connections ("trail" before the change).

Figures 7.7 and 7.8 show screenshots of the trail/connections navigator before and after the changes had been made.

7.6 Support for answering memory questions

This section relates the findings reported in Sections 7.3 and 7.4 to the second goal of this study as introduced in Section 7.1.1, the Digital Parrot's support for answering questions about another person's experiences. These findings give an indication of the Digital Parrot's effectiveness. However, they do not truly show whether the Digital Parrot achieves its design goals because the data collection for the study was not true memory data – it lacked a self-link to the participants.

The section examines several aspects of the participants' use of the Digital Parrot:

1. the overall effectiveness of the approach;
2. the influence of contextual information;
3. the influence of semantic information;
4. the influence of the visualisation type; and
5. the participants' subjective experience in using the Digital Parrot.

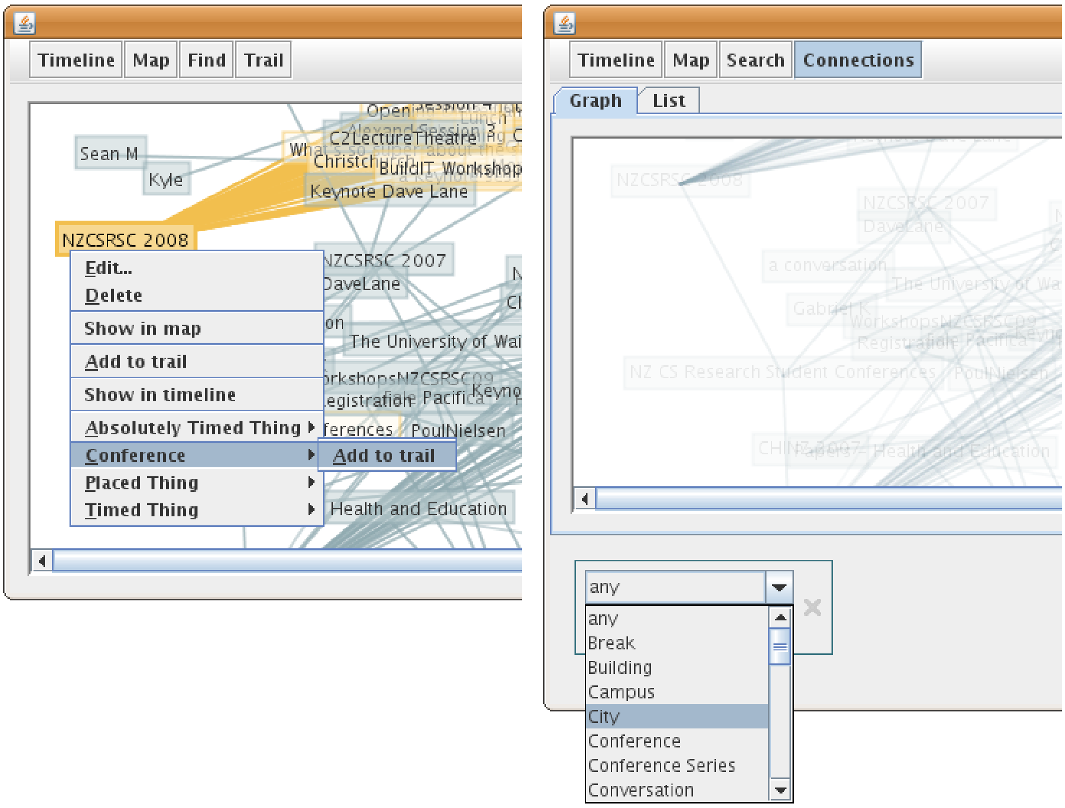

Figure 7.7. Starting a trail/connections chain at the type "Conference", before (left) and after (right) changes

7.6.1 Overall effectiveness of approach

The findings of this study suggest that the approach of combining contextual data with semantic information allows people to successfully answer questions about memories of events and associated information. Almost three quarters of the tasks were answered completely and correctly (29 of the 40 tasks, i. e. 72.5%).

The findings also suggest that the Digital Parrot as the operationalisation of the approach needed further improvements. The usability of the version of the Digital Parrot used in the study was rated below the acceptability threshold on a standard usability measurement tool. The researcher's observations and participants' comments pinpointed the trail navigator, the list visualisation, in particular its interplay with the search,

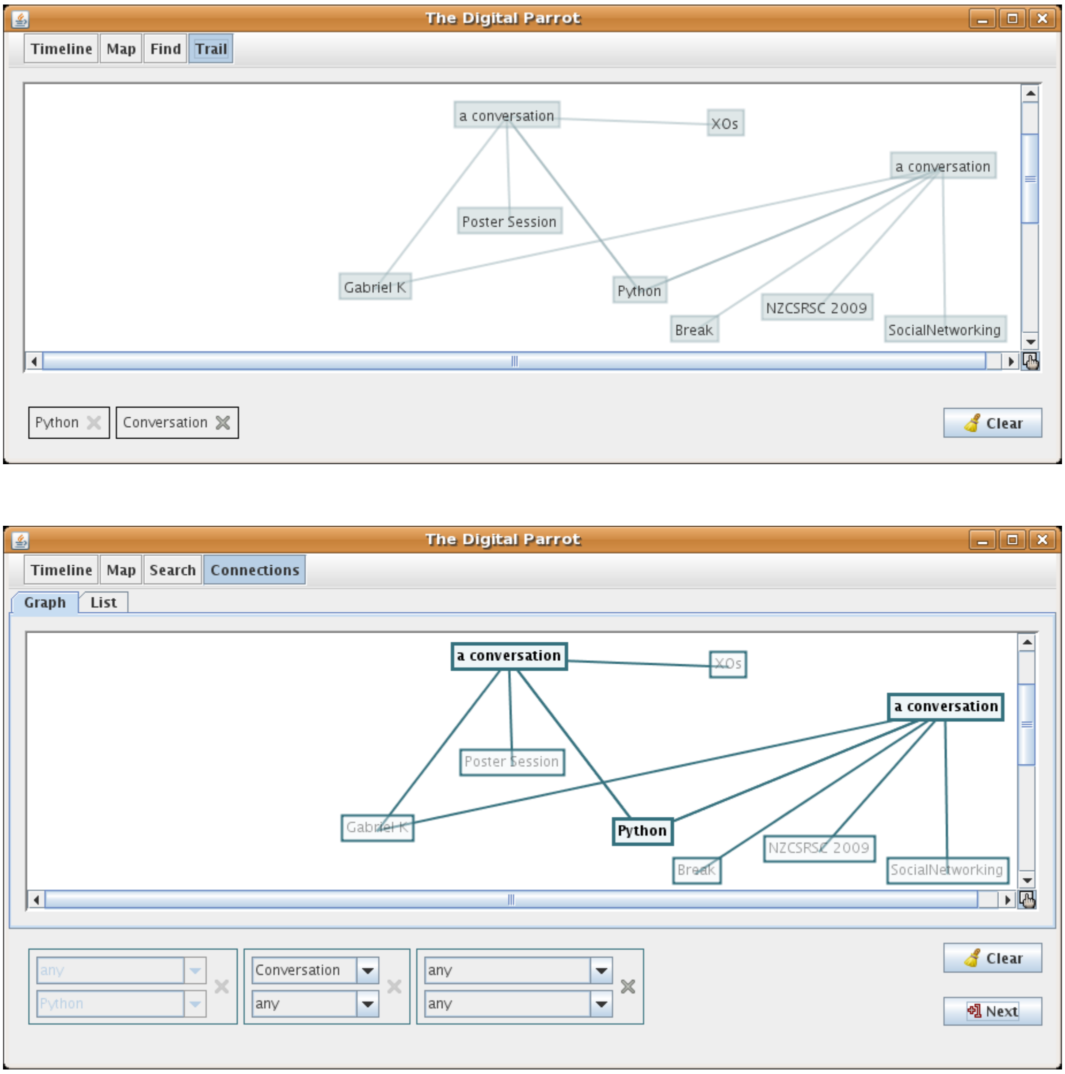

Figure 7.8. Restricting to a trail/chain: items connected to "Conversation" instances connected to the "Python" item, before (top) and after (bottom) changes.

the map navigator and some visualisation aspects of the graph view as the core areas that needed to be addressed.

Subjective and objective measures did not appear to be related; no connection was found between the participants' rating of the Digital Parrot's usability, their confidence in using the system and their rating of the ease of using the system on one hand and the actual effectiveness on the other hand.

7.6.2 Influence of contextual information

Fewer participants than expected used the contextual navigators (Table 7.3). This may be due to the selection of tasks. Another possible reason is that there was not enough data to make restriction by context necessary, or that the question involving temporal context (T4) could be answered using the search because the items concerned ("CHINZ 2007", "CHINZ 2009") contained the temporal context in their name. Similarly, the question involving a location could be answered using a search for the location rather than requiring use of the map. This latter reason may also explain why higher usage of contextual navigators was observed for the list condition – participants were observed to have problems using the search in the list condition and may thus have been more inclined to try alternative means to answer the question.

The contextual navigators were used for the tasks that had been expected to be suitable for them (T2, Section 7.4.1, and T4, Section 7.4.1). This indicates that the participants generally understood for which types of questions the navigators were applicable.

7.6.3 Influence of semantic information

Participants encountered so many problems using the trail navigator that it is difficult to determine the real influence of semantic information. However, all but one participant used the trail navigator for at least one task and all participants who did use it used it for at least two tasks. It was used for almost as many tasks as the search, both overall and on average per participant. Some participants suggested a stronger integration with the search function.

7.6.4 Influence of associations

The influence of the visualisation type gives an indication of the influence of associations. The graph view directly shows connections between information items, while these connections are less straightforward to see and use in the list view.

The list view generally appears to be harder to use than the graph view. Even though the difference of the median SUS score between conditions is not statistically significant, it is reflected in the participants' comments and in the researcher's observations.

Correctness of answers (Table 7.2) is slightly higher in the graph condition than in the list condition. Only in the graph condition did participants miss parts of an answer (T2), while it was only in the list condition that participants struggled to distinguish which of two possible answers was actually correct (T4).

It is not clear whether an observed difference in the number of participants who used contextual navigators in the two study conditions is due to chance or whether it is caused by a difference in the main views.

7.6.5 Subjective experience

Even though participants found the Digital Parrot challenging to use overall, over half of the participants gave a positive answer to the first SUS question, which asked whether they would like to use the Digital Parrot frequently. Some participants were confused about exactly what was being shown in the Digital Parrot. Participants who had shared some of the experiences described in the data file seemed to be more confident in their answers and performed less double-checking.

7.7 Summary

The study described in this chapter had two goals: to detect usability issues in the version of the Digital Parrot used in the study and to examine how the participants used the Digital Parrot to answer memory-related questions.

The study found several usability issues. The most substantial issues found were related to the navigation based on semantic navigation, the list view and, to a lesser extent, the map view and textual search. Subsequently, changes were made to the Digital Parrot's user interface that remedied most of these issues.

The results of the study described in this chapter are promising for the approach taken in this thesis to augment autobiographical memory. Overall, the Digital Parrot allowed the participants to answer questions about another person's experiences. The participants understood the Digital Parrot's user interface components and used them mostly according to expectations.

The Digital Parrot's user interface components were developed based on recommendations derived from Psychology research, namely contextual navigation, navigation based on semantic information and navigation based on associations. The uptake of contextual navigation was somewhat lower than expected, though it remained unclear whether this was because of usability issues or for more fundamental reasons. The user interface for navigation based on semantic information was clearly not designed well in the version of the Digital Parrot used in the study; consequently, no conclusions can be drawn about its effectiveness. A straightforward visualisation of the association structure appeared to be more helpful than a language-based representation. Since some parts of the study were affected by usability issues, further evaluations are necessary in particular with regards to the influence of semantic information.

This study could not, nor attempted to, evaluate the actual effectiveness of the design and implementation of the Digital Parrot. This is because the participants in the study did not actually attempt to remember their own experiences and related information. Some of the observations suggest that the findings differ from those of studies using naturalistic data collections, i. e. the participants' own memories. A study of the Digital Parrot using a naturalistic data collection is described in the following chapter.

Chapter 8

Evaluating the effectiveness of the Digital Parrot

This chapter contributes further to answering the fourth research question, whether the approach to augmenting autobiographical memory introduced in this thesis actually helps people remember. It describes an end-user study which was conducted to evaluate the effectiveness of the Digital Parrot and the underlying conceptual design. In contrast to the evaluation described in the previous chapter, this study investigated how the Digital Parrot aided users in remembering their own experiences and related facts approximately two years after attending an academic conference.

This chapter is structured as follows: Section 8.1 describes the focus employed in designing the study. Sections 8.2 and 8.3 give details about the experimental designs of the experiencing and the remembering phases. Section 8.4 describes the quantitative findings of the study and Section 8.5 describes further observations and participants' comments. Section 8.6 relates the findings of the experiencing phase to the goals of the study. The chapter concludes with a summary in Section 8.7.

8.1 Focus

This section describes the focus employed in designing the study. It outlines the goals of the study, the approach taken and the trade-offs made to overcome the challenges involved in evaluating augmented memory systems.

8.1.1 Goals of the study

The main goal of this study was to answer the fourth research question: Does the Digital Parrot, as the operationalisation of the approach to augmenting autobiographical memory taken in this thesis, actually help people to know about and reconstruct past experiences and related information?

Another goal of the study was to gain insights into the roles of those components of the Digital Parrot that are linked to the three factors to augmenting autobiographical memory stated in the thesis hypothesis: context, semantic information and associations.

The goals of this study are summarised by the following questions:

1. What *strategies* have the study participants established to remember and re-find experiences and related facts associated with attendance of academic conferences without using the Digital Parrot?

2. What is the *effectiveness* of the Digital Parrot in helping users remember past experiences and related information?

3. What influence does *contextual information* have in remembering past experiences and related information using the Digital Parrot?

4. What influence does *semantic information* have in remembering past experiences and related information using the Digital Parrot?

5. What influence do *associations* have on remembering past experiences and related information using the Digital Parrot?

6. What is the study participants' *subjective experience* of using the Digital Parrot for remembering past experiences and related information?

8.1.2 Approach

The study used a two-phase approach. In the first phase, the participants were interviewed about their experiences at a recently attended academic conference. The interviews served to capture the participants' experiences without the use of an augmented autobiographical memory system.

The second phase consisted of an observational laboratory-based study in which the participants were asked to recall experiences and related facts. This phase used the Digital Parrot containing memory data extracted from the interviews and from other information available about the conference (such as the conference website and the published proceedings). The second phase compared unaided remembering to remembering using the Digital Parrot. It used a task-based method in which questions were personalised for each participant but followed a shared set of categories.

8.1.3 Trade-offs

The strengths of the study are its use of a naturalistic data collection and a realistic timespan between making and remembering experiences. A number of trade-off were made to achieve this, as is common with evaluations of augmented memory systems (see Chapter 6).

The first trade-off was to involve a small number of participants (4) who all shared a similar background (Computer Science). The study had fewer participants than typical end-user studies in Human Computer Interaction and similar fields but more participants than many end-user studies of CARPE systems that use the participants' own data.

A smaller trade-off was made with regards to information input into the system, to allow for a meaningful study within the scope and timespan of the work presented in this thesis. The study tasks were not entirely naturalistic, to improve comparability of findings across participants.

As in the other empirical parts of the research presented in this thesis, the study focused on experiences while attending academic conferences.

8.2 Experimental design of the experiencing phase

The first phase of the study captured some of the participant's experiences at an academic conference, similar to the experiencing and revising phases of interacting with an augmented autobiographical memory system as described in Chapter 4. The goal was to collect data for use in the second evaluation phase. The study also allows insights into the kinds of information people might wish to store in such a system when they visit an academic conference.

This section describes the method, participants and procedure of the experiencing phase as well as the types of data captured. It then describes the changes made to the experimental design based on a pilot study and finally the results of the experiencing phase.

8.2.1 Method

This phase consisted of guided open interviews with individual participants. No time constraints were placed on the interviews.

8.2.2 Participants

Graduate students and academic staff members from the Computer Science department at the University of Waikato were invited to participate in this phase of the study. A further selection criterion was that each participant needed to have attended at least one academic conference in the months prior to the interview. Participants were also informed that they should be available for follow-up sessions at a later time. Invitations were sent via departmental e-mail lists.

The phase had five participants. At the time of the interviews, three of these participants were members of academic staff in the Department of Computer Science at the University of Waikato and two were PhD students.

There was no overlap between these five participants and the participants in the study described in the previous chapter.

8.2.3 Procedure

The interviews took place approximately one to three months after the participant had returned from the conference or conferences about which they were interviewed.

At the beginning of their session, the participant was given a brief introduction to the goals and procedure of the study. They were provided with the research consent form and a copy of the research participant's bill of rights and the participant's consent to participate in the interview was obtained. The consent form and bill of rights are given in Appendix C.1.

The researcher then asked the participant to recall and describe those experiences at the conference that they found particularly memorable or important and those that they thought they might wish to remember in the future. The researcher occasionally followed up on points made by the participant. The researcher also prompted the participant occasionally to describe certain types of experiences.

The researcher supplied some paper material for each conference covered in the interview. This was a printed calendar of the month(s) in which the conference took place and a printout of the conference program. The calendar was printed on one A4 sheet per month, leaving room for annotations. The conference days were marked in a different colour. The conference program typically was an abbreviated overview program rather than the full schedule, e. g., showing session titles and times but not necessarily titles and times of individual presentations within a session. The paper material was referred to throughout the session both by the participant and by the researcher.

8.2.4 Types of data collected

The following is a catalogue of typical questions that the researcher asked during an interview, either verbatim or in paraphrased form.

- When and where was the conference? How long was the conference?
- Was it "only" the main conference or for example the conference and some workshops/tutorials?
- Did you meet any new/memorable people, and what do you remember about them? Did you meet up again with people you knew already?

- Did you have any memorable/important conversations?
- What was your role at the conference (for example attendee, presenter, organiser)?
- Were there any memorable/outstanding talks/presentations?
- How did you get to the conference? Was the trip just for the conference or did you go anywhere else?
- Was there a conference dinner or a social program? Is there anything in particular that you remember from that?
- Did you have time for some sightseeing? Is there anything in particular that you remember from that?

Further questions were asked to follow up on a participant's responses. In particular, the researcher typically asked where and when a certain event took place and which people were present.

All interviews were audio recorded, with the participant's permission. Most of the interviews were recorded using a small Internet tablet (Nokia N800) and the remaining interviews were recorded on the researcher's laptop. The device's internal microphone was used in all cases, with the device sitting on a table between the researcher and the participant.

In addition to the audio recordings, the data collected consists of annotations on the paper material made by the participant or the researcher. Most participants annotated the calendar and the conference program to some degree, but none did so extensively.

8.2.5 Pilot study

A pilot study with a single participant was conducted prior to the main study. The pilot study led to a refinement of the catalog of questions asked during the interview. Another change that was made based on the pilot study was the inclusion of the printed material (calendar and conference program) to guide the recollection process.

8.2.6 Results

The five participants were interviewed about their experiences at eight conferences (two participants with one conference, three participants with two conferences each).

The audio recordings of the interviews run on average for thirty minutes per conference. Due to technical problems with the recording equipment, the audio recordings from one of the participants who were interviewed about two conferences could not be used for the second phase.

8.3 Experimental design of the recall phase

In the second phase of the study, participants attempted to remember past experiences and related information with the help of an augmented autobiographical memory system. A task-based user study was conducted in which participants of the experiencing phase (described in the previous section) used the Digital Parrot system to recall experiences from the conference visit(s) they had been interviewed about. This second phase was the central evaluation of the Digital Parrot's effectiveness.

Before the study, information collected in the interviews as well as information about the conference found in other sources was transcribed by the researcher into data files for the Digital Parrot. More details about the transcribed data are given below.

The remainder of this section describes the method, participants, timespans, procedure and tasks of the study, as well as the types of data captured in the study and the changes made to the experimental set-up based on a pilot study.

8.3.1 Method

This study was an observational computer-based laboratory experiment with a task-based between-subjects design. The time that participants could take for the session was not constrained since the study focused on insights into the participants' use of the Digital Parrot rather than on an efficiency analysis.

8.3.2 Participants

Participants were recruited from among the four participants of the first phase with usable audio recordings. All potential participants were approached via e-mail; it was emphasised that participation in the second phase was not compulsory. However, all four participants agreed to take part in the second phase. Of these, one was female and three were male. At the time of the recall session, all four participants were members of academic staff in the Computer Science department at the University of Waikato. The ages of the participants at recall time ranged from 34 to 62 years (median 41 years).

When the participants were approached about taking part in the second phase of the study, they were instructed not to attempt to recall any experiences from the conference visit(s) they had been interviewed about.

8.3.3 Timespans

The timespans involved in the study are shown in Table 8.1. The recall session took place approximately two years after the earliest experience that a participant was asked to recall: the median difference was 720.5 days (1 year and 355.5 days) with an IQR of 38.5 days. The time difference between the interview and the recall session was slightly shorter than two years: the median difference was 671 days (1 year and 306 days) with an IQR of 35.25 days.

The interview had taken place approximately one to three months after the earliest experience that a participant was asked to recall (median 65.5 days, IQR 43.75 days). The main reason for this relatively long gap were practical considerations around participant recruitment. The assumption was made that potentially important experiences (see Section 4.1.1) would still be remembered in sufficient detail after this gap. Conducting the interviews immediately after the conference or even at the end of each conference day would most likely have resulted in more detail; conversely it also might have significantly altered the memories formed by the participants of these experiences. The information collected through the interviews certainly appeared to be a plausible and plausibly exhaustive representation of the type of experiences targeted.

While all participants reported details for at least some of the events they described, the number of events reported and the level of reported

Interval	P1	P2	P3	P4	Median	IQR
Experience–Interview	114	82	38	49	65.5	43.75
Interview–Recall	547	670	672	682	671.0	35.25
Experience–Recall	661	752	710	731	720.5	38.50

Table 8.1. Time in days between experience, interview and recall

detail varied across participants. However, these variations did not appear to have been influenced by the time since the conference; instead, the variation is ascribed by the researcher to factors such as the familiarity of the participant with the interviewer and the participant's comfort level in sharing experiences with another person.

The long gap between experiencing and remembering was designed to give participants sufficient time to forget details of the experiences. Observations made during the study and participants' comments showed that the participants indeed felt that they did not remember much about their conference visit; one participant even stated that they did not remember that they had been interviewed.

The exact timespan involved in this study was selected due to practical considerations and a similar effect may be achievable with shorter gaps. This could make participant recruitment easier. Forgetting curves for autobiographical memory reported in the psychology literature surveyed in Chapter 2 relate to much longer timespans, i. e. decades rather than weeks, months or years. Hence, they offer no guidance on selecting appropriate timespans. However, the academic environment strongly depends on annual cycles (the academic calendar follows a yearly schedule with clearly defined time periods; most academic conferences are held annually or bi-annually), which makes it likely that a gap of at least one year is necessary in this domain.

8.3.4 Procedure

At the beginning of their session, the participant received a brief introduction to the goals and procedure of the study. The participant information sheet and the consent form are shown in Appendix C.2.

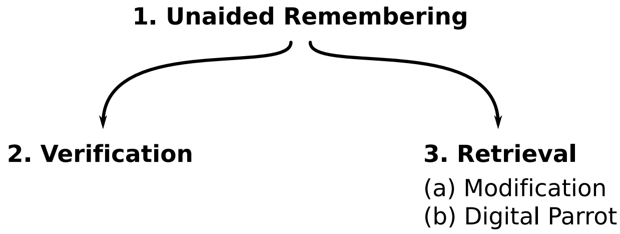

Figure 8.1. The three main stages in the recall phase of the study

After the participant had consented to participate in the study and to have audio and video recordings made of the session, they received a thorough introduction to the system and its features, followed by some practice tasks. This was followed by the three main stages of the study: unaided remembering, verification and retrieval (Figure 8.1). The session was concluded with a brief biographic questionnaire and a discussion between the participant and the researcher. Appendix C.2 shows the material given to each participant for this study: a participant workbook and a sample question sheet as used in the three main stages.

The following describes each of the stages in more detail.

Demonstration The researcher gave a quick overview of the Digital Parrot's main views and navigators. The data shown in the Digital Parrot during this stage was training data and identical to the memory data used in the first evaluation of the Digital Parrot. It is described in detail in Section 7.2.4.

Training Still working with the training data, the participant was first given the opportunity to explore the Digital Parrot's user interface. They were then asked to work through the four retrieval tasks used in the usability study (see Section 7.2.5). The participant was encouraged to ask questions during this stage if any parts of the Digital Parrot's user interface were unclear.

Unaided Remembering The participant was then provided with a list of questions. For each question, the participant was asked to answer it immediately. They were also asked to rate the completeness and the correctness of their answer and to state whether they found the answer satisfactory. The question was then put aside to be used in either the verification or the retrieval stage, depending on the participant's ratings.

Verification The participant was then asked to go through the questions that had been answered satisfactorily in the Unaided Remembering stage. For each question, the participant was asked to verify their answer using the Digital Parrot, which now contained their personalised data. If the answer or their ratings of it changed as a result of using the Digital Parrot, the participant was asked to modify them accordingly.

Retrieval The participant was then asked to go through all remaining questions – i. e. those which had *not* been answered satisfactorily in the Unaided Remembering stage. For each question, the participant was first given the opportunity to modify their answer and their ratings in case these had changed based on other questions/answers seen in the meantime. The participant was then asked to describe verbally how they would normally go about finding the answer to this question and to estimate their chances of success and the time effort for this method. After that, they were asked to use the Digital Parrot to attempt to answer the question. Finally, the participant rated the answer found using the Digital Parrot for its completeness and correctness and stated whether they found this answer satisfactory.

Conclusion The session was concluded with a brief questionnaire collecting information about the background of the participant and a discussion with the participant about their experience of using the Digital Parrot. During the discussion, the participant was encouraged to share any comments they had about the Digital Parrot and other aspects of the session. Additionally, the researcher asked some questions to guide the discussion. These were related to the type of remembering (re-living, knowing, reconstructing or guessing), the degree to which the participant thought the questions were realistic and the participant's use or

Items	P1	P2	P3	P4	Median	IQR
Conferences	5	2	1	3	2.5	1
Conf. with details	2	2	1	1	1.5	1
Sessions	65	52	15	26	39	26
Presentations	65	7	9	76	37	56
Conversations	4	3	4	8	4	0
Topics	5	5	3	7	5	0
Conv. partners	5	3	4	8	4.5	1
People	87	40	23	121	63.5	47
Items with time	181	99	38	126	112.5	27
Items with place	194	112	32	116	114	4
Nodes	344	197	111	370	270.5	147
Statements	1249	909	234	624	766.5	285

Table 8.2. Number of items in memory data, by participant

non-use of the Digital Parrot's components.

8.3.5 Memory data

Information from the interviews in the experiencing phase and from publicly available information about the conferences attended by the participants was transcribed by the researcher into data files for the Digital Parrot. The format of such files is described in Section 5.6.

The memory data for each participant was heavily influenced by the information given in the interview by the participant as well as by publicly available information about the conference. Table 8.2 shows how many items of some important types were transcribed for each of the participants. A number of guidelines were followed when creating the memory data for each participant.

Conversations. All conversations that were clearly recognisable as such were included, along with their topics and conversation partners. Where available, temporal and (more rarely) location information were added to the conversation. All information about conversations, their topics and their conversation partners was based mainly on descriptions by the participant during the interview.

Some additional information was added by the researcher in a few cases based on publicly available information; an example is a case in which a conversation partner's name was added by the researcher after it was identified on the list of attendees based on characteristics described by the participant.

Conference program information. Information about a conference, the conference programs and other people present at the conferences were based on publicly available information about the conference such as the conference program. The variations across participants in the number of sessions, presentations and people are explained by several factors. For example, the conferences attended varied considerably in number of sessions, number of presentations overall as well as presentations per session and number of attendees. The proportion of overall sessions and presentations at a conference that were included varied depending on the ease with which the conference program could be converted semi-automatically into the Digital Parrot's data format. For P4's conference, a list of attendees was available, while the information about people for all other participants in the study was based on the presenters at the conference.

These variations are similar to those that can be expected in a full implementation of the conceptual design for an augmented autobiographical memory introduced in Chapter 4. When information is captured automatically, the level of detail in the captured data is highly dependent on the ease with which information from external sources can be converted into a format that the augmented autobiographical memory system can process.

8.3.6 Tasks

The questions were derived from the information gathered in the interviews and thus personalised for each participant. Each question fell into one of the following categories.

(C1) Conversation partner by topic: find a conversation partner (e. g. their name), given the topic of conversation

Example: "With whom did you speak about *topic* during *conference*?"

All questions in this category required personal knowledge to answer.

(C2) Name of person by characteristic: find the name of a person, given some info about the person (e. g. their affiliation)

Example: "What is the name of the *affiliation* guy at *conference*?"

All questions in this category required personal knowledge to answer.

(C3) Other conversation topics by topic: find additional topics of a conversation, given one topic of the same conversation

This category contained only one question, which required personal knowledge to answer.

(C4) Time of event by characteristic: find the time of an event (e. g. conversation), given a characteristic of the event (e. g. topic)

Example: "When was *conference*?"

All except two questions in this category could be answered using publicly available information. The two remaining questions required personal knowledge to answer.

(C5) Other people by event: find other people present at an event, given a characteristic of the event (e. g. name)

Example: "With whom did you take pictures at *place*?"

All except one question in this category required personal knowledge to answer. The remaining question could be answered using publicly available information.

(C6) Other people by event (advanced): find other people present at an event, where it is non-trivial to find the event for some reason

Example: "Who is the student of *name* who gave a talk at *conference*?"

All questions in this category could be answered using publicly available information.

(C7) *Place of event by characteristic:* find the place of an event (e. g. conference), given a characteristic of the event (e. g. name)

Example: "Where was the *conference* conference dinner?"

All questions in this category could be answered using publicly available information.

(C8) *Characteristic of event by time and type* find characteristics of an event, given the event's time and type. Typically this involved choosing one of several instances of the same type that differed by time.

