

DOI:10.1145/1735223.1735243

Rather than try to capture everything, system design should focus on the psychological basis of human memory.

BY ABIGAIL SELLEN AND STEVE WHITTAKER

Beyond Total Capture: A Constructive Critique of Lifelogging

WHAT IF WE could digitally capture *everything* we do and see? What if we could save every bit of information we touch and record every event we experience? What would such a personal digital archive be like, and how might it affect the way we live? This vision of a complete “lifelog” is the holy grail for many technologists and researchers who consider us to be on the brink of an “e-memory” revolution.

In the past few years, capturing “Memories for Life” has become a U.K. Grand Challenge in Computing (http://www.nesc.ac.uk/esi/events/Grand_Challenges/proposals/), and many research programs today are dedicated to developing technologies to support the archiving of vast amounts of personal data. A 2009

book² by Gordon Bell and Jim Gemmell outlined an enthusiastic view of a future in which technology enables “total recall” of our lives through “total capture” of personally relevant information. Such information includes the paper and digital documents we work on or look at; email, paper mail, and instant messages sent and received; content of telephone conversations; Web sites visited; and charge-card transactions. Also included are data from other everyday activities (such as still images, video, ambient sound and location data). Finally, these personal archives might also be supplemented with environmental measures (such as light intensity and temperature variation) and even internal, biosensor data (such as heart rate and galvanic skin-response measures) reflecting our physical and emotional state.


Constructing such a diverse archive of personal information requires a range of technologies for its capture, management, and storage. Today’s advances in wearable sensors, networking capabilities, and massive increases in digital-storage capacity mean this vision is feasible, fueling enthusiasm for the possibilities offered by the technology itself.

Further impetus comes from speculation about what a comprehensive lifelog might do and how it might change our lives. As outlined in 2006 by Czerwinski et al.,⁹ lifelog-

>> key insights

- **Focusing on “total capture,” current approaches to lifelogging have failed to explore what practical purpose such exhaustive personal digital records might actually serve.**
- **Evaluating new approaches, psychology has emerged as an underexploited resource in defining the nature of human memory and its key processes and weaknesses.**
- **Psychology as design framework could help define the types of memory such systems should support, along with their key interface properties and need to work in synergy with human memory, rather than as its replacement.**

ILLUSTRATION BY GIACOMO MARCHESI



ging could permanently change how we use and share personal data, enabling us to look back over our lives or search through and organize past experiences. It could help us find lost objects, recall names, retrieve details in documents, and review discussions in meetings. It might also offer new ways of sharing data with those we care about or offer up data to interested parties. While there are knotty privacy and security implications (a huge topic, not addressed here), the potential benefits warrant substantial programs of research and development. The vision is indeed compelling. Some proponents have even said that new technologies could give each of us a comprehensive set of “digital memories” to augment or even replace their biological counterparts.²

But how sound are these goals? What benefits would the effort bring us? In reflecting on these questions we see we need a more focused and human-centered research agenda to realize the grand ambition. This, in turn, entails moving away from an obsession with “capturing everything” to a more precise specification of what it means to support human memory, leading to specific systems and concrete design implications.

History

Despite the recent upsurge of interest, lifelogging systems are not new. They can be traced back to Vannevar Bush’s 1945 “Memex” vision (a sort of desk) supporting the archiving, searching, and indexing of personal information stores revolving around documents with which we interact.⁷ Since then, it has inspired many different systems, though mainly in research laboratories. Early forays were close to the original vision, confined to capturing digital objects within an integrated system. Both past¹¹ and present¹² examples are in fact infrastructures for storing collections of heterogeneous digital objects that users generate or encounter, including documents, photos, Web pages, and email. And as infrastructures have developed, so too have the tools for searching, browsing, and retrieving information from the collections. The tools often make use of metadata about the various aspects of an object’s past context, or provenance (such as when and where it


was created, who it came from, and its relationships with other objects), making it easier to re-find information (such as Dumais et al.¹⁰).

