Human–Computer Interaction

Socio-Technical Lifelogging: Deriving Design Principles for a Future Proof Digital Past

Steve Whittaker [a] , Vaiva Kalnikaitė [b] , Daniela Petrelli [c] , Abigail
Sellen [d] , Nicolas Villar [d] , Ofer Bergman [e] , Paul Clough [f] & Jens
Brockmeier [g] [h]

[a] University of California at Santa Cruz

[b] Interactables, Cambridge, United Kingdom

[c] Sheffield Hallam University, United Kingdom

[d] Microsoft Research Cambridge, United Kingdom

[e] Bar Ilan University, Israel

[f] Sheffield University, United Kingdom

[g] Free University Berlin, Germany

[h] University of Manitoba, Canada

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis
Taylor & Francis Group

Socio-Technical Lifelogging: Deriving Design Principles for a Future Proof Digital Past

Steve Whittaker,[1] Vaiva Kalnikaitė,[2] Daniela Petrelli,[3]
Abigail Sellen,[4] Nicolas Villar,[4] Ofer Bergman,[5] Paul Clough,[6]
and Jens Brockmeier[7]

[1]University of California at Santa Cruz
[2]Interactables, Cambridge, United Kingdom
[3]Sheffield Hallam University, United Kingdom
[4]Microsoft Research Cambridge, United Kingdom
[5]Bar Ilan University, Israel
[6]Sheffield University, United Kingdom
[7]Free University Berlin, Germany, and University of Manitoba, Canada

Lifelogging is a technically inspired approach that attempts to address the problem of human forgetting by developing systems that "record everything." Uptake of lifelogging systems has generally been disappointing, however. One reason for this lack of uptake is the *absence of design principles* for developing digital systems to support memory. Synthesizing

Steve Whittaker is a cognitive scientist with an interest in CSCW, computer mediated communication, and memory technologies; he is Professor in the Department of Psychology at University of California at Santa Cruz, United States. **Vaiva Kalnikaitė** is a researcher and interaction designer with interest in human memory and the design of lifelogging technologies; she is a lead user experience designer at Interactables, UK. **Daniela Petrelli** has a multidisciplinary background in computer science, fine art, and social research; her interests include affective memories, empathic design, and digital innovation; she is Reader at the Art & Design Research Centre at Sheffield Hallam University, United Kingdom. **Abigail Sellen** has a background in Cognitive Psychology and a long term interest in evaluating and designing memory technologies; she is a Principal Researcher at Microsoft Research, Cambridge, United Kingdom, and a Special Professor of Interaction at the University of Nottingham, United Kingdom. **Nicolas Villar** is a researcher at the Microsoft Research Lab in Cambridge, United Kingdom; he is particularly interested in the application of embedded systems and digital fabrication as building blocks for designing of interactive objects. **Ofer Bergman** is an information behavior scientist with an interest in human–computer interaction, personal information management and information retrieval; he is a Lecturer at the Department of Information Science at Bar-Ilan University, Israel. **Paul Clough** is a computer scientist researching multilingual search and evaluation of information retrieval systems; he is a Senior Lecturer at the Information School Sheffield, United Kingdom. **Jens Brockmeier** is a Senior Scientist at the Free University Berlin, a Senior Research Fellow at the Centre for Narrative Research of the University of East London, and a Visiting Professor in the Department of Psychology at the University of Manitoba in Winnipeg; his main interest is in autobiographical memory.

CONTENTS

multiple studies, we identify and evaluate 4 new empirically motivated design principles for lifelogging: Selectivity, Embodiment, Synergy, and Reminiscence. We first summarize four empirical studies that motivate the principles, then describe the evaluation of four novel systems built to embody these principles. We show that design principles were generative, leading to the development of new classes of lifelogging system, as well as providing strategic guidance about how those systems should be built. Evaluations suggest support for Selection and Embodiment principles, but more conceptual and technical work is needed to refine the Synergy and Reminiscence principles.

1. THE LIFELOGGING VISION AND THE UPTAKE PARADOX

We are all aware of the fallibility of our memories. Even functioning adults forget or distort names, dates and prospective commitments (Baddeley, Eysenck, & Anderson, 2009; Cohen, 1996; Schacter, Coyle, Fischbach, Mesulam, & Sullivan, 1995). Forgetting was first scientifically studied by Ebbinghaus (1885/1964), and it has been documented across multiple contexts and types of information. But technological advances in networks, sensors, search, and storage could potentially rectify this situation. Bell and Gemmell (2009) outlined future digital technology enabling "total recall" of our entire lives through "total capture" of personally relevant information. Such information includes paper and digital documents we read: e-mail, paper mail, and instant messages; content of telephone conversations; websites visited;

and credit card transactions. Also included are data relating to everyday activities, such as still images, video, ambient sound, and location data. These personal archives might also be supplemented with relevant environmental measures (light intensity and temperature variation) and even biosensor data (heart rate and galvanic skin-response) reflecting our physical and emotional state. The process of digitally archiving one's entire life in this way is known as *lifelogging* (Bell & Gemmell, 2009; Mann, 1997; Sellen & Whittaker, 2010), and it is becoming increasingly feasible. The technical research community has responded to this opportunity by creating different types of lifelogging systems to capture all this information. These research systems fall into three main categories: personal digital repositories, mobile activity capture, and domain specific capture.

Personal Digital Repositories. A key problem with lifelogging is that we encounter digital data in different forms and formats. One common class of lifelogging system provides unifying infrastructures for managing heterogeneous collections of digital objects that users have generated or encountered. Such systems also include tools for searching, and browsing collections. Examples include "Haystack" (Adar, Karger, & Stein, 1999), "Presto" (Dourish, Edwards, LaMarca, & Salisbury, 1999), and "My LifeBits" (Bell & Gemmell, 2009). Metadata about various aspects of a digital object's past context of use (when it was created, who it came from, and relationships to other objects) allows users to search, making refinding easier (Dumais et al., 2003). The same metadata make it possible to view and browse data in flexible ways.

Mobile Activity Capture. A critical limitation of these unified repositories is that they are *linked to the desktop*—omitting useful information about what we are doing in the real world. Other lifelogging systems extend capture into the realm of everyday, mobile life. The earliest such systems involved simple location tracking using networked sensors (Lamming et al., 1994). More recent instantiations can be classified as wearables, portables, or instrumented environments. Wearable systems are based mainly on head-mounted, still, or video cameras (Mann, 1997) or on wearable audio capture devices (Vemuri, Schmandt, Bender, Tellex, & Lassey, 2004). One recent commercial instantiation is SenseCam (Hodges et al., 2006), a wearable digital camera that uses sensors to automatically trigger capture of images of everyday activities. Portable systems largely make use of specialized software on PDAs, notebook computers, or cell phones. Instrumented environments aim to capture activity through installed sensors and local wireless networks, although this approach currently does not scale.

Domain Specific Capture. Finally there are lifelogging systems that capture rich data about interactions in specific domains. Most of these systems have focused on meetings, lectures, or other forms of *conversation* allowing this kind of "organizational knowledge" to be browsed and searched (Whittaker et al., 2008). Early systems recorded and indexed audio and penstrokes (Whittaker, Hyland, & Wiley, 1994). More recent technology-enhanced meeting rooms capture video and audio from

multiple cameras and microphones, whiteboard use, slide presentation, and digital notes (Janin et al., 2006; Renals, Hain, & Bourland, 2007). Complex search and browsing tools support access to this data, including speech summarization (Tucker, Bergman, Ramamoorthy, & Whittaker, 2010; Tucker & Whittaker, 2006) and decision-based summarization (Hsueh & Moore, 2009). Machine learning systems automatically extract key elements from meetings such as action items and decisions (Ehlen, Purves, Niekrasz, Lee, & Peters, 2008). Similar approaches provide access to public lectures (Abowd, 1999).

The existence of these systems implies that lifelogging is technologically possible; however, these systems are not in widespread use. Mann (1997) presented a history of lifelogging, showing some systems are almost 30 years old. Nevertheless, there has been little uptake even within research laboratories (Sellen & Whittaker, 2010). There are many potential reasons for this failure, including the availability, complexity, and cost of the technology for users (Bell & Gemmell, 2009; O'Hara et al., 2006), technical computing challenges (Doherty et al., 2011; O'Hara et al., 2006), and numerous ethical issues (Mayer-Schönberger, 2009). However, one significant barrier to uptake has been a *lack of systematic design principles* for the development of lifelogging systems, specifying why and how to build these systems. We attempt to provide such guidelines in this article.