Example: "What was the topic of the *weekday* keynote at *conference*?"

All questions in this category could be answered using publicly available information.

There were a total of 32 questions, between seven and nine per participant and between one and seven per category. Table 8.3 shows the number of questions per participant per category. The number of questions for each type vary across participants because the questions that could be generated for a participant depended heavily on the information available from their interview and from other material about the conference they had attended. The interviews differed considerably in the level of detail and the types of experiences that were described.

To select the questions, first the researcher listened to the interviews and extracted questions and answers from pieces of information mentioned by the interviewee that were thought to be memorable. This is similar to the question selection process described by Vemuri et al. (2006). The extracted questions were then grouped into the categories listed above. Where there were fewer questions in a category for one

	Question type	P1	P2	P3	P4	Total
C1	conversation partner by topic	–	1	1	1	3
C2	name of person by characteristic	1	–	1	1	3
C3	other conversation topics by topic	–	–	1	–	1
C4	time of event by characteristic	2	2	1	2	7
C5	other people by event	2	2	1	1	6
C6	other people by event (advanced)	2	–	–	1	3
C7	place of event by characteristic	1	1	2	1	5
C8	characteristic of event by time and type	1	1	1	1	4
	Personal	4	3	4	3	14
	Non-Personal	5	4	4	5	18
	Total	9	7	8	8	**32**

Table 8.3. Number of questions for each question type (category and personal/non-personal content), by participant

participant compared to the others, a second attempt was made to extract additional questions of this category from the interviews and from publicly available information.

8.3.7 Types of data collected

The study was conducted in the single-user usability lab at the Computer Science department at the University of Waikato. This lab has equipment for recording audio and video of the participant and for capturing the participant's screen. All participants consented to having audio and video recordings made of their session. However, due to equipment failure, a complete audio recording was made for only one session. The researcher took notes during the sessions, which together with information transcribed from the video recordings and with the information recorded by the participants on paper formed the basis for analysis. Each participant was given the opportunity to review and amend the transcript of their session.

Training answers and answers to the biographical questionnaire were recorded by the participants in a paper-based workbook similar to that used in the usability study. The workbook is shown in Appendix C.2.

Each question to be remembered was printed on its own question sheet that also contained several rating scales and instructions to guide the participant through the stages. A sample question sheet is given in Appendix C.2. Each question sheet contained the question and empty space for free-text answers. Correctness and completeness ratings were given on a six-point Likert-like scale, as was the estimated success of manual retrieval strategies. Satisfaction ratings were given on a binary yes/no scale. Manual retrieval strategies were described verbally by the participant, while the estimate for the time required for the strategies was given in writing.

8.3.8 Pilot Study

A pilot study was conducted with the participant of the pilot study conducted for the experiencing phase. Insights from the pilot study led to several changes in the Digital Parrot's user interface, in the procedure of the study and in the questionnaires.

User Interface The biggest change to the user interface of the Digital Parrot made based on observations during the pilot study was to combine the two types of main view, list and graph view, into a single interface using a tabbed pane. The version of the Digital Parrot used in the study described in the previous chapter had required a restart to change between main view types, which facilitated the between-subjects design of that study. For the study described in the current chapter, it was found that having separate instances of the Digital Parrot running that did not share navigators was confusing for the participants.

Another change was made to the navigator behaviour with regards to highlighting. In the usability study and in the pilot study of the bigger study, all active navigators' highlight filters were combined disjunctively. This meant that those items were highlighted in the main view that were highlighted through *at least one* active navigator. For the main study, this was changed such that only those items were highlighted in the main view that were highlighted by *the most recently active* navigator.

Based on the findings of the study described in the previous chapter and before the pilot study was conducted, the user interface of the connections manager had been changed substantially (see Section 7.5.2). Adding a link to the chain or selecting an instance for a chain link in the version used for the pilot study was done by selecting an item in the main view via a single mouse click. The participant in the pilot study reported that this was very confusing. For the actual study, this operation was changed to require a double-click on the item.

The other changes made focused on appearance, such as improving the formatting of timestamps, and space constraints introduced by the lower screen resolution of the monitor in the usability lab. The map's default size was reduced. Type and instance names in the connection manager's drop-down box were cut off after a length threshold to ensure visibility of several connection links. Some internal types, such as TimedThing and PlacedThing (see Section 5.6.2), were set to be shown less prominently in the user interface. Finally, because the Digital Parrot would have to be re-started with a different data file after the training stage, icons to start the training version and the personalised version of the Digital Parrot were placed on the desktop.

Procedure In the pilot study, questions that the participant had been able to answer satisfactorily were put aside after they had been answered and did not lead to any interaction with the Digital Parrot.

The verification phase was introduced for these questions in the actual study, for two reasons. Firstly, letting the participant use the Digital Parrot to verify answers that they were already satisfied with allowed them to familiarise themselves with their personalised data while working on a presumably easier task. Secondly, this phase made it possible to observe the participants perform tasks that, again presumably, involved purely recognition rather than a mixture of recognition and recall.

Question sheets Based on the participant's comments in the pilot study, some changes were made to the question sheets. These were in addition to the changes made due to the introduction of the verification stage.

The question sheets in the pilot study asked the participant to rate the correctness of their answers and to specify their satisfaction with the answer. The participant in the pilot study commented several times that they found the correctness rating difficult to make – they knew that the answer they had given was incomplete, but they were quite certain that those parts of the answer that they *had* given were correct. To address this issue, an additional rating of the completeness of the answer was introduced in all three main stages of the actual study.

The order of choices for the satisfaction rating was reversed, from yes/no in the pilot study to no/yes in the actual study. This was done to match the direction of the rating scales, with the negative choice on the left and the positive choice on the right.

8.4 Quantitative Findings

This section provides data related to the following questions:

- What and how well did the participants remember without help?
- How would the participants normally answer these questions?
- How did the participants use the Digital Parrot to answer these questions?
- How well did the participants remember with help from the Digital Parrot?

Comparisons are made for these questions across participants as well as across question types (categories and personal/non-personal content).

8.4.1 Unaided remembering

This section first describes the answers produced by the participants during the unaided remembering stage. It then describes how the participants rated these answers.

Answers

During the unaided remembering stage, about 20% of all questions were answered satisfactorily and another 40% of all questions were answered in part. The remaining 40% of questions were not answered at all during the unaided remembering stage. At the beginning of the recall stage

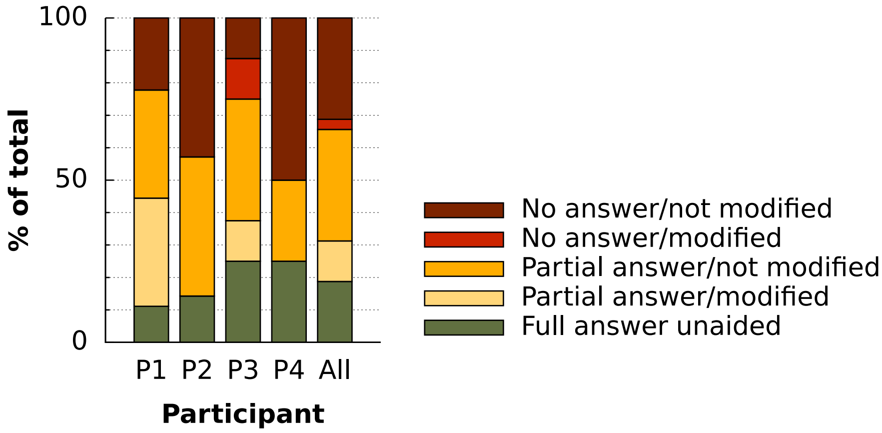

Unaided	Modified?	P1	P2	P3	P4	Total
Full answer	n/a	1	1	2	2	6
Partial answer	Yes	3	–	1	–	4
	No	3	3	3	2	11
No answer	Yes	–	–	1	–	1
	No	2	3	1	4	10
Total		9	7	8	8	**32**

Figure 8.2. How many questions were answered when, without using the Digital Parrot, per participant. The chart shows percentages relative to the total number of questions for this participant.

for each question, participants were given the opportunity to amend or modify their answer before using the Digital Parrot for this question (see Section 8.3.4).

Figure 8.2 shows how many questions were answered fully or in part during the unaided remembering stage and how many were modified at the beginning of the retrieval stage (i. e. before using the Digital Parrot for this question), broken down by participant. Figure 8.3 shows the same data by question type.

Differences across participants Two participants, P1 and P3, remembered on average more questions, fully or in part, compared to the two other participants. Two participants, P2 and P4, never modified any of their initial answers at the beginning of the retrieval stage.

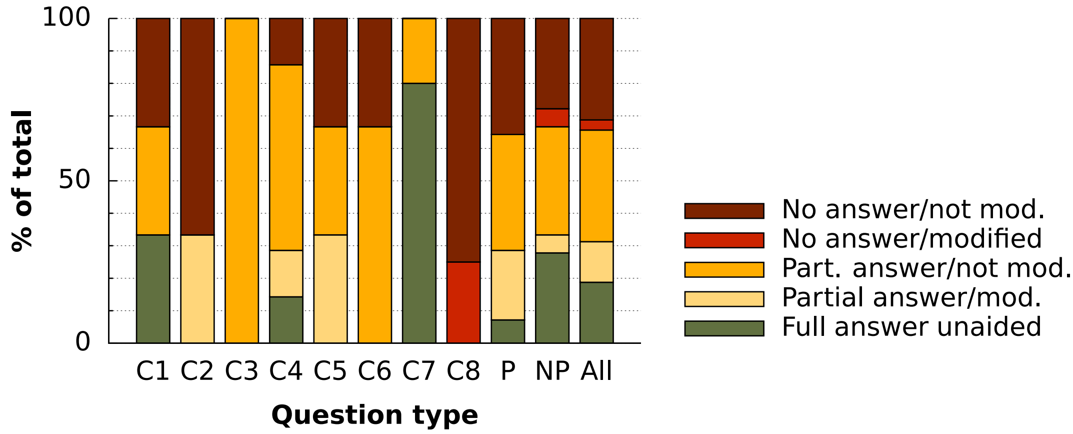

Unaided	Mod.	C1	C2	C3	C4	C5	C6	C7	C8	P	NP
Full answer	n/a	1	–	–	1	–	–	4	–	1	5
Part. answer	Yes	–	1	–	1	2	–	–	–	3	1
	No	1	–	1	4	2	2	1	–	5	6
No answer	Yes	–	–	–	–	–	–	–	1	–	1
	No	1	2	–	1	2	1	–	3	5	5
Total		3	3	1	7	6	3	5	4	14	18

Figure 8.3. How many questions were answered when, without Digital Parrot, by question type (category and personal/non-personal content). The chart shows percentages relative to the total number of questions for this type.

Differences across question types Two categories stand out with below-average numbers of initially answered questions. These are categories C8, “characteristic of event by time and type”, and C2, “name of person by characteristic”. This may be explained for category C8 with the fact that questions in this category typically required the participant to choose from several instances of the same type according to the time they occurred (for example, keynotes by weekday). Often, participants commented on questions in this category with statements such as “I remember that I went to *some* keynotes, but I can’t remember which of them was on what day.” Many of the questions in category C2 had been chosen deliberately to target potential memory failures, for example a name about which the participant had shown some uncertainty in the interview.

Three categories have above-average numbers of initially answered

		Completeness		Correctness		
Unaided	**Mod.?**	initial	mod.	initial	mod.	**Total**
Full answer	n/a	5 [0]		5 [0.75]		6
Partial answer	Yes	3.5 [1.5]	4 [0]	5 [1]	5 [0.25]	4
	No	4 [2]		4 [2.5]		11
No answer	Yes	1 [–]	2 [–]	1 [–]	2 [–]	1
	No	1 [0]		1 [0]		10
Overall		2 [3]	2.5 [3]	4 [4]	4 [4]	**32**

Table 8.4. Median/IQR completeness and correctness ratings of answers before using the Digital Parrot.

questions. These are C3, "other conversation topics by topic", C4, "time of event by characteristic", and C7, "place of event by characteristic". C3 contains only a single question, but the other two categories genuinely stand out. Of these, C7 arguably contained only questions that were easy to answer.

Personal questions were remembered in full less often than the average, while non-personal questions were remembered in full more often than the average. The proportion of questions that were initially not answered at all, however, is very similar for both types.

Ratings

Table 8.4 summarises how the participants rated the completeness and correctness of their answers before using the Digital Parrot. Initial completeness and correctness ratings for questions that were not answered at all initially were not captured for all participants but are assumed to be 1, corresponding to "not certain at all", on both scales.

Participants rated the questions that they answered in full during the unaided remembering stage on average as highly complete and correct. This is not surprising because high ratings on these scales were necessary to consider the question as "answered in full" at this stage.

Completeness and correctness of partial answers during the unaided remembering stage were rated slightly lower than those of full answers, but ratings on average are still on the positive half of the scale. Modifi-

Unaided	Modified?	initially	modified	Total
Full answer	n/a	5	n/a	6
Partial answer	Yes	1	1	4
	No	0	n/a	11
No answer	Yes	0	0	1
	No	1	n/a	10
Total		7	8	**32**

Table 8.5. Number of satisfactory answers before using the Digital Parrot

cations made to these answers at the beginning of the retrieval change led to no change in completeness and correctness ratings.

Questions that were not answered at all initially were all rated with the lowest value on the correctness and completeness scale. Not all of the question sheets in this category contain ratings, because some participants were instructed to leave out the ratings in this case. One participant gave an answer to one of their questions in this category at the beginning of the retrieval stage. This was based on information that they had come across while retrieving the answer to an earlier question. However, the rating for this answer only improved by one step on both scales.

Table 8.5 lists the number of satisfactory answers at each stage before verification or retrieval using the Digital Parrot. Satisfaction ratings were given on a binary yes/no scale.

Almost all full answers given in the unaided remembering stage were considered satisfactory by the participants. One non-satisfactory answer was still grouped into this category by the researcher because the difference between the answer given by the participant and the answer expected by the researcher was merely orthographic in nature.

One of the partial answers was considered satisfactory by the participant but was grouped into the "partial" category by the researcher because it was considered too imprecise by the researcher. One answer was classified as "no answer" by the researcher even though an answer was given and rated as satisfactory by the participant. However, the answer was considered by the researcher as a re-iteration of the question itself.

None of questions that were not answered at all initially were considered satisfactorily answered; as stated above, some participants were instructed not to provide this rating for questions in this category. Modifications to answers at the beginning of the retrieval stage did not lead to changes in satisfaction rating.

Several participants commented that they found it difficult to give satisfaction ratings. Most said that this difficulty was due to the fact that the question lacked a context, which made it particularly difficult to determine how much detail was required in the answer and how much effort they would typically be willing to put into finding an answer.

8.4.2 Retrieval without using the Digital Parrot

This section focuses on the information sources and strategies that the participants described for answering the questions without using the Digital Parrot, as well as on the expected success and expected time effort for these information sources and strategies.

Information sources and strategies

Participants described a range of information sources and strategies that they would normally use to answer the questions. Information sources and strategies for two questions were not recorded. The recorded information sources and strategies are:

Conference website: Three of the four participants said that they would look up information on the conference website. The types of information that the participants expected to find on the conference website ranged from specific facts, such as the conference dates and the name of the conference hotel, to more complex collections of information, such as the conference program and the list of attendees.

Conference proceedings: Three of the four participants said that they would look through the published conference proceedings. All participants here referred to the copy of the proceedings that they brought home from the conference, whether it was printed or electronic (on a USB key). Conference proceedings were specified as the source for details about conference events (e. g. keynotes times and topics) as well for information about conference attendees (by going through the titles and authors of published papers).

E-mail: Three of the four participants said that they would search their collection of e-mail messages to answer some of the questions. Most often this was to determine the names and identities of other people with whom they had communicated prior the conference, for example to organise sightseeing on the day before the conference started.

One participant stated that they would look in their e-mail collection for impersonal information related to the conference as well, such as information that may be in the conference's call for papers.

Notes (digital): One participant uses an idiosyncratic personal knowledge management system which would be the first point of call to answering almost all questions given to this participant. This participant favoured their own system over publicly available information even for impersonal questions.

Ask another person: Three of the four participants said that they would ask another person. Invariably this was to find the answer to questions about events involving other people and the person whom the participant would ask had participated in the event.

Photo: Two participants said that they would look through photographs from the conference, either photographs they had taken themselves or those that had been taken by others and were available online. In both cases this was to determine who else was present at a given event; the participants hoped that everyone present would be shown on the photographs.

Personal document: One participant said that they would retrieve a document from their personal filestore as the first step to answering one of the questions. This document, which had been created by someone other than the participant, was “quite well filed” and thus easily accessible. The participant hoped that it would provide information that could then be used as a starting point for another strategy.

Semi-public website: One participant said that they would look up information on a semi-public website (accessible with a password) to find information about the conference program that was not available directly on the conference’s public website.

From memory: One participant stated that they would attempt to answer one of their questions “just from memory”.

Table 8.6 show the number of distinct strategies and the average number of strategies per question for each participant, for each question type (category and personal/non-personal content) and for all questions together.

Participants named between one and three strategies per question (median = 1, IQR = 1). More than one strategy was named by all participants for at least one question. This was typically a main strategy and a fallback strategy in case the first one did not give the desired result. In one case, the first strategy was used to find a set of possible answers, with the second strategy used to determine which of these was the actual answer to the question. More than one strategy per question was described for 10 out of the 26 questions with known strategies.

Differences across participants One participant, P4, stands out with a lower-than-average number of distinct strategies. Another participant, P3, stands out with a higher-than-average number of strategies per question.

Each participant’s numbers of questions for which they would use each strategy are shown in Table 8.7. The strategies used vary widely across participants. Strategies that are favoured by some participants are not used at all by others, for example the conference website and the conference proceedings. Two participants, P2 and to a lesser degree P4, show a

		Strat. per question		
Part./Type	**Distinct Strat.**	Median	IQR	**Freq**
P1	6	1	0.25	8
P2	4	1	0.75	6
P3	5	2	0	6
P4	3	1	0	6
C1	1	1	0	2
C2	3	2	0.5	3
C3	2	2	0	1
C4	4	1	0	6
C5	4	1	1	6
C6	3	1	0	3
C7	2	2	0	1
C8	3	2	0.25	4
Personal	8	1	1	13
Non-Personal	6	1	1	13
Overall	9	1	1	**26**

Table 8.6. Distinct strategies and number of questions per strategy, by participant and question type (category and personal/non-personal content)

Method	**P1**	**P2**	**P3**	**P4**	**Total**
conference website	3	–	4	2	9
proceedings	–	1	4	4	9
e-mail	2	1	2	–	5
notes (digital)	–	5	–	–	5
ask other person	2	1	1	–	4
photo	1	–	1	–	2
personal document	1	–	–	–	1
semi-public website	1	–	–	–	1
from memory	–	–	–	1	1

Table 8.7. Strategies for manual retrieval, by participant

clear preference for one particular strategy, while the other participants' preferences are spread out across several different strategies.

Differences across question types Even though more distinct strategies were described for some categories, these are simply the categories with

Method	C1	C2	C3	C4	C5	C6	C7	C8	P	NP
conference website	–	–	1	4	–	–	1	3	2	7
proceedings	–	2	1	–	1	1	1	3	4	5
e-mail	–	2	–	1	2	–	–	–	4	1
notes (digital)	1	–	–	1	2	–	–	1	2	3
ask other person	–	–	–	1	2	1	–	–	3	1
photo	–	–	–	–	2	–	–	–	2	–
personal document	–	1	–	–	–	–	–	–	1	–
semi-public website	–	–	–	–	–	1	–	–	–	1
from memory	1	–	–	–	–	–	–	–	1	–

Table 8.8. Strategies for manual retrieval (in number of questions), by question type (category and personal/non-personal content)

more questions overall. The question type does not appear to have an influence on the number of distinct strategies or on the average number of strategies per question.

Table 8.8 shows the number of questions for each strategy by category. The conference website is used much more often to retrieve non-personal information than it is used for personal information; typically, this is to determine the time of an event (C4) or to distinguish between events based on time (C8). Likewise, the proceedings are used more for non-personal than for personal information, again for distinguishing between events based on time (C8) but also to find information about people (C2). The conference website and the proceedings are clearly the preferred methods to distinguish between events based on time, while the time of an event is most commonly looked up on the conference website.

Searching one's e-mail collection, asking other people and looking at photographs are used more often for personal than for non-personal information and in particular to determine other people present at an event (C5). Consulting digital notes (used by only one participant) was divided evenly between personal and non-personal content.

Expected success

Overall, participants expressed high confidence in the success of their approach, but confidence ratings vary considerably across questions (me-

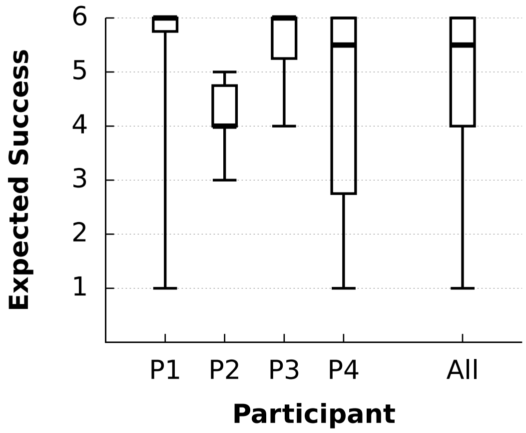

Figure 8.4. Expected success of manual retrieval, by participant. Expected success ratings were given from “not good at all” (1) to “very good” (6).

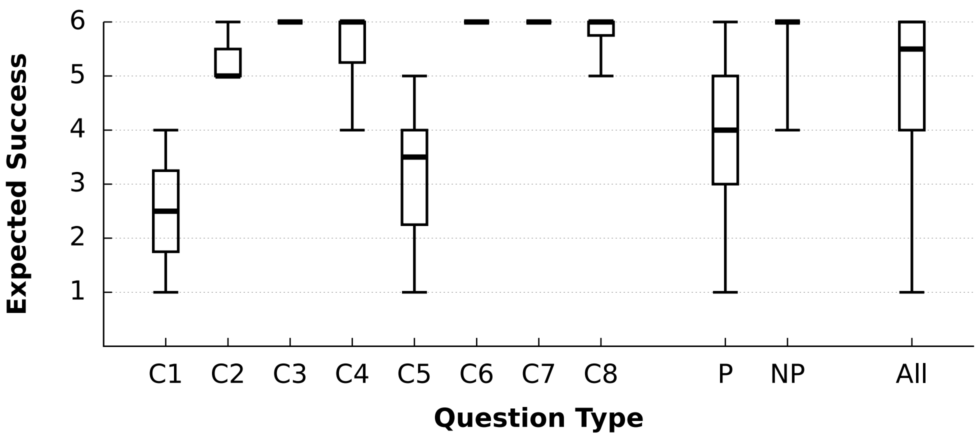

Figure 8.5. Expected success of manual retrieval, by question type (category and personal/non-personal content). Expected success ratings were given from “not good at all” (1) to “very good” (6).

dian = 5.5, IQR = 2).

Figure 8.4 shows boxplots for the expected success ratings by participant and by category; Figure 8.5 shows boxplots for the expected success rating by question type.

Differences across participants. Two participants, P2 and P4, on average expect to be less successful than the others. P4 also shows a much wider dispersion of success ratings than all other participants. The same participant who stands out as having a higher-than-average number or strategies per question, P3, is the only participant who expects to be at least moderately successful in finding an answer to all of their questions.

Differences across question types. The expected success rating is higher on average for non-personal questions than for personal questions. All participants expect to be successful at least to some degree in answering non-personal questions, while some personal questions may very well never be answered at all. Particularly difficult to answer appear to be questions about people: conversation partners (C1), other people present at an event (C5) and, to a lesser degree, the name of a person (C2).

Expected time effort

Participants' estimates for the time required to find the answer for a question using the strategy they had just described vary widely across participants as well as across question types (categories and personal vs non-personal content).

Several participant commented that their estimate for the time it would take to find an answer was actually a combination of two estimates: how long it would take them to find an answer, and how long at most they would look for an answer before giving up.

Differences across participants The expected time effort varies considerably between participants. One participant gave estimates generally in the range of seconds to less than five minutes, two participants gave estimates of a few minutes (under five and under ten minutes, respectively) while the fourth participant gave estimates of at least five minutes, more often about ten to twenty minutes. Table 8.9 shows the fastest and slowest estimate and method for each participant.

Two participants gave their fastest estimate in seconds, while the other

	Fastest		Slowest	
Part.	Speed	Method	Speed	Method
P1	seconds	website	days to weeks or never	ask other person
P2	5 min	notes; e-mail	30 min to 1 day	ask other person
P3	2 min	website; proceedings	1 day	ask other person
P4	30–40 secs	website	never	from memory

Table 8.9. Expected time effort for manual retrieval, by participant

two gave theirs in minutes. The fastest estimated method differs across participants, but looking up information on the conference website was mentioned by three of the four participants as the fastest method.

The slowest estimate was generally given in days and generally was given where the strategy included asking other people present at the event. All participants explained that they person they would ask lived in Europe or North America and that they expected turnaround times of at least half a day to a day because of the timezone difference to New Zealand. Two participants stated that they may never be able to answer some of their questions.

Differences across question types Table 8.10 shows the fastest and the slowest estimate and method for each category. Even though the fastest and slowest estimates are almost identical for personal and non-personal questions, estimated speeds actually differ between these two question types. The "never" estimate listed as slowest estimate for non-personal questions is actually a worst-case estimate in case access to the information source is lost, while answering the question with access to the information source is estimated to take just two or three minutes. The next-slowest estimate for answering a non-personal question is half a day or less when asking another person. All other estimates for answering non-personal questions are at most twenty minutes. In contrast, both "never" estimates for personal questions appeared to be considered as likely outcomes by the participant.

	Fastest		Slowest		Freq
Cat.	Speed	Method	Speed	Method	
C1	10 min	notes	never	from memory	2
C2	30–45 sec	document; e-mail	3–5 min	proceedings	5
C3	< 5 min	website; proceedings	< 5 min	website; proceedings	2
C4	seconds	website	30 min to 1 day	ask other person	7
C5	30 sec	e-mail	days, weeks or never	ask other person	9
C6	2–3 min	semi-public information	never	semi-public information	3
C7	2 min	website; proceedings	2 min	website; proceedings	2
C8	30 sec	website	15–20 min	notes; proceedings	7
P	30 sec	e-mail	never	from memory	19
NP	seconds	website	never	semi-public website	18

Table 8.10. Expected time effort for manual retrieval, by question type (category and personal/non-personal content)

Two of the categories with low estimated success are also among those that take longest to answer: conversation partners (C1) and other people present at an event (C5). Surprisingly, the other slow category, more complex questions about other people present at an event, has a very high estimated success rating. Conversely, finding the name of a person (C2) is expected to be fast even though it has a slightly below estimate of success.

Participants expect to find answers quickly when asked to distinguish between events based on time (C4). Three of the four participants gave estimates for answering questions in this category that are fast compared to the participant's typical speed estimates; all of these questions (four in total) could be answered using publicly available information. Questions about other people at an event (C5) received fast estimates from two participants; these questions (three in total) required personal knowledge to answer. One participant each also estimated that they would be able to answer quickly questions about from two participants for questions about the name of a person (C2; two questions requiring personal knowledge) and the place of an event (C7; two questions answerable based on publicly available information).

8.4.3 Retrieval using the Digital Parrot

This section first describes the participants' strategies in verifying answers and answering questions using the Digital Parrot. It then describes the participants' ratings of the verified and retrieved answers.

Strategies

In describing the participants' interactions with the Digital Parrot while verifying or retrieving the answer to a question, the following component categories are used.

Graph visible: the graph main view was visible for some time while the participant was verifying or retrieving the answer to a question;

Switch to graph: the participant switched the type of main view from list to graph;

Graph features: the participant used features of the graph view, such as selecting an item to view connected items;

List visible: the list main view was visible for some time;

Switch to list: the participant switched the type of main view from graph to list;

List features: the participant used features of the list view, such as selecting an item to see other occurrences of the same item or changing the sort key or sort order of the statements;

Timeline: the participant used the timeline navigator;

Method	P1	P2	P3	P4	Total
Graph visible	7	7	8	7	29
Switch to graph	–	1	1	1	3
Graph features	7	5	3	4	19
List visible	–	1	3	2	6
Switch to list	–	1	2	1	4
List features	–	–	3	–	3
Timeline	2	–	2	2	6
Map	2	–	–	–	2
Search	5	6	7	6	24
Connections	2	5	1	–	8

Table 8.11. Usage of components of the Digital Parrot (in number of questions), by participant

Map: the participant used the map navigator;

Search: the participant used the search function; and

Connections: the participant used the connections navigator.

Table 8.11 shows for how many questions each component was used, broken down by participant. Table 8.12 shows the same data broken down by category. Particularly before working on the first question in the verification stage, clearly exploratory behaviour was shown by some participants. Such behaviour is not considered in the descriptions that follow.

The graph view was clearly preferred over the list view, being visible in every single question. In part this may be influenced by the fact that the graph view was the default on program start-up. However, all participants who switched to the list view switched back to the graph view at some point. Graph features were used for just over half of all questions (55%), while list features were used by only one (P3) of the three participants who used the list.

The search function was used in roughly three quarters of all questions (72%). Not counting differently spelled variants of search terms and repeated searches, 31 searches were performed in total. Table 8.13 shows how many searches involved search terms taken from the question and

Method	C1	C2	C3	C4	C5	C6	C7	C8	P	NP
Graph visible	3	2	1	5	4	2	5	4	11	15
Switch to graph	–	1	–	1	–	–	–	1	1	2
Graph features	2	1	–	2	4	2	3	2	8	8
List visible	–	1	–	2	–	–	1	2	1	5
Switch to list	–	–	–	1	–	–	1	2	–	4
List features	–	1	–	1	–	–	–	1	1	2
Timeline	–	–	–	3	–	–	–	3	–	6
Map	–	–	–	–	–	–	1	–	–	1
Search	2	2	1	5	3	2	3	3	9	12
Connections	1	2	–	1	2	1	–	1	5	3
Overall	3	3	1	7	4	2	5	4	12	17

Table 8.12. Usage of components of the Digital Parrot in number of questions, by question type (category and personal/non-personal content)

how many search terms were ontology classes, or could be classes in a customised vocabulary (such as “lunch”).