However, today’s lifelogging vision extends beyond the mere storage of desktop objects. Just as we can capture and collect personal interactions with documents, we can capture activities away from the computer, out of the office, in the everyday world. Key to it all is that many everyday activities can be captured automatically and comprehensively through digital tools that allow us to not only store important content, but also contextual details of the activities to help access the content

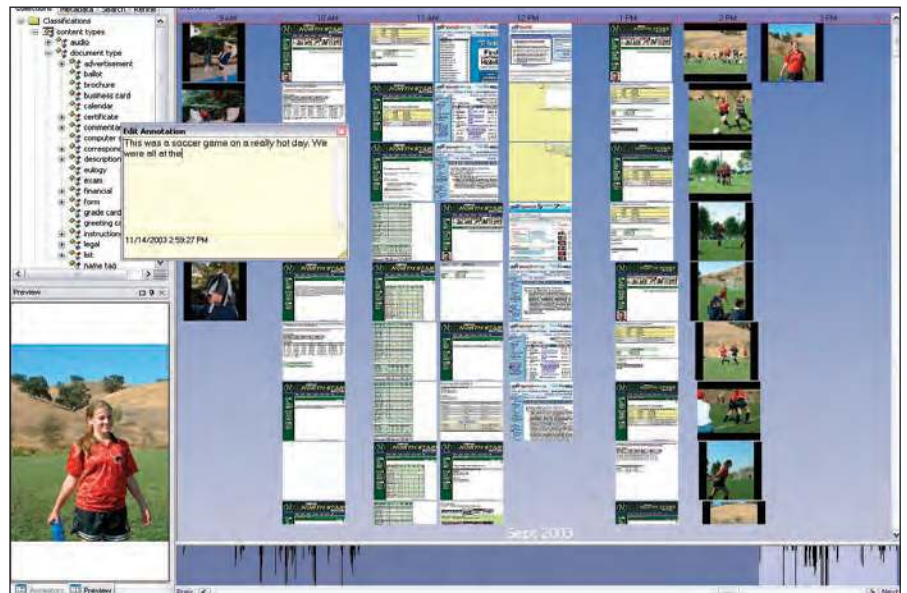
later. Note that we distinguish between lifelogging and other more deliberate activities involving the capture of personal data (such as digital photography and blogging) that involve the *effortful* selective capture and display of digital materials for a particular audience. In contrast, lifelogging seeks to be *effortless* and all-encompassing in terms of data capture.

There are two main classes of lifelogging system:

Total capture. For this one, the target is a complete record of everyday life, capturing as many kinds of data as possible, as continuously as possible. Data types range from documents to



Vannevar Bush's 1945 “Memex” vision.



MyLifeBits by Gordon Bell is a lifetime store of everything, aiming to fulfill the Memex vision.


images to videos to sound to location to ambient temperature, along with light levels and biosensor readings. In early systems (such as Lamming et al.¹⁷), capturing people's locations was a focus, involving users carrying or wearing small devices tracked by localized networked sensors. More recent wearable systems use head-mounted still and video cameras,¹⁵ sensor-triggered still cameras worn around the neck¹⁴ (see the figure here), and audio-capture devices.²⁶ Yet others rely on instrumented environments that capture human activity through sensors or networks, as in MIT's "PlaceLab" (http://architecture.mit.edu/house_n/placelab.html).

Situation-specific capture. These are more limited in scope, aiming to capture rich data in specific domains involving complex information. They can be viewed as a specialized form of lifelogging, aiming to record multiple kinds of data as completely and automatically as possible for specific activities or in particular places where the activity occurs. Most focus on recording activities during meetings, lectures, or other forms of work-related conversation, allowing "organizational knowledge" to be browsed and searched.²⁸ Early systems involved the simple indexing of recorded audio and pen-stroke data. More recent technology-enhanced meeting rooms capture video from multiple cameras and microphones, combining it with whiteboard content, slide capture, and digital pen strokes. Often included is software that automatically summarizes and extracts key events from the data.