The article is structured as follows: We first empirically explore how digital memorabilia and lifelogging tools are currently used, identifying the main problems that users experience with such tools. We synthesise these data to motivate four design principles, Selection, Embodiment, Synergy, and Reflection. We then describe and evaluate new types of systems built according to these principles. The principles were effective in leading to the development of new classes of system. Based on our experience, we outline some of the outstanding problems for research on digital support for memory. Our contribution is the synthesis and evaluation of the design principles. Whereas other work has begun to propose new approaches to designing lifelogging systems (Sellen & Whittaker, 2010), those principles were derived from psychological theory and have yet to be validated in real settings. Instead here we take a socio-technical approach: We derive design principles from studies of existing digital tools and test our principles by developing and evaluating systems with real users.

2. EMPIRICAL MOTIVATIONS FOR DESIGN PRINCIPLES

We conducted four different empirical studies exploring how people use current digital memory technologies. These empirical observations motivate our design principles. We examine how well visual lifelogs support everyday memory and how people manage and access large, personal photo collections. We also compare physical and digital mementos in terms of their evocativeness and salience. Finally we look at

when and why people access digital collections as opposed to relying on their own *unaided memory* (UM). UM refers to situations where people are entirely reliant on internal "in the head" recall. UM contrasts with *mediated memory*, that is, when external digital or analogue artifacts such as a "to-do" list are used to aid recall. UM is what cognitive psychologists typically study when they look at "memory" (Baddeley et al., 2009; Cohen, 1996; Schacter et al., 1995).

2.1. Rich Visual Lifelogs Do Not Radically Enhance Everyday Recall

Our first study (Kalnikaite, Sellen, Whittaker, & Kirk, 2010) explored one of the fundamental premises of lifelogging—that having access to rich records of everyday events should radically enhance autobiographical recall. Lifelogs were captured using SenseCam (Hodges et al., 2006), a wearable digital camera that automatically takes pictures of the wearer's activity when one of its sensors (light, motion, or temperature) is activated (e.g., when the wearer gets up to leave the room). On average this generates about 4,000 images per day for each participant. Captured images are accessed by viewing one's day in "fast forward" mode using simple playback controls.

Theoretical work emphasizes that autobiographical memories are mediated by internal visual images (Conway, 1990; Pillemer, 1998). We therefore expected that having access to personal SenseCam images would assist autobiographical recall, as these images provide a rich visual record of personal activity. Eighteen people wore SenseCam for 2 weeks collecting thousands of images each day. Two months later we asked them to recall as many events as they could for one of their logged days, using the following standard autobiographical memory question: "What did you do, where did you go, and who did you meet on [Monday November 22nd]?" Participants could browse their images for as long as they wanted when answering. We compared such digitally mediated recall of everyday events with UM. To our surprise, having access to a rich set of external images did not promote "total recall." Digitally mediated memory was no better than UM in the total number of events recalled ($M = 0.46, 0.32$), $t(17) = 0.69, p > .05$, after we controlled for the possibility that users were guessing past events simply by studying the images, as opposed to genuinely remembering those events. Participants offered two related explanations why images helped less than expected: First, there were too many images, and second, it was hard to abstract across the huge set of images to "make sense" of their day. Similar findings are reported elsewhere. Sellen et al. (2007) also found that simple image lifelogs do not improve everyday recall in the long term. Section 3.1 discusses how we might better structure lifelog data to overcome these problems.

2.2. Long-Term Retrieval Failure and the Cost of Manually Organizing Archives

Digital photos are one of the most common types of digital memorabilia. Although digital cameras are relatively recent, people have already amassed tens of

thousands of digital photos (Kirk, Sellen, Rother, & Wood, 2006; Whittaker, Bergman, & Clough, 2010). Furthermore, many families with young children regard photos as among their most precious digital assets (Whittaker et al., 2010). Studying how people access digital photo archives should shed light on retrieval from lifelogs. Unlike automatically captured SenseCam images, personal photo archives represent *deliberate attempts* to document our past. Photo collections also cover longer periods than we were able to address in our SenseCam study.

We studied parents with young children, as parents see themselves as "active curators" of family memorabilia (Whittaker et al., 2010). We were interested in whether people can accurately access information from a rich, long-term archive, as well as the methods people use to organize digital photos. We expected people to be effective at organizing and accessing digital photos because these are an extremely valued resource. Indeed, work on accessing *recently* taken photos shows that people are good at retrieving photos taken within the last year (Frohlich, Kuchinsky, Pering, Don, & Ariss, 2002; Kirk et al., 2006).

Finding was much less successful in our own study, however. When we asked 18 people to find digital family pictures taken more than a year ago, participants were often unsuccessful. We first asked participants to name significant family events from more than a year ago. In a subsequent retrieval task, they were asked to show the interviewer digital pictures from these salient past events. For example, having established that someone has young school-age children, we asked people to "Find a photo of your child's first day at school." To avoid participants choosing events that they could easily retrieve, participants were not told about the retrieval task during the initial interview.

In contrast to their expectations, our participants were successful in retrieving pictures in only 61% of retrieval tasks. In the remainder, participants simply could not find pictures of significant family events. Of the 28 unsuccessful retrieval tasks, 21 (75%) were pictures that the participants believed to be stored on their computer (or on CDs) but which they subsequently could not find. The remaining seven were pictures participants initially thought were stored digitally, but during the retrieval process they changed their minds into thinking they were taken with an analogue camera.

Based on participants' comments and behavior during search, we identified two key reasons for their unexpectedly poor retrieval. These echoed the SenseCam findings. One problem was the sheer number of photos kept. The low perceived costs of storing photos meant participants kept many more digital than analogue photos. Participants were unwilling to delete digital photos even when these were near duplicates of others in their collection. Just 8% of photos uploaded from camera to computer were deleted overall. Users did not delete even though they were aware that keeping too many photos affected their ability to find important photos. Other research on digital archives suggests that resistance to deletion is a general phenomenon, occurring in part because people can always think of a context in which they might need that information, "I'll keep that e-mail, just in case" (Bergman, Tucker, Beth-Marom, Cutrell, & Whittaker, 2009; Kirk et al., 2006; Marshall, 2007; Whittaker, 2010).

The resulting large collections were both difficult to manually organize and hard to access. Photos were seldom systematically organized, often being placed in a single folder. Just 16% of participants systematically used subfolders, although active organizers had a statistically higher proportion of successful retrievals than those with more rudimentary organization, $t(16) = 2.38, p < .05$. Folder names were often inconsistent, being mixtures of dates, people, and events and were sometimes vacuous ("mypics"). There were also often duplicate folder titles. One stated reason for this poor organization was that users had little time or motivation to spend organizing. Indeed in many cases we found that participants had *never* revisited photos they had taken years before. As a result when participants came to retrieve they were often unsuccessful. As one participant commented, "[The archive] is a very big mess. I've no idea. . . . It has no logic. . . . I feel full of despair. It's easier to give up on seeing [the pictures] altogether."

2.3. Digital Invisibility: Comparing Digital and Physical Mementos

Another way to understand digital remembering is to compare important physical memorabilia with their digital counterparts. What makes a digital or physical object a significant memento? How does it attain this status, and how does it achieve its evocative effect? What are the differences between physical and digital mementos? We conducted an ethnographic study (Petrelli & Whittaker, 2010; Petrelli, Whittaker, & Brockmeier, 2008) with parents of young children, again chosen as a population who claim to actively "curate" their family's digital memorabilia. We asked 16 parents to give us a "tour" of their homes, identifying objects that invoked important memories and explaining their significance. All of our participants were digitally literate. They all owned digital cameras and computers, and were all experienced users of e-mail and the web. We left it open to participants what types of objects they chose, but given their digital sophistication we expected that many would select digital as well as physical mementos.

The most striking finding was that *only one* of the 159 mementos spontaneously described in the house tour was digital. A single participant described digital maps he created in association with his cycling hobby. This failure to focus on the digital occurred despite the fact that all participants had large digital archives. In follow-up interviews, when explicitly questioned if they had digital as well as physical mementos, there was an initial denial by all participants. The exception was digital photos, which were mentioned by everyone. Eventually with more explicit questions about saved digital videos or e-mails from friends and family, all participants became aware they had substantial collections of digital memorabilia. These were more varied than expected, including digital collages, celebratory Microsoft PowerPoint presentations and recordings of radio broadcasts. Why then were participants' digital collections less salient? One possibility is that digital mementos are simply less *visible* than their physical counterparts. However, invisibility cannot explain why so few digital objects were chosen, as the physical objects that people chose were also often hidden from

view. In 25% of cases the physical mementos people chose were not in direct view and had to be retrieved from the back of a drawer or wardrobe. Another potential reason for failing to choose digital objects is that people were focused on the *distant past*, whereas digital objects tend to be relatively recent. However, physical mementos were also relatively recent; 46% of these objects referred to the last 10 years of the person's life, equivalent to the period covered by digital mementos. The finding that digital mementos are less salient is replicated in other research on digital memorabilia (Kirk & Sellen, 2010; van House, 2009).

