Next-generation personal memory aids

S Vemuri and W Bender

Personal memory assistance is a natural application of ubiquitous computing. Portable computers are decreasing in size, increasing in capability, and the barriers constraining earlier computer-based memory aids are rapidly diminishing. With each engineering advance, a new generation of 'personal memory aids' is enabled. This paper presents iRemember, our prototype wearable 'memory prosthesis'. It allows the wearer to capture and accrue daily experiences (primarily audio) and attempts to remedy a limited set of common memory problems by providing tools to find memory triggers within such collections. This is an early step in an area ripe for growth and controversy. The social and legal implications of ubiquitous recording are discussed and additional memory aids designed to address other memory problems are proposed.

> I get mad at myself when I'm sitting there trying to write, and I want to recall a specific statement, a specific fact, a name, and it doesn't come immediately. I hate to research something that ought to be right there at the press of a little button in my mental computer.
>
> Walter Cronkite [1]

1. Introduction

Imagine being able to reminisce about one's childhood with additional vividness and clarity. Imagine students listening diligently to a lecturer instead of feverish note-taking while maintaining (or even improving) their ability to absorb the material. Imagine being able to remember people's names better. Imagine being able to remember a perfect quote, anecdote, or joke at just the right moment in a conversation. In this paper, we describe how non-invasive, wearable computers can help realise these potentials through assistive memory enhancement.

Human memory is generally poor, prone to error and to manipulation. It fails in unpredictable ways, at the most inconvenient times, and sometimes, with dire consequences. Despite this, most people function by conceding to memory fallibility.

Faulty memory can have societal impacts. For example, memory failures often result in time and resource wastage. Memory errors, both individual and organisational, are often blamed for airline, environmental, and industry disasters [2]. Innocent lives have been ruined as a result of faulty memory conveyed as eyewitness testimony in courts of law [3]. People are poor at remembering decision rationales [4]; in order to learn from your mistakes it is necessary to remember them.

Memory aids have existed for as long as one can remember. Strings tied on fingers, Post-it® notes, and Palm Pilots™ are all examples. These aids can be characterised as either being mechanisms for reminding you to remember or as off-body repositories for the explicit storing of memories for subsequent retrieval.

human memory is generally poor, prone to error and to manipulation

A new generation of memory aids is now possible, facilitated in large part by miniaturisation of computing devices that can store a lifetime's worth of data coupled with retrieval techniques to make such archives a useful repository for memory assistance. It is this new generation of memory aids that is the subject of this paper.

Most current memory aids require active effort to engage the aid in order for the memory to be triggered and few take advantage of the plethora of potentially valuable content and context from everyday experiences that may serve to trigger memories. Such triggers can be derived by capturing and indexing readily available data sources including audio, video, biometrics, physical location, nearby people, electronic correspondence, computer activity, etc. One fortuitous quality of such sources is that they can be captured passively and with relative ease using computers with appropriate sensors.

The idea of recording everything in one's life is not new. One of the earliest proposals, Memex, dates back to 1945 [5]. More than a decade ago, ubiquitous audio-recording systems were built and studied, using desktop computers and workstations [6]. Portability is now making ubiquity possible and it follows that more and more industry, academic, and government groups are investigating the possibility of recording everything in one's life [7—11].

Indeed, ubiquitous recording for the purpose of memory assistance falls within the domain of Weiser's ubiquitous computing vision [12]. But, recording everything is the easy part. The interesting challenge is turning vast repositories of personal recordings into a useful resource while respecting the social, legal, and ethical ramifications of ubiquitous recording.

1.1 Approach

Everyday we forget important things — we are engaged in some task, we forget something, we want to remember, and we are willing to expend some time and effort towards recalling the lost or blocked memory. In certain cases, there are places to look to help resolve the 'memory problem'. If the forgotten item is published somewhere (or otherwise documented as e-mail, a Web page, a blog, a document on your computer, etc), then computer-based search engines can be used to find it.

Frequently, the forgotten item was part of an oral communication — formal exchanges such as meetings and lectures or less-formal conversations such as impromptu chats among friends or colleagues. A large part of our daily lives, even in the workplace, is spent in face-to-face conversations [13], which are unlikely to be recorded. We assert that these oral communications can be valuable sources for helping remedy everyday memory problems.

once the data is captured, the problem is transformed into one better-suited to computers than people

To minimise the chance of missing potentially valuable information, our approach is one of vigilant, but nearly passive documentation. This includes capturing verbatim as much as we can of daily life. The accrual of such data may not be that hard. For example, lectures and meetings are increasingly likely to be recorded and archived. When the environment does not afford such facility, high-capacity, low-cost, portable recording devices can fill in the gaps.

While many things can be captured, we focus primarily, but not exclusively, on audio, since it is integral to face-to-face communication, it is content-rich, there are signal- and speech-processing tools available to index it, and it is easily captured with portable devices. Once the data is captured, the problem of remedying memory problems is transformed into one better-suited to computers than people — searching large collections of data using information retrieval techniques.

In contrast to information retrieval, 'memory retrieval' can benefit from one's remembrance of past-witnessed events to help narrow the search; often we need only find a trigger rather than the original content in order to remember. But, using memory in retrieval tasks can be a 'double-edged sword'. The accurate recollection of a memory can help narrow a search, whereas inaccurate recollection can hurt by leading us astray. Also, culling content and meaning from audio (e.g. speech recognition and audio scene analysis) in support of content-based searches of audio can be problematic due to the inherently noisy and inaccurate analyses, especially under the heterogeneous and acoustically less-than-ideal circumstances typical of everyday life.

The main research question addressed in this paper is how can long-term personal-data archival be effectively used for memory assistance? Some sub-questions are included here.

- Are there certain technologies within the 'memory retrieval' approach more or less useful to the ultimate goal of helping remedy memory problems? A few technologies have already been mentioned (e.g. information retrieval, speech recognition). More will be mentioned in section 2.
- Is the 'memory retrieval' approach more or less useful for particular types of memory problems? The next section discusses human memory and some common memory problems. We focus on a subset of these, but are also interested if the memory retrieval approach inherently favours certain types of memories or problems.
- Can the present solution help people remember in a 'reasonable' amount of time? It is unlikely that one would carefully review every past recording in one's life to remedy a memory problem. But, can the present approach help remedy memory problems within the amount of time people are willing to offer?
- How do people use personal-data archives for their daily tasks? One way to understand how a personal-data archive would be used in daily life is to build an archival tool and to see for what tasks people chose to use it.

In the remainder of this section we review relevant memory research followed by memory aids and ubiquitous recording technologies. In section 2, we describe iRemember, a prototype 'memory prosthesis' we designed that is the basis of the experiments reviewed in section 3. In section 4, we discuss some of the social and legal issues associated with ubiquitous recording. We conclude with a discussion of design principles for memory aids as well as a discussion of potential applications and future directions of this work.

1.2 Memory

Memory research spans a wide area; here we focus on research most relevant to computer-based memory aids. Long-term memories refers to memories that are retained anywhere from about one minute to multiple years. This is different from 'short term' memory, which we use to remember, say, a telephone number between the phone book and the phone. Our focus is on two particular types of long-term memory — episodic and semantic. Episodic refers to our remembrance of events in our lives. Semantic refers to world knowledge (facts, concepts, etc).

With respect to types of memory problems, Schacter developed a taxonomy, 'The Seven Deadly Sins of Memory', highlighting the most common ones [14]. The six involving forgetting and distortion are shown in Table 1. The seventh, 'persistence' — pathological inability to forget — is of less interest to memory-aid designers.

Table 1 Six of the seven 'sins of memory' [14].

Forgetting	Distortion
Transience (memory fading over time)	Misattribution (right memory, wrong source)
Absent-mindedness (shallow processing, forgetting to do things)	Suggestibility (implanting memories, leading questions)
Blocking (memories temporarily unavailable)	Bias (distortions and unconscious influences)

In examining the frequency of memory problems among students, one diary study of everyday forgetting found that most memory failures are the result of failing to perform some action [15]. Subsequent studies [16] explore the frequency of some types of forgetting in workplace settings (Table 2). Relating back to Schacter's taxonomy, 'retrospective memories' can correspond to 'transcience' or 'blocking'. Furthermore, 'prospective memory' and 'action slips' are both forms of 'absent-mindedness'.

Table 2 Eldridge's classification of common memory problems in the workplace [16].