Ten (32%) used search terms that did not appear in the question text; in all but two cases, the participant had given a partial answer in an earlier stage and the terms were taken from this partial answer. In one of the two cases in which no initial answer had been given, the search term was closely associated with a person mentioned in the question. In the other case, the term was taken from the results of a previous search related to the same question. The remaining 21 searches used search terms that were taken from the question text. Nine searches used terms that were, or could be, classes, while the remaining 22 searches used content terms (such as a name).

The connections manager was used in just over a quarter of the questions (27%) and the timeline in a fifth of the questions (20%). The map was used for only one question.

The first interaction strategy was successful for 22 questions, while participants used more than one interaction strategy for ten questions.

Differences across participants One participant did not use the list at all. Another never used the timeline navigator and a third participant

# Searches		Question: Yes	Question: No	**Total**
Class	Yes	8	1	9
	No	13	9	22
Total		21	10	**31**

Table 8.13. Number of searches that used search terms from question/not from question and that were, or could be, classes of information items

never used the connections navigator. The map navigator was used by only one participant. This participant used the map navigator for two questions; in one case, the participant had already answered the question to their satisfaction using a different strategy but wanted to see if they “could have gotten the answer [using the map], too”.

Three participants predominantly used search terms taken from the question, while the remaining participant (P3) almost exclusively used search terms that did not appear in the question (seven out of nine searches). The same participant was the only one who never used a search term that was, or could have been, a class.

Two participants changed strategy once each; one participant changed strategy once for two questions and twice for a third question; one participant changed strategy once for three questions and three times for one question.

Differences across question types. The graph view was used for almost all (89%) of the 29 questions for which components usage is known. The list view was mainly used for non-personal questions (29%, vs 8% for personal questions).

The timeline was used only for non-personal questions and only for distinguishing between events of the same type based on time (C8) and the time of an event (C4). The one time that the map was used was to find the place of an event (C7). The search function was used roughly equally often for personal questions (75%) as for non-personal questions (70%), and particularly often to find the time of an event (C4). The connections manager was used more than twice as often for personal questions (41%) than for non-personal questions (17%).

Two questions about the time of an event (C4) were answered purely from the timeline – a third of all questions for which the timeline was used at all. One question, about the name of a person (C2), was answered purely by using the connections manager. One question, about other people present at an event (C5), was answered using graph features only. One question, about the place of an event (C7), was answered by scrolling through the list of statements.

Search terms that had not been taken from the question were used more frequently than on average for questions about the name of a person (C2) and about the place of an event (C7). Searches for questions about other people at an event (C5) and about the characteristics of an event identified by time (C8) more frequently than on average used terms that were, or could be, classes. The predominance of non-class search terms over class search terms was more pronounced for questions with personal content.

Overall, changes were made from search to connections five times (C2, C5 twice, C6, C8), from connections to search four times (C4, C5 twice, C8), from timeline to search twice (both C4) and from search to timeline once (C2).

Ratings

Table 8.14 compares how the participants rated the completeness and correctness of their answers before and after using the Digital Parrot. The "before" column in this table corresponds to the "modified" or, for non-modified answers, to the "initially" column in Table 8.4. Figure 8.6 shows before-and-after comparisons of completeness ratings by participant, category and question type, while Figure 8.7 shows correctness ratings broken down in the same way.

Use of the Digital Parrot led to a statistically highly significant improvement in both completeness (Mann-Whitney $U = 904.5$, $n_1 = n_2 = 32$, $p < 0.0001$ one-tailed) and correctness (Mann-Whitney $U = 890$, $n_1 = n_2 = 32$, $p < 0.0001$ one-tailed) ratings. Naturally, the biggest increase can

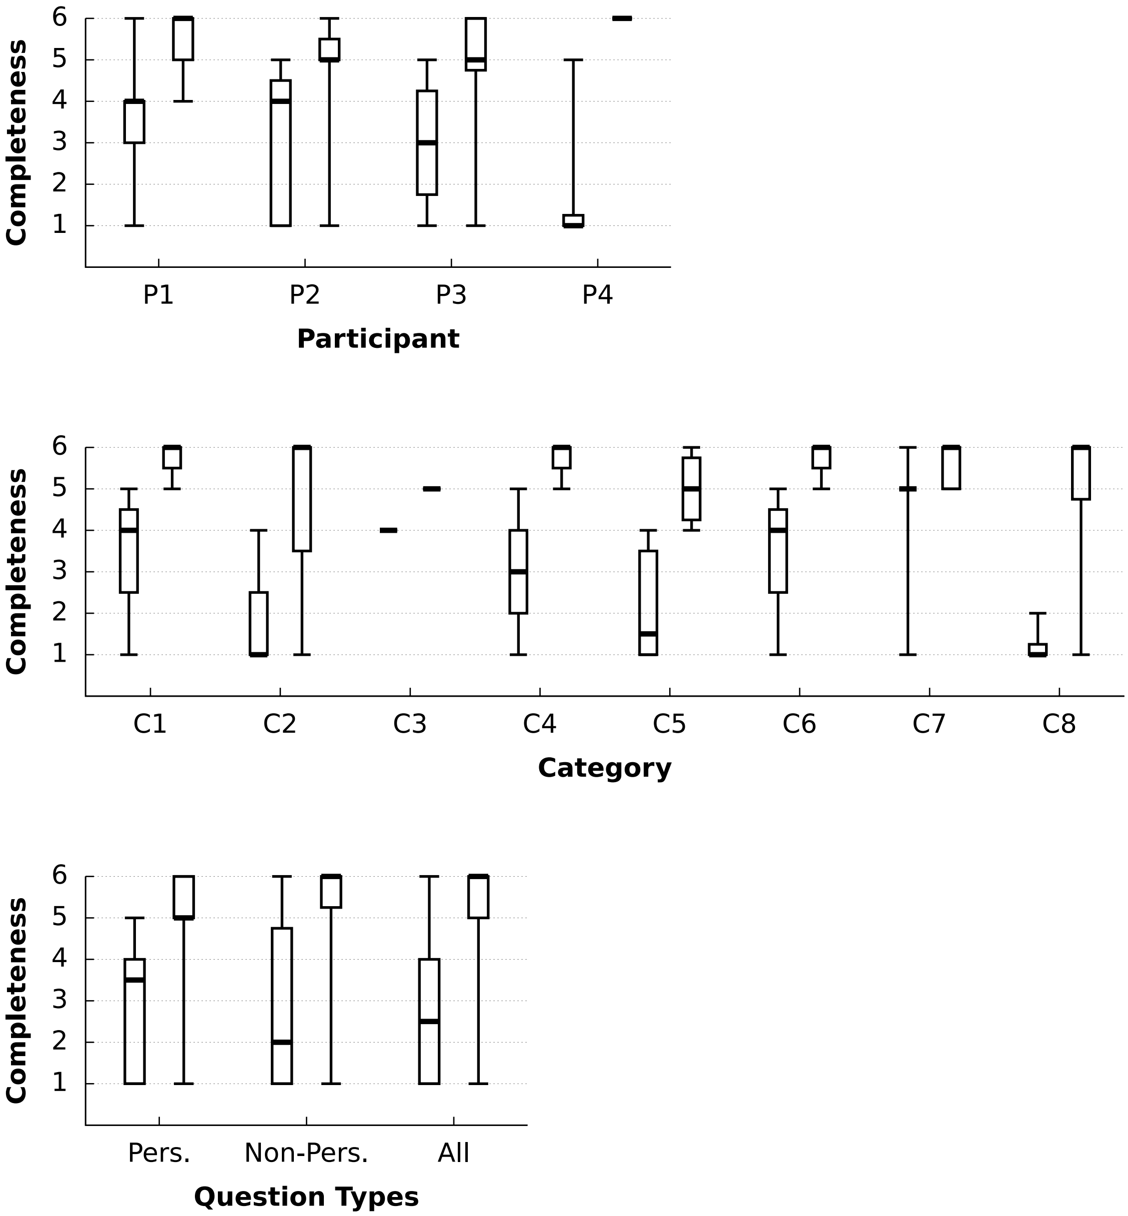

Figure 8.6. Completeness ratings before and after using the Digital Parrot, by participant (top), category (middle) and question type (bottom)

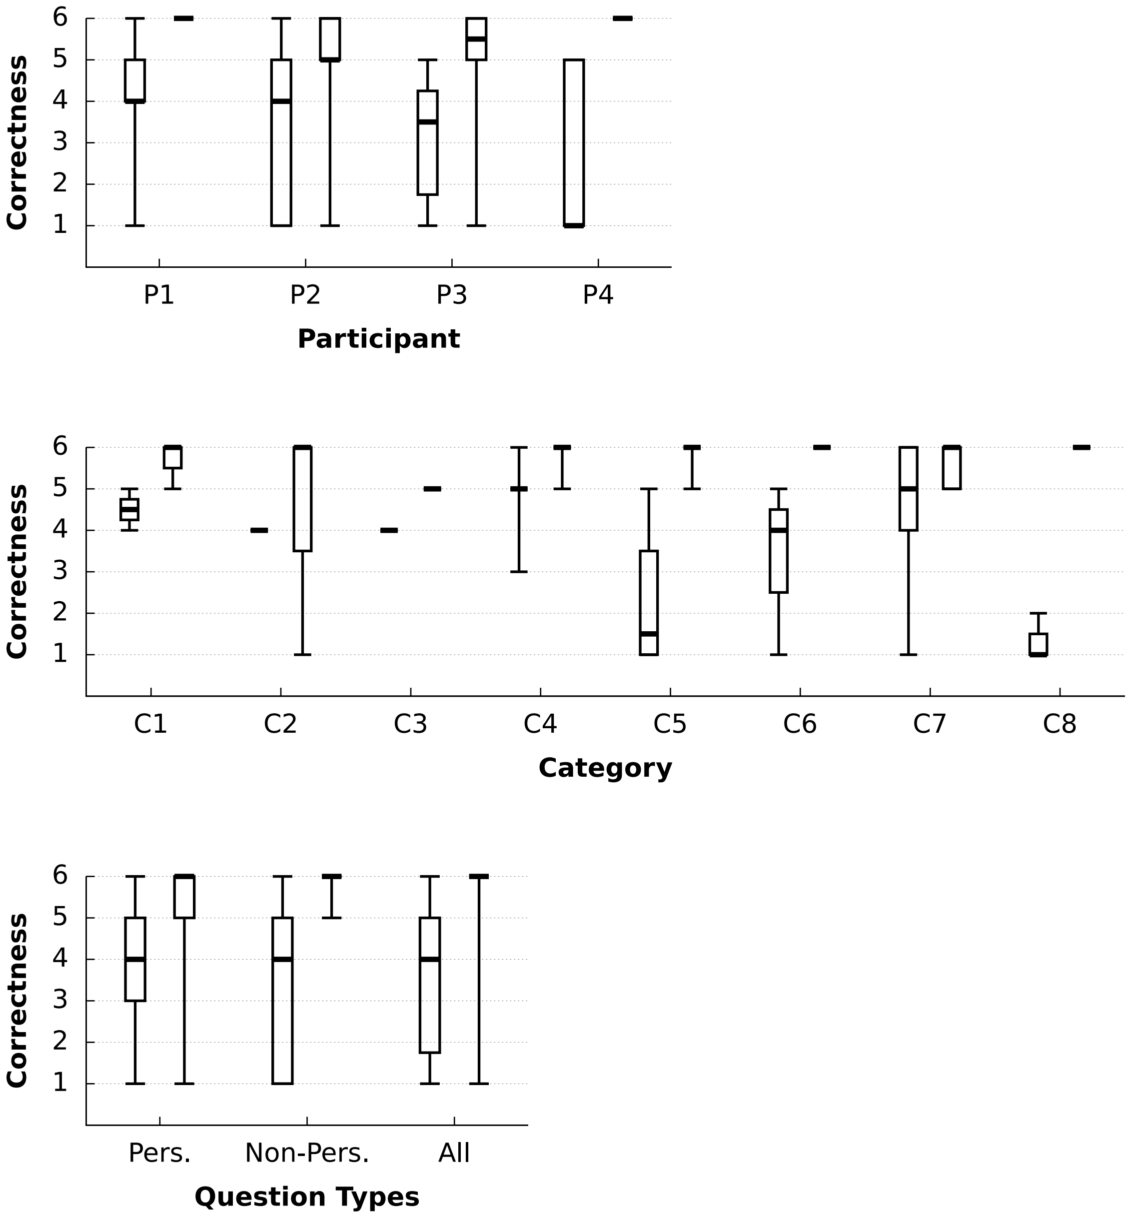

Figure 8.7. Correctness ratings before and after using the Digital Parrot, by participant (top), category (middle) and question type (bottom)

		Completeness		Correctness		
Unaided	**Mod.?**	before	after	before	after	**Total**
Full answer	n/a	5 [0]	6 [0]	5 [0.75]	6 [0]	6
Partial answer	Yes	4 [0]	6 [0.5]	5 [0.25]	6 [0]	4
	No	4 [2]	5 [1]	4 [2.5]	6 [1]	11
No answer	Yes	1 [–]	6 [–]	2 [–]	6 [–]	1
	No	1 [0]	5.5 [1]	1 [0]	6 [0.75]	10
Overall		2.5 [3]	6 [1]	4 [4]	5 [2.25]	**32**

Table 8.14. Median/IQR completeness and correctness ratings of answers before and after using the Digital Parrot.

Unaided	**Modified?**	**before DP**	**with DP**	**Total**
Full answer	n/a	5	6	6
Partial answer	Yes	1	3	4
	No	0	10	11
No answer	Yes	0	1	1
	No	1	8	10
Total		8	28	**32**

Table 8.15. Number of satisfactory answers before and after using the Digital Parrot

be seen for questions that had not been answered at all during the unaided remembering stage. After using the Digital Parrot, most answers were rated as highly complete and absolutely correct, even though there was some dispersion, particularly for the correctness rating.

Table 8.15 gives the number of satisfactory answers, as judged by the participant, before and after using the Digital Parrot. The "before" column corresponds to the "modified" or, for non-modified answers, to the "initially" column in Table 8.5. The number of satisfactory answers more than tripled through use of the Digital Parrot, from 8 to 28 overall.

Only four questions in total could not be answered satisfactorily using the Digital Parrot. In one case, a mistake in the data file led to problems using the connections navigator and the participant gave up trying to find the answer. In the remaining three cases, the answer found using the Digital Parrot was considered incomplete or unreliable by the participant. Incidentally, even though the information in the data file related to each of these questions was based on the interview in the experiencing phase, guesses had been made while creating the file because some information had been omitted in the interview (i. e. the name of a person or the *complete* list of other people present at an event).

Differences across participants There was one unsatisfactory answer each for participants P1 and P2 and two for participant P3. Unsurprisingly, completeness and correctness ratings given by P2 and P3 are slightly lower on average than those given by the other participants and vary more.

Differences across question types There was one unsatisfactory answer in category C2, two in category C5 and one in category C8. Three of these answers were for personal questions and one for a non-personal question. Again this is reflected in the completeness and correctness ratings.

8.5 Observations and participants' comments

This section discusses observations made by the researcher during the sessions and comments made by the study participants.

8.5.1 Types of remembering

The participants' comments show that all forms of remembering, namely re-living, knowing and reconstructing, as well as guessing occurred at some point during each session.

Re-living Often, participants commented on answers they gave themselves and on answers they retrieved with the Digital Parrot in much greater detail than was needed to answer the questions. A lot of these comments indicated some degree of re-living, often incorporating some kind of sensory element. Examples are descriptions of

- the participant's emotional response to the behaviour of someone else present at an event – "*name* was there, they annoyed me on that day";
- additional characteristics of the people whose names were part of the answer to the participant's question – "they looked kind of similar, a bit difficult to tell apart really" and
- additional characteristics of an event, such as the time an event happened when it was the place of the event that had been asked for – "this was before the dinner, a lot of the conference attendees were standing around and taking photos".

Knowing and recognising All six questions that were in the verification stage seemed to have been answered from knowledge. Four of these asked about the place of an event (C7), typically the place of the conference attended. One question asked for a conversation partner (C1) and one for the time of an event (C4). Even though not all of these answers were complete, those parts that were stated by the participants were correct. There were other questions where the participant apparently was certain about the answer without being able to state it: "I know the dinner was at *place* – just what was the name of *place*?"

A counterpart to knowledge is recognition, where someone initially does not know the answer (or not the full answer), but recognises it as correct once they see it. Some of the participants' comments show straightforward recognition: "That's right, I had forgotten that this happened in that year". Others shore more indirect types of recognition: "I knew it was a name with an unusual spelling".

Reconstructing Participants showed several types of reconstructive reasoning. One was to make inferences about the participant's typical behaviour: "I would have done a session with *other person*"; "I always sit with *other person* at the dinner when we're both at the same conference"; "it would have been at the end of *month* – I typically go to conferences during teaching recess". Another was to make inferences based on remembered emotional state: "It would have been towards the end of the conference, after my own talk when I was relaxed and open to go to talks outside my own field."

Interesting was one participant's comment when trying to use the Digital Parrot to retrieve which other people were present at an event: "I have a vague feeling that someone else was there too". This comment shows that meta-cognitive processes do not have to be conscious and rational but can also manifest themselves in feelings.

Guessing Seven partial answers were rated on the negative half of the correctness scale, suggesting that the participants guessed these answers. These answers were spread out evenly across participants, with one or two such answers per participant.

8.5.2 Issues with public information

Several participants expressed awareness that public and semi-public information sources may not stay available forever: "this will probably be difficult in three years or so – who knows if the website is still up then"; "I suspect this information won't be available for much longer".

Another related issue was identified by one participant in relation to searching their e-mail collection to find out who attended an event: "Of course this would only tell me who *said* they'd go, but not who actually came along." Information about the way in which an event was planned to take place is not necessarily a reliable way to determine what actually took place. The same holds true for determining which of the scheduled events were actually attended. Relying on unaided memory may fail here. For example, one participant commented that they didn't think they went to a keynote mentioned in one of their questions. However, the same participant described quite clearly in their interview that they had attended this keynote.

8.5.3 Double-checking of answers

Most participants did not perform any double-checking of their answers. This is particularly notable for questions that required the participant to distinguish between events based on time (C8), but also for a question about other people at an event (C5) that again involved time as a distinguishing element. In these cases, initial answers were often accompanied by comments such as “this would have been on the *weekday*” or “this must have been on the *weekday*”, indicating the presence of some sort of reasoning process that convinced the participant of the correctness of their answer.

One possible explanation is that the participants made assumptions about the data in the Digital Parrot and about the question in the study based on what they remembered telling the researcher during the initial interviews.

8.5.4 Trust in stored information

Participants generally seemed to trust the information in the Digital Parrot, accepting answers found with the Digital Parrot even when they had not been able to answer the question at all before. Some of these answers were clearly recognised, as could be seen from comments such as “That’s right, now I remember this” and “True, I had forgotten that this happened during that conference”. One participant commented that the information in the Digital Parrot was more convincing because it was given in writing, “and I trust things more when they are written down”. Another participant commented that trust in the Digital Parrot was built up when information spotted in the interface cued re-living of experiences and remembered information about the experience matched related information in the Digital Parrot.

During the discussion, participants were reminded of the fact that all information in the Digital Parrot had been entered by the researcher. They confirmed that this did not change their trust in the information – “there are enough details in here that match what I remember to make me confident that you would have got it all right”. In fact, those times where the participant could not find a satisfactory answer using the Digital Parrot, they were quite certain that the answer was incomplete or incorrect.

8.5.5 Task design

During the discussion, the participants were asked how realistic they thought the questions were. Generally, participants indicated that the questions were very realistic. It was not explicitly asked whether the participant had tried find the answer to any of the questions prior to the study. However, some participants explained that some of the questions asked for information that they had initially thought would be important but then later turned out not to be. Most of the time this was information about other people; typically, the participants had spoken to someone to explore the possibility of collaborative research but nothing had come of it in the end: "this could have been important, but then things didn't turn out that way".

Contrary to expectations, questions about the presence of other people at an event were seen as quite important to answer correctly: "I could offend someone if I got this wrong".

8.5.6 Vocabulary

Several participants struggled with aspects of the Conference ontology used in describing the data used in the study. One participant commented that the distinction made by the ontology between sessions and presentations does not match their mental model of these terms. In fact, these terms were used interchangeably at the conference that this participant was interviewed about.

Other participants struggled with the (structurally identical) issue that the ontology distinguishes between a keynote session and the keynote talk held during the session and likewise between a workshop session and a workshop held during the session. This was evident both in the approaches taken by the participants when using the connections manager and also in the participants' expectations of items that should have temporal information available in tooltips – often, temporal information was only available for the enclosing session but not for the individual presentation, workshop or keynote talk.

8.5.7 Suggestions for improvement

Some participants suggested a number of improvements for the user interface, almost all of which were related to the graph view. One participant suggested using different font sizes to distinguish important items (without specifying which types of items would be particularly important). After the session, they shared a "memory map" they had created, which makes use of spatial arrangements of text in different font sizes to capture memories and structure information related to the participant's field of research.

Another suggestion, made by two of the participants, was to include images in the Digital Parrot. People in particular should be associated with photographs, if available. One of the participants also commented that it would be useful to see photographs in their temporal context, i. e. surrounded by photographs taken within a small time interval of the currently viewed photograph. The two participants who made the suggestion both stated that they often take photographs. The other two participants did not mention photographs at all; in fact, one of them even stated that they deliberately did not take any photographs during the conference.

A third suggestion was made on how to possibly improve filtering the graph view. The participant suggested to investigate spatial and multi-criterial filter methods for data organisation, giving the "Dust & magnet" method by Yi et al. (2005) as an example.

8.6 Discussion

The main goal of this study was introduced in Section 8.1.1 as to determine whether the Digital Parrot and the approach taken in this thesis to augmenting autobiographical memory actually helps individuals to know about and reconstruct past experiences and related information. Another goal was to determine how the Digital Parrot's components support users in remembering. The goals were refined into six questions. This section relates the findings to the goals by giving answers to each of the questions.

8.6.1 Established strategies

Examining strategies that the participants have established to remember past experiences and related facts associated with attendance of academic conferences gives a baseline with which augmented autobiographical memory systems can be compared. As described in Section 8.4.2, participants reported a wide range of established strategies to remember such information. Many of the strategies use information sources that are outside the participant's control or that do not directly support the information needs expressed in the questions used in this study.

The participants reported that many of the questions would be answered based on publicly available information such as the conference website and published proceedings. This was the case even for a few questions with personal content, for example to determine the name of a conversation partner. Additionally, questions with personal content would be answered based on e-mail and photographs as well as by asking others. The participants also expressed awareness of the transience of publicly available information such as conference websites and of the fact that published general information does not always reflect their personal experiences (Section 8.5.2).

Overall, the findings described in Section 8.4.2 show that participants were confident that they would find answers to most of the questions asked. They also expected to be able to find answers within a reasonable time span. However, questions with personal content generally took longer to answer and participants acknowledged that some of these questions may very well never be answered at all.

8.6.2 Overall effectiveness of approach

The findings reported in Section 8.4.3 show that the Digital Parrot is generally effective in allowing the retrieval of information about and facts associated with personal experiences. The Digital Parrot allowed the participants to answer questions that they considered very unlikely to be able to answer otherwise, or only with significant time effort. Typically, these questions were concerned with personal information, such as which other people were present at an experience. In addition, the Digital Parrot allowed the participants to retrieve information that they would normally be able to find through other means. Together, this shows that the Digital Parrot is an improvement of the status quo.

As explained earlier in this thesis, augmented memory systems cannot hope to store an experience itself. Rather, they can provide cues that trigger remembering in the user of an augmented memory system. Section 2.2.1 introduced the distinction between re-living, knowing, reconstructing and guessing. For the purposes of the Digital Parrot, it would actually be sufficient if the user were able to know or at least reconstruct past experiences and related information. However, participants in the study described in this chapter reported that while this both was possible, reading the information in the system actually triggered several occasions of re-living (Section 8.5.1). This is particularly interesting in light of related research that has found nonimaginal data to be less likely to trigger re-living (see Section 3.2.2).

8.6.3 Influence of contextual information

The contextual navigators (map and timeline) did not see much use during the study (Section 8.4.3). One explanation given by the participants when asked about this was that “it was all in one place anyway” – they speculated that they would have made more use of temporal and spatial filtering if the system had contained information from more than one timespan or more than one location.

Another explanation is that the map and timeline in their current implementation are not helpful in retrieving memories. The addition of place names to the map, as suggested by participants in the usability study, appeared to be helpful those few times that the map was used in this study; the list of placenames was used to center the map on a particular place every time that the map was used by a participant. However, this did not lead to an increased use of the map navigator.

The fact that one third of the time-related questions was answered using the timeline alone suggests that temporal visualisation of data may need to be made more prominent. An example could be a calendar view or a visualisation of more personal timespans such as the lifetime periods and events associated with autobiographical memory (see Section 2.1.2).

A more far-fetched explanation for the very limited use of temporal navigation is that people typically do not actively seek to restore as much of context as possible of the experience to be remembered. The encoding specificity principle, as introduced in Section 2.2.2 does say that shared context facilitates remembering, but does not predict that people will seek out the context in order to remember. One possible exception to this is behaviour around action slips (where people may retrace their steps in order to remember why they went into a certain room, for example), which is not the type of failure that a system such as the Digital Parrot is targeting.

8.6.4 Influence of semantic information

Use of semantic information, i. e. of the connections navigator, appeared to be a matter of personal preference. Three of the four participants used the connections navigator at least once. One participant used the connections navigator for five of their seven questions, although not always successfully. The connections manager was used more often for personal than for non-personal questions.

The observed frequent changes between the connections navigator and the search function (Section 8.4.3) may indicate the participant's underlying difficulties in determining which of these two component to use for a given question. An integration of semantic information into the search function had already been suggested by participants in the usability study; this is supported by the high frequency with which the participants used search terms that were, or could have been, types.

Altogether this suggests that semantic information can support users in retrieving information from the Digital Parrot but that the current implementation still leaves room for improvement.

The connections manager's user interface, significantly changed based on comments made by participants in the usability study, still presents a number of challenges. Participants seemed to have difficulties understanding how the linear structure of the connections chain maps to the

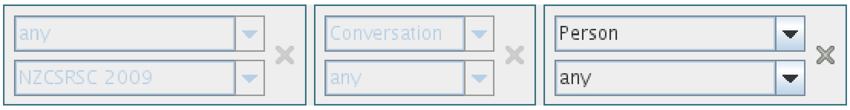

Figure 8.8. Chain example: people in conversations at a conference

resulting sub-graph.

Another problem is that the connections chain essentially needs to be read backwards. A connection chain to find people connected to conversations connected to a conference needs to be built up as "conference instance" – "Conversation type" – "Person type", as shown in Figure 8.8. Even though the same problem was described for an early version of the Feldspar system (Chau et al., 2008a), this issue in the Digital Parrot was discovered only part-way through the study. It was decided not to make modifications to the user interface for the remaining participants to ensure comparability of the findings.

Observations around the ontology and vocabulary (Section 8.5.6) underline how important it is that the ontology terms and structure match the user's language and mental model. The conceptual design underlying the Digital Parrot addresses this by suggesting a personalisable data model. However, the question remains open how well this works for users without a Computer Science background.

8.6.5 Influence of associations

The graph visualisation in the Digital Parrot directly shows the structure of memory items and associations between them, while the list view represents this structure in a more language-based way.

Overall, the participants in the study much preferred the graph visualisation over the list visualisation (Section 8.4.3). Even though all participants found the graph view initially overwhelming (Section 8.5.7), they appeared to become accustomed to it very quickly and preferred it over the list view for almost all questions. One of the participants commented that the graph view "is just like it is in [my] brain".

Participants seemed to like in particular the ability to follow chains of associations. One question was answered using no other part of the user interface. Even though directly connected items are visible in the list view as well, following several steps of associations is much more difficult than in the graph view. Another participant commented enthusiastically on the graph visualisation part of the Digital Parrot's user interface several days after the study.

Two further observations were made with regards to the visualisation type that are independent from the way associations are shown.

Overview versus detail One major advantage of the graph view over the list view is that the graph view makes it easier to switch between overview and detailed view (via zooming and scrolling). Participants often made use of this feature, in particular when determining whether they had seen all relevant results of a search.

However, the fact that zooming or scrolling was necessary to switch between levels of detail shows room for improvement. Scrolling in particular is rather laborious; no participant discovered the panning control and thus scrolling had to be done both horizontally and vertically. A different type of visualisation that integrates detail and overview modes may improve the user experience. Such a visualisation could employ a fish-eye lens, semantic zoom or other techniques that provide a "focus and context" view (Furnas, 2006).

Spatial position of items Going into the study, one of the assumptions was that the spatial position of items is important (see Section 5.3.2). This was one of the reasons for not following suggestions made by a participant in the usability study to always center the graph on the selected item or items (see Section 7.4.2). However, participants in the study described in the current chapter often moved items around when selecting them, which suggests that they did not feel that the position of the item was crucial. No participant was observed going back to a previously encountered item based on its spatial position.

However, it must be noted that the spatial positions of items in this study were randomly assigned. Behaviour may differ when participants continuously interact with their data and they determine the placement of items on the screen themselves.

8.6.6 Subjective experience

Participants' comments throughout the study were generally positive. Critical comments generally were focused on specific parts of the user interface rather than on the overall experience and approach (Section 8.5.7). Some participants expressed delight in remembering experiences from attending the conference about which they were asked in the study.

The participants generally trusted the information that was shown in the Digital Parrot (Sections 8.5.3 and 8.5.4). This suggests that the two-phase method used in this study did not have a negative impact on the study results.