A second unexpected finding was the *nature* of the chosen objects. The most commonly chosen objects were functional. Mundane everyday objects such as a cup, clock, coffee machine, golf tee, pots, cookery book, teapot, children's toys, ladder, calendar, bed, stove, candleholder, and books were all chosen. These accounted for 28% of mementos. Those objects were integrated and used in everyday life, thus daily reinforcing their evocative function.

A third critical observation was the *nature* of recollection. Lifelog applications tend to emphasize factual recall, such as what events occurred when. However, consistent with theoretical autobiographical memory research (Cohen, 1996; Conway, 1990; Pillemer, 1998), when describing valued physical mementos, we observed very different memory processes: people engaged in reminiscence and reflection. A participant here describes a mundane object, but what is being talked about is not a simple recall of factual events. Rather, the narrative is about domestic and work aspects of a specific period in her life, the feelings engendered, and how it relates to her identity:

> Object number one is this mug ... which is actually broken. ... I will never use it again but I can't quite bear to throw it away I feel very emotionally attached to it for some reason. ... I bought it when I was working in London. ... I think [it's the memory of] working in publishing, living in London and going through a sort of fulfilling patch in my career. ... Also I associate it with buying my first house, having it there in the kitchen in my first house. ... So its also an object of continuity because I think I must have had it for ... ooh ... let me think, I've probably had it for nearly 20 years!

2.4. Trade-Offs Between Digital Tools and Unaided Memory

Most lifelogging work tacitly assumes that lifelogs will be users' preferred way to access information, because lifelogs are more accurate and reliable than fallible UM. One unexplored issue is the *trade-offs* leading people to use a lifelogging system versus UM. It is obvious that a well-designed lifelog can help users access information that they may otherwise have forgotten, but there may be good reasons why people sometimes prefer to rely on UM. For example, there is little incentive to use digital tools when one can easily remember information unaided, for example, if the situation requires memory for gist (Baddeley et al., 2009). There may also be *inefficiencies* associated with digital memory tools, leading people to turn to UM where retrieval

is usually rapid. Compare, for example, the ease and rapidity of accessing a familiar contact name from UM, with the effort of retrieving it from a poorly organized e-mail archive.

Kalnikaite and Whittaker (2007) investigated these trade-offs in a lab study of memory for everyday conversations involving 25 participants. They were provided with three different types of memory tool: Dictaphone, Pen and Paper, and a system called ChittyChatty, or they could rely on UM. ChittyChatty worked as follows: it simultaneously recorded and synchronized digital note-taking and audio (see left-hand side of Figure 1). Synchronization meant that digital notes could be used as efficient entry points into the speech recording thus implementing "attentional selection," a technique we describe in more detail in Section 3.1. In contrast, access with Dictaphone was inefficient: requiring serial browsing of the speech recording.

We wanted to understand when and why people relied on digital tools as opposed to UM. First we examined whether digital tools were used only when people felt that their UM was weak. Second, we explored how tool properties influenced usage decisions. Both Dictaphone and ChittyChatty tools were *accurate* lifelog recordings providing a complete verbatim record of the conversation, but they differed in retrieval *efficiency*. In contrast Pen and Paper notes are generally an incomplete record of what was said, but scanning handwritten notes is rapid making access relatively efficient.

Our first finding was that digital tool use is affected by people's *perceptions* of their UM. Before people answered each memory probe, we asked them how confident they were about whether they could remember that information unaided. Digital tool use was strongly negatively correlated with people's confidence in their UM. When people

FIGURE 1. Left-hand side: ChattyWeb device for capture and browser user interface for retrieving speech from handwritten annotations. Right-hand side: PiccyWeb device for capture and browser user interface for retrieving speech from photos. (Color figure available online.)
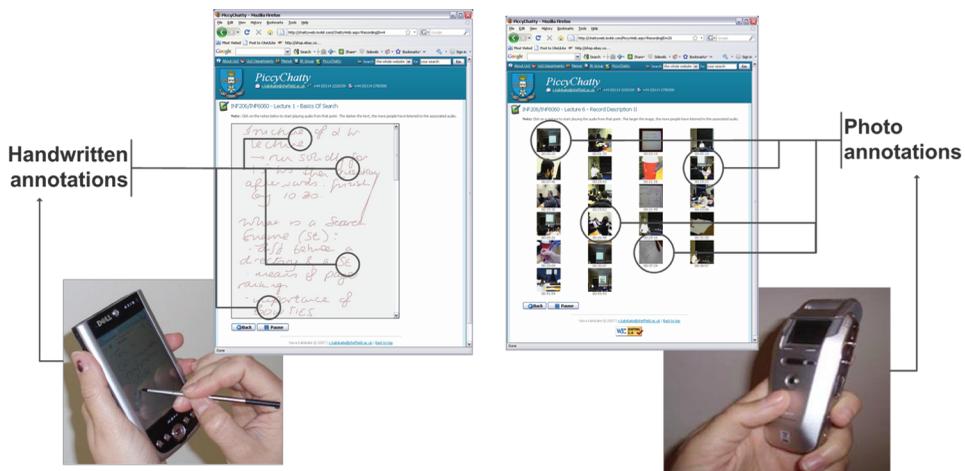
Note. **Clicking on a digital note or photo in the browser accesses the speech occurring when the note or photo was taken.**

were confident they could remember using UM, they were far less likely to use digital tools, $r(1, 899) = -.50, p < .0001$. Thus, even though digital tools like ChittyChatty and Dictaphone provide a complete record, people do not use them when they think that they can remember using UM. Our second finding was that people's choice of retrieval tool was influenced by *efficiency*—which people traded off against *accuracy*. As expected, retrieval was more accurate with Dictaphone and ChittyChatty than Pen and Paper, because both provided verbatim records. However people were statistically more likely to use ChittyChatty than Dictaphone ($F(2, 810) = 8.5, p < .0001$, planned comparison, $p < .05$), because, as users pointed out, ChittyChatty provided more efficient access. Furthermore, they were as likely to use Pen and Paper as Dictaphone (planned comparison $p > .05$), even though Pen and Paper was an incomplete record, again because Dictaphone's inefficiency compromised its usefulness. Overall digital tools were not always preferred; instead, usage was affected by people's perception of UM as well as tool efficiency.

In summary, our four empirical studies reveal a number of problems with existing digital archives and lifelogging systems.

- **Rich Records Do Not Guarantee Effective Recall.** People experience difficulty in accessing information from rich archives when information is not appropriately indexed: we need new methods for user-centered structuring of large archives.
- **Rich Records Impose Organizational Problems.** Archives are large because people are unwilling to delete. Manually organizing large archives is onerous, which can lead to underorganization and poor recall: Lightweight organizational techniques are essential.
- **Digital Invisibility.** Digital archives are generally less salient than their physical counterparts: Digital archives need to be better integrated with everyday life and everyday objects.
- **Beyond Factual Recall.** Memory involves more than simple retrieval of factual information: Digital tools need to also support reminiscence and reflection.
- **Digital Tools Are Not Always Preferred.** People are strategic about when they use digital archives, using them only when they perceive UM to be weak: New systems need to address UM weaknesses.

3. DESIGN PRINCIPLES FOR NEW DIGITAL SYSTEMS TO SUPPORT MEMORY

We derived four design principles from these empirical observations. These are intended to be *generative*—suggesting new types of system. The principles also provide strategic guidance about *how* to build new systems. Our design principles are:

- **Selection:** Capturing a rich archive is insufficient to guarantee recall. Users are loath to delete information, so our approach retains the entire archive, but we

simplify access by making important items more salient. Our selection methods analyze implicit user behaviors, inferring importance by *tracking user attention* and by using *feedback*. Detecting implicit user actions allows importance to be inferred as a side-effect of normal user activities. Selection is lightweight, imposing structure while avoiding manual organization which is onerous for users.

- **Embodiment:** Digital mementos are often overlooked and "invisible" compared with their physical counterparts. We need to integrate them better into people's everyday lives. One way to do this is to *embed* digital archives in significant physical objects.
- **Reminiscence and Reflection:** Memory involves more than veridical recall of simple facts. It often involves reconstructive processes such as reminiscence and reflection, and we need to support these processes in our digital tools.
- **Synergy not Substitution:** Digital tools are used *strategically*. Even when digital tools offer complete records, people do not use them exclusively. Instead people use digital tools to overcome perceived weaknesses in UM. Our systems therefore need to act in synergy with UM, providing support in contexts where UM is weak.