Type		Description	Example
Retrospective memory	47%	Remembering past events or information acquired in the past	Forgetting someone's name, a word, an item on a list, a past event
Prospective memory	29%	Failure to remember to do something	Forgetting to send a letter, forgetting an appointment
Action slips	24%	Very short-term memory failures that cause problems for the actions currently being carried out	Forgetting to check the motor oil level in the car before leaving on a trip

Our research to date has focused primarily on aids for transience and blocking. Two of the convenient qualities of these memory problems include:

- the forgetter is aware of the problem while it is occurring,
- under certain conditions, wishes to remedy it.

Not every instance of a memory problem satisfies this. For example, those experiencing absent-mindedness do not realise they have forgotten until after the fact (e.g. discovering you forgot to buy eggs after leaving the grocery store) and those who experience blocking may be irritated at the failure but may not feel the urgency for an immediate remedy (e.g. forgetting the name of an actor or movie). However, there are some memory problems that become a significant encumbrance if not resolved in a timely fashion (e.g. forgetting the correct highway exit while driving).

With respect to forgetting rates, Wagenaar studied long-term forgetting of salient, personal events [17]. For six years, on a daily basis, he recorded salient events in his life and had an assistant test him on his recall of these events anywhere from months to years later. This illustrates one of the problems with long-term memory experiments on personal data: the large time commitment to both document one's life and to periodically sit for memory tests limits the subject pool.

One interesting result from Wagenaar's study is the value of cueing towards memory retention. As seen in Fig 1 more cues (e.g. who, when, where) lead to better retention. This is of present interest because computers (with appropriate sensors) can collect such cues and use these to attempt to jog one's memory. A particular kind of cue, a landmark event, has been shown to improve time-localisation of salient events [18]. If, for example, the computer can cue you with either public or personal landmarks, we should achieve better time localisation in memory-retrieval tasks.

1.3 Computer-based memory aids

Previous computational memory aids, such as the 'Remembrance Agent' [19], allowed users to write notes on a full-time wearable computer and attempted to serendipitously trigger memories by retrieving relevant notes from the past. 'Forget-Me-Not' [20] used the pager-sized 'ParcTab' to passively collect data such as personal location, encounters with others, workstation activity, file exchange, printing, and telephone calls. The tool attempted to remedy retrospective memory problems by allowing users to review past events when a forgetting incident occurs, and prospective memory problems by providing situation-specific reminders. The 'Audio Notebook' [21] addressed transience and blocking memory problems by recording audio and simultaneously indexing it with the user's hand-written notes. Either data type could act as a memory trigger; the visual-analogue allowed

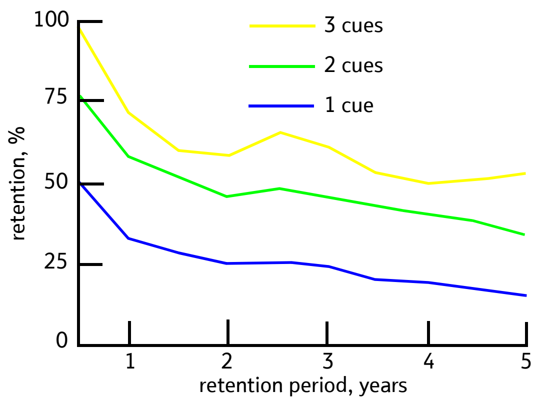

Fig 1 Retention improvements of salient events with additional cueing [17].

easy random-access during audio play back. WearCam [22], LafCam [23], StartleCam [24], and SenseCam [25], all explore a variety of video-based approaches to either capturing a comprehensive video history or to limiting archival to computer-identified salient moments.

A sampling of related efforts involved in collecting digital audio and video archives include classrooms settings [26], meetings [27], voicemail [28], workplace and telephone discussions [6], informal situations [11], and personal histories [7]. Among these and other similar projects in the domain, indexing techniques such as large-vocabulary speech recognition [29], speaker identification [30], face recognition [31], and audio/video scene analysis [32] are used to mine noisy data for search cues. Whereas much attention has been given to improving the accuracy and robustness of these techniques, less has been done on designing tools to help users make sense of noisy data sources and to understand which analyses prove most useful to specific tasks. For example, differences are expected among the search strategies employed by those trying to remember a previously witnessed event versus someone trying to find information within an unfamiliar collection. In the former case, retrieving any information that triggers the memory of the event is sufficient. In the later, finding the exact information is necessary.

2. The 'iRemember' memory prosthesis

2.1 Data sources

One of the first questions during the initial design of iRemember is what data sources should be used as a basis for memory triggering. An ideal data source has maximal memory-triggering value while limiting the CPU and storage demands to within the capabilities of contemporary wearable computers. Furthermore, a design goal is to minimise the effort needed by the user to capture daily experiences. This means, when possible, data should be captured with little effort or passively. Finally, to minimise the chance of missing a potentially valuable memory trigger (at the cost of retaining superfluous ones), nearly continuous recording of daily activity is desired. To these ends, a portable, wearable recording apparatus was constructed based on a commercially available personal digital assistant (PDA), an iPaq.

High on the list of desired data sources was audio, due to the anticipated memory-triggering value of verbatim audio, the desire to capture conversations that occurred in informal settings, the ease of capturing audio using a portable device, and the tractable data-storage requirements. However, for legal and human-subject approval reasons, this data source requires consent from all participants for each recording. Consequently, collecting this data can neither be completely passive nor continuous. Similar to Doppelgänger [33], sources that are captured and archived continuously and passively include the user's location, calendar, e-mail, commonly visited Web sites, and weather.

It should be noted that completely passive data capture may in fact hurt memory recollection as evidenced by the disadvantage of no note-taking among students in classroom situations [34]. The choice to prefer passive data capture is to simplify and reduce forgetfulness in the daily data-capture process. Ironically, forgetting to activate iRemember (i.e. absent-mindedness) is a common problem.

One of the challenges to the design of the prototype is finding ways to manage the large quantity of collected data that is often rife with both noise due to the intrinsic limitations of computational information-extraction from the selected data sources, and irrelevance with respect to a specific information need. The former is more typical of a personal memory assistant and will be addressed in more detail. The latter point is addressed in a manner similar to conventional text corpora searches — keyword searches with ranked retrieval results.

2.2 Audio-based memory retrieval

Speech is rich in meaning, subtlety, and affect. The storage demands are not high and a variety of existing analyses make it possible to construct limited-certainty indexing tools to extract the content [29], speaker identity [30], affect [35] and the conversational structure [36], just to name a few. Our work has focused on using audio recording of conversations as a means for memory assistance.

Conversations may occur in any situation. Previous work has studied ubiquitous audio recording in offices [6] and meetings [27, 37]. One of the goals of the present research was to include the option to capture spontaneous conversations that often occur in hallways, shared laboratory spaces, near water coolers, etc. A portable tool was suitable for this.

A general-purpose memory assistant designed to remedy all of Schacter's 'sins' remains safely in the realm of science fiction and corporate dreams. The current prototype, a memory-retrieval aid, is designed to help remedy memory problems by assisting the retrieval process. Such an aid would be most valuable when a person both recognised that a memory failure has occurred and wishes to recover from the failure. Again, transience and blocking problems would seem to be the most likely benefactors.

retrieving any information that triggers the memory of the event is sufficient

The present approach is to collect, index, and organise data recorded from a variety of sources related to everyday activity, and provide a computer interface to both search and browse the collection. The hope is that some piece of recorded data can be found and act as a trigger for a forgotten memory. This important difference of finding a memory trigger contrasts with the more difficult exact-match metrics of searching from within unfamiliar collections [38].