8.7 Summary

The study described in this chapter was conducted to determine the effectiveness of the design and implementation of the Digital Parrot – whether it allows its users to answer questions about past experiences and information related to these experiences.

The participants in the study were able to use the Digital Parrot to answer such questions significantly more correctly and completely than from unaided memory. Especially for personal content, the Digital Parrot allowed the participants to answer questions that they judged to be difficult or impossible to answer using other means.

The study used naturalistic data derived from interviews with the participants about their experiences. The study itself was conducted approximately two years after the experiences. This study design allowed the examination of the Digital Parrot's effectiveness for cuing the participants' memories. Even though the Digital Parrot was developed to facilitate the simpler types of remembering, knowing about and reconstructing past experiences, it actually allowed the study participants to re-live some of the experiences as well.

The Digital Parrot's user interface components were developed based on recommendations derived from an examination of Cognitive Psychology research, namely contextual navigation, navigation based on semantic information and navigation based on associations between memory items. The uptake of contextual navigation was lower than expected. Navigation based on semantic information was also used less frequently than expected by all but one of the participants; observations suggest that the reasons lie more with the current realisation of the user interface related to this factor than with the factor itself. The visualisation of the information in the system that made associations explicit was much preferred over a representation that was based more strongly on language.

Chapter 9

Summary and Conclusions

This thesis has made the case for a new approach to augmenting autobiographical memory that is grounded in Cognitive Psychology research. Its objective was to help people remember past experiences and related information; it explored the central hypothesis:

> *An interactive software system that combines the context of experiences, semantic information and associations is a suitable means to support individuals in reconstructing autobiographical memories and in retrieving cues and facts related to such memories.*

Our research confirms this hypothesis and shows that the combination of these three factors is suitable to help people remember past experiences and related information and thus to achieve the objective of this thesis. Of these factors, connections between information items (mirroring associations between memory items) were found to be the most effective in accessing personal memories using an augmented memory system. Semantic information was also found to be effective, but open questions remain about its integration in the user interface of an augmented memory system. Contextual information was found to be effective in accessing personal memories using an augmented memory system, but not to the expected degree.

Section 9.1 gives a summary of this thesis and describes the steps that were undertaken to address the research objective and to answer the research questions derived from the objective. Section 9.2 outlines the contributions of this thesis to the field of augmenting autobiographical memory and Section 9.3 answers the research questions. Section 9.4 summarises the roles of the three core factors in the hypothesis by discussing the implications of the two user studies described in this thesis. Section 9.5 discusses the limitations of the work described in thesis. Finally, Section 9.6 gives directions for future work.

9.1 Summary

This thesis analysed research from Cognitive Psychology to derive recommendations and requirements for the design of augmented autobiographical memory systems (Chapter 2). The main outcomes are the statement of the problem in Cognitive Psychology terms and six recommendations for augmenting autobiographical memory, including the identification of three factors that are crucial to remembering autobiographical information: context, semantic information and associations.

An analysis of Computer Science approaches for augmenting autobiographical memory and from related areas revealed both strengths of such approaches that can be built upon as well as shortcomings of such approaches that need to be avoided (Chapter 3). Two main research areas were covered in the analysis which, between them, aim to augment autobiographical memory and make use of the three core factors. A gap was identified in the combination of these aspects, in particular with regards to remembering past experiences and related information.

A new conceptual design was put forward that takes all these considerations into account (Chapter 4). The conceptual design is based on the recommendations from Chapter 2. It bridges the gap identified in existing approaches related to augmenting autobiographical memory and incorporates the strengths of these approaches. It covers the entire lifecycle of memories.

An interactive system, the Digital Parrot, was implemented to realise the remembering aspects of the design (Chapter 5).

The design and implementation of the Digital Parrot were evaluated in two end-user studies. Challenges in evaluating augmented autobiographical systems (Chapter 6) were addressed with a new evaluation technique that is tailored to such systems. The first study used a traditional method to evaluate the Digital Parrot's usability (Chapter 7) and the second used the newly developed technique to evaluate the Digital Parrot's effectiveness (Chapter 8). Together, they show that the Digital Parrot, and more broadly the proposed conceptual design, fulfil their design goal of allowing users to answer questions about personal memories better than using unaided memory or established memory aids.

9.2 Contributions

This section summarises the contributions made by this thesis to the research area of augmenting autobiographical memory using interactive software systems.

Recommendations for augmented autobiographical memory systems

Our examination of research in the field of Cognitive Psychology contributes a thorough understanding and clear definition of the problem of augmenting autobiographical memory. This is important because approaches to related problems in the past have neglected this perspective. The examination yielded six recommendations for the design of augmented autobiographical memory systems. Three of the recommendations clarify choices to be made when designing such a system and three introduce factors that are helpful in remembering experiences: context of an experience, semantic information about items in the system and associations between memory items.

Synthesis of work in two related domains

Our analysis of existing Computer Science approaches enables future research in the area of augmenting autobiographical memory to build on the strengths and avoid the shortcomings of past efforts. The approaches selected for analysis are either designed for personal memories or are designed for related types of information and incorporate the three core factors context, semantic information and associations. Most approaches for personal memories suffer from a narrow focus on technical issues and, as a result, led to the development of interactive systems that do not in fact meet their goals. The analysed systems that incorporate context, semantic information and associations are designed for personal information that is already in digital form rather than for personal memories.

Conceptual design and requirements for an augmented autobiographical memory system

Our conceptual design of a system for supporting autobiographical memory applies the understanding of the problem and the analysis of related approaches to derive a new solution. The design semi-automatically captures context and content data to provide a scaffolding for the user's annotations. It is formalised as a set of requirements for the overall system, the user interface, the capture component and the underlying data model and storage component.

Prototype implementation of the design

Our implementation of the conceptual design, the Digital Parrot, further clarifies the design by giving a realisation of selected aspects. The focus in this implementation is on the retrieval aspect, following the focus of the research presented in this thesis. The Digital Parrot allows the design and its implementation to be evaluated.

Tailored evaluation method

Our evaluation method used for the second user study described in this thesis shows how to evaluate a system designed for personal memories. The method is a synthesis of existing approaches to evaluate systems in this field. It uses two strategies to alleviate the challenges associated with evaluating systems dealing with personal information: a two-phase approach that mirrors experiencing and remembering and a task-based design to allow findings to be compared across participants. These strategies can be used in evaluations of other systems designed for personal memories.

Two evaluations of the approach

Our two user studies evaluated the usability and effectiveness of the design and implementation of the Digital Parrot. They showed that our approach to augmenting autobiographical memory is successful. They also showed that associations, in the form of connections between information items, are useful in augmented autobiographical memory systems and that the roles of context and semantic information require and deserve further investigation.

9.3 Answers to the research questions

This section answers the research questions introduced in Section 1.1.2 that guided the research presented in this thesis. It discusses the implications of these answers for augmenting autobiographical memory.

What does it mean to help someone remember?

The seemingly simple objective of "helping someone remember" is actually quite complex – there are different types of memory and of remembering, as well as different remembering strategies and memory failures.

Memory for experiences and autobiographical facts differs from memory for other types of information. *Remembering* itself takes different shapes: re-living, knowing and reconstructing. People generally have a rich repertoire of *strategies* for attempting to recall and reconstruct information that they do not remember straight away. There are a number of distinct memory *failures*.

The first important insight from our examination of Cognitive Psychology research is the distinction between different types of memory. Even when short-term memory and memory more focused on sensomotoric sequences are disregarded, there are still differences *(a)* between autobiographical memory and memory without a self-link and *(b)* between semantic and episodic memory.

Another important insight is the distinction between different types of remembering; these types are sometimes called true remembering or re-living, knowing and reconstructing. This distinction is important because it is likely that different strategies are helpful in supporting each. This distinction also makes it clear that there are different types of memories that an augmented memory system might deal with. Software systems cannot externalise and store the actual experiences themselves, nor even vivid memories of experiences. Instead, software systems can merely store information that then may act as a cue to trigger re-living. However, an augmented memory system can externalise autobiographical facts that allow the user to know about or reconstruct their past experiences. The remainder of this thesis focuses on support for knowing and reconstructing, not on support for re-living.

In the survey of psychology research, three factors stand out as particularly helpful in cuing recall and enabling reconstructive processes respectively. The first factor is information about the *context* of a past experience, such as the place and time of an experience as well as other people who were present. The second factor are *types* – semantic information in the form of generalisations and abstractions about objects, events and actions that have been derived from past experiences. The third factor are *associations* between memory items, translated into connections between information items.

How can an interactive software system help someone remember?

First answer facet: A description of existing approaches

Information related to experiences, but not experiences themselves, can be captured automatically. This allows an augmented memory system to acquire information. Automatic capture by itself is known not to augment autobiographical memory. Gaps still exist in the knowledge of effective access to captured information and in its effective use to augment autobiographical memory.

There are examples of systems that use one or more of the three core factors (context, semantic information and associations) to organise and retrieve personal information that is already digital. These approaches are not tailored to real-life experiences and autobiographical facts.

Our research bridges a gap between approaches in the area of Capture, Archival and Retrieval of Personal Experiences (CARPE) and approaches in the area of Personal Information Management (PIM).

CARPE systems aim to provide an external repository of a person's experiences or of media files that act as cues for remembering experiences. A typical strategy in this field is to continuously capture life logs consisting of audio, photographs or video as well as context data such as time, the user's location, proximity of other people or the user's interaction with computing devices. Most approaches in this area focus on capturing data rather than on interaction with the system to actually remember past experiences. The efficacy of approaches based on automatic capture alone has recently come under question. Nevertheless, approaches in this area provide examples for methods that an augmented memory system can use to acquire information about its user's experiences.

Several sub-areas of PIM are related to the research presented in this thesis. The first comprises approaches that take into account the user's context. The second comprises approaches that address re-finding of information that was already known to the user at some earlier point. The third comprises approaches that make use of semantic information. PIM approaches take the three core factors and the retrieval aspect into account but are generally concerned with a different type of information than this thesis (i. e. electronic documents but not real-life experiences).

Second answer facet: The introduction of a new system design

A new answer, put forward in this thesis, is a combination of automatic capture and manual annotation together with sophisticated access mechanisms based on context and semantic information as well as a graph visualisation that makes associations between memory items explicit.

Our technique aims to support people in knowing about and reconstructing past experiences and related information. It does so by combining context information with semantic types and connections between information items. Context information and content are captured automatically and provide a scaffolding that the user can then annotate with their own information.

The technique focuses on situations that are semi-structured and potentially important to the user. Semi-structured situations make it more likely that context information and content can be captured automatically. If a situation is potentially important then it is more likely that the user is willing to invest the effort to add their own annotations to the system. Example types of situations are experiences made while travelling or while attending an academic conference.

Third answer facet: The implementation of the conceptual design

This thesis operationalises the conceptual design in the Digital Parrot, a prototype implementation of the retrieval side of an augmented autobiographical memory system. The Digital Parrot allows its users to view and retrieve information items related to past experiences. Navigation is possible based on the temporal and geospatial context of an experience, based on semantic information describing types of information items and of connections between items, as well as on associations. Textual search can also be used.

Information in the Digital Parrot is represented as a collection of statements which each consist of a subject, a predicate and an object. Subject and objects of statements are information items; predicates provide connections between information items. The structure forms a directed graph with information items as nodes and connections as edges. This information is directly visualised as a graph or as a simple list of statements. The graph visualisation makes associations explicit. The set of statements shown can be influenced with navigators based on context and semantic types as well as with textual search.

Time and location are the primary types of context used in the Digital Parrot. The timeline and the map navigator allow to find information items by time and location. They can also be used to narrow down the set of statements shown by excluding items that are outside a given time interval or outside a given geographical area.

The Digital Parrot also allows semantic information, in the form of classes for information items and connections, to be used to influence the set of statements shown to the user. The connections navigator lets the user build subgraph queries by specifying a chain of classes and instances. The set of statements shown is restricted to those nodes and edges that follow the structure of this chain.

How can we determine whether an interactive computer system helps someone remember?

Augmented autobiographical memory systems deal with personal memories. This introduces several complications when evaluating such systems. Time and its passage are important factors in relation to memories, as are rehearsal effects caused by repeated exposure to experiences. If personal memories are to be used in an evaluation, then naturally time must pass between the experience of an event and attempted recall of the event. Generally, evaluations of systems using personal information need to make a trade-off between the degree to which evaluation data and tasks are naturalistic and between comparability of findings.

A two-phase approach that first records events recently experienced by the participants and then derives evaluation data from these records allows the use of personal memories in the evaluation. A task-based design that defines categories of questions allows personalised tasks while ensuring that findings can be compared across participants.

Our evaluation method tailored to autobiographical memory systems consists of an experiencing phase and a recall phase. Experiences are captured in the first phase; guided open interviews and a transcription of experiences by the researcher allow capture without a dedicated system while still maintaining characteristics of the conceptual design. Experiences and related information are remembered in the second phase, both using the augmented autobiographical memory system and with unaided memory.

A task-based approach is used to derive questions. This ensures that comparisons can be made across participants even though each participant works on their own memories and receives tailored questions.

Does our approach help people remember?

Overall the approach taken in this thesis is effective. Participants in a user study were able to improve their answers to questions about two-year-old experiences and autobiographical facts using the Digital Parrot. Unaided answers received higher correctness and completeness ratings after verification using the Digital Parrot. Questions that had been answered only partially unaided and questions that had not been answered at all unaided were answered to the participants' satisfaction using the Digital Parrot.

The positive effect of the Digital Parrot is more pronounced the more personal the question is. While participants described a range of strategies to find unremembered answers elsewhere, many of these depend on the availability of information that is outside the participants' control and most fail completely for remembering experiences.

The graph visualisation was seen as useful by participants in a user study, clearly superior to the list visualisation of the same underlying data. Textual search and, to a lesser extent, navigation based on semantic types and worked well for most participants. Surprisingly, contextual navigation was not as effective as expected.

The goals of the studies were to determine the usability (first study) and effectiveness of the Digital Parrot and to investigate the roles of the Digital Parrot's components (second study).

The usability study found that participants could successfully use the Digital Parrot to answer questions about experiences in an example data set. Participants in the graph view condition found the system easier to use than those in the list view condition. Minor usability issues were discovered with both main views and with the map navigator. Major usability issues were discovered with the user interface for the navigation based on semantic types. The navigator based on semantic types was not used as often as expected; however, this was assumed to be due to the problems with its user interface. Both contextual navigators, the map and the timeline, were also not used as often as expected.

The second study found that participants could successfully use the Digital Parrot to answer questions about their own experiences. Answers verified or retrieved using the Digital Parrot were rated as more complete and correct than unaided answers. The graph view was preferred to the list view by all participants. Navigation based on semantic types was found useful by the participants, as was textual search. Again, both contextual navigators were used less often than expected, suggesting that the issues with these navigators are fundamental in nature and not just caused by details of the user interface or by the data used in the evaluations.

9.4 Implications

This section discusses the implications of the findings of this thesis for the area of augmenting autobiographical memories. It does so by discussing the combined findings of both evaluations conducted of the Digital Parrot with regards to recommendations for augmented autobiographical memory systems made in Section 2.4.3.

The first three of the six recommendations mainly serve to distinguish characteristics of the problem space. The remaining three recommendations describe factors that should be present in an augmented autobiographical memory system. These factors are the context of an experience, semantic information and associations between memory items. This section addresses each of the three factors in turn.

9.4.1 Context

Both evaluations of the Digital Parrot found that contextual navigation was not used as often for accessing information items in the system as recommendation R4 suggests. As discussed in Sections 7.6.2 and 8.6.3, this lack of uptake of the time-based and the location-based navigator in the Digital Parrot can be explained in several ways. The reason could be

1. the data used in the studies;

2. design decisions made in the implementation of the contextual navigators in the Digital Parrot; or
3. a more fundamental mismatch between expectations of user behaviour and actual user behaviour.

The first explanation was put forward by a participant in the second evaluation: there may not have been enough variation in the temporal and geospatial context of the experiences with which the evaluation was concerned. However, uptake of contextual navigation was just as low in the first evaluation, whose data showed a higher degree of variation in the context of experiences. This suggests that either an even higher degree of variation in context is required to make contextual navigation useful, or that this is not the reason for the observed behaviour.

Another explanation is that the lack of uptake of contextual navigation was caused by details in the user interfaces for the Digital Parrot's contextual navigation components. Based on observations made in the evaluations, several alternative treatments of context can be suggested. The representation of temporal and geospatial context may play an important role. The addition of placenames to the map navigator appeared to be helpful. Similar approaches to representing temporal context with semantic labels may improve the time navigator.

A third explanation is that people may not be aware, and thus not consciously take advantage, of overlaps between the context of an experience and their context when remembering the experience. However, such an overlap of context is known to help with remembering based on research results from Cognitive Psychology.

To summarise, it is unknown whether the low usage of contextual access to memory items in the evaluations of the Digital Parrot is caused by the data used in the studies, by the way in which such access was realised in the Digital Parrot or by more fundamental reasons. If the third explanation proves correct, this will have an impact far beyond the Digital Parrot and require a significant change of direction in the development of user interfaces to access personal memories.

9.4.2 Semantic information

The evaluations of the Digital Parrot show that semantic information is promising as a means to allow people to access their memories. However, both the initial and the current user interface of the Digital Parrot clearly do not make use of the full potential of semantic information.

In both evaluations, the Digital Parrot's interface component that allows access to information items using semantic information was used by only a small proportion of the participants. These participants appeared to find it easy to build up a mental model of the component and were mainly successful to use it for answering memory-related questions. In contrast, the majority of participants in both evaluations appeared to find the component confusing. The findings of both studies suggest that a stronger integration of semantic information with textual search may help overcome these issues.

Further research is necessary to determine the true influence of semantic information in augmenting autobiographical memory and to find suitable user interface and interaction techniques for such information. Our analysis of existing systems in Chapter 3 showed that semantic information for augmenting autobiographical memory is the least explored factor to date. Our findings, particularly of the second study, show that this factor has potential and thus that further research in this area is merited.

9.4.3 Associations

Both evaluations of the Digital Parrot showed that a straightforward graph visualisation of associations between memory items was seen as helpful by the participants. Those participants in the first study who used an alternative visualisation that was more strongly based on language appeared to find the Digital Parrot somewhat harder to use; when shown the graph visualisation, these participants expressed their preference for the graph visualisation.

A stronger effect was seen in the second study, in which the participants could choose freely which visualisation to use at any stage of their interaction with the Digital Parrot. All participants chose to use the graph visualisation most of the time.

One effect of the direct representation of associations was that selecting an item in the graph also caused its neighbours to be shown more prominently. This was especially the case in the version of the Digital Parrot used in the second study. Reading the neighbours of an information item allowed the participants to confirm that they had found the correct item; it provided item-level context and thus helped to transform the task from one of retrieval to one of recognition.

Most participants in the second study perceived the graph visualisation chosen as a suitable representation of the structure of their memories. However, all participants in both studies have a Computer Science background and are familiar with graphs and their visualisations. Such a visualisation may not be as suitable for users without a background in Computer Science or similar areas.

Our findings suggest that associations – item-level context – are helpful in augmenting autobiographical memory. Consequently, even if further research should show that the graph visualisation is not suitable for users with backgrounds different from that of the participants in our studies, alternative user interface and interaction techniques should be developed that make use of associations.

9.5 Limitations

We made choices of direction in the research presented in this thesis that led to a number of limitations.

Design choices. Recommendation R1 lists three types of remembering that an augmented autobiographical memory system can support: reliving, knowing and reconstructing. We chose to focus exclusively on support for knowing and reconstructing throughout the research presented in the thesis.

Recommendation R3 lists three stages of the remembering process at which an augmented autobiographical memory system can support its users memory: while memories of an experience are formed, when the user attempts to remember an experience, and in-between even when the user is not explicitly trying to remember the experience. Our analysis of related approaches includes systems for all three stages; our conceptual architecture for augmented autobiographical memory systems address all three stages but focuses on the former two. The Digital Parrot and both evaluations focus exclusively on support when the user attempts to remember an experience.

Implementation choices. Our conceptual design of an augmented autobiographical memory system includes support for a variety of information representations including text, photographs, audio and video. Our implementation of the Digital Parrot represents all information in the system in textual form; consequently, only this type of representation was included in the evaluations.

Early designs of the Digital Parrot's user interface envisaged a strong integration of the visualisation of an information item with the visualisation of its context. It also envisaged a seamless interaction between viewing the context of items and using context to control which items were being shown. In implementing the Digital Parrot, we deviated from these early designs for reasons of scope; however, evaluation findings suggest that an implementation closer to the original designs may better support users of the system.

Evaluation method. The study described in Chapter 8 investigates remembering experiences and related information two years after the original experience. Even though it uses a longer timespan between experience and recall than most other studies with similar goals, this study is not truly longitudinal because it does not take repeated samples over time. The gap between the time the experience was made and the time it was captured was relatively long (one to three months). The memory data used in both studies is naturalistic but the studies were still laboratory-based.

The study described in Chapter 8 uses a task-based design to enable comparison of findings across participants. This design draws on that introduced by Elsweiler and Ruthven (2007). However, while the task categories employed by Elsweiler and Ruthven were derived from a diary study and other qualitative investigations into people's e-mail behaviour, the task categories in the study described in Chapter 8 were chosen by the researchers in part based from the data available and in part based on personal experience with the domain.

The evaluations were limited in the number and choice of participants and in the choice of domain. In particular the second study had a small number of participants; all participants in both studies had a Computer Science background. Both studies focused exclusively on the domain of academic conferences.

9.6 Future Work

The work presented in this thesis is a significant contribution in itself. However, there are a number of promising points that could be pursued to extend our work further and to address its limitations. These points fall into three major areas:

- following up on observations made in the user studies;
- conducting further user studies to generalise the findings discussed in the previous section; and
- exploring areas of the conceptual design that were not pursued in the Digital Parrot.

The remainder of this section outlines each area in turn.

9.6.1 Following up on observations made in user studies

This section discusses possible extensions of the work presented in this thesis that follow up on observations made in both evaluations of the Digital Parrot, as well as on suggestions for improvement made by participants in these studies.

Exploring the use of context

The observed lack of use of contextual navigation in the studies may have been caused by a wide range of reasons (see Section 9.4.1). To date it is unknown which, if any, of the possible explanations we give are correct. Further user studies, involving data with greater temporal and spatial dispersion, will be a first step to confirm or rule out the first possible explanation. Modification of the Digital Parrot, followed by further user studies, will lead to confirmation or rejection of the second and third explanation.

Modifications to Digital Parrot's user interface can start with changes to the contextual navigators. The addition of a list of placenames to the map navigator appeared to have a positive impact on the use of the map navigator in the second study (see Section 8.6.3). The timeline could be modified in a similar way, for example using temporal landmarks (Ringel et al., 2003), activities as timespans (Krishnan and Jones, 2005) or hierarchies of temporal representations (André et al., 2007) of personal and public events.

Besides modifications made to the contextual navigators, the interplay between the Digital Parrot's main view and the contextual navigators can also be modified. A promising direction is to follow experiential principles (Jain, 2003) as used in early designs for the Digital Parrot. This type of user interface will allow the user to see the temporal and spatial context of items in the main view much more easily than is currently the case, which may have the same benefits as the immediacy with which directly connected items are currently visible in the graph view. Such modifications of the Digital Parrot's user interface will require the resolution of conceptual design issues (see Section 5.3.4).

Both options that involve further software development will also need to involve further user studies to determine the effectiveness of the modifications made.

Exploring the use of semantic information

The evaluation results indicate that while semantic information is useful to remember experiences and related information, the current version of the Digital Parrot's user interface does not use the full potential offered by semantic information. Issues with the current interface for semantic filtering most likely indicate a significant mismatch between the implementation and users' mental models (see Section 8.6.4).

To gain further insights into the role of semantic information in augmenting autobiographical memory, alternative user interfaces for semantic filtering can be explored. One avenue to pursue is to integrate semantic filtering more strongly with textual search. Another option is to investigate further transfer of user interface elements from systems in the area of semantic PIM (e. g. Karger et al., 2005; Chau et al., 2008a).

Exploring the use of associations

The evaluation results indicate that a direct visualisation of the connections between information items was seen as beneficial, if initially overwhelming. Insights from existing research into visualisation method for large graphs (e. g. Wattenberg, 2006; Ham and Wattenberg, 2008) can guide the development of visualisation methods that retain the benefits of showing an information item's in-system context while scaling up to the amounts of data that is likely to be collected during everyday use of an augmented memory system over a long time period.

9.6.2 Conducting further user studies

This section describes which further user studies will be helpful to corroborate the findings of our user studies.

Different study designs

A truly longitudinal study, in which measurements of the participants' use of the Digital Parrot are taken repeatedly over time, will further validate the findings of our study. A longitudinal study design can take into account both learning times associated with a new type of user interface and rehearsal aspects from interaction with memory data during the study period. It will allow for the observation of more naturalistic interaction with the system.

Such a truly longitudinal study will be particularly valuable if it is conducted in a naturalistic setting rather than in the controlled environment of a laboratory. Challenges that will need to be overcome are: the difficulty to balance participants' expectations about the stability of software with the realities of developing research prototypes; the difficulty to recruit participants for studies involving long timeframes and personal information; and the difficulty to genuinely evaluate systems that require participants to deviate from their established ways to manage personal information (see Section 6.1.1).

A longitudinal and naturalistic study design will allow for the exploration of alternative types of measures, for example evaluating the success of a system to support someone's memory by determining "how other people perceive[d] the [user's] memory" (Vemuri, 2004, p. 117).

Confirm information needs

An investigation into people's actual, rather than supposed, information needs around remembering experiences from conferences will validate the choice of tasks in our studies. It will enable further evaluations using tasks based on the results of such an investigation.

9.6.3 Further implementation of the conceptual design

This section outlines parts of the conceptual design that were left largely unexplored in this thesis and that can be implemented and evaluated further.

Experiencing and revising phases

Of the three interaction phases covered by our conceptual design, only the remembering phase has been the focus of the current version of the Digital Parrot and its evaluations conducted to date. Assumptions were made about the availability of information for automatic capture and the willingness and ability of users to revise and add to the information in the system.

To validate these assumptions, the work presented in this thesis is currently being extended (Allah, 2011). Central topics in this extension are automatic capture of information, mechanisms for marking a particular moment for later annotation and the types of automatically captured information that are most beneficial in the revision phase.

An interesting direction to pursue is the integration of the Digital Parrot with recent approaches that automatically add semantic annotations to the types of data captured in CARPE systems (Byrne et al., 2009).

Other aspects that can be explored in this context are the integration of the Digital Parrot with mobile devices and with traditional PIM tools, as well as the support of information representations other than text in the Digital Parrot's user interface.

Finally, an implementation of the experiencing and revising aspects of the conceptual design will allow a full evaluation of the design. This will be particularly valuable if combined with the truly longitudinal approach outlined in the previous section.

Other application domains

Both the conceptual design introduced in this thesis and the Digital Parrot make no assumptions about the application domain other than targeting situations that *(a)* are semi-structured and *(b)* the user of such a system is likely to wish to remember later (see Section 4.1.1). However, the empirical part of this thesis has focused exclusively on the domain of academic conferences, for example in the assumptions made about the types of information available for automatic capture and the information needs related to this domain.

Examples for similar types of situations are visits to trade conventions, scientific fieldwork and travel. Events in these domains differ from academic conferences in the amount of publicly available information, in the degree of inherent structure and in the level of importance of spatial context.

Conducting further evaluations in different application domains will serve two purposes. It will confirm to which degree the findings in this thesis can be transferred to other domains. It will also determine whether the differences between domains lead to differences in interaction patterns, particularly in the use of contextual navigation. This will be particularly interesting if these further evaluations include participants from backgrounds other than Computer Science.

Appendix A

Data Model

This appendix lists the OWL ontologies used in the Digital Parrot and the RDF data used to produce the screenshots in Section 5.3. Refer to Section 5.6 for descriptions of the ontologies. All OWL and RDF is presented in Notation 3 format (Berners-Lee and Connolly, 2008).

Some of the ontologies in this appendix refer to classes and properties defined in the Friend-of-a-Friend (FOAF) project[1] and in the UMBEL ontology project[2].

A.1 Ontologies

Section A.1 lists the two ontologies that are explicitly referred to in the Digital Parrot's code. The Context ontology defines vocabulary to add temporal and geospatial context to information items. The Digital Parrot ontology provides the means to specify which classes and properties are shown in Digital Parrot's user interface.

