

Memory Prosthesis

by Sunil Vemuri

Thesis Proposal for the degree of Doctor of Philosophy
at the
Massachusetts Institute of Technology

March 2002

Thesis Advisor

Walter Bender
Senior Scientist
MIT Media Laboratory

Thesis Reader

Chris Schmandt
Principal Research Scientist
MIT Media Laboratory

Thesis Reader

Rosalind Picard
Associate Professor
MIT Media Laboratory

Abstract

This proposal describes a *memory prosthesis*: a novel device designed to help people with everyday memory problems. This proposed portable device's assistance begins by recording everyday conversations (audio and participants' biometrics) and identifying salient events in those conversations based on audio and biometric analysis. Next, it helps the biological encoding of a memory by allowing the wearer to review salient events soon after the event. Finally, it allows the wearer to actively seek a forgotten memory when a forgetting incident occurs.

To this end, this research includes a set of experiments designed to identify the most salient parts and effective memory triggers of an event with the goal of designing computational detectors for these triggers. Furthermore, it includes experiments in which subjects are given prosthetics designed to integrate with their daily work lives along with a suite of interfaces to review and find memories. Additional goals of this research include developing an understanding of how strengthened individual memories can improve group memories and the social effects of ubiquitous memory prosthetics and audio recording.

1. Introduction

Human memory is generally poor, prone to error, and to manipulation. It fails in unpredictable ways and at the most inconvenient times. Despite that, people function by either conceding memory fallibility or by relying on tools to improve their memories. Memory prosthetics have existed for as long as one can remember. Strings on fingers, Post-it Notes, and Palm Pilots are all examples of tools people use to help remember. However, all of these require the memorizer to make active effort to engage the prosthetic in order for the memory to be triggered. Furthermore, these tools only help with limited classes of forgetfulness.

In this paper I propose to develop a novel memory prosthetic: one that is designed to minimize the active human effort, while providing a range of targeted memory assistances. Memory assistance will be provided by passively recording audio and biometric data from all conversations, analyzing the data to identify salient parts, and then using these salient parts to either trigger a past forgotten event or strengthen one's biological memory of an event by improving the initial encoding of that event. Although research exists that explores the relationship of audio and memory in laboratory situations [Wilding and Rugg 1997] and long-term memory in real-world situations [Wagenaar 1986], no work presently exists that explores the role of audio and biometrics in real-world long-term memory. A significant goal of this thesis will be to use the prosthetic to take a first step in this direction by collecting quantitative data on memory improvement through prosthetics and qualitative data on usage of prosthetics in practice.

The proposed research will quantify how effectively audio can trigger past memories and help strengthen newly formed memories. This includes producing a set of algorithms for identifying salient events from conversations. Another goal is to identify optimal retrieval practices using a prosthetic when a forgetting incident occurs. Together, this body of work will lead to the interaction design of a prosthetic under specific work practices. Finally, this research will examine the effect of ubiquitous prosthetics on both the social norms and legal protections of both the wearer and the community.

1.1 Outline

A quick read of this paper should include Section 2 (Contributions), sections 3.1 (Taxonomy of Memory Problems), and Section 4 (Research Plan). The rest of this section provides a roadmap for a detailed reading of this paper.

The proposal begins by detailing the contributions of this research. Next, it provides pertinent background in human memory from a cognitive psychological perspective. This includes discussions of previous work done on forgetting, long-term memory in real-world situations, and the role of audio in memory recall. It goes on to examine previous work on computational human-memory prosthetics.

Section 4 describes the specific research proposal including the parts of human memory that will be tested and how the experiments will be conducted. This will include a description of the conference-room recording setting that will establish a control from which we will have a better understanding of both the role of audio and biometrics in real-world situations, but will also explore how memory prosthetics can assist with organizational memory by testing how a group's collective memory can be augmented with memory prosthetics.

This section also describes the wearable memory prosthetic that will be used to evaluate the role of memory prosthetics in the uncontrolled situation of the daily lives of the subjects. In addition to collecting everyday conversations, this device will introduce the ability for subjects to retrieve memories at their will instead of only at an experimenter's behest. These tests will also include an examination of the appropriate interface necessary for memory retrieval and which categories of memory failure can be best served by a prosthetic, and which cannot.

2. Contributions

Quantitative analyses of how effectively audio can both trigger past memories and strengthen newly formed memories is the most explicit contribution. This includes collecting data from real-world situations, analysis of what classes of events are better suited for memory triggering, and a strategy for optimally triggering such memories in practice. Considering no research to date has examined the role of audio in long-term memory recall in real-world situations, empirical results from these studies are expected to add to cognitive psychological understanding of human memory.

Different interfaces for retrieving memories will be developed and evaluated to determine what type of interface works best for different classes of memory failure. To date, no research has collected large amounts of conversational data and studied interfaces for accessing such a corpus. There have been many designs and studies of interfaces for accessing text, audio, conversations, etc., but none have studied how to access archives of one's own past.

The proposed research is intrinsically controversial in that it envisions a world in which memory prosthetics are ubiquitous. This has the potential to save a lot of "bathwater along with all the babies." Conversations including those with ill-formed thoughts, private personal opinions, or publicly inappropriate rants will be preserved. The author posits this future world as inevitable. This research also hopes to provide a better understanding of the social implications of memory prosthetics. This includes a qualitative analysis of what social settings prosthetics are appropriate, how prosthetics affect conversations, and what are the legal implications of prosthetics?