2.3 Prototype

A working system, iRemember, has been built and is currently in use at the Media Lab. This includes recording devices, data storage, indexing tools, search engines, and user interfaces to facilitate memory retrieval. Figure 2 shows the prototype

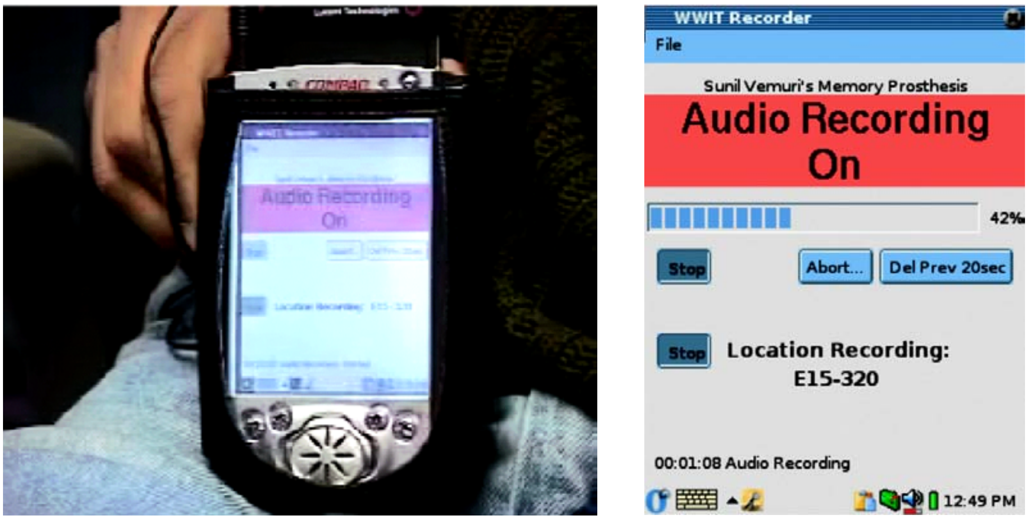

Fig 2 Wearable recording device used for daily recording of audio and physical location (left) . Close-up of the screen is shown on the right.

wearable recording device. It records audio and physical location of the wearer. Location is determined via proximity to stationary IEEE802.11b wireless-network base-stations. For data security and user-comforting reasons, the current implementation is limited to recording within the Media Lab. When the security blockades are disabled, the system is known to work on any of the numerous, open, wireless networks available in the Boston area.

Such networks are commonplace in many locations around the world. The wireless network also serves as a conduit to offload data from the limited-capacity recording device to a high-capacity server. This server also has the necessary computational power to perform data analysis and indexing. Finally, the server records data sources such as e-mail, calendar, local weather, and user-selected Web sites, which are then made available through the retrieval interface.

Although many data sources are recorded, initial evaluations have focused on audio as the primary means for memory triggering. This includes integrating audio analysis, indexing, and searching tools to help find relevant data. In particular, IBM's ViaVoice® large-vocabulary speech recognition system [39] was selected as the initial data analysis and indexing technique. This enables keyword searching across all audio recordings (included as part of the interfaces shown on Figs 3 and 4) and allows for a visual accompaniment during audio playback (Fig 5). Details of these are now discussed.

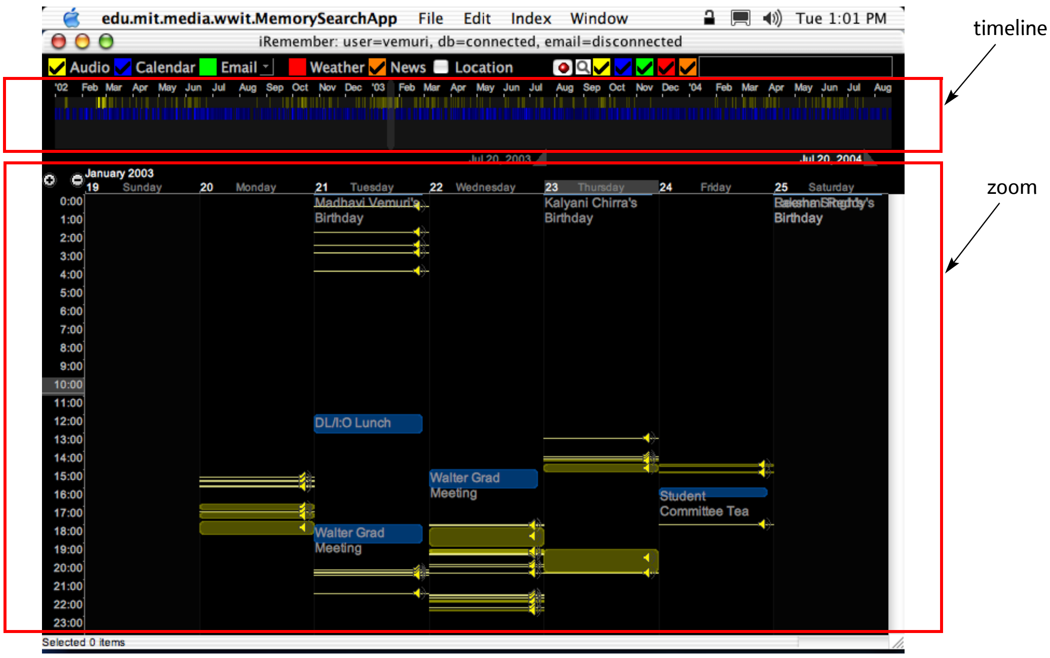

Fig 3 Visual interface for browsing and searching through all recordings. This view shows a multi-year timeline and a zoomed-in view of one week.

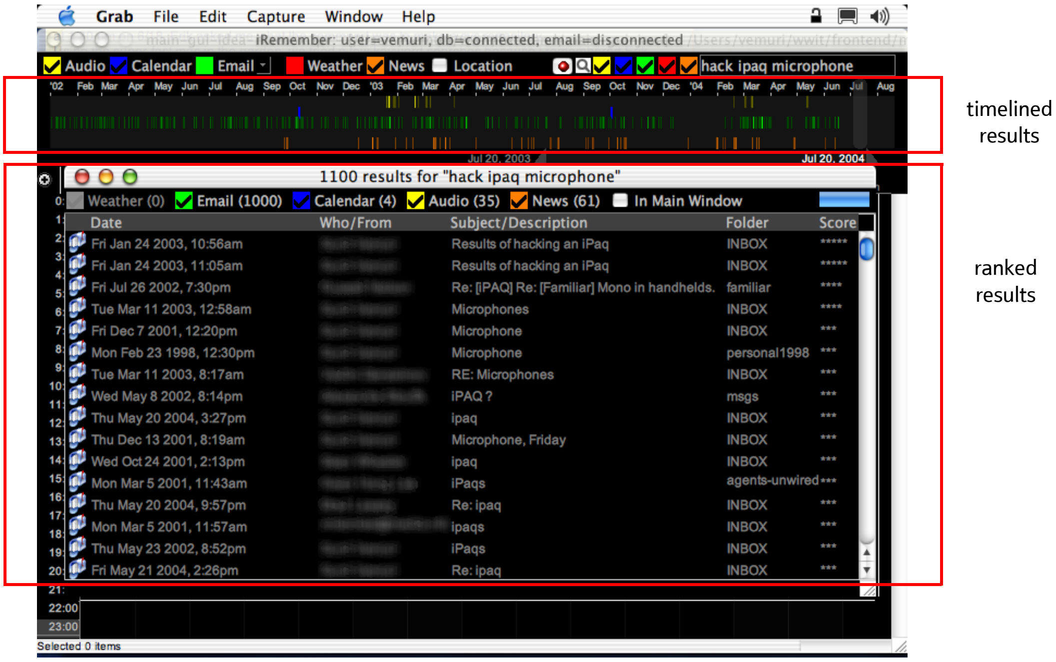

Fig 4 Search results are shown simultaneously as a ranked list and on the multi-year timeline.