A.1.1 Context ontology

```
@prefix : <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
   TimeAndPlace/2008/11/TimeAndPlace.owl#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
```

[1] http://www.foaf-project.org/, vocabulary at http://www.mindswap.org/2003/owl/foaf

[2] http://umbel.org/, vocabulary at http://umbel.org/umbel/sc/

```
<http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/TimeAndPlace
   /2008/11/TimeAndPlace.owl> a owl:Ontology;
  rdfs:comment "Vocabulary for describing things in relation to
     time and place."@en .

:AbsolutelyPlacedThing     a owl:Class;
  rdfs:comment "Something that is anchored in space (via geo
     coordinates)."@en;
  rdfs:label "Absolutely Placed Thing"@en;
  owl:equivalentClass
    [ a owl:Class ;
      owl:intersectionOf
        (
           [ a owl:Restriction ;
             owl:onProperty :lat ;
             owl:minCardinality "1"^^xsd:nonNegativeInteger ]
           [ a owl:Restriction;
             owl:onProperty :long ;
             owl:minCardinality "1"^^xsd:nonNegativeInteger ]
         )
       ];
     rdfs:subClassOf :PlacedThing .

:AbsolutelyTimedThing a owl:Class;
  rdfs:comment "Something that is anchored in time (via timestamps)
     ."@en;
  rdfs:label "Absolutely Timed Thing"@en;
  owl:equivalentClass
    [ a owl:Class;
      owl:intersectionOf
        (
           [ a owl:Restriction;
             owl:onProperty :startsAt ;
             owl:minCardinality "1"^^xsd:nonNegativeInteger ]
           [ a owl:Restriction;
             owl:onProperty :endsAt ;
             owl:minCardinality "1"^^xsd:nonNegativeInteger ]
         )
       ];
  rdfs:subClassOf :TimedThing .

:PlacedThing a owl:Class;
  rdfs:comment "Something that describes a place."@en;
  rdfs:label "Placed Thing"@en .
```

```
:TimedThing a owl:Class;
  rdfs:comment "Something that is anchored in time."@en;
  rdfs:label "Timed Thing"@en .

:CoordPrecisionValueType a owl:Class;
  rdfs:comment "Value partition type for the coordinatesPrecision
     property. Describes the precision of a PlacedThing’s
     coordinates. So eg Rome might be described by a set of lat/
     long coordinates and CityPrecision precision."@en;
  rdfs:label "Precision"@en;
  owl:equivalentClass
  [ a owl:Class;
    owl:oneOf
      ( :RoomPrecision :BuildingPrecision :BlockPrecision
         :SuburbPrecision :CityPrecision :SmallCountryPrecision
         :MediumCountryPrecision :LargeCountryPrecision
         :ContinentPrecision)
  ] .

:RoomPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a room"@en;
  rdfs:comment "A precision of the size of an average room. Up to a
      few metres."@en .

:BuildingPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a building"@en;
  rdfs:comment "A precision of the size of an average building. Up
     to a few tens of metres."@en .

:BlockPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a block of buildings on a street"
     @en;
  rdfs:comment "A precision of the size of an average street block.
      Up to a few hundreds of metres."@en .

:SuburbPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a suburb"@en;
  rdfs:comment "A precision of the size of an average suburb. Up to
      a few kilometres."@en .

:CityPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a city"@en;
  rdfs:comment "A precision of the size of an average city. Up to a
      few tens of kilometres."@en .
```

```
:SmallCountryPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a small country"@en;
  rdfs:comment "A precision of the size of a small country or
      district. Up to a few hundreds of kilometres."@en .

:MediumCountryPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a medium-sized country"@en;
  rdfs:comment "A precision of the size of a medium-sized country
      or state. Up to a few thousands of kilometres."@en .

:LargeCountryPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a large country"@en;
  rdfs:comment "A precision of the size of a large country. Up to
      ca 5,000 kilometres."@en .

:ContinentPrecision a :CoordPrecisionValueType;
  rdfs:label "precision of about a continent"@en;
  rdfs:comment "A precision of the size of a continent. More than
      ca 5,000 kilometres."@en .

:after a owl:ObjectProperty, owl:TransitiveProperty;
  rdfs:domain :TimedThing;
  rdfs:range :TimedThing .

:before a owl:ObjectProperty, owl:TransitiveProperty;
  rdfs:domain :TimedThing;
  rdfs:range :TimedThing .

:encloses a owl:ObjectProperty, owl:TransitiveProperty;
  rdfs:domain :PlacedThing;
  rdfs:range :PlacedThing;
  owl:inverseOf :locatedIn .

:endsAt a owl:DatatypeProperty;
  rdfs:domain :TimedThing;
  rdfs:range xsd:dateTime;
  rdfs:label "ends at"@en .

:locatedIn a owl:ObjectProperty, owl:TransitiveProperty;
  rdfs:domain :PlacedThing;
  rdfs:range :PlacedThing;
  rdfs:label "is in"@en .

:during a owl:ObjectProperty,
  owl:TransitiveProperty;
```

```
  rdfs:domain :TimedThing;
  rdfs:range  :TimedThing;
  rdfs:label "is during"@en .

:spans a owl:ObjectProperty, owl:TransitiveProperty;
  rdfs:domain :TimedThing;
  rdfs:range :TimedThing;
  owl:inverseOf :during .

:startsAt a owl:DatatypeProperty;
  rdfs:domain :TimedThing;
  rdfs:range xsd:dateTime;
  rdfs:label "starts at"@en .

:lat a owl:DatatypeProperty;
  rdfs:domain :PlacedThing;
  rdfs:range xsd:string;
  rdfs:label "latitude"@en .

:long a owl:DatatypeProperty;
  rdfs:domain :PlacedThing;
  rdfs:range xsd:string;
  rdfs:label "longitude"@en .

:coordPrecision a owl:ObjectProperty;
  rdfs:domain :PlacedThing;
  rdfs:range :CoordPrecisionValueType;
  rdfs:label "precision of the coordinates"@en .
```

A.1.2 Digital Parrot Ontology

```
@prefix : <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
   DigitalParrot/2009/02/DigitalParrot.owl#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .

<http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/DigitalParrot
   /2009/02/DigitalParrot.owl> a owl:Ontology .

:showThisType    a owl:AnnotationProperty;
    rdfs:comment "types that are annotated with this property are
       shown in the Digital Parrot’s user interface. The value of
       this property must be one of the string literal ’primary’
       and ’secondary’, to indicate a primary or secondary type,
```

```
      respectively. Primary types may be shown more prominently
      than secondary types."@en .
```

A.2 Example data

This section contains example memory data that was used to produce the general user interface screenshots in Section 5.3. The same data was also used in the first user study (Chapter 7) and in the training phase of the remembering phase of the second user study (Chapter 8).

A.2.1 Conference memories

The data in this section describes the some of the researcher's memories related to attending five conferences. All places and times are based on the actual conference programs, which are currently, or were at the time of the conference, publicly available. All personal names that cannot be deduced from publicly available information have been altered for publication in this thesis. Real names, including full last names where known, were used in the studies.

```
@prefix : <file:res/data/nzcsrsc.n3#> .
@prefix conf: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms
   /Conferences/2008/11/Conferences.owl#> .
@prefix interact: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/
   Terms/Interaction/2008/11/Interaction.owl#> .
@prefix timeplace: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/
   Terms/TimeAndPlace/2008/11/TimeAndPlace.owl#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix parrot: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/
   Terms/DigitalParrot/2009/02/DigitalParrot.owl#> .

<file:res/data/nzcsrsc.n3> a owl:Ontology ;
  owl:imports <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms
     /Conferences/2008/11/Conferences.owl> ,
```

```
    <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
       Interaction/2008/11/Interaction.owl> ,
    <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/Memories
       /2008/11/Memories.owl> .

# Conferences

:NZCSRSC a conf:ConferenceSeries ;
  rdfs:label "NZ CS Research Student Conferences" .

:CHINZ a conf:ConferenceSeries ;
  rdfs:label "NZ CHI Conferences" .

:NZCSRSC07 a conf:Conference ;
  timeplace:locatedIn :WaikatoUni ;
  rdfs:label "NZCSRSC 2007" ;
  conf:partOfSeries :NZCSRSC ;
  timeplace:startsAt "2007-04-10T15:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2007-04-13T14:00:00+13:00"^^xsd:dateTime .

:CHINZ07 a conf:Conference ;
  timeplace:locatedIn :WaikatoUni ;
  rdfs:label "CHINZ 2007" ;
  conf:partOfSeries :CHINZ ;
  timeplace:startsAt "2007-07-03T09:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2007-07-04T15:00:00+13:00"^^xsd:dateTime .

:NZCSRSC08 a conf:Conference ;
  timeplace:locatedIn :CanterburyUni ;
  rdfs:label "NZCSRSC 2008" ;
  conf:partOfSeries :NZCSRSC ;
  timeplace:startsAt "2008-04-14T13:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-17T14:00:00+13:00"^^xsd:dateTime .

:NZCSRSC09 a conf:Conference ;
  timeplace:locatedIn :AucklandUni ;
```

```
  rdfs:label "NZCSRSC 2009" ;
  conf:partOfSeries :NZCSRSC ;
  timeplace:startsAt "2009-04-06T15:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-09T13:45:00+13:00"^^xsd:dateTime .

:CHINZ2009 a conf:Conference ;
  timeplace:locatedIn :AucklandUni ;
  rdfs:label "CHINZ 2009" ;
  conf:partOfSeries :CHINZ ;
  timeplace:startsAt "2009-07-06T10:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-07T18:00:00+13:00"^^xsd:dateTime .

:NZCSRSC10 a conf:Conference ;
  rdfs:label "NZCSRSC 2010" ;
  conf:partOfSeries :NZCSRSC .

# Places

:City rdfs:subClassOf timeplace:PlacedThing ;
      parrot:showThisType "primary"^^xsd:string .
:Campus rdfs:subClassOf timeplace:PlacedThing ;
      parrot:showThisType "primary"^^xsd:string .
:Building rdfs:subClassOf timeplace:PlacedThing ;
      parrot:showThisType "primary"^^xsd:string .
:Room rdfs:subClassOf timeplace:PlacedThing ;
      parrot:showThisType "primary"^^xsd:string .

:Auckland a :City ;
  rdfs:label "Auckland"@en ;
  timeplace:coordPrecision timeplace:CityPrecision ;
  timeplace:lat "-36.847384"^^xsd:string ;
  timeplace:long "174.765733"^^xsd:string .

:Hamilton a :City ;
  rdfs:label "Hamilton"@en ;
  timeplace:coordPrecision timeplace:CityPrecision ;
```

```
  timeplace:lat "-37.78752"^^xsd:string ;
  timeplace:long "175.282386"^^xsd:string .

:Christchurch a :City ;
  rdfs:label "Christchurch"@en ;
  timeplace:coordPrecision timeplace:CityPrecision ;
  timeplace:lat "-43.531432"^^xsd:string ;
  timeplace:long "172.636628"^^xsd:string .

:WaikatoUni a :Campus ;
  rdfs:label "The University of Waikato"@en ;
  timeplace:coordPrecision timeplace:BlockPrecision ;
  timeplace:lat "-37.788209"^^xsd:string;
  timeplace:long "175.318435"^^xsd:string ;
  timeplace:locatedIn :Hamilton .

:CanterburyUni a :Campus ;
  rdfs:label "Canterbury University"@en ;
  timeplace:coordPrecision timeplace:BlockPrecision ;
  timeplace:lat "-43.52"^^xsd:string ;
  timeplace:long "172.58"^^xsd:string ;
  timeplace:locatedIn :Christchurch .

:AucklandUni a :Campus ;
  rdfs:label "Auckland University"@en ;
  timeplace:coordPrecision timeplace:BlockPrecision ;
  timeplace:lat "-36.85167"^^xsd:string ;
  timeplace:long "174.769574"^^xsd:string ;
  timeplace:locatedIn :Auckland .

:OwenGlennBuilding a :Building ;
  rdfs:label "Owen G. Glenn Building"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision;
  timeplace:lat "-36.853208"^^xsd:string ;
  timeplace:long "174.771111"^^xsd:string;
  timeplace:locatedIn :AucklandUni .
```

```
:MaraeAuckland a :Building ;
  rdfs:label "Marae"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision ;
  timeplace:lat "-36.85146"^^xsd:string ;
  timeplace:long "174.772789"^^xsd:string ;
  timeplace:locatedIn :AucklandUni .

:FalePacifica a :Building ;
  rdfs:label "Fale Pacifica"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision ;
  timeplace:lat "-36.852233"^^xsd:string ;
  timeplace:long "174.772152"^^xsd:string ;
  timeplace:locatedIn :AucklandUni .

:WhareKai a :Building ;
  rdfs:label "Whare Kai"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision ;
  timeplace:lat "-36.851526"^^xsd:string ;
  timeplace:long "174.773069"^^xsd:string ;
  timeplace:locatedIn :MaraeAuckland .

:BishopJulius a :Building ;
  rdfs:label "Bishop Julius"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision ;
  timeplace:lat "-43.524344"^^xsd:string ;
  timeplace:long "172.573371"^^xsd:string ;
  timeplace:locatedIn :CanterburyUni .

:SpicersRestaurant a :Room ;
  rdfs:label "Spicers Restaurant"@en ;
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :OwenGlennBuilding .

:NgaioMarshTheatre a :Room ;
  rdfs:label "Ngaio Marsh Theatre"@en ;
```

```
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :UCSABuilding .

:C1LectureTheatre a :Room ;
  rdfs:label "C1 Lecture Theatre"@en ;
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :CLectureBuilding .

:C3LectureTheatre a :Room ;
  rdfs:label "C3 Lecture Theatre"@en ;
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :CLectureBuilding .

:CLectureBuilding a :Building ;
  rdfs:label "Lecture Theatre"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision ;
  timeplace:lat "-43.523128"^^xsd:string ;
  timeplace:long "172.583714"^^xsd:string ;
  timeplace:locatedIn :CanterburyUni .

:UCSABallRoom a :Room ;
  rdfs:label "UCSA Ball Room"@en ;
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :UCSABuilding .

:UCSAShelleyCommonRoom a :Room ;
  rdfs:label "UCSA Common Room (Shelley)"@en ;
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :UCSABuilding .

:UCSAUpperCommonRoom a :Room ;
  rdfs:label "UCSA Upper Common Room"@en ;
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :UCSABuilding .

:UCSAFoundry a :Room ;
```

```
  rdfs:label "UCSA Foundry"@en ;
  timeplace:coordPrecision timeplace:RoomPrecision ;
  timeplace:locatedIn :UCSABuilding .

:UCSABuilding a :Building ;
  rdfs:label "UCSA Building"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision ;
  timeplace:lat "-43.523932"^^xsd:string ;
  timeplace:long "172.578971"^^xsd:string ;
  timeplace:locatedIn :CanterburyUni .

:HITLab a :Room ;
  rdfs:label "HIT Lab"@en ;
  timeplace:coordPrecision timeplace:BuildingPrecision ;
  timeplace:lat "-43.522088"^^xsd:string ;
  timeplace:long "172.582823"^^xsd:string ;
  timeplace:locatedIn :CanterburyUni .

# Conference sessions

:RegistrationCHINZ09 a conf:SocialEvent ;
  rdfs:label "Registration"@en ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:during :CHINZ2009 ;
  timeplace:startsAt "2009-07-06T10:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T10:45:00+13:00"^^xsd:dateTime .

:WelcomeCHINZ09 a conf:SocialEvent ;
  rdfs:label "Welcome"@en ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:during :CHINZ2009 ;
  timeplace:startsAt "2009-07-06T10:45:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T11:00:00+13:00"^^xsd:dateTime .

:KeynoteSessionMondayCHINZ09 a conf:PresentationSession ;
  rdfs:label "Keynote Session"@en ;
```

```
  conf:sessionIn :CHINZ2009 ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:startsAt "2009-07-06T11:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T12:30:00+13:00"^^xsd:dateTime .

:KeynoteMondayCHINZ09 a conf:KeynoteTalk ;
  rdfs:label "Opening Keynote"@en ;
  conf:inSession :KeynoteSessionMondayCHINZ09 ;
  conf:hasPresenter :StephenBrewster .

:LunchMondayCHINZ09 a conf:SocialEvent ;
  rdfs:label "Lunch"@en ;
  timeplace:during :CHINZ2009 ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:startsAt "2009-07-06T12:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T13:30:00+13:00"^^xsd:dateTime .

:HealthAndEducationSessionCHINZ09 a conf:PresentationSession ;
  rdfs:label "Papers - Health and Education"@en ;
  conf:sessionIn :CHINZ2009 ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:startsAt "2009-07-06T11:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T15:00:00+13:00"^^xsd:dateTime .

:AfternoonTeaMondayCHINZ09 a conf:SocialEvent ;
  rdfs:label "Afternoon Tea"@en ;
  timeplace:during :CHINZ2009 ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:startsAt "2009-07-06T15:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T15:30:00+13:00"^^xsd:dateTime .

:UsabilitySessionCHINZ09 a conf:PresentationSession ;
  rdfs:label "Papers - Usability"@en ;
  conf:sessionIn :CHINZ2009 ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:startsAt "2009-07-06T15:30:00+13:00"^^xsd:dateTime ;
```

```
  timeplace:endsAt   "2009-07-06T17:00:00+13:00"^^xsd:dateTime .

:DrinksCHINZ09 a conf:SocialEvent ;
  rdfs:label "Drinks"@en ;
  timeplace:during :CHINZ2009 ;
  timeplace:locatedIn :SpicersRestaurant ;
  timeplace:startsAt "2009-07-06T17:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T18:00:00+13:00"^^xsd:dateTime .

:DinnerCHINZ09 a conf:SocialEvent ;
  rdfs:label "Dinner"@en ;
  timeplace:during :CHINZ2009 ;
  timeplace:locatedIn :SpicersRestaurant ;
  timeplace:startsAt "2009-07-06T18:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-07-06T20:00:00+13:00"^^xsd:dateTime .

# Sessions NZCSRSC 08

:CheckInNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Check in"@en ;
  timeplace:locatedIn :BishopJulius ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-14T13:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-14T14:45:00+13:00"^^xsd:dateTime .

:AssembleNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Assemble"@en ;
  timeplace:locatedIn :CanterburyUni ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-14T15:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-14T15:15:00+13:00"^^xsd:dateTime ;
  timeplace:lat "-43.523545"^^xsd:string ;
  timeplace:long "172.582319"^^xsd:string .

:OpeningNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Opening Ceremony"@en ;
```

```
  timeplace:during :NZCSRSC08 ;
  timeplace:locatedIn :C2LectureTheatre ;
  timeplace:startsAt "2008-04-14T15:15:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-14T15:45:00+13:00"^^xsd:dateTime .

:SupercomputerSessionNZCSRSC08 a conf:PresentationSession ;
  rdfs:label "What’s so super about the supercomputer?"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C2LectureTheatre ;
  timeplace:startsAt "2008-04-14T16:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-14T18:00:00+13:00"^^xsd:dateTime .

:OpeningDinnerNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Opening Dinner and after dinner presentation"@en ;
  timeplace:during :NZCSRSC08 ;
  timeplace:locatedIn :UCSABallRoom ;
  timeplace:startsAt "2008-04-14T18:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-14T21:00:00+13:00"^^xsd:dateTime .

:BreakfastTuesdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Breakfast"@en ;
  timeplace:during :NZCSRSC08 ;
  timeplace:locatedIn :BishopJulius ;
  timeplace:startsAt "2008-04-15T07:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T08:30:00+13:00"^^xsd:dateTime .

:Session1NZCSRSC08 a conf:PresentationSession ;
  rdfs:label "Session 1"@en ;
  conf:sessionIn :NZCSRSC08;
  timeplace:locatedIn :C3LectureTheatre ;
  timeplace:startsAt "2008-04-15T09:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T11:10:00+13:00"^^xsd:dateTime .

:MorningTeaTuesdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Morning tea"@en ;
  timeplace:during :NZCSRSC08 ;
```

```
  timeplace:locatedIn :CBlock ;
  timeplace:startsAt "2008-04-15T11:10:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T11:30:00+13:00"^^xsd:dateTime .

:KeynoteTuesdayNZCSRSC08 a conf:KeynoteTalk ;
  rdfs:label "Keynote Edwin Dando"@en ;
  conf:inSession :KeynoteSessionTuesdayNZCSRSC08 ;
  conf:hasPresenter :EdwinDando .

:KeynoteSessionTuesdayNZCSRSC08 a conf:PresentationSession ;
  rdfs:label "a keynote session" ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C3LectureTheatre ;
  timeplace:startsAt "2008-04-15T11:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T12:30:00+13:00"^^xsd:dateTime .

:PhotoSessionNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Photo session"@en ;
  timeplace:locatedIn :CanterburyUni ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-15T12:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T13:00:00+13:00"^^xsd:dateTime .

:LunchTuesdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Lunch"@en ;
  timeplace:locatedIn :CBlock ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-15T13:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T14:00:00+13:00"^^xsd:dateTime .

:ChalkNTalk a conf:PresentationSession ;
  rdfs:label "Chalk’n’talk sessions"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :CBlock ;
  timeplace:startsAt "2008-04-15T13:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T14:00:00+13:00"^^xsd:dateTime .
```

```
:Session2NZCSRSC08 a conf:PresentationSession ;
  rdfs:label "Session 2"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C3LectureTheatre ;
  timeplace:startsAt "2008-04-15T14:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T15:50:00+13:00"^^xsd:dateTime .

:PosterSessionNZCSRSC08 a conf:PosterSession ;
  rdfs:label "Poster Session"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :CBlock ;
  timeplace:startsAt "2008-04-15T15:50:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T18:00:00+13:00"^^xsd:dateTime .

:OrionHealthNightNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Orion Health Social Night"@en ;
  timeplace:locatedIn :CanterburyUni ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-15T18:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T21:45:00+13:00"^^xsd:dateTime .

:BBQDinnerNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "BBQ Dinner"@en ;
  timeplace:locatedIn :UCSAFoundry ;
  timeplace:during :OrionHealthNightNZCSRSC08 .

:ComedyNight a conf:SocialEvent ;
  rdfs:label "BBQ Dinner"@en ;
  timeplace:locatedIn :NgaioMarshTheatre ;
  timeplace:during :OrionHealthNightNZCSRSC08 ;
  timeplace:startsAt "2008-04-15T20:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-15T21:45:00+13:00"^^xsd:dateTime .

:BreakfastWednesdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Breakfast"@en ;
```

```
  timeplace:locatedIn :BishopJulius ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-16T07:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T08:30:00+13:00"^^xsd:dateTime .

:Session3NZCSRSC08 a conf:PresentationSession ;
  rdfs:label "Session 3"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C2LectureTheatre ;
  timeplace:startsAt "2008-04-16T09:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T11:10:00+13:00"^^xsd:dateTime .

:MorningTeaWednesdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Morning tea"@en ;
  timeplace:locatedIn :CBlock ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-16T11:10:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T11:30:00+13:00"^^xsd:dateTime .

:KeynoteSessionWednesdayNZCSRSC08 a conf:PresentationSession ;
  rdfs:label "a keynote session" ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C2LectureTheatre ;
  timeplace:startsAt "2008-04-16T11:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T12:30:00+13:00"^^xsd:dateTime .

:KeynoteWednesdayNZCSRSC08 a conf:KeynoteTalk ;
  rdfs:label "Keynote Dave Lane" ;
  conf:inSession :KeynoteSessionWednesdayNZCSRSC08 ;
  conf:hasPresenter :DaveLane .

:LunchWednesdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Lunch"@en ;
  timeplace:locatedIn :CBlock ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-16T12:30:00+13:00"^^xsd:dateTime ;
```

```
  timeplace:endsAt   "2008-04-16T13:30:00+13:00"^^xsd:dateTime .

:Session4NZCSRSC08 a conf:PresentationSession ;
  rdfs:label "Session 4"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C2LectureTheatre ;
  timeplace:startsAt "2008-04-16T13:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T15:20:00+13:00"^^xsd:dateTime .

:AfternoonTeaWednesdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Afternoon tea"@en ;
  timeplace:locatedIn :CBlock ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-16T15:20:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T15:40:00+13:00"^^xsd:dateTime .

:WorkshopsNZCSRSC08 a conf:PresentationSession ;
  rdfs:label "BuildIT Workshop sessions"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :CBlock ;
  timeplace:startsAt "2008-04-16T15:40:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T18:00:00+13:00"^^xsd:dateTime .

:IndustryEventNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "NZi3 Industry Event"@en ;
  timeplace:during :NZCSRSC08 ;
  timeplace:locatedIn :CanterburyUni ;
  timeplace:startsAt "2008-04-16T18:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-16T21:00:00+13:00"^^xsd:dateTime .

:IndustryDinnerNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Industry dinner"@en ;
  timeplace:during :IndustryEventNZCSRSC08 ;
  timeplace:locatedIn :UCSAUpperCommonRoom ;
  timeplace:startsAt "2008-04-16T18:00:00+13:00"^^xsd:dateTime .
```

```
:IndustryMixerNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Industry mixer"@en ;
  timeplace:during :IndustryEventNZCSRSC08 ;
  timeplace:locatedIn :HITLab ;
  timeplace:startsAt "2008-04-16T19:00:00+13:00"^^xsd:dateTime .

:BreakfastThursdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Breakfast"@en ;
  timeplace:during :NZCSRSC08 ;
  timeplace:locatedIn :BishopJulius ;
  timeplace:startsAt "2008-04-17T07:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-17T08:30:00+13:00"^^xsd:dateTime .

:Session5NZCSRSC08 a conf:PresentationSession ;
  rdfs:label "Session 5"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :C3LectureTheatre ;
  timeplace:startsAt "2008-04-17T09:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-17T10:30:00+13:00"^^xsd:dateTime .

:MorningTeaThursdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Morning tea"@en ;
  timeplace:locatedIn :CBlock ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-17T10:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-17T11:00:00+13:00"^^xsd:dateTime .

:KeynoteSessionThursdayNZCSRSC08 a conf:PresentationSession ;
  rdfs:label "a keynote session"@en ;
  conf:sessionIn :NZCSRSC08 ;
  timeplace:locatedIn :UCSAShelleyCommonRoom ;
  timeplace:startsAt "2008-04-17T11:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-17T12:00:00+13:00"^^xsd:dateTime .

:KeynoteThursdayNZCSRSC08 a conf:KeynoteTalk ;
  rdfs:label "Keynote David Park"@en ;
```

```
  conf:inSession :KeynoteSessionThursdayNZCSRSC08 ;
  conf:hasPresenter :DavidPark .

:ClosingNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Awards and Closing"@en ;
  timeplace:locatedIn :UCSAShelleyCommonRoom ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-17T12:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-17T12:30:00+13:00"^^xsd:dateTime .

:LunchThursdayNZCSRSC08 a conf:SocialEvent ;
  rdfs:label "Lunch"@en ;
  timeplace:locatedIn :UCSAShelleyCommonRoom ;
  timeplace:during :NZCSRSC08 ;
  timeplace:startsAt "2008-04-17T12:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2008-04-17T14:00:00+13:00"^^xsd:dateTime .

# Sessions NZCSRSC 09

:RegistrationNZCSRSC09 a conf:SocialEvent ;
  rdfs:label "Registration"@en ;
  timeplace:locatedIn :OwenGlennBuilding ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-06T15:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-06T16:30:00+13:00"^^xsd:dateTime .

:PowhiriNZCSRSC09 a conf:SocialEvent ;
  rdfs:label "Powhiri"@en ;
  timeplace:locatedIn :MaraeAuckland ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-06T16:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-06T17:30:00+13:00"^^xsd:dateTime .

:OpeningDinnerNZCSRSC09 a conf:SocialEvent ;
  rdfs:label "Endace Opening Dinner"@en ;
  timeplace:locatedIn :FalePacifica ;
```

```
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-06T18:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-06T20:00:00+13:00"^^xsd:dateTime .

:BreakfastTuesdayNZCSRSC09 a conf:SocialEvent ;
  rdfs:label "Breakfast"@en ;
  timeplace:locatedIn :WhareKai ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-07T07:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T08:30:00+13:00"^^xsd:dateTime .

:Presentations1NZCSRSC09 a conf:PresentationSession ;
  rdfs:label "Presentation session 1"@en ;
  conf:sessionIn :NZCSRSC09 ;
  timeplace:startsAt "2009-04-07T09:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T10:50:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding .

:MorningBreakTuesdayNZCSRSC09 a conf:Break ;
  rdfs:label "Break"@en ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-07T10:50:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T11:20:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :WhareKai .

:KeynoteTuesdayNZCSRSC09 a conf:KeynoteTalk ;
  rdfs:label "Keynote 1"@en ;
  conf:hasPresenter :AlanBlackwell ;
  conf:inSession :KeynoteSessionTuesdayNZCSRSC09 .

:KeynoteSessionTuesdayNZCSRSC09 a conf:PresentationSession ;
  rdfs:label "a keynote session" ;
  timeplace:startsAt "2009-04-07T11:20:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T12:20:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding ;
  conf:sessionIn :NZCSRSC09 .
```

```
:ConferencePhotoNZCSRSC09 a conf:Break ;
  rdfs:label "Conference photo"@en ;
  timeplace:startsAt "2009-04-07T12:20:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T12:40:00+13:00"^^xsd:dateTime ;
  timeplace:during :NZCSRSC09 .

:LunchNZCSRSC09 a conf:Break ;
  rdfs:label "Lunch"@en ;
  timeplace:startsAt "2009-04-07T12:40:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T13:40:00+13:00"^^xsd:dateTime ;
  timeplace:during :NZCSRSC09 ;
  timeplace:locatedIn :WhareKai .

:Presentations2NZCSRSC09 a conf:PresentationSession ;
  rdfs:label "Presentation session 2"@en ;
  conf:sessionIn :NZCSRSC09 ;
  timeplace:startsAt "2009-04-07T13:40:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T15:30:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding .

:PostersNZCSRSC09 a conf:PosterSession ;
  rdfs:label "Poster session"@en ;
  conf:sessionIn :NZCSRSC09 ;
  timeplace:startsAt "2009-04-07T15:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T17:00:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :WhareKai .

:DinnerCruise a conf:SocialEvent ;
  rdfs:label "Orion Health Social Night"@en ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-07T18:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-07T21:00:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :AucklandHarbour .

:BreakfastWednesdayNZCSRSC09 a conf:SocialEvent ;
```

```
  rdfs:label "Breakfast"@en ;
  timeplace:locatedIn :WhareKai ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-08T07:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T08:30:00+13:00"^^xsd:dateTime .

:Presentations3NZCSRSC09 a conf:PresentationSession ;
  rdfs:label "Presentation session 3"@en ;
  conf:sessionIn :NZCSRSC09 ;
  timeplace:startsAt "2009-04-08T09:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T10:50:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding .

:MorningBreakWednesdayNZCSRSC09 a conf:Break ;
  rdfs:label "Break"@en ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-08T10:50:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T11:20:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :WhareKai .

:KeynoteWednesdayNZCSRSC09 a conf:KeynoteTalk ;
  rdfs:label "Keynote 2"@en ;
  conf:hasPresenter :JPLewis ;
  conf:inSession :KeynoteSessionWednesdayNZCSRSC09 .

:KeynoteSessionWednesdayNZCSRSC09 a conf:PresentationSession ;
  rdfs:label "a keynote session" ;
  timeplace:startsAt "2009-04-08T11:20:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T12:20:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding ;
  conf:sessionIn :NZCSRSC09 .

:DiscussionGroups a conf:Session ;
  rdfs:label "Discussion groups and lunch"@en ;
  timeplace:startsAt "2009-04-08T12:20:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T13:50:00+13:00"^^xsd:dateTime ;
```

```
  conf:sessionIn :NZCSRSC09 ;
  timeplace:locatedIn :WhareKai .

:Presentations4NZCSRSC09 a conf:PresentationSession ;
  rdfs:label "Presentation session 4"@en ;
  conf:sessionIn :NZCSRSC09 ;
  timeplace:startsAt "2009-04-08T13:50:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T15:30:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding .

:AfternoonBreakWednesdayNZCSRSC09 a conf:Break ;
  rdfs:label "Break"@en ;
  timeplace:startsAt "2009-04-08T15:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T16:00:00+13:00"^^xsd:dateTime ;
  timeplace:during :NZCSRSC09 ;
  timeplace:locatedIn :WhareKai .

:WorkshopsNZCSRSC09 a conf:Session ;
  conf:sessionIn :NZCSRSC09 ;
  timeplace:startsAt "2009-04-08T16:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-08T17:30:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :WhareKai .

:IndustryDinner a conf:SocialEvent ;
  rdfs:label "Industry dinner"@en ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-08T18:00:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :Scala .

:BreakfastThursdayNZCSRSC09 a conf:SocialEvent ;
  rdfs:label "Breakfast"@en ;
  timeplace:locatedIn :WhareKai ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-09T07:30:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-09T08:30:00+13:00"^^xsd:dateTime .
```

```
:Presentations5NZCSRSC09 a conf:PresentationSession ;
  rdfs:label "Presentation session 5"@en ;
  conf:sessionIn :NZCSRSC09 ;
  timeplace:startsAt "2009-04-09T09:00:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-09T10:50:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding .


:MorningBreakThursdayNZCSRSC09 a conf:Break ;
  rdfs:label "Break"@en ;
  timeplace:during :NZCSRSC09 ;
  timeplace:startsAt "2009-04-09T10:50:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-09T11:20:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :WhareKai .

:KeynoteThursdayNZCSRSC09 a conf:KeynoteTalk ;
  rdfs:label "Keynote 3"@en ;
  conf:hasPresenter :PoulNielsen ;
  conf:inSession :KeynoteSessionThursdayNZCSRSC09 .

:KeynoteSessionThursdayNZCSRSC09 a conf:PresentationSession ;
  rdfs:label "a keynote session" ;
  timeplace:startsAt "2009-04-09T11:20:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-09T12:20:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :OwenGlennBuilding ;
  conf:sessionIn :NZCSRSC09 .

:ClosingLunchNZCSRSC09 a conf:SocialEvent ;
  rdfs:label "Closing lunch and awards ceremony" ;
  timeplace:startsAt "2009-04-09T12:20:00+13:00"^^xsd:dateTime ;
  timeplace:endsAt   "2009-04-09T13:45:00+13:00"^^xsd:dateTime ;
  timeplace:locatedIn :WhareKai ;
  timeplace:during :NZCSRSC09 .

# Conversations
```

```
:TalkingAboutXOsAtPosterSession a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :PosterSessionNZCSRSC08 ;
  interact:hasTopic :XOs,
    :Python ;
  interact:hasConversationPartner :Gabriel .

:TalkingAboutScubaBeforePowhiri a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:before :PowhiriNZCSRSC09 ;
  interact:hasTopic :ScubaDiving ;
  interact:hasConversationPartner :Alexander,
    :Gabriel .

:TalkingToRubyBeforeDinner a interact:Conversation ;
  rdfs:label "a conversation" ;
  interact:hasConversationPartner :Ruby ;
  timeplace:before :OpeningDinnerNZCSRSC09 .

:TalkingAtDinnerFirstNight a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :OpeningDinnerNZCSRSC09;
  interact:hasConversationPartner :MozillaRobert .

:TalkingToAmyAtBreakfast a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :BreakfastTuesdayNZCSRSC09 ;
  interact:hasConversationPartner :Amy ;
  interact:hasTopic :AnitaBorgScholarship .

:TalkingToChristopherKyleAtBreakfast a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :BreakfastTuesdayNZCSRSC09 ;
  interact:hasConversationPartner :Christopher,
    :KyleWellington ;
```

```
  interact:hasTopic :NZCSRSC10 .

:TalkingToAmySeanAtBreakfast a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :BreakfastTuesdayNZCSRSC09 ;
  interact:hasConversationPartner :Amy ,
    :Sean ;
  interact:hasTopic :AmysResearch,
    :Web2Dot0 .

:TalkingToLiamInBreak a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :MorningBreakTuesdayNZCSRSC09 ;
  interact:hasTopic :Music ,
    :NZCSRSC10 ;
  interact:hasConversationPartner :Liam .

:TalkingToGabrielInBreakTuesday a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :MorningBreakTuesdayNZCSRSC09 ;
  interact:hasTopic :Python,
    :SocialNetworking ;
  interact:hasConversationPartner :Gabriel .

:TalkingToRubyRegistrationCHINZ09 a interact:Conversation ;
  rdfs:label "a conversation" ;
  timeplace:during :RegistrationCHINZ09 ;
  interact:hasTopic :NZCSRSC09, :CHINZ2009 ;
  interact:hasConversationPartner :Ruby .

# People

:Alexander a foaf:Person ;
  rdfs:label "Alexander L"@en .

:Amy a foaf:Person ;
```

```
  rdfs:label "Amy T"@en .

:Christopher a foaf:Person ;
  rdfs:label "Christopher A"@en .

:Gabriel a foaf:Person ;
  rdfs:label "Gabriel K"@en .

:Liam a foaf:Person ;
  rdfs:label "Liam W"@en .

:KyleWellington a foaf:Person ;
  rdfs:label "Kyle"@en .

:MozillaRobert a foaf:Person ;
  rdfs:label "Robert O"@en;
  foaf:fundedBy :MozillaNZ .

:Ruby a foaf:Person ;
  rdfs:label "Ruby B"@en .

:Sean a foaf:Person ;
  rdfs:label "Sean M"@en .

# Topics

:ScubaDiving a interact:Topic ;
  rdfs:label "Scuba diving"@en .

# Other

:MozillaNZ a foaf:Organization ;
  rdfs:label "Mozilla NZ"@en .

:AucklandHarbour a timeplace:PlacedThing ;
  rdfs:label "Auckland Harbour"@en ;
```

```
  timeplace:coordPrecision timeplace:SuburbPrecision ;
  timeplace:locatedIn :Auckland .

:Scala a timeplace:PlacedThing ;
  rdfs:label "Scala"@en ;
  timeplace:locatedIn :AucklandUni ;
  timeplace:coordPrecision timeplace:RoomPrecision .
```