We now elucidate these principles, describing four novel digital tools built according to the principles; their evaluations provide evidence for the utility of the principles. In each case, we describe how we map from abstract principle to specific implementation and how we evaluated each resulting system.

3.1. Selection

Identifying Important Events Using Attentional Selection

As we have seen, people experience difficulty in accessing rich digital archives. Furthermore, users find it cognitively difficult and time-consuming to manually organize their collections. Together these suggest that we need lightweight methods for structuring information to support selection. We therefore analyze implicit user activity to detect important events within lifelog records and make important events more salient.

Our first selection technique exploits *user attention* to identify important events in the lifelog. Following arguments from functional cognition (Anderson & Schooler, 1994; Conway, 1990), attentional selection assumes that the events that people will want to remember are those they *pay careful attention to*. But how can we computationally determine what people are attending to? To do this we record simple behavioral correlates of attention. For example, people demonstrate attention to critical aspects of a meeting by taking careful notes during those phases. At a family wedding, they reveal their focus of attention by taking photos of important people or events. Our systems log these attentional behaviors, as well as recording the entire event. The systems then identify those events in the recorded archive that *temporally correlate* with attentional behaviors. We therefore facilitate selection by increasing the salience of those points

in a meeting archive where notes were taken, or the parts of the wedding archive when photos were taken.

Attentional selection exploits *implicit* behaviors. Unlike folder creation or tagging, users do not have to engage in complex deliberations to categorize information for future retrieval. Instead, important parts of the archive are selected in a lightweight way as a side-effect of implicit actions that the user already carries out, such as taking notes or photos. Note- and photo-taking (Figure 1) are examples of *preparatory* acts where people *anticipate* future retrieval demands as the event unfolds. However we can also track and exploit attention *after the event*, for example, during retrieval. Again, functional cognition would argue that events that are *frequently reaccessed* from the archive can be inferred to be important. Our second technique employs this implicit feedback for selection, making frequently accessed events more salient (Kalnikaite & Whittaker, 2008).

An alternative to Selection might be to permanently *delete* some of the archive, thus reducing its complexity. Others have argued that deletion would not only be more efficient but also more comprehensible and ethical for users (Bannon & Kuutti, 1996; Mayer-Schönberger, 2009; O'Hara et al., 2006). However, we have already presented strong evidence that users are highly unwilling to delete information (Kirk et al., 2006; Marshall, 2007; Whittaker et al., 2010) making deletion an implausible solution. Another alternative to Selection might be to *autonomously* index archives, using techniques that do not require manual user intervention such as event detection or image analysis (Doherty et al., 2011; Doherty & Smeaton, 2008). Though potentially promising, the benefits of these techniques for users remain to be demonstrated, and we return to these in Section 4.

We now describe systems that have been built to support Selection and their evaluation. In choosing a domain, we also adhered to the Synergy not Substitution principle. We wanted to develop applications where UM was weak. Education is a context where UM demand is particularly high. In many pedagogic situations there is a need to master and remember complex novel information delivered verbally. Students experience major problems in determining what is critical (and hence important to remember) while processing complex new information (Brown, 1987; Bransford, Brown, & Cocking, 1999). Digital records might therefore free students from the pressures of "not missing anything important" while trying to comprehend novel ideas or contribute to class discussion (Abowd, 1999).

Preparatory Attentional Selection System

We built two preparatory attentional selection systems (Figure 1) that detect implicit user actions to determine important events (Kalnikaite & Whittaker, 2010). The first, ChattyWeb, like the ChittyChatty system described in Section 2.4, exploits pen-based digital notes to detect important events for future retrieval. Users take notes that are synchronized with the underlying speech record, which is recorded onto the local device. Clicking on a note in the browser retrieves the speech occurring when the note was taken. The second attentional selection tool PiccyWeb, allows students

take photos of important events, for example, a key slide in a lecture or critical notes on a whiteboard. The photos are again synchronized with the speech record, so that clicking on a photo in the browser retrieves what was being said when the photo was taken. In both cases, the attentional index provides efficient entry points into the speech record.

We evaluated the tools in field trials and a lab study with 98 students. ChattyWeb and PiccyWeb were used to record and annotate lectures over 2 years of an information retrieval course. Recordings were made available to students after each lecture allowing them to reaccess course materials to revise and prepare coursework. We logged access behaviors and determined how system usage affected student performance such as course grades. There were few differences between Note (ChattyWeb) and Photo (PiccyWeb) versions of the system, so we combine the results for the two versions here.

Both versions were used extensively by students, who were highly enthusiastic about them. Spontaneous usage in Figure 2 shows the total number of accesses for all lecture recordings during each week of the 1st year of the course. Thus, in Week 13, there were 193 accesses. There is a clear relation between active system usage and coursework evaluations. Coursework evaluations occurred in Weeks 3, 8, 10, and 13, and on each occasion, we saw increased system use. These spontaneous data provide suggestive evidence that the system was being actively used to prepare for coursework evaluations. To test system utility more rigorously we conducted a linear regression where we regressed system usage and student characteristic variables against learning outcome (final student grade). The overall model was highly significant ($R^2 = 0.744$, $p < .001$). Number of system sessions was a significant predictor of final grade ($t = 2.12$, $p < .05$) when student ability and native language were controlled for, suggesting that system use improved learning. One student described how having the

FIGURE 2. Combined ChattyWeb and PiccyWeb accesses over the duration of the course.
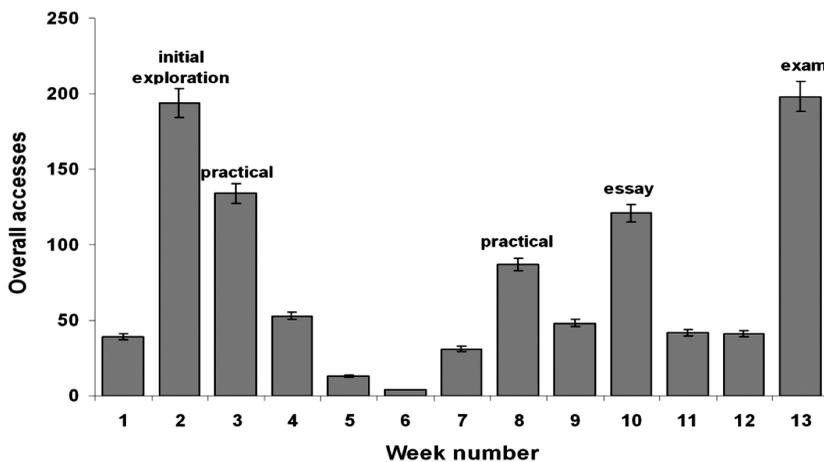
Note. Tool use is associated with course evaluations, in Weeks 3, 8, 10, and 13.

FIGURE 3. Quiz accuracy scores for different retrieval strategies: The combined digital tools (ChattyWeb and PiccyWeb) are the most effective retrieval technique. (Color figure available online.)
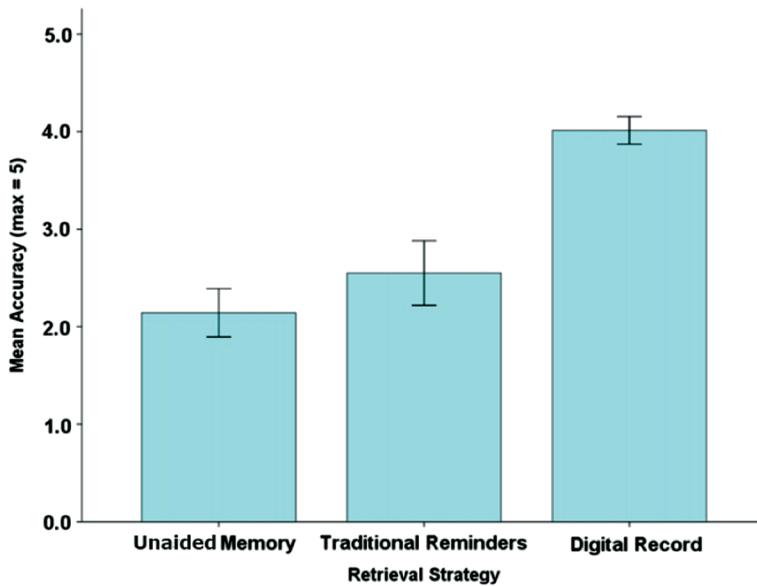
multimedia recording allowed her to browse lectures in an active manner to trigger recall: "[PiccyWeb] gives you the chance to move around in the lecture, as well as jog your memory by seeing the pictures of the lecture at different points."