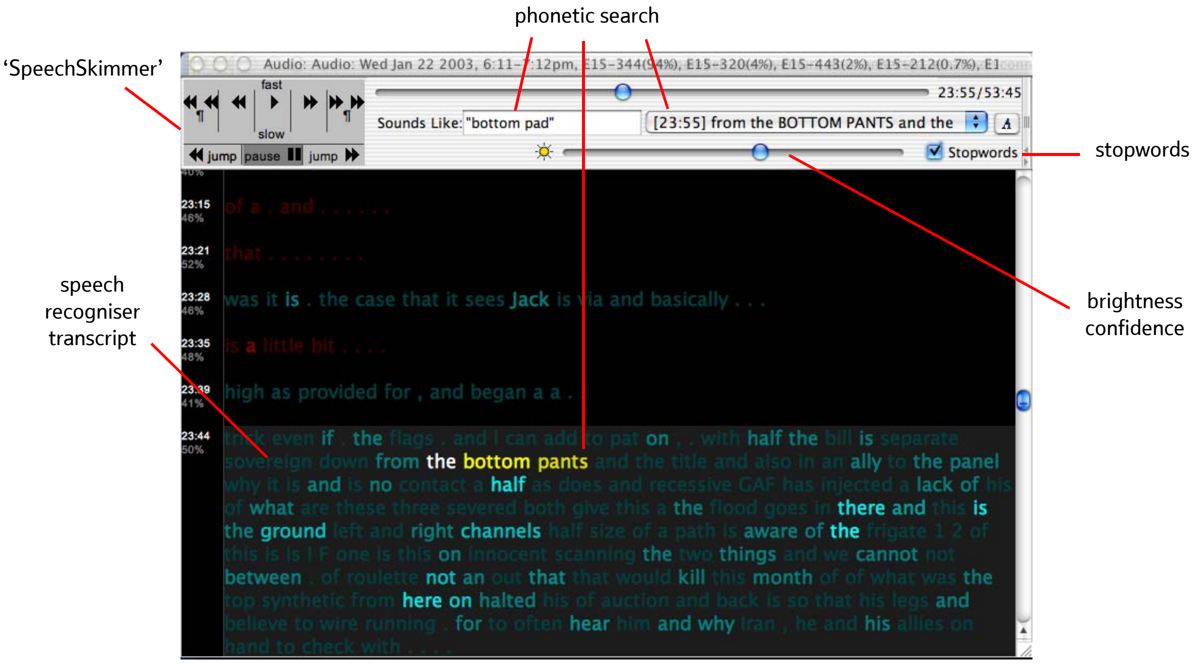

Fig 5 Interface for browsing, searching, and playing recordings accompanied by error-laden speech-recogniser-generated transcripts. Phonetic 'sounds-like' keyword matches are rendered in yellow text. Brightness of words based on recogniser-assigned confidence. The text colours represent different speakers.

The desktop- and laptop-computer-based retrieval interface for browsing and searching through recordings is shown in Fig 3. A multi-year timeline is shown on top with a zoomed-in view below; days are on the x-axis and time-of-day is on the y-axis. Audio recordings are shown as yellow rectangles and calendar entries are blue. A keyword-search feature (upper right), based on the Lucene open-source search engine [40], is available. Users may conduct simultaneous searches on any combination of the previously mentioned data sources (including audio); search results are presented both as a rank-ordered list and as tick marks on the timeline. Presenting results on the timeline allows users to focus on date ranges with a high density of hits or see results relative to a landmark event.

Results from a sample query are shown in Fig 4. The results in the figure are based on an actual forgetting situation experienced by one of the authors. He was trying to remember details of an esoteric hardware modification done to the iPaq; the answer could only be found in a conversation. The search produced an e-mail that quickly narrowed the investigation to a few days' worth of audio (instead of a few years' worth).

Figure 5 shows the interface to play and 'view' a single audio recording. As mentioned before, the transcript provides a visual accompaniment to the audio. Notably in this interface is a phonetic or 'sounds like' search feature. Partially based on Wechsler et al's algorithm [41] and using CMUDict [42], it allows users to identify misrecognised, out-of-vocabulary, and morphological variant words by locating phoneme-sequence similarity. The interface includes a few additional features meant to further help users focus their reading and listening attention within a recording. This includes the sophisticated 'SpeechSkimmer' [43] audio skimming tool and options to render text brighter or dimmer based on recogniser-reported confidence or if the word is a common English 'stopword' (e.g. 'a', 'an', 'the').

The visualisations and search features are admittedly crutches meant to support an error-prone technology. Unfortunately, there is little near-term hope that speech recognition will achieve high-accuracy in the poorly microphoned, heterogeneous, noisy settings typically experienced by iRemember users. Witbrock [44] suggests general-purpose audio-information retrieval tasks can still be performed despite high recognition error rates. Indeed, speech recognition has been used successfully for a similar, voicemail retrieval task [28], a calendar-scheduling task [45], and broadcast-news retrieval tasks [38]. But, these benefit from limiting vocabularies and acoustic conditions, which in turn improves recognition accuracy. Such benefits are not expected for iRemember users.

3. Evaluations

With these tools and infrastructure in hand, experiments of the iRemember memory-retrieval aid are being conducted to better understand how a long-term personal-data archive can act as a memory-triggering resource. To date, studies have focused on memory-retrieval tasks while being mindful of user interface issues.

In this section we present summaries of three studies. The first, a study of memory retrieval within the context of public events, was conducted to establish a baseline for the efficacy of our tools. The second experiment validates a technique for reducing the time necessary to browse through audio, which is a critical component of our approach to remedying memory problems via searching audio archives. The third study examines memory retrieval within the context of long-term personal data. We conclude the section with discussions of our experiences with ubiquitous recording and portable, as opposed to desktop-computer-based, memory retrieval.

3.1 Memory retrieval of public events

Short of a large-scale deployment, it is unlikely that an adequately controlled evaluation over a general populace can be performed. Qualitative results from smaller-scale trials are instructive, but alternates that approximate the circumstances of iRemember users in their daily lives can give valuable quantitative data. To this end, memory tests based on content from semi-public recorded events (e.g. conferences) have been performed. This includes providing (or for control groups, not providing) tools like the one described in section 2 to assist with the memory-retrieval tasks.

Experimental details are found in Vemuri et al [46] and summarised here. Subjects were given a question-answering task based on facts presented at conference seminars one month in the past. When subjects attempted to answer these questions without any assistance, transience (63%) and blocking (11%) memory problems were common, and misattribution errors (26%) were unexpectedly high. Subjects were then given a memory-retrieval aid similar to the one shown in Figs 3—5 and asked to remedy some of these memory problems based on events they witnessed (i.e. memory retrieval). For comparison, they were also asked to find answers to questions based on talks they did not witness (i.e. information retrieval). When using the aid, subjects committed no mistakes, and the amount of guessing decreased. It is speculated — and somewhat supported by subject interviews — that the availability of unimpeachable evidence (e.g. verbatim recordings) discourages guessing. Comparing memory and information retrieval performance, subjects successfully found answers in 82% of the question-answering attempts of witnessed events and 55% of unwitnessed events. Not surprisingly, unaided memory recall took less time than aided. But, when using the aid, subject search time was often limited to roughly four minutes whether or not an answer was found. This may indicate a potential upper-bound design-constraint for memory-retrieval tasks.

In addition to the quantitative measures, subject problem-solving strategies were observed. Two challenges of audio-based information-retrieval include finding the correct recording within a collection and choosing the correct section to play back within the recording. In most observed cases, subjects used calendar-navigation along with their remembrance of the past to identify the correct recording. Speech recognition, despite poor-quality transcripts, was more useful when subjects attempted to localise and select fragments of individual recordings to play back. In particular, visually skimming confidence-based, brightness-coded, error-laden transcripts and performing phonetic searching were commonly used and cited as useful localisation techniques. Subjects familiar with speech recognition technology often tuned their keyword choices and search strategies based on the expected limits of speech recognition by choosing words anticipated to be 'in-vocabulary'.

Despite the infrequent use of keyword-searching to identify candidate recordings in the collection, it is expected that this will still be useful when searching larger archives or archives spanning longer time spans. The collections used in the evaluation (59 recordings, covering 14 hours over 3 days) are dissimilar to the archive collected by the iRemember-wearing author. The latter shows a sparser recording pattern, includes many more recordings, covers a longer period of time, and has less calendar-annotated structure. Second, keyword-searching was commonly used in the control condition in which subjects attempted to answer questions about events they did not witness. These results may not generalise to searching unfamiliar collections, but there may be implications to organisational-memory applications. For example, absentees could have improved ways of quickly finding information within a recorded meeting. This example notwithstanding, memory fades over time and it is anticipated that the search process on events in the distant past will resemble that which is experienced by non-witnesses.

While information-retrieval of text corpora is now common practice among computer users (especially those using Web-based search engines), the novel and somewhat esoteric features in the memory-retrieval aid led to a certain degree of subject adjustment. Some subjects reported self-improvement within about 15 minutes of searching: 'I think I'm getting better at figuring out how to search the audio just in terms of thinking about things that might work.'

3.2 'Speed listening'

One observation from the study presented in section 3.1 is that considerable time is spent simply listening to audio to find the relevant bit. Reducing audio-listening time would noticeably reduce the overall memory-retrieval search time. It is well established that comprehension of audio decreases as speech is accelerated (using time-compression) [47]. We hypothesised that having a transcript associated with the audio would either increase comprehension at a given playback speed or enable an increased playback speed for a given comprehension rate, thereby reducing the time needed to listen. To study this, we conducted an experiment where we demonstrate that combining speech recognition transcripts and time-compressed audio can produce such time reductions without sacrificing user comprehension. Details are presented in Vemuri et al [48] and summarised below.