Defining the Benefits: Five Rs

Surprisingly, many lifelogging systems lack an explicit description of potential value for users, focusing instead on technical challenges (such as data capture and retrieval mechanisms). When claims are made about potential benefits, they tend to be ill-defined. Nevertheless, it is important to clarify what these claims might be. Here, we outline potential benefits for memory by describing the ways such systems might support "the five Rs," or the activities we call recollecting, reminiscing, retrieving, reflecting, and remembering intentions:

Recollecting. Technology could help



Surprisingly, many lifelogging systems lack an explicit description of potential value for users, focusing instead on technical challenges.



us mentally "re-live" specific life experiences, thinking back in detail to past personal experiences (often called "episodic" memories²⁵). Remembering aspects of a past experience can serve many practical purposes; examples include locating lost physical objects by mentally retracing our steps, recollecting faces and names by recalling when and where we met someone, or remembering the details of what was discussed in a meeting.

Reminiscing. As a special case of recollection, lifelogs could also help users re-live past experiences for emotional or sentimental reasons. This can be done by individuals or socially in groups, as an (often pleasurable) end in itself; examples are watching home videos and flipping through photo albums with friends and family.

Retrieving. Lifelogs promise to help us retrieve specific digital information we've encountered over the years (such as documents, email, and Web pages). Retrieval can include elements of recollection; for example, retrieving a document might require remembering the details of when we wrote it, when we last used it, or where we put it. Alternatively, retrieval might depend on inferential reasoning (such as trying to deduce keywords in a document or thinking about the document's other likely properties, like type and size). Pondering information properties need not involve recollection of past experiences at all, as long as other ways are available for finding the desired information.

Reflecting. Lifelogs might support a more abstract representation of personal data to facilitate reflection on, and reviewing of, past experience. Reflection might include examining patterns of past experiences (such as about one's behavior over time). Such patterns might provide useful information about our general level of physical activity or emotional states in different situations, allowing us to relate it to other data about, say, our health. Alternatively, reflection might involve looking at one's past experiences from different angles or perspectives. Here, the value is not in re-living past events (as in recollecting) but in seeing things anew and framing the past differently.¹³ The value is less about memory per se than it is about

learning and self-identity.

Remembering intentions. Another type of activity concerns remembering prospective events in one's life ("prospective memory"), as opposed to the things that have happened in the past. Our everyday activities require that we constantly defer actions and plan future activities; examples include remembering to run errands, take medication, and show up for appointments.

A careful analysis of the lifelogging literature for proposed user value suggests a general focus on the processes of recollection and retrieval, but these benefits are usually implied rather than explicit.

Evaluation

Do lifelogging systems deliver these benefits? One way to identify evidence of utility is through systematic evaluation (such as focused lab studies and studies of longer-term real-world use). However, such evaluations have yet to provide overwhelming evidence of effectiveness, and extended usage studies show little impact outside research labs. Worse, few lifelogging systems are in widespread use, even within the research laboratories that developed them.

In practice, total-capture systems have been used by only a small number of people (often those with direct investment in the technology), and lifelogging infrastructures^{11,12} are not in widespread use. While Gordon Bell of Microsoft² and Steve Mann of the University of Toronto¹⁸ have both "lived the vision" by recording many aspects of their everyday lives, they are unusual in the extreme extent of their engagement with lifelogging. Otherwise, there are few instances of full-fledged use of total capture.

More controlled system evaluations are also not encouraging in terms of demonstrable utility. The SenseCam (see the figure), which captures a series of still images based on movement and changes in the intensity of light and heat, has been shown to support the recollection of everyday experience, as well as retrieval of information about past events.²³ However, the same study showed that the capacity for these images to help people recollect their past experience rapidly decreased after only three months, casting doubt on whether such devices can support longer-term recollection.

Most evaluations described in the literature have examined situation-

specific capture systems. For example, FiloChat is an application that allows users to access spoken information from meetings via personal digital notes. Both lab experiments and a field trial at Hewlett-Packard Labs in the early 1990s demonstrated FiloChat's superiority in supporting retrieval of meeting information compared with traditional techniques (such as pen-and-paper notes and Dictaphone recordings). An active user group reported positive reactions to the system because it allowed them to generate accurate meeting minutes. Similar findings were reported in other field studies of related meeting-capture technologies; see Whittaker et al.²⁸ for a review. But despite such early positive results, more recent research should make us skeptical, suggesting that records may be less useful than we might first think. For example, lecture recordings don't significantly improve student grades,¹ and evaluations of meeting-capture systems have shown little uptake of sophisticated records.²⁰