A.2.2 Annotated types

This section shows the type annotations for the conference memories in the previous section. They were used by the Digital Parrot to determine which classes and properties to show prominently, less prominently or not at all in the user interface.

```
@prefix : <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
   DigitalParrot/2009/02/DigitalParrot.owl#> .
@prefix conf: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms
   /Conferences/2008/11/Conferences.owl#> .
@prefix int: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
   Interaction/2008/11/Interaction.owl#> .
@prefix mem: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
   Memories/2008/11/Memories.owl#> .
@prefix timeplace: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/
   Terms/TimeAndPlace/2008/11/TimeAndPlace.owl#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

conf:Break                  :showThisType "primary"^^xsd:string .
conf:Conference             :showThisType "primary"^^xsd:string .
conf:ConferenceSeries       :showThisType "primary"^^xsd:string .
conf:Demo                   :showThisType "primary"^^xsd:string .
conf:DemoSession            :showThisType "primary"^^xsd:string .
conf:KeynoteTalk            :showThisType "primary"^^xsd:string .
conf:Poster                 :showThisType "primary"^^xsd:string .
conf:PosterPresentation     :showThisType "primary"^^xsd:string .
conf:PosterSession          :showThisType "primary"^^xsd:string .
conf:Presentation           :showThisType "primary"^^xsd:string .
```

```
conf:PresentationSession :showThisType "primary"^^xsd:string .
conf:Session             :showThisType "primary"^^xsd:string .
conf:SocialEvent         :showThisType "primary"^^xsd:string .
conf:Talk                :showThisType "primary"^^xsd:string .
conf:Track               :showThisType "primary"^^xsd:string .
conf:Tutorial            :showThisType "primary"^^xsd:string .
conf:TutorialSession     :showThisType "primary"^^xsd:string .
conf:Workshop            :showThisType "primary"^^xsd:string .
conf:WorkshopSession     :showThisType "primary"^^xsd:string .

conf:chairs           :showThisType "primary"^^xsd:string .
conf:hasChair         :showThisType "primary"^^xsd:string .
conf:hasInstance      :showThisType "primary"^^xsd:string .
conf:hasPresentation  :showThisType "primary"^^xsd:string .
conf:hasPresenter     :showThisType "primary"^^xsd:string .
conf:hasSession       :showThisType "primary"^^xsd:string .
conf:hasTrack         :showThisType "primary"^^xsd:string .
conf:inSession        :showThisType "primary"^^xsd:string .
conf:partOfSeries     :showThisType "primary"^^xsd:string .
conf:presents         :showThisType "primary"^^xsd:string .
conf:sessionIn        :showThisType "primary"^^xsd:string .
conf:trackIn          :showThisType "primary"^^xsd:string .

int:Conversation  :showThisType "primary"^^xsd:string .
int:FirstMeeting  :showThisType "primary"^^xsd:string .
int:Introduction  :showThisType "primary"^^xsd:string .
int:Topic         :showThisType "primary"^^xsd:string .

int:conversationPartnerIn  :showThisType "primary"^^xsd:string .
int:firstMet               :showThisType "primary"^^xsd:string .
int:hasConversationPartner :showThisType "primary"^^xsd:string .
int:hasTopic               :showThisType "primary"^^xsd:string .
int:introducedAt           :showThisType "primary"^^xsd:string .
int:introducer             :showThisType "primary"^^xsd:string .
int:newContact             :showThisType "primary"^^xsd:string .
int:topicIn                :showThisType "primary"^^xsd:string .
```

```
mem:Event             :showThisType "primary"^^xsd:string .
mem:Experience        :showThisType "primary"^^xsd:string .
mem:LifetimePeriod    :showThisType "primary"^^xsd:string .
timeplace:AbsolutelyPlacedThing
  :showThisType "secondary"^^xsd:string .
timeplace:AbsolutelyTimedThing
  :showThisType "secondary"^^xsd:string .
timeplace:PlacedThing :showThisType "primary"^^xsd:string .
timeplace:TimedThing  :showThisType "primary"^^xsd:string .

timeplace:after     :showThisType "primary"^^xsd:string .
timeplace:before    :showThisType "primary"^^xsd:string .
timeplace:during    :showThisType "primary"^^xsd:string .
timeplace:encloses  :showThisType "primary"^^xsd:string .
timeplace:endsAt    :showThisType "secondary"^^xsd:string .
timeplace:lat       :showThisType "secondary"^^xsd:string .
timeplace:locatedIn :showThisType "primary"^^xsd:string .
timeplace:long      :showThisType "secondary"^^xsd:string .
timeplace:spans     :showThisType "primary"^^xsd:string .
timeplace:startsAt  :showThisType "secondary"^^xsd:string .

<http://xmlns.com/foaf/0.1/Person>
  :showThisType "primary"^^xsd:string .
```

A.3 Additional ontologies

This section contains additional ontologies used in the research described in this thesis, most notably in the user studies.

These ontologies define vocabulary related to academic conferences (Section A.3.1) and to people's interactions with each other (Section A.3.2).

A.3.1 Conferences

```
@prefix : <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
   Conferences/2008/11/Conferences.owl#> .
```

```
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix timeplace: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/
    Terms/TimeAndPlace/2008/11/TimeAndPlace.owl#> .
@prefix umbel: <http://umbel.org/umbel/sc/> .

<http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/Conferences
    /2008/11/Conferences.owl> a owl:Ontology;
  rdfs:comment "Vocabulary for describing academic conferences."
      @en;
  owl:imports <http://www.mindswap.org/2003/owl/foaf>,
    <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
       TimeAndPlace/2008/11/TimeAndPlace.owl> .

:Break a owl:Class;
  rdfs:comment "A break in the conference program"@en;
  rdfs:subClassOf timeplace:TimedThing;
  owl:disjointWith :Conference, :ConferenceSeries, :Poster,
      :Presentation, :Session, :SocialEvent, :Track;
  rdfs:label "Break"@en .

:Conference a owl:Class;
  rdfs:comment "An individual conference, for example HCI 2008."
      @en;
  rdfs:subClassOf timeplace:TimedThing,
    timeplace:PlacedThing,
    umbel:ScientificConference;
  owl:disjointWith :Break, :ConferenceSeries, :Poster,
      :Presentation, :Session, :SocialEvent, :Track;
  rdfs:label "Conference"@en .

:ConferenceSeries a owl:Class;
  rdfs:comment "A series of conferences, for example the series of
      HCI conferences."@en;
  owl:disjointWith :Break, :Conference, :Poster, :Presentation,
      :Session, :SocialEvent, :Track;
  rdfs:label "Conference Series"@en .

:Demo a owl:Class;
  rdfs:comment "A demonstration, usually of a piece of software."
      @en;
  rdfs:subClassOf :Presentation;
  rdfs:label "Demo"@en .
```

```
:DemoSession a owl:Class;
  rdfs:comment "A session with demos."@en;
  rdfs:subClassOf :Session,
    [ a owl:Restriction ;
      owl:onProperty :hasPresentation ;
      owl:allValuesFrom :Demo
    ];
  owl:disjointWith :PosterSession,
    :PresentationSession;
  rdfs:label "Demo Session"@en .

:KeynoteTalk a owl:Class;
  rdfs:comment "A keynote talk."@en;
  rdfs:subClassOf :Presentation;
  rdfs:label "Keynote Talk"@en .

:Poster a owl:Class;
  rdfs:comment "A poster as can be seen during a poster session.
     Note that this describes the actual poster; walk-throughs by
     the presenter are encoded as instances of PosterPresentation.
     "@en;
  owl:disjointWith :Break, :Conference, :ConferenceSeries,
     :Presentation, :Session, :SocialEvent, :Track;
  rdfs:label "Poster"@en .

:PosterPresentation a owl:Class;
  rdfs:comment "A walk-through for a poster in a poster session."
     @en;
  rdfs:subClassOf :Presentation;
  rdfs:label "Poster Presentation"@en .

:PosterSession a owl:Class;
  rdfs:comment "A poster session. Note that this has instances of
     PosterPresentation as its hasPresentation instances -- the
     Poster is used to describe the actual poster that is being
     shown."@en;
  rdfs:subClassOf :Session,
    [ a owl:Restriction ;
      owl:onProperty :hasPresentation ;
      owl:allValuesFrom :PosterPresentation
    ];
  owl:disjointWith :DemoSession,
    :PresentationSession;
  rdfs:label "Poster Session"@en .
```

```
:Presentation a owl:Class;
  rdfs:comment "Something that happens in a session. Note that a
     typical paper presentation is a Talk, not a Presentation."
     @en;
  owl:disjointWith :Break, :Conference, :ConferenceSeries, :Poster,
      :Session, :SocialEvent, :Track;
  rdfs:subClassOf timeplace:TimedThing,
    timeplace:PlacedThing,
    umbel:ScientificPresentation;
  rdfs:label "Presentation"@en .

:PresentationSession a owl:Class;
  rdfs:comment "A session with talks, as opposed to a demo session
     or a poster session."@en;
  rdfs:subClassOf :Session,
    [ a owl:Restriction ;
      owl:onProperty :hasPresentation ;
      owl:allValuesFrom :Talk
    ];
  owl:disjointWith :DemoSession,
    :PosterSession;
  rdfs:label "Presentation Session"@en .

:Session a owl:Class;
  rdfs:comment "A session. Most conferences come divided into
     sessions."@en;
  owl:disjointWith :Break, :Conference, :ConferenceSeries, :Poster,
      :Presentation, :SocialEvent, :Track;
  rdfs:subClassOf timeplace:TimedThing,
    timeplace:PlacedThing,
    umbel:ConferenceSession,
    [ a owl:Restriction ;
      owl:onProperty :hasPresentation ;
      owl:someValuesFrom :Presentation
    ];
  rdfs:label "Session"@en .

:SocialEvent a owl:Class;
  rdfs:comment "A social event, for example the conference dinner."
     @en;
  owl:disjointWith :Break, :Conference, :ConferenceSeries, :Poster,
      :Presentation, :Session, :Track;
  rdfs:subClassOf timeplace:TimedThing,
    timeplace:PlacedThing;
  rdfs:label "Social Event"@en .
```

```
:Talk a owl:Class;
  rdfs:comment "A talk, usually with slides and by a single
     speaker."@en;
  rdfs:subClassOf :Presentation;
  rdfs:label "Talk"@en .

:Track a owl:Class;
  rdfs:comment "A track, for conferences that have more than one
     parallel track."@en;
  owl:disjointWith :Break, :Conference, :ConferenceSeries, :Poster,
      :Presentation, :SocialEvent, :Session;
  rdfs:label "Track"@en .

:Tutorial a owl:Class;
  rdfs:comment "A tutorial."@en;
  rdfs:subClassOf :Presentation;
  rdfs:label "Tutorial"@en .

:TutorialSession a owl:Class;
  rdfs:comment "A session with tutorials."@en;
  rdfs:subClassOf :Session,
    [ a owl:Restriction ;
      owl:onProperty :hasPresentation ;
      owl:allValuesFrom :Tutorial
    ];
  rdfs:label "Tutorial Session"@en .

:Workshop a owl:Class;
  rdfs:comment "A workshop. Typically, workshops are made up of
     talks, but they might have other things too (keynotes etc)."
     @en;
  rdfs:subClassOf :Session;
  rdfs:label "Workshop"@en .

:WorkshopSession a owl:Class;
  rdfs:comment "A session consisting of workshops."@en;
  rdfs:subClassOf :Session;
  rdfs:label "Workshop Session"@en .

:chairs a owl:ObjectProperty;
  rdfs:comment "The act of chairing a session or conference."@en;
  rdfs:domain foaf:Person;
  rdfs:range  [
    a owl:Class;
```

```
    owl:unionOf  ( :Session :Conference )
  ];
  owl:inverseOf :hasChair;
  rdfs:label "chairs"@en .

:hasChair a owl:ObjectProperty;
  rdfs:comment "To link to the chair of a session or conference."
     @en;
  rdfs:domain  [
     a owl:Class;
     owl:unionOf  ( :Conference :Session )
  ];
  rdfs:range foaf:Person;
  owl:inverseOf :chairs;
  rdfs:label "has chair"@en .

:hasInstance a owl:ObjectProperty;
  rdfs:comment "To link a conference series to its instances."@en;
  rdfs:domain :ConferenceSeries;
  rdfs:range :Conference;
  owl:inverseOf :partOfSeries;
  rdfs:label "has instance"@en .

:hasPresentation a owl:ObjectProperty;
  rdfs:comment "To link a session to its presentations. This would
     typically be used with an instance of a _subclass_ of Session
      and an instance of a _subclass_ of Presentation."@en;
  rdfs:domain :Session;
  rdfs:range :Presentation;
  owl:inverseOf :inSession;
  rdfs:label "has presentation"@en .

:hasPresenter a owl:ObjectProperty;
  rdfs:comment "To link a presentation to its presenter(s)."@en;
  rdfs:domain :Presentation;
  rdfs:range foaf:Person;
  owl:inverseOf :presents;
  rdfs:label "has presenter"@en .

:inSession a owl:ObjectProperty;
  rdfs:comment "To link a presentation to its session. Also see
     comment on the inverse property."@en;
  rdfs:domain :Presentation;
  rdfs:range :Session;
  owl:inverseOf :hasPresentation;
```

```
  rdfs:label "in session"@en .

:partOfSeries a owl:ObjectProperty;
  rdfs:comment "Denotes that a conference is part of a conference
     series. For example, NZCSRSC 2007 was a conference in the
     NZCSRSC series."@en;
  rdfs:domain :Conference;
  rdfs:range :ConferenceSeries;
  owl:inverseOf :hasInstance;
  rdfs:label "part of series"@en .

:presents a owl:ObjectProperty;
  rdfs:comment "To link a presenter to his/her presentation(s)."
     @en;
  rdfs:domain foaf:Person;
  rdfs:range :Presentation;
  owl:inverseOf :hasPresenter;
  rdfs:label "presents"@en .

:hasTrack a owl:InverseFunctionalProperty;
  rdfs:comment "To link something to its track(s). Domain would
     typically be a conference but might be a workshop."@en;
  rdfs:range :Track;
  rdfs:label "has track"@en .

:trackIn a owl:ObjectProperty, owl:FunctionalProperty;
  rdfs:comment "To link a track to its bigger event. Range would
     typically be a conference but might be a workshop."@en;
  rdfs:domain :Track;
  rdfs:label "track in"@en .

:hasSession a owl:InverseFunctionalProperty;
  rdfs:comment "To link something to its session(s). Domain would
     typically be a conference or track but might be a workshop."
     @en;
  rdfs:range :Session;
  rdfs:label "has session"@en .

:sessionIn a owl:ObjectProperty, owl:FunctionalProperty;
  rdfs:comment "To link a session to its bigger event. Range would
     typically be a conference or track but might be a workshop."
     @en;
  rdfs:domain :Session;
  rdfs:label "session in"@en .
```

A.3.2 Interaction

```
@prefix : <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/
    Interaction/2008/11/Interaction.owl#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix memories: <http://parrot.resnet.scms.waikato.ac.nz/Parrot/
    Terms/Memories/2008/11/Memories.owl#> .
@prefix umbel: <http://umbel.org/umbel/sc/> .

<http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/Interaction
    /2008/11/Interaction.owl> a owl:Ontology;
  rdfs:comment "Vocabulary to describe interactions between self
      and other people."@en;
  owl:imports <http://www.mindswap.org/2003/owl/foaf>,
  <http://parrot.resnet.scms.waikato.ac.nz/Parrot/Terms/Memories
      /2008/11/Memories.owl> .

:Conversation a owl:Class;
  rdfs:subClassOf memories:Event,
    [ a owl:Restriction ;
      owl:onProperty :hasConversationPartner ;
      owl:someValuesFrom foaf:Person
    ],
    [ a owl:Restriction ;
      owl:onProperty :hasTopic ;
      owl:someValuesFrom :Topic
    ],
    umbel:Conversation;
  rdfs:comment "A conversation between self and at least one other
      person."@en;
  rdfs:label "Conversation"@en;
  owl:disjointWith :FirstMeeting,
    :Topic .

:FirstMeeting a owl:Class;
  rdfs:subClassOf memories:Event,
    [ a owl:Restriction ;
      owl:onProperty :newContact ;
      owl:someValuesFrom foaf:Person
    ];
  rdfs:comment "The occasion when one first meets another person."
      @en;
  rdfs:label "First Meeting"@en;
```

```
  owl:disjointWith :Conversation,
    :Topic .

:Introduction a owl:Class;
  rdfs:subClassOf :FirstMeeting,
    [ a owl:Restriction ;
      owl:onProperty :introducer ;
      owl:someValuesFrom foaf:Person
    ];
  rdfs:comment "The occasion when one is introduced to another
     person by a third party."@en;
  rdfs:label "Introduction"@en .

:Topic a owl:Class;
  rdfs:comment "A topic as dicussed in a conversation. To be used
     like a tag."@en;
  rdfs:label "Topic"@en;
  rdfs:subClassOf umbel:Communication_Topic;
  owl:disjointWith :Conversation,
    :FirstMeeting .

:conversationPartnerIn a owl:ObjectProperty;
  rdfs:comment "To link a person to the conversations that one had
     with this person"@en;
  rdfs:domain foaf:Person;
  rdfs:range :Conversation;
  owl:inverseOf :hasConversationPartner;
  rdfs:label "takes part in"@en .

:hasConversationPartner a owl:ObjectProperty;
  rdfs:comment "To link a conversation to the people one spoke
     with."@en;
  rdfs:domain :Conversation;
  rdfs:range foaf:Person;
  owl:inverseOf :conversationPartnerIn;
  rdfs:label "with"@en .

:hasTopic a owl:ObjectProperty;
  rdfs:comment "To describe the topic of a conversation."@en;
  rdfs:domain :Conversation;
  rdfs:range :Topic;
  owl:inverseOf :topicIn;
  rdfs:label "is about"@en .

:introducedAt a owl:ObjectProperty,
```

```
    owl:FunctionalProperty;
  rdfs:subPropertyOf :firstMet;
  rdfs:comment "To link a person to the occasion when one was
     introduced to her/him"@en;
  rdfs:domain foaf:Person;
  rdfs:range :Introduction;
  rdfs:label "was introduced at"@en .

:introducer a owl:ObjectProperty;
  rdfs:comment "To link an introduction to the person(s) who did
     the introducing"@en;
  rdfs:domain :Introduction;
  rdfs:range foaf:Person;
  rdfs:label "was introduced by"@en .

:firstMet a owl:ObjectProperty,
    owl:FunctionalProperty;
  rdfs:comment "To describe since when one knows a person."@en;
  rdfs:domain foaf:Person;
  rdfs:range :FirstMeeting;
  owl:inverseOf :newContact;
  rdfs:label "was first met"@en .

:newContact a owl:ObjectProperty,
    owl:InverseFunctionalProperty;
  rdfs:comment "To describe who it is one met"@en;
  rdfs:domain :FirstMeeting;
  rdfs:range foaf:Person;
  owl:inverseOf :firstMet;
  rdfs:label "was the first meeting with"@en .

:topicIn a owl:ObjectProperty;
  rdfs:comment "To link a topic to the conversations in which it
     occurs."@en;
  rdfs:domain :Topic;
  rdfs:range :Conversation;
  owl:inverseOf :hasTopic;
  rdfs:label "is a topic in"@en .
```

Appendix B

Material for the usability study

This appendix contains material given to the participants in the study that evaluated the Digital Parrot's usability, described in Chapter 7:

- the approval letter from the Human Research Ethics Committee of the School of Computing and Mathematical Sciences at the University of Waikato, dated 27 July 2009;
- the Participant Information Sheet, which outlines the study goals and procedure as well as the participant's rights;
- the Research Consent Form, which each participant signed at the beginning of their session;
- the Participant Workbook, which contains instructions, the study tasks and the SUS questionnaire as well as the biographical questionnaire.

Computing and Mathematical Sciences
Rorohiko me ngā Pūtaiao Pāngarau
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4322
www.scms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

27 July 2009

Andrea Schweer
C/- Department of Computer Science
THE UNIVERSITY OF WAIKATO

Dear Andrea

Request for approval to conduct a research project study using human participants

I have considered your request for approval to conduct a research study as part of your PhD research on Context-Aware Augmented Memory, the purpose being to gain some insights about the usability of the software application you are developing to detect any serious flaws that might impact on the results of a bigger study later this year.

The procedure described in your request is acceptable. I note your statement that participants anonymity will be strictly maintained and that all information gathered will be used for statistical analysis only and no names or other identifying characteristics will be stated in the final or any other reports.

The research participants' instruction sheet and consent forms comply with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to conduct the study.

Yours sincerely

Dave Nichols
Human Research Ethics Committee
School of Computing and Mathematical Sciences

Participant Information Sheet

Ethics Committee, School of Computing and Mathematical Sciences

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Project Title
Context-Aware Augmented Memory: Usability Test of the Digital Parrot

Purpose
This research is conducted as part of my PhD research on Context-Aware Augmented Memory. I'm developing a software application and wish to gain some insights about its usability.

What is this research project about?
In my PhD project, I am investigating augmented memory systems – computer systems which support autobiographic memory, the part of long-term memory that stores memories about a person's life. My research explores which components are needed for a successful augmented memory system.

What will you have to do and how long will it take?
After you have signed the consent form, I'll ask you to perform a few tasks: finding answers to a few questions using the software that I'm developing. I will watch you work on the task and take notes, and I'll also ask you to speak about what you are doing as you are working on the tasks. After that, I'd like you to answer a few questions about your experience and about your background. I might ask a few questions to follow up on observations I made while you were working on the tasks. Altogether this will take about an hour.

What will happen to the information collected?
The information collected will be used by the researcher to write parts of her PhD dissertation. It is likely that articles and presentations will be the outcome of the research. The researcher will keep the paper-based materials as well as transcriptions of the notes and questionnaire responses but will treat them with the strictest confidentiality. The paper-based and electronic materials will be archived once the study has been completed, but in such a way that they can't be linked back to the individual participant. No participants will be named in the publications and every effort will be made to disguise their identity.

Declaration to participants
If you take part in the study, you have the right to:

- Refuse to answer any particular question, and to withdraw from the study before analysis has commenced on the data.
- Ask any further questions about the study that occurs to you during your participation.
- Be given access to a summary of findings from the study when it is concluded.

Who's responsible?
If you have any questions or concerns about the project, either now or in the future, please feel free to contact either:

Researcher:
Andrea Schweer
Department of Computer Science, The University of Waikato
Room FG 2.01
schweer@cs.waikato.ac.nz

Supervisors:
Annika Hinze
Department of Computer Science, The University of Waikato
hinze@cs.waikato.ac.nz

Steve Jones
Department of Computer Science, The University of Waikato
stevej@cs.waikato.ac.nz

Research Consent Form

Ethics Committee, School of Computing and Mathematical Sciences

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Context-Aware Augmented Memory: Usability Test of the Digital Parrot

Consent Form for Participants

I have read the **Participant Information Sheet** for this study and have had the details of the study explained to me. My questions about the study have been answered to my satisfaction, and I understand that I may ask further questions at any time.

I also understand that I am free to withdraw from the study at any time before analysis has commenced on the data, or to decline to answer any particular questions in the study. I understand I can withdraw any information I have provided up until the researcher has commenced analysis on my data. I agree to provide information to the researchers under the conditions of confidentiality set out on the **Participant Information Sheet**.

I agree to participate in this study under the conditions set out in the **Participant Information Sheet**.