Several experimental conditions were established; these are described below and summarised in Table 3. Firstly, perfect transcripts (Perfect) were compared to audio only (Audio-only), but perfect transcripts are expensive to construct because they must be human generated. A less-expensive way of generating transcripts is using a speech recogniser (Reco-uniform). The problem with this approach is that the transcripts are intrinsically error-prone. So, we added a condition in which word brightness was rendered proportional to 'phrase score' (Reco-vary), IBM ViaVoice's metric for recognition confidence [39]. The expectation was that subjects would focus attention on high-confidence words, which will, in turn, improve overall comprehension. We maintained uniform word brightness for all of the other conditions. Finally, because large-vocabulary speech recognition can suffer from very high word error rate (WER), we added a final condition (Incorrect) in which the transcripts were intentionally all wrong. We hoped to learn about comprehension penalties incurred with high WER transcripts.

The results of the experiment are shown in Fig 6. Speech-recognition-generated transcripts, even with high WER, are an inexpensive way to improve audio browsing (by either increasing comprehension or reducing listening time). Subject comprehension was slightly higher when using 'Reco-vary' transcripts versus 'Reco-uniform', but the significance was weak. Surprisingly, high WER does not seem to hurt performance: speech recognition can be used with no penalty.

Table 3 Experimental conditions for 'speed listening'.

Transcript style	Description
Perfect	Human-constructed 'perfect' transcript with uniform word brightness
Reco-vary	Speech-recogniser-generated transcript with word brightness proportional to 'phrase score'
Reco-uniform	Speech-recogniser-generated transcript uniform word brightness
Audio only	No transcript; audio only
Incorrect	Completely incorrect transcript with uniform word brightness

3.3 Memory retrieval on long-term personal data

We tested the memory-retrieval technique in the conference setting (section 3.1), and we developed a speed-listening technique (section 3.2). But how will these techniques work with conversational audio spanning several years?

To study this, like Wagenaar [17], one of the authors documented parts of his life for two years using the recording device described in section 2. Unlike Wagenaar, he tested his colleagues' memories of past conversations instead of his own memory. When they could not remember, they could use the memory-retrieval tools. (One subject described the task as analogous to having someone create a Trivial Pursuit®-like game, but the category is always about you.) Questions were chosen by identifying topics and facts that might be in the subjects' long-term memory, but they likely forgot. The focus was on transience memory problems.

The procedure started with the investigator giving the subject a single 'main' question along with a pre-questionnaire. Subjects were not allowed to use the memory-retrieval software yet. The pre-questionnaire included questions such as: Do you know the answer? Do you remember having this conversation? Do you remember where, when, who, or anything else about the conversation? Based on pre-questionnaire answers, if a memory problem is observed,

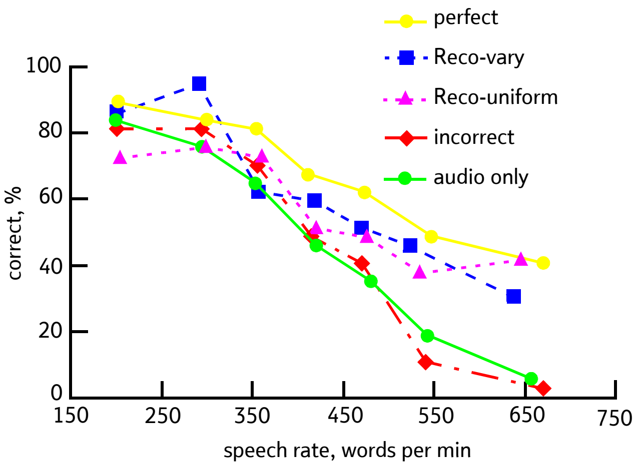

Fig 6 Question-answering performance versus speech rate for the five transcript styles.

subjects were then asked to use the memory aid (which included verbatim audio recordings of two years of past conversations) to attempt to remedy the memory problem. An attempt was considered successful if the subject found the answer and unsuccessful if they did not find the answer or submitted an incorrect guess. Three subjects spent four to five hours (spread over multiple sessions) answering questions. Each attempted 18—20 questions and spent 5—30 minutes per iteration (i.e. answering the 'main' question, using software, and answering pre- and post-questionnaires).

landmarks help us achieve better time localisation in memory-retrieval tasks

In all, 56 questions were asked. In 7 cases, subjects already knew the answer and thus did not have a memory problem. In 27 cases, subjects remembered having the conversation, but not the answer to the question. In 22 cases, subjects did not remember having the conversation, let alone the answer. This does not necessarily mean the memory of that conversation was not lurking somewhere in the subject's long-term memory; it just meant that the question did not trigger a remembrance of the conversation.

In 36 (64%) cases, the subjects were successful in remembering either on their own (7 cases, 13%) or with help from the software (29 cases, 52%). As seen in Fig 7, most of the non-success cases were more than one year in the past. Subjects were more successful (68% versus 44% success rate) when they were able to benefit from using temporal memory cues, i.e. when they were able to correctly remember the approximate timeframe of the conversation. By properly constraining the time bounds, the information retrieval task is facilitated. Two of the subjects (subjects A and C) tended to have narrower time-bounds and could thereby focus their searches better (75% success rate). The third subject (subject B) had a lower success rate (31%), which can be partially attributed to wider time bounds and misattribution problems, both of which exasperated the retrieval task.

Based on the pre- and post-interviews, there was no evidence across subjects suggesting different remembrance of having the conversations. However, there were some between-subject differences. When they remembered having the conversation, subjects A and C could cite specifics of the past conversations and the surrounding circumstances (e.g. who, where, what else). Subject B's descriptions were more general and included references to multiple conversations on the asked-about topics or conversations on the topics with people outside the study. This might suggest that subject B's memories of these topics have become more consolidated as semantic memories whereas subjects A and C are retained in episodic memory. The blurring of episodic details is not uncommon as memories become consolidated in semantic memory. This may further explain subject B's lower success rate and illustrate a limitation of the memory-retrieval approach for these types of memories.

The mean time subjects spent question-answering using the software when successful was under six minutes. The mean time when unsuccessful was almost 12 minutes — there is a time penalty associated with failure.

Like the conference study presented in section 3.1, we were also interested in subject strategies and anecdotes as they performed memory retrievals. General observations from the present study include:

- within-recording localisation strategies, using the interface in Fig 5, were similar to those in the conference study (section 3.1),
- collection-wide audio search was preferred, unlike the conference study where calendar-navigation was the primary choice,
- landmarks helped with time-localisation when used, but subjects used simple mechanisms like calendar, e-mail, Web search to find landmarks,
- accurate speech recognition makes the task easier,
- less-vivid remembrance made it harder — this was based on subjects' answers in the pre-questionnaire indicating

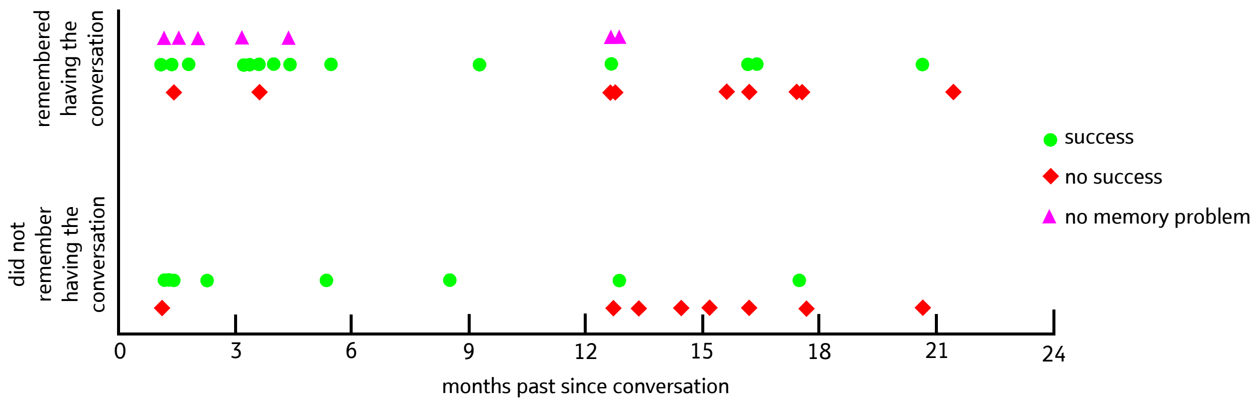

Fig 7 Memory-retrieval results over two years. Top row shows attempts when subjects remembered having conversation. Bottom row shows attempts when subjects did not remember having conversation.

vivid details of the conversation, even if they did not remember the answer to the specific question,

- misattribution, both in answering the 'main' question and in answering questions related to the circumstances surrounding the conversation, led subjects astray, e.g. looking in the wrong time period.