These naturalistic data offered evidence of when and how the system was used. We next quantified system benefits under more controlled conditions. For 35 students, we administered class quizzes testing material from previous lectures. We compared scores when students answered quizzes using: (a) digital tools (ChattyWeb and PiccyWeb); (b) traditional educational tools, that is, handouts and personal lecture notes; and (c) UM. Figure 3 shows the differences in retrieval accuracy. Students who used digital tools performed better than those who relied on UM or traditional tools. An analysis of variance showed an effect for tool used, $F(2, 47) = 8.62, p < .001$, with post hoc tests showing digital records outperforming UM ($p < .0001$) and traditional tools ($p < .02$).

Feedback-Based Selection

The previous study showed clear benefits for *preparatory attentional selection*. However, there are limitations to this technique. Students pointed out difficulties in judging important information to index during the lecture, observing that sometimes it was only *afterward* that they realized which aspects of the lecture were important. We therefore explored a second selection principle that utilized *post hoc feedback*. Rather than relying purely on preparatory indices we exploited people's implicit evaluations

of those indices at retrieval, when people were revising class materials. We assumed that a strong indicator of the utility of an index was how frequently that index was accessed for retrieval. We provided students with implicit visual feedback about which annotations (and hence which parts of the lecture) were accessed most often during retrieval. Feedback was combined across the entire class. The feedback interfaces for pictures and notes are shown in Figure 4. Each time a student uses a photo index to access part of the lecture, we make that photo *more salient* by enlarging it relative to other photos. For digital notes we increase salience by highlighting them in bold and in red (Kalnikaite & Whittaker, 2008).

We again evaluated feedback-based selection in a naturalistic classroom setting. We conducted a field trial with 25 students where we provided lecture recordings that the students could access in different ways. We created four different versions of the system. Two versions included feedback about which indices had been accessed most often: Picture feedback and Notes feedback (Figure 4). The other two versions (basic ChattyWeb and PiccyWeb systems) provided the same set of annotations but without feedback information. Thus, in the basic versions, all photos were the same size and handwritten notes were presented at the same level of grayscale and color salience.

There was almost no use of the basic systems but significantly greater use of the two feedback systems, indicating benefits for feedback (M session lengths = 10.9 vs. 462.1 s, $z = -3.2$, $p < .001$). We also evaluated the utility of feedback *within* each lecture distinguishing between indices that have previously been accessed from those that have not. We observed a global shift to using feedback. In the second half of our trial, *all* uses of photos exploited indices that had previously been accessed. Again we looked at how system use related to student grades. There was a significant correlation between coursework grades and number of sessions using the feedback

FIGURE 4. Feedback-based selection—Implicit feedback from prior access signals the importance of different indices. (Color figure available online.)
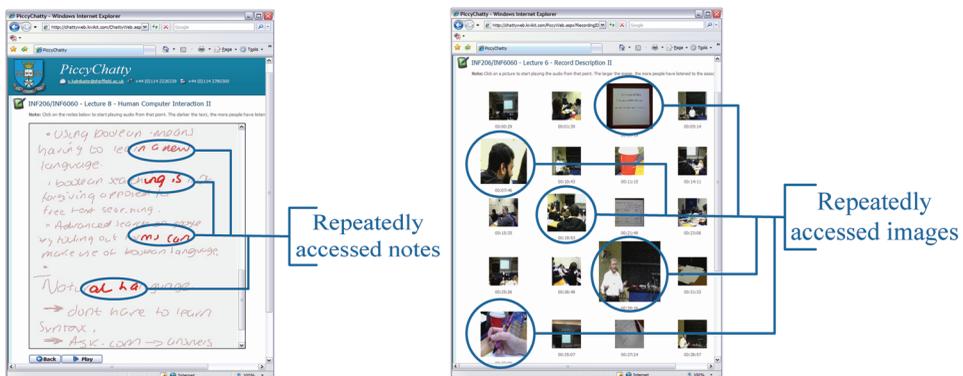
Note. **Photos and notes accessed more often are made more visually salient, being bolder or larger.**

systems, $r(23) = .51$, $p < .05$. However, there was no correlation between grades and number of basic sessions, $r(23) = 0.21$, $p > .05$. Thus systems providing feedback improved grades, whereas the basic system did not, suggesting benefits of feedback for student learning.

3.2. Embodiment: Combining the Digital and Physical

We all keep physical as well as digital mementos, but the two are currently treated very differently. Physical mementos are frequently displayed around homes often taking the form of mundane functional objects that remind people of significant past events. In contrast, digital mementos tend to be poorly integrated with people's everyday lives, making them invisible (Petrelli et al., 2008). In our next study, we wanted to see whether embodying digital archives in significant physical objects would make those archives more salient. Embodiment has been explored as a general interaction paradigm (Dourish, 2001; Ishii, 2008), but it has not generally been applied to memory tools. One notable exception is Hoven and Eggen (2005, 2008), who created augmented souvenirs by linking digital information to physical mementos and explored the effects of different cue types on recall (Hoven & Eggen, 2009). We extend their approach, exploring the embodiment of digital "sonic souvenirs"—sound-only family recordings (Petrelli et al., 2010). Our design embodied these digital sound souvenirs in a familiar functional domestic object: a radio (Figure 5). Our design was intended to fit seamlessly into the home by repurposing a familiar object to access digital mementos. By using a radio for embodiment, we maintained the evocativeness and ambiguity of sound, at the same time allowing for natural exploration of the sound collection.

The Family Memory Radio (Figure 5) reflected insights from a prior field study (Dib, Petrelli, & Whittaker, 2010) in which 10 families recorded "sonic souvenirs"

FIGURE 5. Three siblings interacting with the Family Memory Radio. (Color figure available online.)

(audio-only mementos) of their summer holiday. Participants recorded between nine and 197 different sounds, including mock interviews, family conversations, giggles, pseudo radio shows, commentary about what they were doing (waiting in an airport, having breakfast), family arguments, ambient sounds both natural (animals, water) and human (volleyball match, murder mystery game), and created sounds (bubbles blown with a straw in water, the creak of a door). A few participants recorded verbal diaries or more abstract reflections, for example, favorite parts of a holiday.

The design goal was to embed these digital memorabilia in the radio so as to facilitate a shared listening experience while reminding people about their digital archive. We wanted minimal controls while still mimicking a radio, so content was organized into a few channels for different recording types: human voices versus ambient sounds, family favourites plus an all-inclusive channel organized by time.

We evaluated Family Memory Radio with the families who participated in the initial sonic souvenirs field study. One year after making their recordings, we invited families to use the Radio for revisiting the sounds they collected the previous summer, and 23 people from six families took part. We compared their Family Memory experiences with using a PC or a Dictaphone to access the same sonic memorabilia. We observed families interacting, listening, reminiscing, and playing their sounds. After about 30 min of self-discovery, 10 open questions were posed to investigate: feelings on relistening to their sounds, perceptions of the interaction, the aesthetic of the radio, and its projected use in family life. Observed behaviors and comments made during the interaction were used to stimulate discussion and further elicit participants' views.

Listening to sounds on the FM Radio engendered extensive laughter and family jokes. One defining characteristic was the *social interaction* focused on the device. This included frequent conflicts between siblings over the controls (see Figure 5) usually ending up in shared laughter and animated discussion. The Radio afforded a level of collective interaction that neither the Dictaphone nor the PC allowed. Participants commented on how much better their experience was with the Radio: "With [the Dictaphone] you have to pass it around and lean on it." The Radio was also a more democratic way of accessing mementos than a PC: "The files are on my laptop and [the kids] don't have easy access to it." The radio therefore seemed to overcome perceived invisibility barriers with current ways of accessing digital memorabilia. In addition to collective listening, families became deeply involved in conversations about the original events. They discussed when a certain sound was recorded and aspects of the holiday unrelated to the clip. They explored the different channels and their collections extensively. All families commented on the excellent quality of the audio and how vivid the experience was, "It's incredible! It seems like having him in this room!"