Other research confirms the view that digital archives may be generally less valuable than people would hope. Even when—contrary to lifelogging principles—we deliberately choose to save digital memorabilia, we seldom access them. Studies by Petrelli and Whittaker¹⁹ in 2010 of digital family memorabilia (such as photos, videos, scanned images, and email) showed that digital archives are rarely accessed. In the same study, when people with large collections of digital memorabilia were asked to select objects of mnemonic significance in their homes, less than 2% of the objects they reported were digital.¹⁹ Other work has found that users with collections of thousands of digital photos never access the majority of them.²⁷

While this lack of interest in digital data doesn't imply that all such archives have little value, it raises questions about their utility in remembering or reviewing the past. One might surmise that we simply lack effective techniques for accessing the archives. However, most situation-specific capture systems for meetings and lectures include sophisticated access tools (such as multimedia search and bespoke browsing interfaces), suggesting



Microsoft SenseCam captures a series of still images triggered by changes in activity (such as movement) and includes sensors that capture other kinds of data (such as ambient light levels and temperature); courtesy Microsoft Research Cambridge.

that having poor access is likely not a general explanation. Other researchers have argued that new access tools (such as desktop search) will facilitate exploitation of digital archives.²¹ However, research at the University of Sheffield in 2008 on desktop search suggests it is used only infrequently. More important, there is no consistent evidence that improving the quality of search leads to increased use of search tools.³ These results indicate that desktop search is not a “silver bullet” leading to effective access to and exploitation of our personal digital archives.


In general, these findings imply that archival data may be less valuable than the considerable effort expended on these systems would justify.

Design Principles


How can we prevent creating underused infrastructures or only proof-of-concept demonstrators? Needed is a new conceptual framework that is better focused on the functions lifelogging technologies could serve. Despite the memory terminology used in lifelogging work, little attention seems to focus on human memory and how it operates. Psychological studies of memory are largely ignored, even though they provide relevant concepts and results that lead directly to new design principles:

Strategically targeting the weaknesses of human memory. Total-capture systems are indiscriminate, assuming that all kinds of data are equally valuable and the more data the better. The argument often goes that we should capture “as much as we can” because we never know what we might need to remember in the future. But this “just-in-case” principle has two weaknesses: First, we can never capture absolutely everything, so choices must indeed be made when designing and building systems; for example, different kinds of data require different kinds of sensors or capture devices, adding complexity for the people using and building the systems. Second, capturing vast arrays of data might overwhelm end users maintaining and retrieving valuable information from large archives; it also ignores the burden huge amounts of data impose on system designers and developers.

Previous research provides future guidance; for example, psychology re-



Collections of digital data can serve as cues to trigger autobiographical memory about past events but are not memories in themselves.



search provides a deeper understanding of the most frequent and critical kinds of memory problems people have, allowing system designers to focus on areas of true value to users. This means that, rather than the overambitious goal of “logging everything,” creators of lifelogging systems would be better off directing their efforts at the kinds of data people find most valuable and the issues they find most problematic. In addition to the problem of the transience of memory (implicitly addressed by much lifelogging technology), people are subject to myriad other distortions and inaccuracies in memory (such as absentmindedness, blocking, misattribution, suggestibility, bias, and persistence).²² And while almost all lifelogging applications focus on supporting people’s past (retrospective) memory, strong evidence indicates that people have greater difficulty remembering what they intend to do in the future (prospective memory).²⁴ Other memory studies have identified specific groups (such as Alzheimer’s patients and aging populations) with debilitating memory conditions and demonstrated how visual-recording technologies (such as SenseCam) can be of help.⁴

Not “capturing experience” but designing effective retrieval cues. The language used by lifelog proponents sometimes conflates cueing with experiential capture. This distinction may seem obvious but is worth restating. Collections of digital data (such as sets of digital photos and sounds) can serve as cues to trigger autobiographical memory about past events but are not memories in themselves or in any way facsimiles of personal experience. Following this principle, we are thus better able to address the precise mechanisms by which cues help memory. For example, metadata can help cue retrieval of lost files by, say, providing contextual information about who wrote a document and when it was written. Alternatively, information in the digital archive may itself serve to cue a forgotten memory (such as when a stored digital photo prompts reminiscence about a previously forgotten incident).