In summary, we demonstrate that the memory-retrieval technique can work on long-term collections of personal data. Unlike other audio-search tasks, the difficult part of this task is finding the correct recording. Previous studies indicate that people perform searching tasks better with landmarks [49]. Though we found evidence of this here too, the landmark tools provided were infrequently used. $N = 3$ is not conclusive and more data and subjects are needed to strengthen the results and also to design more-controlled studies. But we are a step closer towards using personal-data archives with information retrieval tools as a memory aid.

3.4 Experiences with ubiquitous recording

As mentioned before, one of the authors (a student) has used iRemember to record a variety of day-to-day experiences over the past two years. In addition, another student in the department (not directly involved in the research project) is also a frequent audio-recorder of select conversations and uses a variety of computer-based audio recording tools, including iRemember. Both record with different frequencies and for different purposes, the former collects many recordings in anticipation of transience problems and the latter limits recordings to student-advisor discussions to prevent both transience and absent-mindedness. However, both tend to make use of the recordings for the very same purpose — reviewing student-advisor recordings while under time-pressure of a paper-writing deadline. Also, the latter student, when using a non-iRemember recorder, adopts a space-saving deleting strategy upon task completion. This last point has implications towards the upcoming discussions on privacy and future directions.

unlike other audio-search tasks, the difficult part of this task is finding the correct recording

3.5 Transition to portable retrieval

A desktop personal computer was used for the experiments described above due in large part to the computational and visual demands of the searching, browsing, and skimming capabilities. The evaluations have indicated that a subset of the features is adequate for most memory-retrieval tasks and many of these have been implemented on the iRemember PDA. Pilot studies have demonstrated subjects can successfully perform memory-retrieval tasks using the PDA. However, the computational limits of current PDAs still require a server-based approach. In particular, large-vocabulary, automatic speech recognition and phonetic searching are just beyond the capability (CPU speed) of current PDAs, but this is expected to change within a few years.

4. Social and legal issues

Brin posits ubiquitous recording is inevitable [50]. Yet, ubiquitous recording devices, especially ones like iRemember, raise some obvious privacy concerns. What permissions are needed to record? Who will have access to the data? What social conventions are needed for such devices? What legal protections are available? Common use is not expected until these points are adequately addressed.

what social and legal conventions are needed for such devices?

Many countries have laws requiring consent before audio recordings can take place. Within the USA, there is some state-to-state variation with respect to both the setting (public versus private) and how many people must consent (one or all) [51]. While these standards describe what is legal, social conventions prescribe what is appropriate. The author who regularly wears iRemember has observed various reactions to the device. Consent is required and always requested prior to any recording; the social greeting includes a somewhat awkward request (e.g. 'Hello. Good to see you. May I record this?'). Despite this, when the device is off, some conversation partners (who prefer not to be recorded) still ask for verification before speaking freely. Others who have used the device found it uncomfortable to ask conversation partners for permission to record. In particular, those in supervisory roles report more hesitation and discomfort asking permission of their subordinates as compared to asking peers. Though people occasionally decline requests to be recorded, it is not known if there are instances of accepted requests in which the person truly preferred not to be recorded, but agreed out of a sense of co-operation for a fellow researcher or other unspecified reasons. These reports are based on experiences of university researchers in a technology-friendly environment — the Media Lab. More studies are needed to understand what social conventions are appropriate among a more general population.

Assuming such conventions are possible, there is another reason for caution. Once recordings are made, they can be searched and seized via court-approved warrants. The Fourth and Fifth Amendments to the US Constitution describe protections against searches, seizures, and self-incrimination, but these are not expected to extend to memory prostheses. From a legal perspective, a close cousin to iRemember is a personal diary. US courts have addressed protection of personal diaries, and their current position is diaries can be searched and seized [52]. Hence, it is unlikely that the less-private iRemember and similar devices would be afforded more protection.

Encrypting data is an option. Hiding it is another, and the safest place might be inside one's body. While US courts have not set an absolute limit to what can be seized, the standard for extracting things from inside one's body is higher than things outside [53].

5. Future directions

Current research has only scratched the surface of memory assistances afforded by ubiquitous computing. Lessons learned from iRemember, while promising, are limited to a small set of memory assistances and a restricted population. A rich future is in store.

The authors feel solving the social and legal implications are more challenging than the engineering ones. In this section, proposals for future memory aid research are presented. These attempt to address additional memory problems, use additional data sources, reduce the social and legal barriers, and address memory problems among the elderly.

5.1 Memory-encoding aids

Instead of remedying a memory failure, can a memory aid reduce the chance of the failure from happening in the first place by improving biological memory-encoding? One theory suggests humans encode memories in summary form. Details are filled in as more time, effort, and attention are spent on a particular event. The more details one encodes, the more association paths there will be to the memory. The amount of detail one remembers about a particular event depends both on one's ability to find a stored memory and how much effort is dedicated to the initial encoding process. The 'depth-of-processing' theory suggests that more processing done during memory encoding leaves a stronger memory trace for future retrieval [54, 55].

> recordings can be searched and seized via court-approved warrants

All of this suggests that memory-strengthening efforts soon after an event may mitigate the need for future memory assistance. Computer-based memory aids can potentially facilitate this biological improvement by encouraging well-focused memory-strengthening exercises at the time of biological memory encoding. As mentioned earlier, the applicability of iRemember is limited to transience and blocking problems in which people both recognise the problem is occurring and wish to remedy it. This limitation is not expected for memory-encoding aids.

Ebbinghaus' forgetting curve [56] shows that people forget over 50% after the first hour and 66% after the first day. This suggests a memory-encoding aid would be most effective within the first 24 hours of an event. In addition to formation of new memories, maintaining existing memories can also benefit from periodic rehearsal. A single test of a previous memory, even after years, has significant impact on its subsequent retrievability [57]. This suggests a memory aid that occasionally elucidated past relevant events (even if only once every few years) would help improve retention of such events.

Despite its potential benefits, computer-assisted encoding might interfere with the normal, biological selection of memorable events. Forgetting not only has to do with the amount of time that passes between event and retrieval, but how much activity occurs between event and retrieval [58]. Hence, it could be argued that encouraging memory-strengthening exercises have the adverse effect of artificially amplifying the strength of computer-selected salient memories in lieu of an individual's biological selection process. It remains to be seen if this selection-bias will have such an effect.

One vision of a memory-encoding aid is as follows: as events transpire in one's life, the aid would try to identify moments most worthy of later remembrance (or let users identify these), record those bits, and before one's biological memory has faded, present a summary of the salient parts. This presentation would hopefully inspire rehearsal and consequently, help strengthen the biological memory of the event. This ideal is very difficult to achieve. In fact, this vision is loaded with challenging sub-problems. Two that stick out include the identification of salient events and subsequent summarisation. Interestingly and perhaps conveniently, these are also challenges in the design of the memory-retrieval aid.

5.2 Limited archival

It is worth considering the trade-off between remedying potential memory problems that require long-term archival and the benefits of limiting archival. The reported scenarios in which iRemember was put to use (i.e. reviewing student-advisor conversations for a pending deadline) can succeed without long-term archival. Task completion may be an opportune time for data purge, or at least, extraction of salient parts and deletion of the remains. A memory-encoding aid intrinsically does not require long-term archival. The necessity for persistence implied by a lifetime archive, at least for certain memory assistances, is not clear.

A limited- or non-archival strategy has additional benefits. Firstly, a restricted search space may improve search experiences by reducing the time to find answers (or abandon fruitless search paths). Secondly, an illicit data intrusion would have limited ramifications. Thirdly, if a user were embroiled in legal struggles in which recordings were subpoenaed, an established deletion strategy may avoid allegations of illegal destruction of evidence. Finally, conversation partners might be more willing to be recorded if there is an agreed destruction policy.

5.3 Non-audio data

iRemember focuses on audio as a memory-triggering resource. However, with the data-collection infrastructure in place, adding new data sources is now a straightforward process. Doppelgänger [33] is an open architecture for adding heterogeneous sensors to user-modeling systems. A similar approach can be used for gathering resources for memory triggering.

Previous work at the Media Lab to build technology to remedy 'visual amnesia' [22], a natural addition to the current system, would provide users with a visual-capture capability. Not surprisingly, among those who have tried our wearable device for short periods of time, one to five days, the most commonly requested capability is visual recording. There is an active research community investigating general-purpose video retrieval [59] and it is expected that a visual record, even without any content analysis, can trigger memories. Formal evaluations would be needed to better understand the impact of visual triggering and if image quality and image content analysis play a significant role.