On various occasions people pointed out parallels with physical mementos. With two exceptions, participants liked the embodiment in an old-fashioned radio. Just as with physical mementos (Petrelli et al., 2008), people saw the radio as a prompt for conversation. "I can see visitors asking about it. It would make a good conversation

point." The embodied character was also seen as being good for reminding. Unlike the Dictaphone or mementos on the PC, participants thought that the physical presence of the Radio would remind them about their digital sounds (Hoven & Eggen, 2005), addressing the invisibility problem with digital collections. Participants were confident that the device would not end up forgotten and unused in a drawer "like so many digital gadgets we have." Such reminding could prompt more recordings of sonic mementos: "[While listening to the sounds] I regret I did not record more this year. I suppose it is a matter of remembering that we can."

3.3. Reflection and Reminiscence

The majority of lifelogging applications focus on verbatim recall of facts, whereas significant physical mementos engender reminiscence and reflection. Our prior work (Petrelli et al., 2008) showed the importance of the home as a focus for physical mementos and reminiscence. Other work has also demonstrated the importance of people and places for cueing reminiscence (Cohen, 1996; Linton, 1982; Wagenaar, 1994). Our next application, MemoryLane (Kalnikate & Whittaker, 2011), allows users to organize and reflect on digital mementos in the context of home, people, and places. MemoryLane is a *digital* environment that mimics important properties of the physical environment (Figure 6). Rather than supporting veridical recall, our tool is intended to provide an evocative set of familiar cues allowing participants to *reminisce* and *reflect* on their lives.

There are three contextual views in Memory Lane. Each can be populated with digital representations of personal mementos: (a) *home* (Figure 6, left-hand side)—a schematic of a house and garden; (b) *places*—maps of potential locations; and (c) *people* (Figure 6, right-hand side)—a photo frame for images of significant people. Although

FIGURE 6. MemoryLane interface. (Color figure available online.)
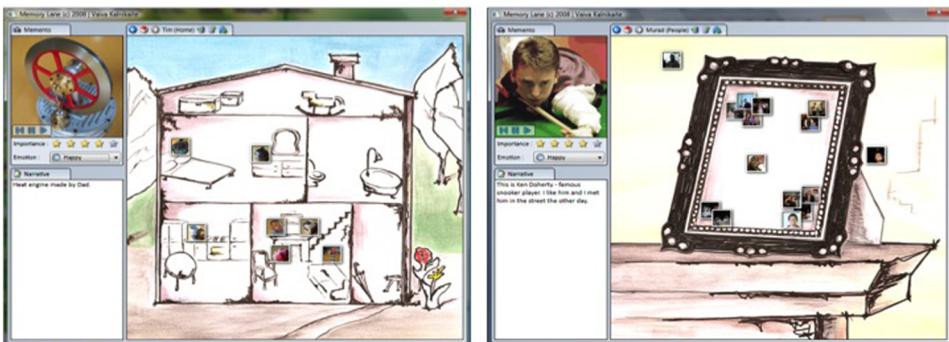
Note. **The left-hand side shows the home view: Small images in the home represent different mementos. The right-hand side shows the people view with significant people organized in and around a photo frame. Enlarged images on the left-hand side of each view show metadata about the currently selected memento.**

the home is technically also a place, we deliberately created a privileged home view, because prior work has shown its centrality for memento collection (Csikszentmihalyi & Rochberg-Halton, 1981). All mementos are represented as pictures with optional additional metadata, including audio and text, allowing people to add narratives or other information, such as affective ratings.

We asked 31 people to capture their physical mementos (using a combination of digital cameras and audio-recording tools). Participants then organized mementos in MemoryLane. For each memento, we asked them to add (a) how important it was to them, (b) how they felt about it, and (c) a narrative about the memories it evoked. We recorded which mementos were augmented with narratives and judgments of importance and emotion. After participants captured and organized their mementos, we interviewed them about how they had used the system and how they saw using it longer term.

Overall, participants actively captured, organized, augmented, and reflected on 356 mementos; 163 of these were home objects, 112 were of places, and 81 were of people. They were actively added narratives to 93% of these objects. As we anticipated, most (76%) of the captured mementos were judged to be important to participants. Interviews and surveys indicated the tool was effective in supporting reflection. Our participants stated that it allowed their everyday lives to become more visible, and this visibility allowed for observation and tracking of personal change, for example, relationships with other people or moving house. MemoryLane also emulates key properties of physical mementos. Most captured mementos featured home objects, replicating studies of the physical world (Csikszentmihalyi & Rochberg-Halton 1981; Petrelli et al., 2008). Reflections about home objects or locations revealed that these were often intended to trigger people, and mementos in the people area engendered the strongest emotional reactions. As with prior work on photos (Chalfen, 1987; King, 1986), participants tended to capture mementos that engendered positive affect, although we did find sadness surfaced with locations associated with deceased relatives.

Consistent with our design goals, participants envisaged using MemoryLane for reflecting on long-term aspects of their lives. One participant talked about using it to reflect on significant life events such as moving house:

> I would definitely like to use [MemoryLane] for a longer period of time. If we ever sell our house and move, it would be great to add mementos of our life in the current house (well, the good times anyway!). It would be nice to look back on these in time, after we'd moved house.

This long-term focus was also reflected in the objects chosen. Participants were strategic about the mementos that they captured, focusing on things they might otherwise forget:

> I tried to record things which I wouldn't necessarily remember without Memo-ryLane in the long term—so generally people who I see every day I didn't add, but little objects which I might not have in the future I did. I also added things

which have a bit of a story attached to them and detailed that with them (e.g.,
the Festival of Britain coin), so if my own memory fails me I will always have the
MemoryLane and stories to remind me.

In the same way they strategically recorded images of people they were likely to forget
(acquaintances or colleagues) and were less likely to include family or significant others.

At the same time, other mementos clearly evoked strong long-term relationships.
For instance, one participant captured an image of her old piano originally purchased
by her grandfather and passed on in the family. It was the only remaining physical
link to that aspect of her past, and the piano image in MemoryLane served to directly
remind her of him. Another participant originally placed the only picture she had of
her extended family together in the people view of MemoryLane, but later decided
to move it to the home view because that was the place where the framed picture
was located in her real house. Place mementos also reflected this mix of habitual
and unique. Some places reflected long-term habits: One person chose a favorite pub
visited for occasional Sunday lunches with someone special; another chose a nostalgic
panoramic spot in the hills—a favorite place visited with a late best friend. Places
from abroad were mainly reminders of holidays and long trips, for example, a trip
to Machu Picchu as part of traveling across South America or teaching for a year
in Tanzania. These are significant life events associated with locations that people
wanted to preserve and narrate for future reflection and reminiscence.

Our multimedia design also allowed people to attach narratives to digital me-
mentos, making them more evocative than simple physical mementos. People were
clear about the benefits of such accompanying narratives in reevoking memories that
they might otherwise have forgotten: "I think this would be really useful when looking
back at a memento that maybe happened some time ago and had been forgotten. [The
narrative] would bring back the memory straight away."

4. DISCUSSION AND CONCLUSIONS

We began by detailing four design principles for the digital tools to support
memory: Structure, Embodiment, Synergy, and Reminiscence. We now evaluate the
success of our efforts and suggest future research issues.

Selection. We have made important progress in developing new lightweight
methods to impose structure on archives through selection. Our approach preserves
the entire archive but uses the implicit user actions of attention and feedback to
identify important aspects of that archive. However, there are alternatives ways
to impose structure involving simple contextual metadata, unsupervised or supervised
machine learning (Doherty et al., 2011; Doherty & Smeaton, 2008). Many researchers
have pointed to the potential utility of metadata, such as location data to automatically
geotag lifelogs. Our own research shows that providing location indices to SenseCam
data helps recall (Kalnikaite et al., 2010). However, locations need to vary significantly

if they are to provide distinct cues. For example, many family photos are taken in the home and many SenseCam images at work, reducing the distinctiveness of location information. Others are now capturing additional contextual information such as heart rate or motion, but future research needs to examine how such context data might be profitably exploited. This research might be usefully informed by psychological work documenting how cue type affects what is remembered (Hoven & Eggen, 2009; Linton, 1982; Wagenaar, 1994).

Another way to structure archives is to have users actively label archives to support machine learning, for example, face recognition in photos. Aside from the considerable burden of labeling for users, there are practical issues of coverage. To be useful, users have to label broadly, including infrequently encountered people who are more likely to be forgotten. One solution might be to use social feedback (Kalnikaite & Whittaker, 2008). It may be that people will allow close friends and family to tag personal photos or videos, but designs will have to be sensitive to the clear privacy issues this involves (Millen et al., 2007).

Once we have richer event logs of people's lives, it might also be possible to use unsupervised machine learning to identify the objects people encounter, the activities they are engaged in, and whether these events are unusual. Although promising algorithmic work has begun in this area (Doherty et al., 2011), we are some way from demonstrating the benefits of these techniques to users. However, if successful, event analyses could provide rich semantic indices to improve access to digital archives. They also have the benefit of supporting retrieval for events that people did not anticipate having to remember.