There is again highly relevant psychology research detailing how different cues (such as time, place, people, and events) trigger autobiographical memories, suggesting (in contrast to

the design of many lifelogging user interfaces) that place, events, and people are stronger cues than time.⁵ Other research has found that cues can trigger inferences about what *must* have happened in someone's life, rather than genuine recollection²³; for example, a photo showing us having lunch with friends may cause us to truly remember the event or simply lead us to conclude we must have been there. Finally, while recollecting the past is highly dependent on the kind of cues presented to people, for prospective memory (memory for future events and intentions), the important issue is not so much the type of cue but rather when a cue is delivered, allowing an intention to be remembered at the right time and place.²⁴ This observation suggests that the capture of data (such as location or other contextual cues) might be used to trigger reminders rather than provide content to be remembered.


These observations run counter to much of the rhetoric surrounding lifelogging, where such phrases as “digital memories” and “experience capture” are often used. They show instead the importance of understanding the precise relationship between the cues we are able to capture through lifelogging technologies and the memory experiences they trigger, with clear implications for how we might design improved systems.

Support for the Five Rs. Most lifelogging proponents presume that the systems deliver benefits without being specific about what the benefits might be, assuming a one-system-suits-all approach. As we outlined earlier in the section on lifelogging benefits, the psychological literature distinguishes many different types of memory, each involving different retrieval processes. Thus, it is not simply a question of what the system captures but determining *how* such a system would be used. This determination depends largely on the type of memory being targeted or, more generally, the kind of benefit system designers hope to deliver to end users.

Many total-capture systems implicitly seem to address *recollection*, or remembering past personal experiences. It is well known in the psychological literature that there are strong connections between these autobiographical memories and visual images.⁸ This



Systems supporting reminiscence and reflection have received far less attention than those supporting recollection and retrieval.



suggests that the interfaces for such systems should focus on images as the backbone of their design. In contrast, systems for *retrieval* need not be concerned with recollection, but rather with efficient ways of searching through large heterogeneous collections of data and so provide access to metadata that might support effective search. If system designers decide to support *reminiscence*, other kinds of factors become important (such as optimizing the sharing of data with others). Systems for *reflection* might be different still where abstraction is important, offering flexible and novel methods for viewing personal data in ways that might surprise, provoke, or educate users. Finally, designing systems to support *remembering intentions* need to focus on delivering timely cues in appropriate contexts if they are to provide effective reminders.

Applying such memory taxonomies is vital for designing effective systems. First, they clarify the aspects of memory the systems are trying to support; without such clarification it is difficult to know whether the systems succeed. Second, understanding the relevant psychological processes allows designers to create systems to better support the processes. Third, taxonomies can suggest new directions for exploration; for example, systems supporting reminiscence and reflection have received far less attention than those supporting recollection and retrieval.

Offloading and metamemory. Much of the lifelogging approach is motivated by the observation that human memory is fallible. Lifelog proponents argue we need to remove the memory burden from humans, offloading it to reliable and comprehensive external digital stores. These claims need careful scrutiny, as we calculate the costs associated with digital archives, even if they make it possible to store vast amounts of data. For example, the effort required to capture, create, and maintain some kinds of data can be prohibitive. Moreover, accessing data can be inefficient compared with exploiting “organic” human memory.

How, when, and why people exploit external memory tools has been studied extensively in the psychological subfield of metamemory,⁶ which addresses people's understanding of the strengths and weaknesses of their

own memories and the strategies they employ to overcome the weaknesses. Kalnikaite and Whittaker¹⁶ looked at the interplay between organic memory and metamemory in determining when people choose to access digital conversational records, showing that even when a digital record is accurate and complete, users do not rely on it if they feel they can remember the information unaided. The decision to use a digital-memory aid also depends on the efficiency with which a memory aid can be accessed. Indeed, the study found that efficiency of access sometimes overrides accuracy, with subjects being willing to settle for less-than-perfect accuracy, as long as the method is quick and easy to use.