StartleCam demonstrates how biometrics can be used to identify instances of salient events [24]. Similar devices for continuous biometric recording for personal health monitoring have been studied [8] and commercially available biometric sensors are becoming readily available. While the technical hurdles for continuous biometric recording are diminishing, there is limited evidence relating memory to biometrics. What little is known is related to memorability of fear response in animals and the role of affect in the formation of 'flashbulb memories'.

The fear studies are similar to Pavlov's classical conditioning experiments. LeDoux showed that when presented with a fearful stimulus, animals quickly learn the conditions under which the stimulus occurs [60]. In terms of biometrics, blood pressure, galvanic skin response, and physical agitation are perturbed by presentation of the fearful stimulus. If, during training, the fearful stimulus is presented in conjunction with a conditioned stimulus (e.g. an audio tone, a light, etc), the same physiological reaction will result when the animal is presented with only the conditioned stimulus. In effect, these studies have quantified the intuitive notion that a memory of a fearful experience can elucidate a biometric and emotional reaction without actually re-enacting the experience.

Flashbulb memories refer to memories of events of such significance that observers can recall their circumstances (who, what, where, when, emotional response, etc) with abnormally high detail, clarity, and durability [61]. Many people claim to experience this phenomenon with news events of high significance (e.g. JFK assassination, Space Shuttle Challenger disaster, 9/11 terrorist attacks). Theories on how these memories are formed differ in the mechanisms involved; yet they are consistent in that all suggest a relationship between elevated affective state and memory.

This limited evidence leaves a thirst for more data regarding any possible relationship between biometrics and human memory. While it is premature to assert much about this interaction, biometric sensor technology has reached the point in which hypothesis-driven, controlled experiments can be conducted. Such work is needed to clarify this murky picture.

5.4 Memory and aging

This research project focuses on healthy people in academic and knowledge-worker settings. But, there is potential for impact on a much broader range of individuals, groups, and settings. Normal, healthy individuals suffer decreased memory facility due to aging and US demographics suggests increasing numbers of the 'baby boomer' generation are reaching ages where decreased memory facility has more pronounced effects on daily life. The elderly feel a great deal of anxiety as a result of memory loss and that anxiety is a good predictor for future cognitive decline [62]. While there may not be medicinal means to prevent the cognitive decline, having better memory aids may relieve some anxiety. In fact, the elderly, while showing decreased performance in laboratory memory tests compared to their younger counterparts, often perform adequately if not better on daily life tasks due to more efficient and diligent use of memory aids [63]. There may be near-term potential for improved quality of life for larger segments of an aging population through better and more accessible memory aids. Related to these goals, an exploration studying family histories illustrated how archived news media can help trigger memories among the elderly [64]. More personal, family-oriented archives may be of similar value.

5.5 Infrastructure-assisted capture

Public and semi-public environments are already filled with devices recording a variety of day-to-day activity. Buildings, roadways, and automatic bank teller machines are often equipped with video surveillance and these have proven useful in solving crimes. In Australia, taxicabs have continuously recording cameras (both interior and exterior); some police cars in the USA are similarly equipped. Broadcast and sports media outlets vigilantly document publicly witnessed newsworthy events. It is becoming commonplace for universities to broadcast and archive lectures, seminars, and classroom discussions through 'distance learning' programmes. Similar videoconferencing systems are used for more private communications in a variety of settings. Use of personal digital cameras (still and motion) has risen sharply over the past few years, especially via mobile telephones. Many Web-based cameras provide continuous feeds of public locations. We and the events we witness are recorded far more often than we think [50]. Coverage is bound to increase.

solving the technology problems is less challenging than the social, legal and ethical ones

The responsibility of memory-aid recording need not rest solely in the hands of each individual. The previous examples illustrate existing infrastructure already capturing a valuable and largely unused memory-triggering source. The technical hurdles to integrate with the existing memory prosthesis are not large. For example, a user control to identify interesting moments along with a mechanism to collect or tag data from nearby recording apparatuses would suffice. As with the memory prosthesis, the engineering hurdles are overshadowed by the social, ethical, and legal ones. Less-restricted access to such surveillance data may draw ire; Brin discusses some scenarios in detail along with cautions against hypocrisy on data-sharing stances [50]. Projects allowing people to avoid public-surveillance might avert some controversy [65].

6. Conclusions

Advances in computer miniaturisation and data analyses have enabled a new generation of personal memory aids. A prototype 'memory prosthesis' has been presented with the eventual goal of helping to alleviate some common memory problems by creating a searchable, personal archive of everyday experiences (with initial focus on audio recordings). This includes memory-retrieval capabilities allowing users to browse, search, and listen to audio and associated error-laden speech-recogniser-generated transcripts of past, recorded events. Results are promising and give authors hope that the approach is viable to solve a limited set of memory problems.

But, solving the technology problems of building a ubiquitous recording and memory-retrieval system is less challenging than the social, legal, and ethical ones. Wide acceptance is not anticipated until those are resolved. Perspectives on future directions for this field were discussed. This includes alternate memory aids to address additional memory problems, using non-audio data types including those from a variety of existing sources, addressing memory problems unique to healthy, older individuals, and design alternatives that may minimise social, legal, and ethical concerns.

Acknowledgements

The authors wish to thank Chris Schmandt, Rosalind Picard, Philip DeCamp, Stefanie Tellex, Brad Lassey, Charles Nesson, and all of our experimental subjects. This work was sponsored in part by the 'information:organized' research consortium at the MIT Media Lab.

References

1 Sabbag R: 'The reluctant retiree', Boston Globe Magazine (November 2003).

2 Kletz T: 'Lessons from disaster: how organisations have no memory and accidents recur', Institution of Chemical Engineers, Rugby, Warwickshire, UK (1993).

3 Loftus E F: 'Eyewitness testimony', Harvard Univ Press, Cambridge, Massachusetts (1996).

4 Walsh J P and Ungson G R: 'Organizational Memory', Academy of Management Review, 16, No 1, pp 57—91 (January 1991).

5 Bush V: 'As we may think', Atlantic Monthly, 76, No 1, pp 101—108 (July 1945).

6 Hindus D and Schmandt C: 'Ubiquitous audio: capturing spontaneous collaboration', Proc CSCW '92, pp 210—217 (1992).

7 Gemmell J, Bell G, Lueder R, Drucker S and Wong C: 'MyLifeBits: fulfilling the Memex vision', Proc ACM Multimedia '02, Juan-les-Pins, France, pp 235—238 (2002).

8 Gerasimov V: 'Every Sign of Life', PhD thesis, MIT program in Media Arts and Sciences (2003).

9 LifeBlog, Nokia — http://www.nokia.com/nokia/0,54628,00.html

10 LifeLog — http://www.darpa.mil/ipto/programs/lifelog/

11 Lin W and Hauptmann A: 'A wearable digital library of personal conversations', JCDL 2002, pp 277—278 (2002).

12 Weiser M: 'The computer for the twenty-first century', Scientific American, pp 94—10 (September 1991).

13 Reder S and Schwab R G: 'The temporal structure of cooperative activity', in Proc CSCW 1990, pp 303—316 (1990).

14 Schacter D L: 'The Seven Sins of Memory: Insights from Psychology and Cognitive Neuroscience', American Psychologist, 54, No 3, pp 182—203 (1999).

15 Terry S S: 'Everyday forgetting: Data from a diary study', Psychological Reports, 62, pp 299—303 (1988).

16 Eldridge M, Sellen A and Bekerian D: 'Memory problems at work: their range, frequency, and severity', Technical Report EPC-1992-129, Rank Xerox Research Centre (1992).

17 Wagenaar W A: 'My memory: a study of autobiographical memory over six years', in Cognitive Psychology, 18, pp 225—252 (1986).

18 Loftus E F and Marburger W: 'Since the eruption of Mt St Helens, did anyone beat you up? Improving the accuracy of retrospective reports with landmark events', Memory and Cognition, 11, pp 114—120 (1983).

19 Rhodes B and Maes P: 'Just-in-time information retrieval agents', in IBM Systems Journal special issue on the MIT Media Laboratory, 39, Nos 3 and 4, pp 685—704 (2000).