Embodiment. Digital archives are often "imprisoned" in people's hard drives. Embodiment attempts to imbue existing physical objects with compelling digital interactivity involving these archives. This area presents some significant research challenges. Our studies of physical mementos (Petrelli et al., 2008) showed that these were used to evoke memory in three main ways: (a) facilitating social narratives and sharing of experience, (b) acting as reminders in frequently used everyday objects, and (c) private immersion in rich collections of emotionally evocative objects. Family Memory Radio explicitly attempted to combine the first two of these by supporting shared narratives around a familiar object. Private recollection is often triggered by unexpected reencounters with mementos, and the Radio has potential for casual rediscovery in large sound collections. Other work addressed immersive memory experiences with "memory boxes" (Frohlich & Murphy, 2000; Stevens, Abowd, Truong, & Vollmer, 2003), or around augmented artifacts (Hoven & Eggen, 2005, 2008), but techniques for supporting each of these separate functions still need refining. Although our FM Radio results were promising, two major challenges remain. First we need better *empirical understanding* of the functions of physical memorabilia (Hoven & Eggen, 2005, 2008; Kirk & Sellen, 2010), allowing digital designs to mimic those functions. Second, there are significant *design challenges* in building engaging augmented objects needed to implement this (Hoven & Eggen, 2008; Petrelli, Villar, Kalnikaité, Dib, & Whittaker, 2010).

Synergy not Substitution. Lifelogging approaches implicitly assume that digital archives are preferred to UM. Our empirical results suggest otherwise, indicating that digital tools are used only when UM is perceived as being weak (Kalnikaite & Whittaker, 2007). We need better accounts about when and where lifelogging will bring genuine benefits. In this paper we demonstrated clear utility for lifelogs in learning situations, where there are high cognitive demands and a strong requirement for accurate recall. However, we need more research to identify other application areas where UM is weak. One promising candidate is *prospective memory*. Although almost all Lifelogging applications focus on people's past (retrospective memory), there is good evidence that people have greater difficulty in remembering *what they intend to do* (Baddeley et al., 2009, Cohen, 1996). Psychological studies can give us a better understanding of when UM is fallible, allowing us to focus on areas of true value to users. A variant of the Synergy Principle is to identify *entire user populations* such as the elderly who could benefit from memory support. One very successful application of SenseCam has been to supporting Alzheimer's patients (Berry et al., 2009). There are many opportunities in the health monitoring and neuroscience areas for this class of application (Barnard et al., 2011), but we still need to identify other areas where access to extensive personal archives will be of value.

Reminiscence and Reflection. With a few exceptions, the primary focus of lifelogging has been on building systems to support factual recollection. However, not only did our empirical research point to the importance of reminiscence and reflection, we were also able to build systems to support these (Kalnikaite & Whittaker, 2011). Here again there are research challenges. Recent work has begun to explore these other forms of memory (Harper et al., 2008; Kirk & Sellen, 2010; Peesapati et al., 2010). Overall, however, much less is known about these aspects of remembering (Pillemer, 1998), making it more difficult to design and evaluate systems that are intended to support these processes. Furthermore, reminiscence and reflection are long-term processes, and it may be that we need to build lasting digital archives to explore these issues in more depth.

4.1. Conclusion

In conclusion, this article has attempted to shift the prevailing technology-centric focus in digital memory tools. Moving beyond technologies that simply rely on "recording everything," our socio-technical approach identified empirically motivated system design principles: Selection, Embodiment, Synergy, and Reminiscence. The principles were effective in leading to the development of new classes of system, as well as showing that systems adhering to the principles were successful in supporting memory. Evaluation of those systems suggests support for Selection and Embodiment principles, but more conceptual and technical work is needed to refine concepts of Synergy and Reminiscence.

NOTES

Authors' Present Addresses. Steve Whittaker, University of California at Santa Cruz, 1156 High St., Santa Cruz, CA 95064. Email: swhittak@ucsc.edu. Vaiva Kalnikaitė, Inter-actables, Studio 151, 23 King Street, Cambridge, CB1 1AH, United Kingdom Email: vaiva@interactables.com. Daniela Petrelli, Sheffield Hallam University, Howard Street, Sheffield, S1 1WB, United Kingdom. Email: d.petrelli@shu.ac.uk. Abigail Sellen, Microsoft Research Cambridge, 7 J J Thomson Ave, Cambridge CB3 0FB, United Kingdom. Email: asellen@microsoft.com. Nicolas Villar, Microsoft Research Cambridge, 7 J J Thomson Avenue, Cambridge CB3 0FB, United Kingdom. Email: nvillar@microsoft.com. Ofer Bergman, Bar Ilan University, Ramat-Gan 52900, Israel. Email: oferbergman@googlemail.com. Paul Clough, Sheffield University, 211 Portobello, Sheffield, S1 4DP, United Kingdom. Email: p.d.clough@shef.ac.uk. Jens Brockmeier, University of Manitoba, Winnipeg MB R3T 2N2, Canada. Email: brockmei@cc.umanitoba.ca.

REFERENCES

Abowd, G. D. (1999). Classroom 2000: An experiment with the instrumentation for a living educational environment. *IBM Systems Journal, 38*, 508–530.

Adar, E., Karger, D., & Stein, L. A. (1999). Haystack: Per-user information environments. *Proceedings of Information and Knowledge Management*. New York, NY: IEEE.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2*, 396–408.

Baddeley, A., Eysenck, M., & Anderson, M. (2009). *Memory*. Hove, UK: Psychology Press.

Bannon, L., & Kuutti, K. (1996). Shifting perspectives on organizational memory: From storage to active remembering. *Proceedings of HICSS–29*. New York, NY: IEEE.

Barnard, P. J., Murphy, F. C., Carthery-Goulart, M. T., Ramponi, C., & Clare, L. (2011). Exploring the basis and boundary conditions of SenseCam-facilitated recollection. *Memory, 19*, 758–767.

Bell, G., & Gemmell, J. (2009). *Total recall*. New York, NY: Dutton.

Bergman, O., Tucker, S., Beth-Marom, R., Cutrell, E., & Whittaker, S. (2009). It's not that important: Demoting personal information of low subjective importance using GrayArea. *Proceedings of CHI 2009*. New York, NY: ACM.

Berry, E., Hampshire, A., Rowe, J., Hodges, S., Kapur, N., Watson, P., . . . Owen, A. (2009). The neural basis of effective memory therapy in a patient with limbic encephalitis. *Journal of Neurology, Neurosurgery, and Psychiatry, 80*, 1202–1205.

Bransford, J., Brown, A. L., & Cocking, R. R. (1999). *How people learn: Brain, mind, experience and school*. Washington, DC: National Academy Press.

Brown, A. 1987. *Metacognition, executive control, self-regulation, and other more mysterious mechanisms*. Mahwah, NJ: Erlbaum.

Chalfen, R. (1987). *Snapshot Versions of Life*. Bowling Green, OH: The Popular Press/The University of Wisconsin Press.

Cohen, G. (1996). *Memory in the real world*. London, UK: Psychology Press.

Conway, M. (1990). *Autobiographical memory*. Milton Keynes, UK: Open University Press.

Csikszentmihalyi, M., & Rochberg-Halton, E. (1981). *The meaning of things—Domestic symbols and the self*. Cambridge, UK: Cambridge University Press.

Doherty, A., Caprani, N., O'Conaire, C., Kalnikaite, V., Gurrin, C., O'Connor, N., & Smeaton, A. (2011). Passively recognising human activities through lifelogging. *Computers in Human Behavior, 27*, 1948–1958.

Doherty, A., & Smeaton, A. F. (2008). Automatically segmenting lifelog data into events. *Proceedings of the 2008 Ninth International Workshop on Image Analysis for Multimedia Interactive Services*. New York, NY: IEEE.

Dourish, P. (2001). *Where the action is*. Cambridge, MA: MIT Press.

Dourish, P., Edwards, W. K., LaMarca, A., & Salisbury, M. (1999). Presto: An experimental architecture for fluid interactive document spaces. *ACM Transactions on Computer–Human Interaction, 6*, 133–161.

Dib, L., Petrelli, D., & Whittaker, S. (2010). Sonic souvenirs: Exploring the paradoxes of recorded sound for family remembering. *Proceedings of CSCW 2010*. New York, NY: ACM.

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. *Proceedings of SIGIR 2003*. New York, NY: ACM.