Signed: ____

Name: ____

Date: ____

Researcher:
Andrea Schweer
Department of Computer Science, The University of Waikato
Room FG 2.01
schweer@cs.waikato.ac.nz

Supervisors:
Annika Hinze
Department of Computer Science, The University of Waikato
hinze@cs.waikato.ac.nz
Steve Jones
Department of Computer Science, The University of Waikato
stevej@cs.waikato.ac.nz

Participant's Workbook

Context-Aware Augmented Memory
Usability Test of the Digital Parrot

Participant number: ___________

Study conducted by:

Andrea Schweer
Department of Computer Science
The University of Waikato
schweer@cs.waikato.ac.nz

October/November 2009

Instructions

First of all, thank you very much for agreeing to take part in this study!

The Digital Parrot

The Digital Parrot is a repository of memories and of information associated with memories. The version that you will use today already contains some memories. These are some of my memories, collected at several conferences that I attended.

The Digital Parrot lets you navigate among the stored memories and associated information. There are four tools: the timeline, the map, textual search and trail navigation. I'll give you a quick demo of each tool before you start; please let me know when you are ready.

Structure of today's study

On the following pages you'll find a few tasks. In each task, please try to use the memories stored in the Digital Parrot answer the question that's provided. In some tasks, there may be no answer or there may be more than one correct answer. Go through the tasks in your own time, moving on to the next task when you wish.

As you go through the tasks, please say aloud whatever you are looking at, thinking, doing and feeling. I will take notes and occasionally prompt you to keep talking, if necessary.

After you have completed all tasks, I will ask you to fill in a short questionnaire about the Digital Parrot and your background. I might then ask you a few additional questions to follow up on my observations.

Again, thank you for helping us test the Digital Parrot. Turn the page and begin with the first task when you are ready. Finally, please remember: we're testing the Digital Parrot's performance, not yours!

Tasks

Task 1

To whom did I talk about scuba diving? Write their name(s) into the space below.

Task 2

Which conferences did I attend in Auckland? Write the conference name(s) into the space below.

Task 3

At which conference(s) did I speak to someone about Python during the poster session? Write the conference name(s) into the space below.

Task 4

In which place was the NZ CHI conference in 2007? Write the place name into the space below.

A few questions about the Digital Parrot

Please check the "○" that reflects your immediate response to each statement. Don't think too long about each statement. Make sure you respond to every statement. If you don't know how to respond to a question, simply check the middle "○".

I think that I would like to use the Digital Parrot frequently.

○ ○ ○ ○ ○
strongly disagree strongly agree

I found the Digital Parrot unnecessarily complex.

○ ○ ○ ○ ○
strongly disagree strongly agree

I thought the Digital Parrot was easy to use.

○ ○ ○ ○ ○
strongly disagree strongly agree

I think that I would need the support of a technical person to be able to use the Digital Parrot.

○ ○ ○ ○ ○
strongly disagree strongly agree

I found the various functions in the Digital Parrot were well integrated.

○ ○ ○ ○ ○
strongly disagree strongly agree

I thought there was too much inconsistency in the Digital Parrot.

○ ○ ○ ○ ○
strongly disagree strongly agree

I imagine that most people would learn to use the Digital Parrot very quickly.

○ ○ ○ ○ ○
strongly disagree strongly agree

I found the Digital Parrot very awkward to use.

○ ○ ○ ○ ○

strongly disagree — strongly agree

I felt very confident using the Digital Parrot.

○ ○ ○ ○ ○

strongly disagree — strongly agree

I needed to learn a lot of things before I could get going with the Digital Parrot.

○ ○ ○ ○ ○

strongly disagree — strongly agree

A few questions about yourself

Please circle your gender. male female

How old are you? _______ years

Please indicate your occupation.

- ○ CS research student
- ○ member of CS academic staff
- ○ research student in other department (please specify):
- ○ member of academic staff in other department (please specify):
- ○ other (please specify):

Please tick the conferences that you attended.

- NZCSRSC
  - ○ 2007 in Hamilton
  - ○ 2008 in Christchurch
  - ○ 2009 in Auckland
  - ○ none of these
- CHINZ
  - ○ 2007 in Hamilton
  - ○ 2009 in Auckland
  - ○ none of these

That's it – thank you!

Appendix C

Material for the effectiveness study

This appendix contains material given to the participants in the experiencing phase and in the recall phase of the evaluation of the Digital Parrot's effectiveness, described in Chapter 8.

The participant information material differs slightly between the two study phases because of changes in policy regarding experimentation involving human participants.

C.1 Experiencing Phase

The documents shown for the experiencing phase are:

- the approval letter from the Human Research Ethics Committee of the School of Computing and Mathematical Sciences at the University of Waikato, dated 1 August 2008.
- the Research Consent Form for the experiencing phase, which outlines the study goals and procedure and contains the agreement form that each participant signed at the beginning of their first interview; and
- the Research Participants' Bill of Rights that accompanied the Participant Information Sheet.

School of Computing &
Mathematical Sciences
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4021
www.scms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

1st August 2008

Andrea Schweer
C/- Department of Computer Science
University of Waikato

Dear Andrea

Request for approval to perform several experiments involving human participants

I have considered your request for approval to perform several experiments involving human participants during July to December this year, consisting of semi-structured interviews in which the study participants talk about academic conferences which they attended shortly before the respective interview.

The purpose of the interviews is to assist you in building a computer system that helps people store and recall information about their own past experiences and you wish to collect samples of the type of information that people would wish to store in such a system when they visit an academic conference, to aid you in your PhD project.

The procedure described in your request is acceptable. I note your statements that confidentiality and participant anonymity will be strictly maintained, all information gathered will be used for statistical analysis only and no names or other identifying characteristics will be stated in the final or any other reports. Data will be securely stored and all data containing personal information in non-anoymised form will be destroyed at the end of the research project.

The research participants' Bill of Rights and the Research Consent form comply with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to undertake the experiment.

Yours faithfully

M Mayo

Mike Mayo
Department of Computer Science
Human Research Ethics Committee
School of Computing and Mathematical Sciences

Research Consent Form

The University of Waikato
School of Computing and Mathematical Sciences

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

Research Project Title

Context-Aware Augmented Memory – Interview series about conference visits

Experiment Purpose

For my PhD project, I am building a computer system that helps people store and recall information about their own past experiences. With this interview series, I wish to collect samples of the type of information that people would wish to store in such a system when they visit an academic conference.

Participant Recruitment and Selection

Graduate students and academic staff members from the University of Waikato will be recruited for this experiment. Participants need to have visited, or plan to visit, at least one academic conference during June to December 2008 and should be available for follow-up sessions early in 2009.

Procedure

This session will require about 10 to 30 minutes of your time for each conference that you agree to be interviewed about.

Each interview will be for one particular conference that you visited shortly before the interview. The researcher will ask you some questions about your conference visit and may follow up on your answers with further questions. Where feasible, the researcher will bring a printout of the conference schedule to the interview to help you remember your conference visit.

Data Collection

The researcher will make an audio recording of each interview. The researcher may also take notes during the interview. In those interviews where a printed copy of the conference program is available, the researcher will keep this printout including any notes written on it by the participant or the researcher during the interview.

Data Archiving/Destruction

Data will be kept securely stored by the researcher. All data containing personal information in non-anonymised form will be destroyed at the end of the research project.

Confidentiality

Confidentiality and participant anonymity will be strictly maintained. All information gathered will be used for statistical analysis only and no names or other identifying characteristics will be stated in the final or any other reports.

Likelihood of Discomfort

There is no likelihood of discomfort or risk associated with participation.

Researcher

Andrea Schweer is working on her doctorate in the Computer Science Department at the University of Waikato, New Zealand. This study will contribute to her research on Context-Aware Augmented Memory Systems. Her supervisor is Dr. Annika Hinze.

Andrea can be contacted in room G 2.06 of the School of Computing and Mathematical Sciences building at the University of Waikato. Her phone number is 838 4466 ext. 6011 and her e-mail address is schweer@cs.waikato.ac.nz.

Finding out about Results

The participants can find out the results of the study by contacting the researcher after April 30, 2009.

2

Agreement

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive you legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. You are free to not answer specific items or questions in interviews or on questionnaires. You are free to withdraw from the study at any time without penalty. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact the researcher.

Participant Date

Researcher Date

A copy of this consent form has been given to you to keep for your records and reference.

Research Participant's Bill of Rights

The University of Waikato
School of Computing and Mathematical Sciences

The following is a list of your rights if you participate in a research project organised within the School of Computing and Mathematical Sciences at the University of Waikato.

As a research participant, you have the right:

- To be treated with respect and dignity in every phase of the research.
- To be fully and clearly informed of all aspects of the research prior to becoming involved in it.
- To enter into clear, informed, and written agreement with the researcher prior to becoming involved in the activity. You should sense no pressure, explicit or otherwise, to sign this contract.
- To choose explicitly whether or not you will become involved in the research under the clearly stated provision that refusal to participate or the choice to withdraw during the activity can be made at any time without penalty to you.
- To be treated with honesty, integrity, openness, and straightforwardness in all phases of the research, including a guarantee that you will not unknowingly be deceived during the course of the research.
- To receive something in return for your time and energy.
- To demand proof that an independent and competent ethical review of human rights and protections associated with the research has been successfully completed.
- To demand complete personal confidentiality and privacy in any reports of the research unless you have explicitly negotiated otherwise.
- To expect that your personal welfare is protected and promoted in all phases of the research, including knowing that no harm will come to you.
- To be informed of the results of the research study in a language you understand.
- To be offered a range of research studies or experiences from which to select, if the research is part of fulfilling your educational or employment goals.

The contents of this bill were prepared by the University of Calgary who examined all of the relevant Ethical Standards from the Canadian Psychological Association's Code of Ethics for Psychologists, 1991 and rewrote these to be of relevance to research participants.

The complete CPA Ethical Code can be found in: Canadian Psychological Association, Companion manual for the Canadian Code of Ethics for Psychologists (1992).

C.2 Recall Phase

The documents shown for the experiencing phase are:

- the approval letter from the Human Research Ethics Committee of the School of Computing and Mathematical Sciences at the University of Waikato, dated 10 February 2010;
- the Participant Information Sheet for the recall phase, which outlines the study goals and procedure as well as the participant's rights;
- the Consent Form for Participants, which each participant signed at the beginning of their session;
- the Participant's Workbook, which contains instructions, the practice tasks and instructions for the three main stages of the study as well as the biographical questionnaire; and
- a sample Question Sheet.

Computing and Mathematical Sciences
Rorohiko me ngā Pūtaiao Pāngarau
The University of Waikato
Private Bag 3105
Hamilton
New Zealand

Phone +64 7 838 4322
www.scms.waikato.ac.nz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

10 February 2010

Andrea Schweer
C/- Department of Computer Science
THE UNIVERSITY OF WAIKATO

Dear Andrea

Request for approval to conduct a research project study using human participants

I have considered your request for approval to conduct a research study as part of your PhD research on Context-Aware Augmented Memory – Recall Phase, the purpose being to gain insight into the usefulness of the Digital Parrot software system.

The procedure described in your request is acceptable. I note your statement that participants anonymity will be strictly maintained and that all information published about this study anonymised. Potentially sensitive information in the data used for experiements will be obscured in all publications. Video recordings, inlcuding excerpts, still pictures and audio, will never be published.

The research participants' instruction sheet and consent forms comply with the requirements of the University's human research ethics policies and procedures.

I therefore approve your application to conduct the study.

Yours sincerely

Dave Nichols
Human Research Ethics Committee
School of Computing and Mathematical Sciences

Participant Information Sheet
Ethics Committee, School of Computing and Mathematical Sciences

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Participant Information Sheet

Project Title

Context-Aware Augmented Memory: Recall Phase

Purpose

This research or related activity is conducted as part of my PhD research on Context-Aware Augmented Memory.

What is this research project about?

This research is to investigate approaches for augmenting autobiographical memory – that is, for improving people's memory of the events of their own lives.

The goal of this study is to gain insight into the usefulness of the Digital Parrot, a software system that I have developed; into the relative usefulness of its components, into the retrieval strategies used by the participants and into the participants' perception of the system.

What will you have to do and how long will it take?

You will first receive a thorough introduction to the system and its features, followed by some practice tasks.

You will then be provided with a list of questions. For each question, I will first ask you whether you think you can answer it straight away. If not, then this question will be put aside. Once this has been completed for all questions, I will ask you to go back to those questions that were put aside in the first phase. For each of these questions, I will first ask you whether you know the answer now (the reason behind this is that you may have gotten an idea of the answer from other questions/answers in the meantime). If you still don't know the answer, I will ask you to describe how you would normally go about finding the answer to this question. I will then ask you to use the Digital Parrot to attempt to answer the question.

Once the list of questions has been worked through, I will likely follow this up with some questions to clarify observations. I will then ask you to reflect on your experiences when using the Digital Parrot, with a particular focus on your opinions on the usefulness of the Digital Parrot and on the relative usefulness of its components.

The total duration of a session, including set-up, training and the exit interview, is expected to be no more than 2.5 hours. There will be opportunities to take breaks throughout the session.

I will take notes throughout the session. If you consent, I will also make a video recording of the session.

What will happen to the information collected?

In publications

All information published about this study will be anonymised. Any quotes from participants will only ever be published in forms that don't allow them to be linked back to an individual. Information about the participants' backgrounds (e.g., occupation, gender) will only be published in aggregated form. Potentially sensitive information in the data used for the experiments, such as the names of conferences attended by participants and the names of people with whom the participants interacted, will be obscured in all publications.

1

Participant Information Sheet
Ethics Committee, School of Computing and Mathematical Sciences

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Notes and transcripts

All paper material (observer's notes, participants' workbooks) will be deposited in the SCMS data archive once it has been transcribed into electronic form.

All electronic transcriptions of notes and questionnaire responses will be stored in a portion of the researcher's university file space, which is protected by a password. This data will be deleted from the researcher's file space once data analysis has been completed. A copy will be deposited in the SCMS data archive.

Video recordings

Video recordings will be made of those participants who consent to it. These video recordings will never be published (this includes excerpts, still pictures and audio taken from the recordings).

All video recordings will be stored on DVDs. Only the researcher and possibly the supervisors will have access to these recordings. The DVDs will be deposited in the SCMS data archive once data analysis has been completed.

Destruction

All information placed in the SCMS data archive will be marked with a destruction date of 30th June 2015, five years after the anticipated submission date of my PhD dissertation.

Declaration to participants

If you take part in the study, you have the right to:

- Refuse to answer any particular question, and to withdraw from the study within one week after the session.
- Ask any further questions about the study that occurs to you during your participation.
- Be given access to a summary of findings from the study when it is concluded.

Who's responsible?

If you have any questions or concerns about the project, either now or in the future, please feel free to contact either the researcher or the researcher's supervisors.

Researcher:
Andrea Schweer
schweer@cs.waikato.ac.nz

Supervisors:
Annika Hinze
hinze@cs.waikato.ac.nz

Steve Jones
stevej@cs.waikato.ac.nz

2

Research Consent Form
Ethics Committee, School of Computing and Mathematical Sciences

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Consent Form for Participants

I have read the **Participant Information Sheet** for this study and have had the details of the study explained to me. My questions about the study have been answered to my satisfaction, and I understand that I may ask further questions at any time.

I also understand that I am free to withdraw from the study within one week after the session, or to decline to answer any particular questions in the study. I understand I can withdraw any information I have provided up until the researcher has commenced analysis on my data. I agree to provide information to the researchers under the conditions of confidentiality set out on the **Participant Information Sheet**.

I agree to participate in this study under the conditions set out in the **Participant Information Sheet**.

Signed: ______________________________

Name: ______________________________

Date: ______________________________

I agree / do not agree to my session to be video recorded.

Signed: ______________________________

Name: ______________________________

Date: ______________________________

Researcher's name and contact information:
Andrea Schweer
schweer@cs.waikato.ac.nz

Supervisors' names and contact information:
Annika Hinze
hinze@cs.waikato.ac.nz
Steve Jones
stevej@cs.waikato.ac.nz

3

Participant's Workbook

Context-Aware Augmented Memory:
Recall Phase

Participant Number:

Study conducted by:
Andrea Schweer
Department of Computer Science
The University of Waikato
schweer@cs.waikato.ac.nz

Mai/June 2010

Instructions

First of all, thank you very much for agreeing to take part in this study!

Structure of today's study

1. An introduction to the Digital Parrot and its features. You'll also get some time to explore the system and there will be a few practice tasks.

2. The main study. I will give you a version of the Digital Parrot containing information from our interview(s) in 2008 and a list of related questions. I will ask you to try and answer these questions first without and then with the Digital Parrot.

   You will get more detailed instructions about this step later on and I will guide you through this part of the study.

3. A very brief questionnaire with some demographic questions.

4. Finally, I will invite you to reflect on your experiences when using the Digital Parrot. We'll focus in particular on your opinions on the usefulness of the Digital Parrot and on the relative usefulness of its components.

Again, thank you for helping me with my research. Turn the page when you are ready. Finally, please remember: this study is testing the Digital Parrot, not you! You're actually helping me *more* with everything you *don't* remember.

The Digital Parrot

The Digital Parrot is a repository of memories and of information associated with memories.

The version of the Digital Parrot that you will use for this part of the study already contains some memories. These are some of my memories, collected at several conferences that I attended.

The Digital Parrot lets you navigate among the stored memories and associated information. There are four tools: the timeline, the map, textual search and connections navigation. There are also two different main views: a graph view and a list view.

Demonstration

Please let me know when you're ready and I will demonstrate the Digital Parrot's features.

Exploration

After the demonstration, you'll have a few minutes to familiarise yourself with the Digital Parrot. Feel free to ask questions.

Practice tasks

Once you feel comfortable enough, please move on to the practice tasks. Please read each task carefully and then attempt to solve it using the Digital Parrot. Write the requested information into the space provided.

5

Practice task 1

To whom did I talk about scuba diving? Write **their name(s)** into the space below.

Practice task 2

Which conference(s) have I attended in Auckland? Write **the conference name(s)** into the space below.

6

Practice task 3

At which conference(s) did I speak to someone about Python during the poster session? Write **the conference name(s)** into the space below.

Practice task 4

In which place was the New Zealand HCI conference in 2007? Write **the place name** into the space below.

Conference memories

The version of the Digital Parrot that you will use now already contains some memories. I have transcribed these from what you told me when I interviewed you about your conference visit(s) in 2008. I am also providing you with some questions, one per sheet.

First phase

Please go through these questions one at a time in the order provided. For each question, please follow the workflow on the facing page. Once you have worked through all the questions, we will move on to the second phase.

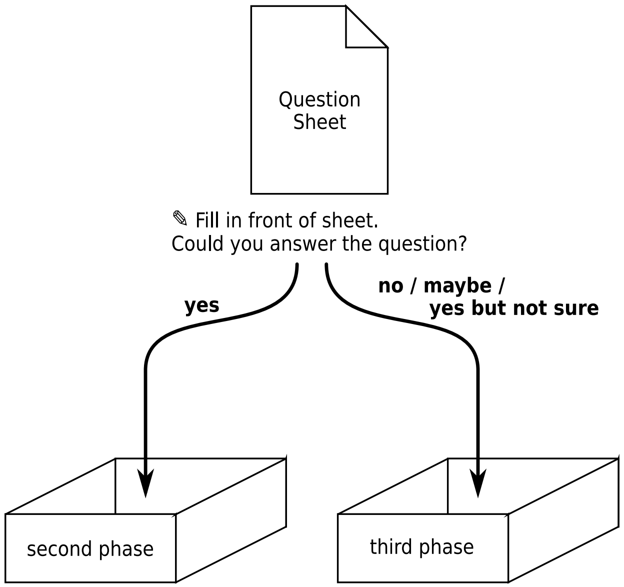

Second phase

We will now go back to the questions from the “second phase” box. Please go through these questions one at a time. For each question, please follow the workflow on the facing page. Once you have worked through all the questions, we will move on to the final phase.

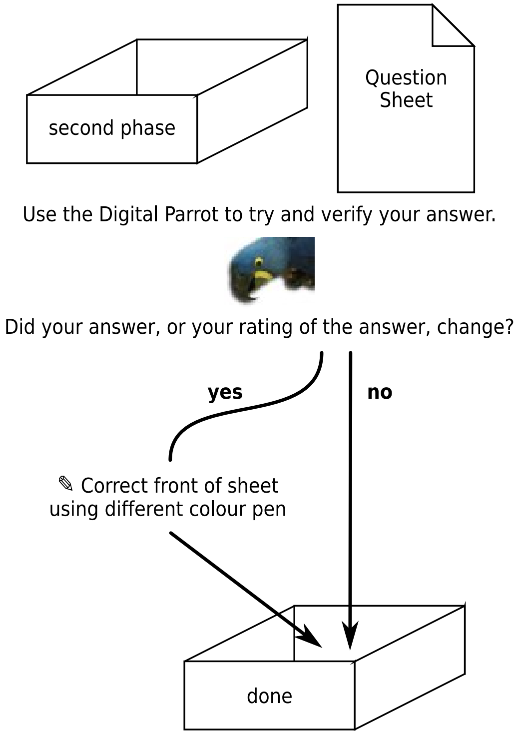

Third phase

We will now go back to the questions from the “third phase” box. Please go through these questions one at a time. For each question, please follow the workflow on the facing page.

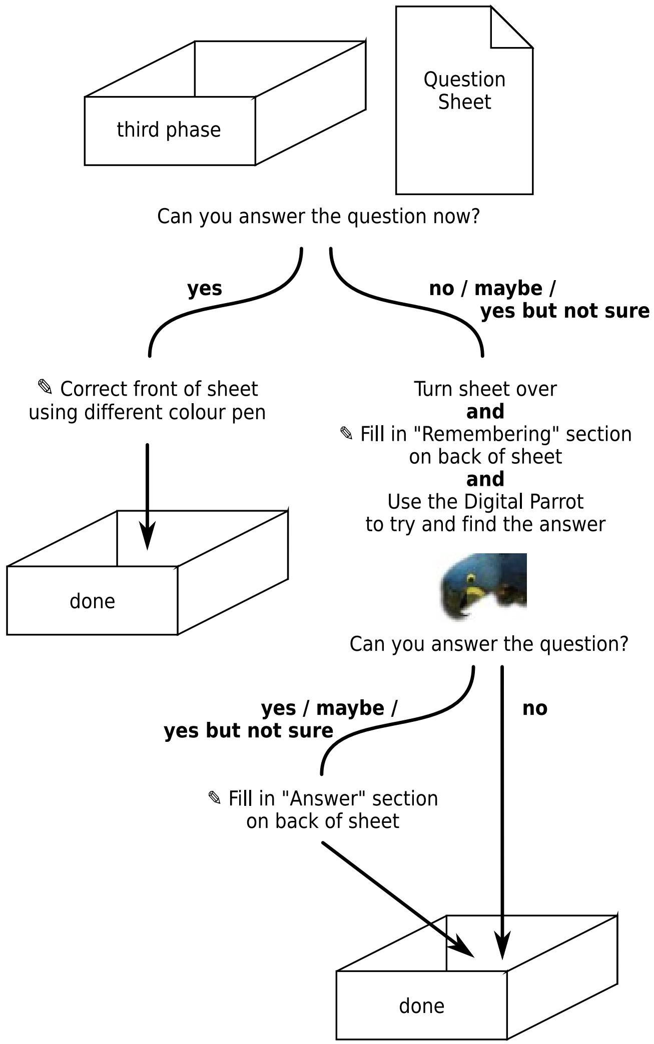

Some questions about you

Please circle your gender: female male

Please indicate your age:

What is your main occupation?

- ○ member of academic staff
- ○ PhD student
- ○ Other (please indicate):

Discussion

I’m likely to have some questions for you to follow up on my observations.

This is it – thank you very much for taking part in this study!

Digital Parrot User Study • Participant <participantID> • Question <ID>

Question

<text>

Answer

What would your answer be? Guess if you wish.

How certain are you that this answer is complete?

not certain at all ◌ ◌ ◌ ◌ ◌ ◌ *absolutely certain*

How certain are you that your (partial) answer is correct?

not certain at all ◌ ◌ ◌ ◌ ◌ ◌ *absolutely certain*

Would you be satisfied with this answer if it were important?

◌ no ◌ yes

Question

<text>

Remembering

How would you normally go about trying to find an answer?

Briefly describe to me (you don't need to write down anything) how you would normally go about trying to answer the question, assuming that it is important and without using the Digital Parrot.

What do you think are your chances of success this way?

not good at all ◌ ◌ ◌ ◌ ◌ ◌ *very good*

How quickly do you think you would find an answer this way?

Now please try to find an answer using the Digital Parrot.

Answer

What would your answer be? Guess if you wish.

How certain are you that this answer is complete?

not certain at all ◌ ◌ ◌ ◌ ◌ ◌ *absolutely certain*

How certain are you that your (partial) answer is correct?

not certain at all ◌ ◌ ◌ ◌ ◌ ◌ *absolutely certain*

Would you be satisfied with this answer if it were important?

◌ no ◌ yes

Bibliography

Adar, E., Karger, D., Stein, L. A. (1999). Haystack: Per-user information environments. In *CIKM '99: Proceedings of the 8th International Conference on Information and Knowledge Management (Kansas City, MO, USA, November 2–6, 1999)*, pp. 413–422. New York, NY, USA: ACM. doi:10.1145/319950.323231.

Ahmed, M., Hoang, H. H., Karim, S., Khusro, S., Lanzenberger, M., Latif, K., Michlmayr, E., Mustofa, K., Nguyen, H. T., Rauber, A., Schatten, A., Tho, M. N., Tjoa, A. M. (2004). SemanticLIFE – A framework for managing information of a human lifetime. In *iiWAS2004: Proceedings of 6th International Conference on Information Integration and Web-based Applications & Services (Jakarta, Indonesia, September 27–29, 2004)*, pp. 687–696. Vienna, Austria: OCG.

Aizawa, K., Hori, T., Kawasaki, S., Ishikawa, T. (2004a). Capture and efficient retrieval of life log. In *Proceedings of the Pervasive 2004 Workshop on Memory and Sharing of Experiences (Vienna, Austria, April 20, 2004)*, pp. 15–20. Retrieved from http://www.ii.ist.i.kyoto-u.ac.jp/~sumi/pervasive04/.

Aizawa, K., Tancharoen, D., Kawasaki, S., Yamasaki, T. (2004b). Efficient retrieval of life log based on context and content. In *CARPE '04: Proceedings of the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (New York, NY, USA, October 15, 2004)*, pp. 22–31. New York, NY, USA: ACM. doi:10.1145/1026653.1026656.

Allah, J. S. M. (2011). *Capturing Situational Context in an Augmented Memory System*. Master's thesis, Department of Computer Science, The University of Waikato, Hamilton, New Zealand. Under examination.

André, P., Wilson, M. L., Russell, A., Smith, D. A., Owens, A., schraefel, m. (2007). Continuum: Designing timelines for hierarchies, relationships and scale. In *UIST '07: Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology (Newport, Rhode Island, USA, October 7–10, 2007)*, pp. 101–110. New York, NY: ACM. doi:10.1145/1294211.1294229.

Bangor, A., Kortum, P., Miller, J. (2009). Determining what individual SUS scores mean: Adding an adjective rating scale. *Journal of Usability Studies*, *4*(3), 114–123. Retrieved from http://www.upassoc.org/upa_publications/jus/2009may/bangor1.html.

Barclay, C. R. (1986). Schematization of autobiographical memory. In Rubin, D. C. (Ed.), *Autobiographical Memory*, chap. 6, pp. 82–99. Cambridge, UK: Cambridge University Press. ISBN 0-521-36850-2.

Barreau, D. K. (1995). Context as a factor in personal information management systems. *Journal of the American Society for Information Science and Technology*, *46*(5), 327–339. doi:10.1002/(SICI)1097-4571(199506)46:5<327::AID-ASI4>3.0.CO;2-C.

Barsalou, L. W. (1988). The content and organization of autobiographical memories. In Neisser, U., Winograd, E. (Eds.), *Remembering Reconsidered*, vol. 2 of *Emory Symposia in Cognition*, chap. 8, pp. 193–243. Cambridge University Press, 1st paperback edn. ISBN 0-521-48500-2.

Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F., Stein, L. A. (2004). OWL web ontology language reference. W3C Recommendation 10 February 2004. Retrieved from http://www.w3.org/TR/2004/REC-owl-ref-20040210/.

Bell, G. (2001). A personal digital store. *Communications of the ACM*, *44*(1), 86–91. doi:10.1145/357489.357513.

Bergman, O., Beyth-Marom, R., Nachmias, R., Gradovitch, N., Whittaker, S. (2008). Improved search engines and navigation preference in personal information management. *ACM Transactions on Information Systems, 26*(4), [Online, article no. 20]. doi:10.1145/1402256.1402259.

Berners-Lee, T., Connolly, D. (2008). Notation 3 (N3): A readable RDF syntax. W3C Team Submission 14 January 2008. Retrieved from http://www.w3.org/TeamSubmission/2008/SUBM-n3-20080114/.

Bernstein, M., Van Kleek, M., Karger, D., schraefel, m. (2008a). Information scraps: How and why information eludes our personal information management tools. *ACM Transactions on Information Systems, 26*(4), [Online, article no. 24]. doi:10.1145/1402256.1402263.

Bernstein, M., Van Kleek, M., schraefel, m., Karger, D. R. (2008b). Evolution and evaluation of an information scrap manager. Paper presented at PIM 2008: CHI 2008 Workshop on Personal Information Management (Florence, Italy, April 5–6, 2008). Retrieved from http://pim2008.ethz.ch/acceptedpapers.php.

Bloch, D. (2007). *Aristotle on Memory and Recollection: Text, Translation, Interpretation, and Reception in Western Scholasticism*, vol. 110 of *Philosophia Antiqua*. Leiden, The Netherlands: Brill Academic Publishers. ISBN 9789047421436.

Blum, M., Pentland, A. S., Tröster, G. (2006). InSense: Interest-based life logging. *IEEE MultiMedia, 13*(4), 40–48.

Blunschi, L., Dittrich, J.-P., Girard, O. R., Karakashian, S. K., Salles, M. A. V. (2007). A dataspace odyssey: The iMeMex personal dataspace management system. In *CIDR 2007: Proceedings of the 3rd Biennial Conference on Innovative Data Systems Research (Asilomar, CA, USA, January 7–10, 2007)*, pp. 114–119. Retrieved from http://www.cidrdb.org/cidr2007/.