20 Lamming M, Brown P, Carter P, Eldridge M, Flynn M, Louie P, Robinson P and Sellen A: 'The design of a human memory prosthesis', The Computer Journal, 37, No 3, pp 153—163 (1994).

21 Stifelman L, Arons B and Schmandt C: 'The audio notebook: paper and pen interaction with structured speech', in Proceedings of the SIG-CHI on Human Factors in Computing Systems, pp 182—189. (2001).

22 Mann S: 'Wearable Tetherless Computer-Mediated Reality: WearCam as a wearable face-recogniser, and other applications for the disabled', in Proceedings of AAAI Fall Symposium on Developing Assistive Technology for People with Disabilities, (November 1996).

23 Lockerd A and Mueller F: 'LAFCam: Leveraging Affective Feedback Camcorder', in Proc CHI 2002 (2002).

24 Healey J and Picard R: 'StartleCam: a cybernetic wearable camera', in Proc of the Second International Symposium on Wearable Computers (1998).

25 Williams L: 'SenseCam', — http://research.microsoft.com/hwsystems/

26 Abowd G D: 'Classroom 2000: an experiment with the instrumentation of a living educational environment', IBM Systems Journal, 38, No 4, pp 508—530 (1999).

27 Morgan N, Baron D, Edwards J, Ellis D, Gelbart D, Janin A, Pfau T, Shriberg E and Stolcke A: 'The Meeting Project at ICSI', Human Language Technologies Conference (2001).

28 Whittaker S, Hirschberg J, Amento B, Stark L, Bacchiani M, Isenhour P, Stead L, Zamchick G and Rosenberg A: 'SCANMail: a voicemail interface that makes speech browsable, readable and searchable', Proc CHI 2002, pp 275—282 (2002).

29 Rabiner L and Juang B: 'Fundamentals of Speech Recognition', Englewood Cliffs, NJ, Prentice-Hall (1993).

30 Campbell Jr J P: 'Speaker recognition: a tutorial', Proc IEEE, 85, pp 1436—1462 (September 1997).

31 Chellappa R, Wilson C L and Sirohey S: 'Human and machine recognition of faces: a survey', in Proceedings of the IEEE, 83, No 5, pp 705—740 (1995).

32 Dimitrova N: 'Multimedia content analysis and indexing for filtering and retrieval applications', in Informing Science, Special Issue on Multimedia Informing Technologies, 2, No 4 (1999).

33 Orwant J: 'Heterogenous learning in the Doppelgänger user modeling system', User Modeling and User-Adapted Interaction, 4, No 2, pp 107—130 (1995).

34 Monty M L: 'Issues for Supporting Notetaking and Note Using in the Computer Environment', PhD thesis, Department of Psychology, University of California, San Diego, CA (1990).

35 Fernandez R: 'A Computational Model for the Automatic Recognition of Affect in Speech', PhD thesis, MIT program in Media Arts and Sciences (February 2004).

36 Heeman P A, Byron D and Allen J F: 'Identifying discourse markers in spoken dialog', Proc AAAI 1998 Spring Symposium on Applying Machine Learning to Discourse Processing (1998).

37 Moran T P, Palen L, Harrison S, Chiu P, Kimber D, Minneman S, van Melle W and Zellweger P: 'I'll get that off the audio: A case study of salvaging multimedia meeting records', Proc of CHI'97 (1997).

38 Garofolo J, Auzanne C and Voorhees E: 'The TREC spoken document retrieval track: a success story', in Proc TREC, 8, pp 107—130 (1999).

39 ViaVoice — http://www-3.ibm.com/software/speech/

40 Lucene — http://jakarta.apache.org/lucene/

41 Wechsler M, Munteanu E and Schäuble P: 'New techniques for open-vocabulary spoken document retrieval', Proc SIGIR 1998, pp 20—27 (1998).

42 CMU Pronouncing Dictionary (cmudict0.6d) — http://www.speech.cs.cmu.edu/cgi-bin/cmudict

43 Arons B: 'SpeechSkimmer: interactively skimming recorded speech', Proc UIST 1993, pp 187—196 (1993).

44 Witbrock M: 'Lycos', (1999) — http://infonortics.com/searchengines/boston1999/witbrock/index.htm

45 Wong B A, Starner T E and McGuire R M: 'Towards conversational speech recognition for a wearable computer based appointment scheduling agent', GVU Tech Report GIT-GVU-02-17 (2002).

46 Vemuri S, Schmandt C, Bender W, Tellex S and Lassey B: 'An audio-based personal memory aid', To appear in Proc of Ubicomp 2004 (2004).

47 Foulke W and Sticht T G: 'Review of research on the intelligibility and comprehension of accelerated speech', Psychological Bulletin, 72, pp 50—62 (1969).

48 Vemuri S, DeCamp P, Bender W and Schmandt C: 'Improving speech playback using time-compression and speech recognition', in Proc CHI 2004, pp 295—302 (2004).

49 Ringel M, Cutrell E, Dumais S T and Horvitz E: 'Milestones in time: the value of landmarks in retrieving information from personal stores', Proc Interact 2003 (2003).

50 Brin D: 'Transparent Society', Addison-Wesley (1998).

51 'Hidden Cameras, Hidden Microphones: At the Crossroads of Journalism, Ethics, and the Law', — http://www.rtnda.org/resources /hiddencamera/allstates.html

52 Johnson C: 'Privacy lost: the supreme court's failure to secure privacy in that which is most private — personal diaries', McGeorge Law Rev, 129 (2001).

53 Rogers M G: 'Bodily intrusion in search of evidence: a study in Fourth Amendment decision-making', 62 Ind LJ, pp 1181 (1987).

54 Baddeley A D: 'Essentials of Human Memory', Psychology Press Ltd (1999).

55 Craik F I M and Lockert R S: 'Levels of processing: A framework for memory research', in Journal of Verbal Learning and Verbal Behaviour, 11, pp 671—684 (1972).

56 Ebbinghaus H: 'Memory: a contribution to experimental psychology', trans by Ruber H A and Bussenius C E, New York, Dover (original work published 1885) (1964).

57 Linton M: 'Memory for real-world events', in Norman D A and Rumelhart D E (Eds): 'Explorations in cognition', (Chapter 14) San Francisco, Freeman (1975).

58 Baddeley A D and Hitch G J: 'Recency re-examined', in Dornic S (Ed): 'Attention and performance', Hillsdale, NJ, Lawrence Erlbaum Associates Inc, pp 647—667 (1977).

59 Smeaton A F and Over P: 'The TREC-2002 Video Track Report', Proc TREC, 11 (November 2002).

60 LeDoux J E: 'Emotion, memory, and the brain', Scientific American, pp 50—57 (June 1994).

61 Finkenauer C, Luminet O, Gisle L, El-Ahmadi A, Van Der Linden M and Philippot P: 'Flashbulb memories and the underlying mechanisms of their formation: Towards an emotional-integrative model', Memory & Cognition, 26, No 3, pp 516—531 (1998).

62 Sinoff G and Werner P: 'Anxiety disorder and accompanying subjective memory loss in the elderly as a predictor of future cognitive decline', Int J Geriatr Psychiatry, 18, No 10, pp 951—959 (Oct 2003).

63 Balota D A, Dolan P O and Ducheck J M: 'Memory changes in healthy older adults', in Tulving E and Craik F I M (Eds): 'The Oxford Handbook of Memory', Oxford University Press, pp 395—409 (2000).

64 Hadis M: 'From generation to generation: family stories, computers and genealogy', Masters Thesis, MIT Media Arts and Sciences (2002).

65 Levin T Y, Frohne U and Weibel P (Eds): 'CTRL[SPACE]: Rhetorics of Surveillance from Bentham to Big Brother', The MIT Press (2002).

PhD candidate Sunil Vemuri is studying ways to help individuals and groups remember through the use of 'memory prostheses'.

This work investigates the role of audio in human memory recall, with the goal of designing wearable devices that help to trigger past memories.

Before coming to the Media Lab, he studied artificial intelligence at Stanford University and cybernetics at UCLA.

Walter Bender is executive director of the MIT Media Laboratory, a senior research scientist, director of the Electronic Publishing group, and a member of the Laboratory's information:organized consortium. He also directs the Gray Matters special interest group, which focuses on technology's impact on the aging population. A founding member of the Media Lab, he studies new information technologies, building upon the interactive styles associated with existing media and extending them into domains where a computer is incorporated into the interaction. He has participated in much of the pioneering research in the field of electronic publishing and personalised interactive multimedia. He received his BA from Harvard University, and his MS from MIT.