Ebbinghaus, H. (1964). *Memory: A contribution to experimental psychology*. New York, NY: Dover. (Original work cited 1885)

Ehlen, P., Purves, M., Niekrasz, J., Lee, K., & Peters, S. (2008). Meeting adjourned: Off-line learning interfaces for automatic meeting understanding. *Proceedings of IUI 2008*. New York, NY: ACM.

Frohlich, D., Kuchinsky, A., Pering, C., Don, A., & Ariss, S. (2002). Requirements for photoware. *Proceedings of CSCW 2002*. New York, NY: ACM.

Frohlich, D., & Murphy, R. (2000). The memory box. *Personal and Ubiquitous Computing, 4*, 238–240.

Harper, R., Randall, D., Smyth, N., Evans, C., Heledd, L., & Moore, R. (2008). The past is a different place: They do things differently there. *Proceedings of DIS 2008* (pp. 271–280). New York, NY: ACM.

Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., . . . Wood, K., (2006). SenseCam: A retrospective memory aid. *Proceedings of UbiComp 2006*. Amsterdam, the Netherlands: Kluwer.

Hoven, E. van den, & Eggen, B. (2005). Personal souvenirs as Ambient Intelligent objects. *Proceedings of Soc-EUSAI 2005*. Grenoble, France.

Hoven, E. van den, & Eggen, B. (2008). Informing augmented memory system design through autobiographical memory theory. *Personal and Ubiquitous Computing, 12*, 433–443.

Hoven, E. van den, & Eggen, B. (2009). The effect of cue media on recollections. *Human Technology, 5*, 47–67.

Hsueh, P-Y., & Moore, J. (2009). Improving meeting summarization by focusing on user needs: A task-oriented evaluation. *Proceedings of IUI*. New York, NY: ACM.

Ishii, H. (2008). The tangible user interface and its evolution. *Commun. ACM, 51*, 32–36.

Janin, A., Stolcke, A., Anguera, X., Boakye, K., Çetin, O., Frankel, J., & Zheng, J. (2006). The ICSI-SRI Spring 2006 meeting recognition system. *Proceedings of Multimodal Interaction and Machine Learning*. New York, NY: ACM.

Kalnikaite, V., Sellen, A., Whittaker, S., & Kirk, D. (2010). Now let me see where I was: Understanding how lifelogs mediate memory. *Proceedings of CHI 2010*. New York, NY: ACM Press.

Kalnikaite, V., & Whittaker, S. (2008). Social summarization: does social feedback improve access to speech data? *Proceedings of CSCW 2008*. New York, NY: ACM Press.

Kalnikaite, V., & Whittaker, S. (2007). Software or wetware? Discovering when and why people use digital prosthetic memory. *Proceedings of CHI 2007*. New York, NY: ACM.

Kalnikaite, V., & Whittaker, S. (2010). Beyond being there? Evaluating augmented digital records. *The International Journal of Human–Computer Studies, 69*, 627–640.

Kalnikaite, V., & Whittaker, S. (2011). A saunter down memory lane: Digital reflection on personal mementos. *The International Journal of Human–Computer Studies, 69*, 298–310.

King, G. (1986). *Say "Cheese." The snapshot as art and social history*. London: Collins.

Kirk, D. S., & Sellen, A. (2010). On human remains: Values and practice in the home archiving of cherished objects. *ACM Transactions on Computer–Human Interaction, 17*, 1–43.

Kirk, D., Sellen, A., Rother C., & Wood, K. (2006). Understanding photowork. *Proceedings of CHI 2006*. New York, NY: ACM Press.

Lamming, M., Brown, P., Carter, K., Eldridge, M., Flynn, M., Louie, G., . . . Sellen, A. (1994). The design of a human memory prosthesis. *Computer Journal, 37*, 153–163.

Linton, M. (1982). Transformation of memory in everyday life. In U. E. Neisser (Ed.), *Memory observed: Remembering in natural context* (pp. 107–118). San Francisco, CA: Freeman.

Mann, S. (1997). An historical account of the "WearComp" and "WearCam" projects developed for "personal imaging." *IEEE International Symposium on Wearable Computers*. New York, NY: IEEE.

Marshall, C. (2007). How people manage personal information over a lifetime. In W. Jones & J. Teevan (Eds.), *Personal information management* (pp. 57–75). Seattle: University of Washington Press.

Mayer-Schönberger, V. (2009). *Delete: The virtue of forgetting in the digital age*. Princeton, NJ: Princeton University Press.

Millen, D. R., Yang, M., Whittaker, S., & Feinberg, J. (2007). Social bookmarking and exploratory search. In *Proceedings of ECSCW 2007*. Amsterdam, the Netherlands: Kluwer.

O'Hara, K., Morris, R., Shadbolt, N., Hitch, G. J., Hall, W., & Beagrie, N. (2006). Memories for life: a review of the science and technology. *Journal of the Royal Society Interface, 3*, 351–365.

Peesapati, S. T., Schwanda, V., Schultz, J., Lepage, M., Jeong, S., & Cosley, D. (2010). Pensieve: Supporting everyday reminiscence. *Proceedings of CHI 2010*. New York, NY: ACM.

Petrelli, D., Villar, N., Kalnikaité, V., Dib, L., & Whittaker, S. (2010). FM radio: Family interplay with sonic mementos. *Proceedings of CHI 2010*. New York, NY: ACM.

Petrelli, D., & Whittaker, S. (2010). Family memories in the home: contrasting physical and digital mementos. *Personal and Ubiquitous Computing, 14*, 153–169.

Petrelli, D., Whittaker, S., & Brockmeier, J. (2008). Autotopography: What can physical mementos tell us about digital memories? *Proceedings of CHI 2008*. New York, NY: ACM.

Pillemer, D. (1998). *Momentous events, vivid memories: How unforgettable moments help us understand the meaning of our lives*. Cambridge, MA: Harvard University Press.

Renals, S., Hain, T., & Bourlard, H. (2007). Recognition and interpretation of meetings: The AMI and AMIDA projects. *Proceedings of IEEE ASRU*. New York, NY: IEEE.

Schacter, D. L., Coyle, J. T., Fischbach, G. D., Mesulam, M. M., & Sullivan, L. E. (Eds.). (1995). *Memory distortion: How minds, brains and societies reconstruct the past*. Cambridge, MA: Harvard University Press.

Sellen, A., Fogg, A., Aitken, M., Hodges, S., Rother, C., & Wood, K. (2007). Do life-logging technologies support memory for the past? *Proceedings of CHI 2007*. New York, NY: ACM.

Sellen, A., & Whittaker, S. (2010). Beyond total capture: A constructive critique of lifelogging. *Communications of the ACM, 53*(5), 70–77.

Stevens, M. M., Abowd, G. D., Truong, K. N., & Vollmer, F. (2003). Getting into the Living Memory Box: Family archives & holistic design. *Personal and Ubiquitous Computing, 7*, 210–216.

Tucker, S., Bergman, O., Ramamoorthy, A., & Whittaker, S. (2010). Catchup: A useful application of time-travel in meetings. *Proceedings of CSCW 2010.* New York, NY: ACM.

Tucker, S., & Whittaker, S. (2006). Time is of the essence: An evaluation of temporal compression algorithms. *Proceedings of CHI 2006.* New York, NY: ACM.

Van House, N. A. (2009). Collocated photo sharing, story-telling, and the performance of self. *The International Journal of Human–Computer Studies, 67*, 1073–1086.

Vemuri, S., Schmandt, C., Bender, W., Tellex, S., & Lassey, B., (2004). An audio-based personal memory aid. *Proceedings of UbiComp 2004.* Amsterdam, the Netherlands: Kluwer.

Wagenaar, W. A. (1994). Is memory self-serving? In U. Neisser & R. Fivush (Eds.), The remembering self: Construction and accuracy in the self-narrative (pp. 191–204). Cambridge, UK: Cambridge University Press.

Whittaker, S. (2010). Digital curation. *Annual Review of Information Science and Technology*, 1–33.

Whittaker, S., Bergman, O., & Clough, P. (2010). Easy on that trigger dad: A study of long-term family photo retrieval. *Personal and Ubiquitous Computing, 14*, 31–43.

Whittaker, S., Hyland, P., & Wiley, M. (1994). Filochat: Handwritten notes provide access to recorded conversations. *Proceedings of CHI 1994*. New York, NY: ACM Press.

Whittaker, S., Tucker, S., Swampillai, K., and Laban, R. (2008). Design and Evaluation of Systems to Support Interaction Capture and Retrieval. *Personal and Ubiquitous Computing.* Retrieved from http://dx.doi.org/10.1007/s00779-007-0146-3