Brewer, W. F. (1986). What is autobiographical memory? In Rubin, D. C. (Ed.), *Autobiographical Memory*, chap. 3, pp. 25–49. Cambridge, UK: Cambridge University Press. ISBN 0-521-36850-2.

Brooke, J. (1996). SUS: A "quick and dirty" usability scale. In Jordan, P. W., Thomas, B., McClelland, I. L., Weerdmeester, B. A. (Eds.), *Usability evaluation in industry*, chap. 21, pp. 189–194. CRC Press. ISBN 0748404600.

Brown, S. C., Craik, F. I. M. (2000). Encoding and retrieval of information. In Tulving, E., Craik, F. I. M. (Eds.), *The Oxford Handbook of Memory*, chap. 6, pp. 93–107. Oxford, UK: Oxford University Press. ISBN 0-19-512265-8.

Bush, V. (1945). As we may think. *The Atlantic Monthly*, *176*(1), 101–108.

Byrne, D., Doherty, A. R., Snoek, C. G. M., Jones, G. J. F., Smeaton, A. F. (2009). Everyday concept detection in visual lifelogs: validation, relationships and trends. *Multimedia Tools and Applications*, *49*(1), 119–144. doi:10.1007/s11042-009-0403-8.

Cai, Y., Dong, X. L., Halevy, A., Liu, J. M., Madhavan, J. (2005). Personal information management with SEMEX. In *SIGMOD '05: Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data (Baltimore, MD, USA, June 13–16, 2005)*, pp. 921–923. New York, NY, USA: ACM. doi:10.1145/1066157.1066289.

Capra, R. G., Pérez-Quiñones, M. A. (2003). Re-finding found things: An exploratory study of how users re-find information. arXiv preprint cs/0310011v1. Retrieved from http://arxiv.org/abs/cs.HC/0310011.

Capra, R. G., Pérez-Quiñones, M. A. (2005). Using web search engines to find and refind information. *Computer*, *38*(10), 36–42. doi:10.1109/MC.2005.355.

Catarci, T., Dong, L., Halevy, A., Poggi, A. (2007). Structure everything. In Jones and Teevan (2007b), chap. 7, pp. 108–126.

Chau, D. H., Myers, B., Faulring, A. (2008a). Feldspar: A system for finding information by association. Paper presented at PIM 2008: CHI 2008 Workshop on Personal Information Management (Florence, Italy, April

5–6, 2008). Retrieved from http://pim2008.ethz.ch/acceptedpapers.php.

Chau, D. H., Myers, B., Faulring, A. (2008b). What to do when search fails: Finding information by association. In *CHI '08: Proceeding of the 26th Annual SIGCHI Conference on Human Factors in Computing Systems (Florence, Italy, April 5–10, 2008)*, pp. 999–1008. New York, NY, USA: ACM. doi:10.1145/1357054.1357208.

Clarkson, B., Mase, K., Pentland, A. (2001). The Familiar: A living diary and companion. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems (Seattle, WA, USA, March 31 – April 5, 2001)*, pp. 271–272. New York, NY, USA: ACM. doi:10.1145/634067.634228.

Cohen, G. (1996). *Memory in the real world*. Hove, UK: Psychology Press, 2nd edn. ISBN 0-86377-729-5.

Conway, M., Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological Review*, *107*(2), 261–288. doi:10.1037/0033-295X.107.2.261.

Conway, M. A., Pleydell-Pearce, C. W., Whitecross, S. E. (2001). The neuroanatomy of autobiographical memory: A slow cortical potential study of autobiographical memory retrieval. *Journal of Memory and Language*, *45*(3), 493–524. doi:10.1006/jmla.2001.2781.

Conway, M. A., Rubin, D. C. (1993). The structure of autobiographical memory. In Collins, A. F., Gathercole, S. E., Conway, M. A., Morris, P. E. (Eds.), *Theories of Memory*, vol. 1, chap. 4, pp. 103–137. Hillsdale, NJ, USA: Erlbaum. ISBN 978-0-86377-290-0.

Cooper, A. (2004). *The Inmates are Running the Asylum: Why High Tech Products Drive Us Crazy and How to Restore the Sanity*. Indianapolis, IN, USA: Sams Publishing. ISBN 0-672-326141-0.

Cutrell, E. (2006). Search user interfaces for PIM. Paper presented at PIM 2006: SIGIR 2006 Workshop on Personal Information Management (Seattle, WA, USA, August 10–11, 2006). Retrieved from http://pim.ischool.washington.edu/pim06/posters.php.

Cutrell, E., Dumais, S. T., Teevan, J. (2006a). Searching to eliminate personal information management. *Communications of the ACM, 49*(1), 58–64. doi:10.1145/1107458.1107492.

Cutrell, E., Robbins, D., Dumais, S., Sarin, R. (2006b). Fast, flexible filtering with phlat. In *CHI '06: Proceedings of the 24th Annual SIGCHI Conference on Human Factors in Computing Systems (Montréal, Canada, April 22–27, 2006)*, pp. 261–270. New York, NY, USA: ACM. doi:10.1145/1124772.1124812.

Dey, A. K., Salber, D., Abowd, G. D., Futakawa, M. (1999). The conference assistant: Combining context-awareness with wearable computing. In *ISWC 1999: Proceedings of the 3rd International Symposium on Wearable Computers (San Francisco, CA, USA, October 18–19, 1999)*, pp. 21–28. Los Alamitos, CA, USA: IEEE Computer Society. doi:10.1109/ISWC.1999.806639.

Dickie, C., Vertegaal, R., Fono, D., Sohn, C., Chen, D., Cheng, D., Shell, J. S., Aoudeh, O. (2004). Augmenting and sharing memory with eyeBlog. In *CARPE '04: Proceedings of the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (New York, NY, USA, October 15, 2004)*, pp. 105–109. New York, NY, USA: ACM. doi:10.1145/1026653.1026670.

Dittrich, J., Salles, M. A. V., Blunschi, L. (2008). Managing personal information using iTrails. Paper presented at PIM 2008: CHI 2008 Workshop on Personal Information Management (Florence, Italy, April 5–6, 2008). Retrieved from http://pim2008.ethz.ch/acceptedpapers.php.

Dittrich, J., Salles, M. A. V., Blunschi, L. (2009). iMeMex: From search to information integration and back. *IEEE Data Engineering Bulletin, 32*(2), 28–35. Retrieved from http://sites.computer.org/debull/A09june/issue1.htm.

Dittrich, J.-P., Salles, M. A. V., Kossmann, D., Blunschi, L. (2005). iMeMex: Escapes from the personal information jungle. In *VLDB 2005: Proceedings of the 31st International Conference on Very Large Data*

Bases (Trondheim, Norway, August 30 – September 2, 2005), pp. 1306–1309. VLDB Endowment. ISBN 1-59593-154-6. Retrieved from http://portal.acm.org/citation.cfm?id=1083759.

Dong, X., Halevy, A. (2005). A platform for personal information management and integration. In *CIDR '05: Proceedings of the 2nd Biennial Conference on Innovative Data Systems Research (Asilomar, CA, USA, January 4–7, 2005)*, pp. 119–130. Retrieved from http://www.cidrdb.org/cidr2005/.

Dong, X., Halevy, A., Nemes, E., Sigundsson, S. B., Domingos, P. (2004). SEMEX: Toward on-the-fly personal information integration. In *IIWeb-2004: Proceedings of the 2004 VLDB Workshop on Information Integration on the Web (Toronto, Canada, August 30, 2004)*. Retrieved from http://data.cs.washington.edu/papers/semex_iiweb.pdf.

Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R., Robbins, D. C. (2003). Stuff I've Seen: A system for personal information retrieval and re-use. In *SIGIR '03: Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (Toronto, Canada, July 28 – August 1, 2003)*, pp. 72–79. New York, NY, USA: ACM. doi:10.1145/860435.860451.

Elsweiler, D., Baillie, M., Ruthven, I. (2008). Exploring memory in email refinding. *ACM Transactions on Information Systems*, *26*(4), [Online, article no. 21]. doi:10.1145/1402256.1402260.

Elsweiler, D., Baillie, M., Ruthven, I. (2009). On understanding the relationship between recollection and refinding. *Journal of Digital Information*, *10*(5), [Online]. Retrieved from http://journals.tdl.org/jodi/article/view/436.

Elsweiler, D., Ruthven, I. (2007). Towards task-based personal information management evaluations. In *SIGIR '07: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 23–30. New York, NY, USA: ACM. doi:10.1145/1277741.1277748.

Elsweiler, D., Ruthven, I., Jones, C. (2005). Dealing with fragmented recollection of context in information management. In *CIR-05: Proceedings of the Context-Based Information Retrieval Workshop at CONTEXT-05 (Paris, France, July 5, 2005)*, no. 151 in CEUR Workshop Proceedings. Retrieved from http://www.ceur-ws.org/Vol-151/.

Elsweiler, D., Ruthven, I., Jones, C. (2007). Towards memory supporting personal information management tools. *Journal of the American Society for Information Science and Technology*, *58*(7), 924–946. doi: 10.1002/asi.20570.

Eysenck, M. W., Keane, M. T. (2005). *Cognitive Psychology: A Student's Handbook*. Hove, UK: Psychology Press, 5th edn. ISBN 1-84169-359-6.

Fertig, S., Freeman, E., Gelernter, D. (1996a). "Finding and reminding" reconsidered. *SIGCHI Bulletin*, *28*(1), 66–69. doi:10.1145/249170.249187.

Fertig, S., Freeman, E., Gelernter, D. (1996b). Lifestreams: An alternative to the desktop metaphor. In *CHI '96: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Vancouver, Canada, April 13–18, 1996)*, pp. 410–411. New York, NY, USA: ACM. doi:10.1145/257089.257404.

Freeman, E., Fertig, S. (1995). Lifestreams: Organizing your electronic life. In *Papers from the AAAI 1995 Fall Symposium: AI Applications in Knowledge Navigation and Retrieval*, pp. 38–44. Menlo Park, CA, USA: AAAI Press.

Furnas, G. W. (2006). A fisheye follow-up: Further reflections on focus + context. In *CHI '06: Proceedings of the 24th Annual SIGCHI Conference on Human Factors in Computing Systems (Montréal, Canada, April 22–27, 2006)*, pp. 999–1008. New York, NY, USA: ACM. doi: 10.1145/1124772.1124921.

Gardiner, J. M., Richardson-Klavehn, A. (2000). Remembering and knowing. In Tulving, E., Craik, F. I. M. (Eds.), *The Oxford Handbook of*

Memory, chap. 15, pp. 229–244. Oxford, UK: Oxford University Press. ISBN 0-19-512265-8.

Gemmell, J., Aris, A., Lueder, R. (2005). Telling stories with MyLifeBits. In *ICME 2005: Proceedings of the IEEE International Conference on Multimedia and Expo (Amsterdam, The Netherlands, July 6–9, 2005)*, pp. 1536–1539. Los Alamitos, CA, USA: IEEE Computer Society. doi: 10.1109/ICME.2005.1521726.

Gemmell, J., Bell, G., Lueder, R. (2006). MyLifeBits: A personal database for everything. *Communications of the ACM*, *49*(1), 88–95. doi:10.1145/1107458.1107460.

Gemmell, J., Bell, G., Lueder, R., Drucker, S., Wong, C. (2002). MyLifeBits: Fulfilling the Memex vision. In *MULTIMEDIA '02: Proceedings of the 10th ACM international conference on Multimedia (Juan-les-Pins, France, December 1–6, 2002)*, pp. 235–238. New York, NY, USA: ACM. doi:10.1145/641007.641053.

Gemmell, J., Sundaram, H. (2006). Guest editors' introduction: Capture, archival, and retrieval of personal experience. *IEEE MultiMedia*, *13*(4), 18–19. doi:10.1109/MMUL.2006.82.

Gwizdka, J. (2006). Finding to keep and organize: Personal information collections as context. Paper presented at PIM 2006: SIGIR 2006 Workshop on Personal Information Management (Seattle, WA, USA, August 10–11, 2006). Retrieved from http://pim.ischool.washington.edu/pim06/posters.php.

Ham, F. V., Wattenberg, M. (2008). Centrality based visualization of small world graphs. *Computer Graphics Forum*, *27*(3), 975–982. doi: 10.1111/j.1467-8659.2008.01232.x.

Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., Wood, K. (2006). SenseCam: A retrospective memory aid. In *UbiComp 2006: Proceedings of the 8th International Conference on Ubiquitous Computing (Orange County, CA,* *USA, September 17–21, 2006)*, vol. 4206 of *Lecture Notes in Computer Science*, pp. 177–193. Berlin/Heidelberg, Germany: Springer. doi:10.1007/11853565_11.

van den Hoven, E. (2004). Exploring graspable cues for everyday recollecting. In *Proceedings of the Pervasive 2004 Workshop on Memory and Sharing of Experiences (Vienna, Austria, April 20, 2004)*, pp. 21–24. Retrieved from http://www.ii.ist.i.kyoto-u.ac.jp/~sumi/pervasive04/.

van den Hoven, E., Eggen, B. (2007). Informing augmented memory system design through autobiographical memory theory. *Personal and Ubiquitous Computing*, *12*(6), 433–443. doi:10.1007/s00779-007-0177-9.

van den Hoven, E., Eggen, B. (2009). The effect of cue media on recollections. *Human Technology*, *5*(1), 47–67. Retrieved from http://www.humantechnology.jyu.fi/.

Huynh, D., Karger, D., Quan, D. (2002). Haystack: A platform for creating, organizing and visualizing information using RDF. In *Proceedings of the International Workshop on the Semantic Web at WWW2002 (Honolulu, HI, USA, May 7, 2002)*. Retrieved from http://semanticweb2002.aifb.uni-karlsruhe.de/programme.htm.

Jain, R. (2003). Experiential computing. *Communications of the ACM*, *46*(7), 48–55. doi:10.1145/792704.792729.

Jones, W. (2007a). How people keep and organize personal information. In Jones and Teevan (2007b), chap. 3, pp. 35–56.

Jones, W. (2007b). Personal information management. *Annual Review of Information Science and Technology*, *41*(1), 453–504. doi:10.1002/aris.2007.1440410117.

Jones, W., Teevan, J. (2007a). Introduction. In Jones and Teevan (2007b), chap. 1, pp. 3–20.

Jones, W., Teevan, J. (Eds.) (2007b). *Personal Information Management*. Seattle, WA, USA: University of Washington Press. ISBN 978-0-295-98737-8.

Kalnikaitė, V., Sellen, A., Whittaker, S., Kirk, D. (2010). Now let me see where I was: Understanding how lifelogs mediate memory. In *CHI '10: Proceedings of the 28th International Conference on Human Factors in Computing Systems (Atlanta, GA, USA, April 10–15, 2010)*, pp. 2045–2054. New York, NY, USA: ACM. doi:10.1145/1753326.1753638.

Kalnikaitė, V., Whittaker, S. (2007). Software or wetware? Discovering when and why people use digital prosthetic memory. In *CHI '07: Proceedings of the 25th Annual SIGCHI Conference on Human Factors in Computing Systems (San Jose, CA, USA, April 28 – May 3, 2007)*, pp. 71–80. New York, NY, USA: ACM. doi:10.1145/1240624.1240635.

Kalnikaitė, V., Whittaker, S. (2008). Cueing digital memory: How and why do digital notes help us remember? In *Proceedings of the 22nd British HCI Group Annual Conference on HCI 2008 (Liverpool, UK, September 1–5, 2008)*, pp. 153–161. Swinton, UK: British Computer Society. Retrieved from http://portal.acm.org/citation.cfm?id=1531536.

Karger, D. R. (2007). Unify everything: It's all the same to me. In Jones and Teevan (2007b), chap. 8, pp. 127–152.

Karger, D. R., Bakshi, K., Huynh, D., Quan, D., Sinha, V. (2005). Haystack: A general purpose information management tool for end users of semistructured data. In *CIDR '05: Proceedings of the 2nd Biennial Conference on Innovative Data Systems Research (Asilomar, CA, USA, January 4–7, 2005)*, pp. 13–27. Retrieved from http://www.cidrdb.org/cidr2005/.

Kelly, D. (2006). Evaluating personal information management behaviors and tools. *Communications of the ACM*, *49*(1), 84–86. doi:10.1145/1107458.1107497.

Kelly, D., Teevan, J. (2007). Understanding what works: Evaluating PIM tools. In Jones and Teevan (2007b), chap. 11, pp. 190–204.

Kim, P., Gargi, U., Jain, R. (2005). Event-based multimedia chronicling systems. In *CARPE '05: Proceedings of the 2nd ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (Singapore, November 11, 2005)*, pp. 1–12. New York, NY, USA: ACM. doi:10.1145/1099083.1099085.

Kim, P., Podlaseck, M., Pingali, G. (2004). Personal chronicling tools for enhancing information archival and collaboration in enterprises. In *CARPE '04: Proceedings of the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (New York, NY, USA, October 15, 2004)*, pp. 56–65. New York, NY, USA: ACM. doi:10.1145/1026653.1026662.

Kolodner, J. L. (1983). Reconstructive memory: A computer model. *Cognitive Science*, *7*(4), 281–328.

Krishnan, A., Jones, S. (2005). TimeSpace: Activity-based temporal visualisation of personal information spaces. *Personal and Ubiquitous Computing*, *9*(1), 46–65. doi:10.1007/s00779-004-0291-x.

Lamming, M., Brown, P., Carter, K., Eldridge, M., Flynn, M., Louie, G., Robinson, P., Sellen, A. (1994). The design of a human memory prosthesis. *The Computer Journal*, *37*(3), 153–163. doi:10.1093/comjnl/37.3.153.

Lamming, M., Flynn, M. (1994). "Forget-me-not": Intimate computing in support of human memory. In *Proceedings of FRIEND21, 1994 International Symposium on Next Generation Human Interface (Tokyo, Japan, February 2–4, 1994)*. Retrieved from http://www.lamming.com/mik/Papers/fmn.pdf.

Latif, K., Mustofa, K., Tjoa, A. M. (2006). An approach for a personal information management system for photos of a lifetime by exploiting semantics. In *DEXA '06: Proceedings of the 17th International Conference on Database and Expert Systems Applications (Krakow, Poland, September 4–8, 2006)*, no. 4080 in Lecture Notes in Computer Science, pp. 467–477. Berlin/Heidelberg, Germany: Springer. doi:10.1007/11827405_46.

Latif, K., Tjoa, A. M. (2006). Combining context ontology and landmarks for personal information management. In *ICOCI 2006: Proceedings of the International Conference on Computing and Informatics (Kuala Lumpur, Malaysia, June 6–8, 2006)*. Los Alamitos, CA, USA: IEEE Press. doi:10.1109/ICOCI.2006.5276629.

Linton, M. (1986). Ways of searching and the contents of memory. In Rubin, D. C. (Ed.), *Autobiographical Memory*, chap. 4, pp. 50–67. Cambridge, UK: Cambridge University Press. ISBN 0-521-36850-2.

Mustofa, K., Tjoa, A. M. (2006). Building ontology-based annotations for personal information management systems – The SemanticLIFE design and implementation case. In *ICOCI 2006: Proceedings of the International Conference on Computing and Informatics (Kuala Lumpur, Malaysia, June 6–8, 2006)*. Los Alamitos, CA, USA: IEEE Press. doi:10.1109/ICOCI.2006.5276630.

Norman, D. A. (1992). The Teddy. In *Turn signals are the facial expression of automobiles*, chap. 6. Addison Wesley. Out of print. Retrieved from http://jnd.org/dn.mss/chapter_6_the_teddy.html.

Peesapati, S. T., Schwanda, V., Schultz, J., Lepage, M., Jeong, S.-y., Cosley, D. (2010). Pensieve: Supporting everyday reminiscence. In *CHI '10: Proceedings of the 28th International Conference on Human Factors in Computing Systems (Atlanta, GA, USA, April 10–15, 2010)*, pp. 2027–2036. New York, NY, USA: ACM. doi:10.1145/1753326.1753635.

Petrelli, D., van den Hoven, E., Whittaker, S. (2009). Making history: Intentional capture of future memories. In *PIM 2009: Proceedings of the ASIS&T 2009 Workshop on Personal Information Management (Vancouver, Canada, November 7–8, 2009)*, pp. 1723–1732. New York, NY, USA: ACM. doi:10.1145/1518701.1518966.

Petrelli, D., Whittaker, S. (2010). Family memories in the home: Contrasting physical and digital mementos. *Personal and Ubiquitous Computing*, *14*(2), 153–169. doi:10.1007/s00779-009-0279-7.

Pinzon, J. C. C., Singh, R. (2005). Designing an experiential annotation system for personal multimedia information management. In *IASTED-HCI: Proceedings of the IASTED International Conference on Human-Computer Interaction (Phoenix, AZ, USA, November 14–16, 2005)*, pp. 234–239. Calgary, Canada: ACTA Press.

Reiser, B. J., Black, J. B., Kalamarides, P. (1986). Strategic memory search processes. In Rubin, D. C. (Ed.), *Autobiographical Memory*, chap. 7, pp. 100–121. Cambridge, UK: Cambridge University Press. ISBN 0-521-36850-2.

Rhodes, B. (1997). The wearable remembrance agent: A system for augmented memory. *Personal and Ubiquitous Computing*, *1*(4), 218–224. doi:10.1007/BF01682024.

Rhodes, B. (2003). Using physical context for just-in-time information retrieval. *IEEE Transactions on Computers*, *52*(8), 1011–1014. doi:10.1109/TC.2003.1223636.

Rhodes, B. J., Starner, T. (1996). Remembrance agent: A continuously running automated information retrieval system. In *Papers from the 1996 AAAI Spring Symposium: Acquisition, Learning and Demonstration: Automating Tasks for Users*, pp. 487–495. Menlo Park, CA, USA: AAAI Press.

Ringel, M., Cutrell, E., Dumais, S., Horvitz, E. (2003). Milestones in time: The value of landmarks in retrieving information from personal stores. In *Interact 2003: Proceedings of the 9th IFIP TC13 International Conference on Human-Computer Interaction (Zurich, Switzerland, September 1–5, 2003)*, pp. 184–191. Amsterdam, The Netherlands: IOS Press. ISBN 978-1-58603-363-7.

Russell, D. M., Lawrence, S. (2007). Search everything. In Jones and Teevan (2007b), chap. 9, pp. 153–166.

Sauermann, L., Grimnes, G. A., Kiesel, M., Fluit, C., Maus, H., Heim, D., Nadeem, D., Horak, B., Dengel, A. (2006). Semantic desktop 2.0: The gnowsis experience. In *ISWC 2006: Proceedings of the 5th International Semantic Web Conference (Athens, GA, USA, November 5–9, 2006)*, vol. 4273 of *Lecture Notes in Computer Science*, pp. 887–900. Berlin/Heidelberg, Germany: Springer. doi:10.1007/11926078_64.

Sauermann, L., Heim, D. (2010). Evaluating long-term use of the Gnowsis semantic desktop for PIM. In *Proceedings of the 7th international Conference on the Semantic Web (Karlsruhe, Germany, October 26–30, 2008)*, vol. 5318 of *Lecture Notes in Computer Science*, pp. 467–482. Berlin/Heidelberg, Germany: Springer. doi:10.1007/978-3-540-88564-1_30.

Schacter, D. L. (1999). The seven sins of memory: Insights from psychology and cognitive neuroscience. *American Psychologist*, *54*(3), 182–203. doi:10.1037/0003-066X.54.3.182.

Schweer, A., Hinze, A. (2007a). Augmented memory for conference attendees. In *Proceedings of NZCSRSC '07, the Fifth New Zealand Computer Science Research Student Conference*.

Schweer, A., Hinze, A. (2007b). Challenges in interface and interaction design for context-aware augmented memory systems. In *CHINZ 2007: Proceedings of the 8th Annual Conference of the NZ ACM Special Interest Group on Human-Computer Interaction (Hamilton, New Zealand, July 01–04, 2007*. New York, NY, USA: ACM. doi:10.1145/1278960.1278968.

Schweer, A., Hinze, A. (2007c). The Digital Parrot: Combining context-awareness and semantics to augment memory. Paper presented at MeMos, the Workshop on Supporting Human Memory with Interactive Systems, 21st British HCI Group Annual Conference on HCI 2008 (Lancaster, UK, September 4, 2007). Retrieved from http://diuf.unifr.ch/people/lalanned/MeMos07/files/MeMos07.pdf.

Schweer, A., Jones, S., Hinze, A. (2009). Trails of experiences: Navigating personal memories. In *CHINZ 2009: Proceedings of the 10th Annual Conference of the NZ ACM Special Interest Group on Human-Computer Interaction (Auckland, New Zealand, July 6–7, 2009)*, pp. 105–106. New York, NY, USA: ACM. doi:10.1145/1577782.1577801.

Sellen, A. J., Fogg, A., Aitken, M., Hodges, S., Rother, C., Wood, K. (2007). Do life-logging technologies support memory for the past? An experimental study using SenseCam. In *CHI '07: Proceedings of the 25th Annual SIGCHI Conference on Human Factors in Computing Systems (San Jose, CA, USA, April 28 – May 3, 2007)*, pp. 81–90. New York, NY, USA: ACM. doi:10.1145/1240624.1240636.

Sellen, A. J., Whittaker, S. (2010). Beyond total capture: A constructive critique of lifelogging. *Communications of the ACM*, *53*(5), 70–77. doi: 10.1145/1735223.1735243.

de Silva, G. C., Yamasaki, T., Aizawa, K. (2006). Interactive experience retrieval for a ubiquitous home. In *CARPE '06: Proceedings of the 3rd ACM workshop on Continuous Archival and Retrieval of Personal Experiences (Santa Barbara, CA, USA, October 27, 2006)*, pp. 45–48. New York, NY, USA: ACM. doi:10.1145/1178657.1178667.

Singh, R., Knickmeyer, R., Gupta, P., Jain, R. (2004). Designing experiential environments for management of personal multimedia. In *MULTIMEDIA '04: Proceedings of the 12th annual ACM international conference on Multimedia (New York, NY, USA, October 10–16, 2004)*, pp. 496–499. New York, NY, USA: ACM. doi:10.1145/1027527.1027647.

Ståhl, A., Höök, K., Svensson, M., Taylor, A. S., Combetto, M. (2009). Experiencing the affective diary. *Personal and Ubiquitous Computing*, *13*(5), 365–378. doi:10.1007/s00779-008-0202-7.

Teevan, J., Alvarado, C., Ackerman, M. S., Karger, D. R. (2004). The perfect search engine is not enough: A study of orienteering behavior in directed search. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems (Vienna, Austria, April 24–29,* *2004)*, pp. 415–422. New York, NY, USA: ACM. doi:10.1145/985692.985745.

Truong, K. N., Hayes, G. R. (2009). Ubiquitous computing for capture and access. *Foundations and Trends in Human-Computer Interaction*, *2*(2), 95–171. doi:10.1561/1100000014.

Van Kleek, M., Shrobe, H. E. (2007). A practical activity capture framework for personal, lifetime user modeling. In *Proceedings of the 11th International Conference on User Modeling (Corfu, Greece, July 25–29, 2007)*, vol. 4511 of *Lecture Notes in Artificial Intelligence*, pp. 298–302. Berlin/Heidelberg, Germany: Springer. doi:10.1007/978-3-540-73078-1_33.

Veith, R. H. (2006). Memex at 60: Internet or iPod? *Journal of the American Society for Information Science and Technology*, *57*(9), 1233–1242. doi:10.1002/asi.20415.

Vemuri, S. (2004). *Personal long-term memory aids*. Ph.D. thesis, MIT Media Lab, Massachusetts Institute of Technology, Cambridge, MA, USA. Retrieved from http://hdl.handle.net/1721.1/30242.

Vemuri, S., Bender, W. (2004). Next-generation personal memory aids. *BT Technology Journal*, *22*(4), 125–138. doi:10.1023/B:BTTJ.0000047591.29175.89.

Vemuri, S., Schmandt, C., Bender, W. (2006). iRemember: A personal, long-term memory prosthesis. In *CARPE '06: Proceedings of the 3rd ACM workshop on Continuous Archival and Retrieval of Personal Experiences (Santa Barbara, CA, USA, October 27, 2006)*, pp. 65–74. New York, NY, USA: ACM. doi:10.1145/1178657.1178670.

Vemuri, S., Schmandt, C., Bender, W., Tellex, S., Lassey, B. (2004). An audio-based personal memory aid. In *UbiComp 2004: Proceedings of the 6th International Conference on Ubiquitous Computing (Nottingham, UK, September 7–10, 2004)*, vol. 3205 of *Lecture Notes in Computer Science*, pp. 400–417. Berlin/Heidelberg, Germany: Springer. doi:10.1007/978-3-540-30119-6_24.

Wattenberg, M. (2006). Visual exploration of multivariate graphs. In *CHI '06: Proceedings of the 24th Annual SIGCHI Conference on Human Factors in Computing Systems (Montréal, Canada, April 22–27, 2006)*, pp. 811–819. New York, NY, USA: ACM. doi:10.1145/1124772.1124891.

Weiser, M. (1991). The computer for the 21st century. *Scientific American*, *265*(3), 94–110.

Yeh, R., Liao, C., Klemmer, S., Guimbretière, F., Lee, B., Kakaradov, B., Stamberger, J., Paepcke, A. (2006). ButterflyNet: A mobile capture and access system for field biology research. In *CHI '06: Proceedings of the 24th Annual SIGCHI Conference on Human Factors in Computing Systems (Montréal, Canada, April 22–27, 2006)*, pp. 571–580. New York, NY, USA: ACM. doi:10.1145/1124772.1124859.

Yi, J. S., Melton, R., Stasko, J., Jacko, J. A. (2005). Dust & magnet: multivariate information visualization using a magnet metaphor. *Information Visualization*, *4*(4), 239–256. doi:10.1057/palgrave.ivs.9500099.