Lifelogging applications must be better at analyzing these trade-offs. When should people use efficient but fallible organic memory instead of less efficient but potentially more accurate digital records? Rather than seeing lifelogs as replacing memory, system designers would be better off viewing them as working in synergy with organic memory.

Lessons and Research Questions

For lifelogging research, prior work offers four key insights:

Selectivity, not total capture. Rather than unfocused efforts to “capture everything,” system designers should channel their efforts more fruitfully by identifying the situations where human memory is poor or targeting the things users most want to remember. These situations are where the systems would provide their greatest utility. Furthermore, people may sometimes want to *forget*, a view anathema to the lifelogging philosophy.

Cues not capture. System designers must be explicit about the fact that these systems do not “capture experiences” but instead provide cues that might trigger different kinds of memories. This is not a simple matter of infelicity in language, pointing instead to the need to better understand cueing processes to build systems that genuinely support user requirements for memory support.

Memory refers to a complex, multi-faceted set of concepts. There are different types of memory, and system designers must clarify the aspects of memory they are targeting (such as recollection,

reminiscence, retrieval, reflection, and remembering intentions). Greater clarity should produce systems that better fit their intended purpose and support the user’s relevant cognitive processes. Unless system designers know precisely what their systems are intended to do, they can’t determine whether their designs are successful.

Synergy not substitution. Much of the rhetoric on behalf of lifelogging assumes the systems will replace human memories. However, digital records are used only when they are efficient to access and when users feel their own memory is unreliable. Rather than try to replace human memory with digital systems, system designers should look to capitalize on the strengths of human memory and help overcome its weaknesses. Lifelogging design must therefore be engineered to work in synergy with users’ own memories.

Incorporating the psychology of memory into the design of novel mnemonic technologies opens up exciting possibilities for ways to augment and support human endeavors. In light of the diversity of human memory and how it plays out in everyday life, we’ve sought to outline a design space that exploits the strong body of psychological knowledge we already have in hand. **□**

References

1. Abowd, G. Classroom 2000: An experiment with the instrumentation of a living educational environment. *IBM Systems Journal (Special Issue on Pervasive Computing)* 38, 4 (Oct. 1999), 508–530.
2. Bell, G. and Gemmell, J. *Total Recall: How the E-Memory Revolution Will Change Everything*. Dutton, New York, 2009.
3. Bergman, O., Beyth-Marom, R., Nachmias R., Gradovitch, N., and Whittaker S. Improved search engines and navigation preference in personal information management. *ACM Transactions on Office Information Systems* 26, 4 (Sept. 2008), 1–24.
4. Berry, E., Kapur, N., Williams, L., Hodges, S., Watson, P., Smyth, G., Srinivasan, J., Smith, R., Wilson, B., and Wood, K. The use of a wearable camera: SenseCam as a pictorial diary to improve autobiographical memory in a patient with limbic encephalitis. *Neuropsychological Rehabilitation* 17, 4/5 (2007), 582–681.
5. Brewer, W. Memory for randomly sampled autobiographical events. In *Remembering Reconsidered: Ecological and Traditional Approaches to the Study of Memory*, U. Neisser and E. Winograd, Eds. Cambridge University Press, Cambridge, U.K., 1988, 21–90.
6. Brown, A. Metacognition, executive control, self-regulation, and other mysterious mechanisms. In *Metacognition, Motivation, and Understanding*, F.E. Weinert and R.H. Kluwe, Eds. Lawrence Erlbaum Associates, Hillsdale, NJ, 1987, 65–116.
7. Bush, V. As we may think. *Atlantic Monthly* 176, 1 (July 1945), 101–108.
8. Conway, M. *Autobiographical Memory*. Open University Press, Milton Keynes, U.K., 1990.
9. Czerwinski, M., Gage, D., Gemmell, J., Marshall, C., Perez-Quinones, M., Skeels, M., and Catarci, T. Digital memories in an era of ubiquitous computing and abundant storage. *Commun. ACM* 49, 1 (Jan. 2006), 44–50.

10. Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin R., and Robbins D. Stuff I’ve seen: A system for personal information retrieval and re-use. In *Proceedings of the Conference of the Special Interest Group on Information Retrieval* (Toronto, 2003). ACM Press, New York, 2003, 72–79.
11. Freeman, E. and Fertig, S. Lifestreams: Organizing your electronic life. In *Proceedings of the AAAI Fall Symposium on AI Applications in Knowledge Navigation and Retrieval* (Cambridge, MA, Nov.). AAAI Press, Menlo Park, CA, 1995.
12. Gemmell, J., Bell, G., and Lueder, R. MyLifeBits: A personal database for everything. *Commun. ACM* 49, 1 (Jan. 2006), 88–95.
13. Harper, R., Randall, D., Smyth, N., Evans, C., Heledd, L., and Moore, R. The past is a different place: They do things differently there. In *Proceedings of the Conference on Designing Interactive Systems*. ACM Press, New York, 2008, 271–280.
14. Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., and Wood, K. SenseCam: A retrospective memory aid. In *Proceedings of Ubicomp* (Orange County, CA, Sept.). Springer, Berlin, 2006, 177–193.
15. Hori, T. and Aizawa, K. Context-based video retrieval system for the lifelog applications. In *Proceedings of the ACM Workshop on Multimedia Information Retrieval* (Berkeley, CA, Nov.). ACM Press, New York, 2003, 31–38.
16. Kalnikaite, V. and Whittaker, S. Software or wetware? Discovering when and why people use digital prosthetic memory. In *Proceedings of the Conference on Human Factors in Computing Systems* (San Jose, CA, Apr.–May). ACM Press, New York, 2007, 71–80.
17. Lamming, M., Brown, P., Carter, K., Eldridge, M., Flynn, M., Louie, G., Robinson, P., and Sellen, A. The design of a human memory prosthesis. *Computer Journal* 37, 3 (Jan. 1994), 153–163.
18. Mann, S. Wearable computing: A first step toward personal imaging. *Computer* 30, 2 (Feb. 1997), 25–32.
19. Petrelli, D. and Whittaker, S. Family memories in the home: Contrasting physical and digital mementos. *Personal and Ubiquitous Computing* 14, 2, (Feb. 2010), 153–169.
20. Richter, H., Miller, C., Abowd, G., and Funk, H. Tagging knowledge acquisition to facilitate knowledge traceability. *International Journal on Software Engineering and Knowledge Engineering* 14, 1 (Feb. 2004), 3–19.
21. Russell, D.M. and Lawrence, S. Search everything. In *Personal Information Management*, W. Jones and J. Teevan, Eds. University of Washington Press, Seattle, 2007.
22. Schacter, D. *The Seven Sins of Memory: How the Mind Forgets and Remembers*. Houghton Mifflin, New York, 2001.
23. Sellen, A., Fogg, A., Hodges, S., Rother, C., and Wood, K. Do lifelogging technologies support memory for the past? An experimental study using SenseCam. In *Proceedings of the Conference on Human Factors in Computing Systems* (San Jose, CA, Apr.–May). ACM Press, New York, 2007, 81–90.
24. Sellen, A., Louie, G., Harris, J., and Wilkins, A. What brings intentions to mind? An in situ study of prospective memory. *Memory* 5, 4 (July 1997), 483–507.
25. Tulving, E. *Elements of Episodic Memory*. Oxford University Press, New York, 1983.
26. Vermuri, S., Schmandt, C., Bender, W., Tellex, S., and Lassey, B. An audio-based personal memory aid. In *Proceedings of Ubicomp* (Nottingham, U.K., Sept.). Springer, Berlin, 2004, 400–417.
27. Whittaker, S., Bergman, O., and Clough, P. Easy on that trigger dad: A study of long-term family photo retrieval. *Personal and Ubiquitous Computing* 14, 1 (Jan. 2010), 31–43.
28. Whittaker, S., Tucker, S., Swappillai, K., and Laban, R. Design and evaluation of systems to support interaction capture and retrieval. *Personal and Ubiquitous Computing* 12, 3 (Mar. 2008), 197–221.

Abigail J. Sellen (asellen@microsoft.com) is a principal researcher at Microsoft Research in Cambridge, U.K., and Special Professor of Interaction at the University of Nottingham, U.K.

Steve Whittaker (whittaker@almaden.ibm.com) is a research scientist at IBM Almaden Research Center, San Jose, CA.