Section 4 details the specific hypotheses of this research. The key claims are described in summary here:

With regard to memory triggering, the exact audio of a past event will serve as a better trigger than a textual or computer spoken representation. Furthermore, certain classes of audio and sounds are expected to serve as better triggers than others. Humans are expected to be the better at selecting these optimal triggers compared to computational techniques. Finally, collective recollection of memories by groups will be better than individuals separately.

Daily-worn prosthetics are expected to cause an initial social awkwardness that will pass for some and remain for others. The prosthetic wearers when given a choice of memory-retrieval interfaces (laptop, PDA, speech-only) are expected to perform best with a graphical laptop interface, but use of the PDA interface more frequently in everyday practice. Finally, wearable-prosthetic users, when provided a summary of recent salient events, are expected to have overall improved biological memories compared to those who are not presented with summaries.

3. Background

Memory research spans a wide area, and a full treatment of this subject is beyond the scope of this paper. This section focuses on research most relevant to memory prosthetics. It begins by identifying a taxonomy of memory problems that will be used throughout the paper and identifies which among these problems are most common in workplace settings. Next, it discusses forgetting in real-world situations, expands on the different types of memory failures, and illustrates which ones have the potential benefit of prosthetics. Previous attempts at memory prosthetics are discussed in Section 3.3.

The specific design of the memory prosthetic will include a large vocabulary automatic speech recognition. The final background section will present information from both the speech-recognition and information-retrieval communities suggesting the feasibility of using this technology for the purpose of memory triggering.

3.1 Taxonomy of Memory Problems

Before getting into the specifics of forgetting, a useful place to start is Schacter's taxonomy of memory problems, or the "Seven Deadly Sins of Memory" [Schacter 1999]. He specifically broke down memory problems as follows:

Forgetting

1. Transience (memory fading over time)
2. Absent-mindedness (shallow processing, forgetting to do things)
3. Blocking (memories temporarily unavailable)

Distortion

4. Misattribution (right memory, wrong source)
5. Suggestibility (implanting memories, leading questions)
6. Bias (distortions and unconscious influences)
7. Persistence (pathological inability to forget)

The remainder of this proposal will use this taxonomy to position specific pieces of prior work and to indicate how the present proposed research fits into the greater body of memory research. The emphasis will be on the first three sins as those are most relevant to the present work on prosthetics. This is not to suggest that memory prosthetics cannot address sins 4-6. Rather, as will be shown in the "Frequency of Memory Failures" section, sins 1-3 are simply more common in the context in which the present research will be performed. Sin #7 will not be mentioned further as the present research aims to help people remember; it is not expected that prosthetics will help people forget.

3.1.1 Forgetting

Textbook treatments of forgetting typically start with the seminal work of Ebbinghaus whose experimental designs and “Forgetting Curve” (Figure 1) set the standard for decades of subsequent memory research. Ebbinghaus restricted his attention to “nonsense syllables.” In the context of memory prosthetics, this research is limiting because it only addresses nonsense syllables and the maximum duration of 31 days is far shorter than the lifetime of memory failures we expect in a daily-wear prosthetic.




Figure 1: Ebbinghaus' Forgetting Curve [Ebbinghaus 1885]

In addition to the fading of memory over time, there are distinctions on how memory fades based on the type of memory. For example, people significantly forget headline news and public events over the years but younger people are more capable of remembering both recent and distant events [Warrington and Sanders 1971]. Regarding names and portraits of classmates, even after 30 years, people are good at recognition but are poor at recall [Barrick et al. 1975]. Knowledge of foreign languages has been shown to fade quickly, but is then retained at a slowly-degrading “permastore” level (Figure 2) [Barrick and Phelps 1987]. Furthermore, these results suggest that a deeper initial learning is still retained even after 50 years. Motor skills involving more continuous motion (e.g., swimming, steering a car) are more easily forgotten than discrete skills having recognizable beginnings and ends (e.g., throwing, striking a match) [Baddeley, 1999][Schmidt and Lee, 1999]. Finally, memory retention depends on the type of material being remembered [Conway et al. 1991]. These examples suggest that memory recollection techniques are not universal and a memory prosthesis would need to be task specific.

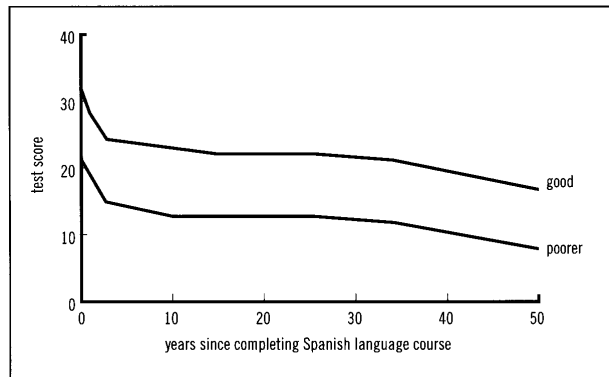


Figure 2: Bahrck and Phelps graphs on forgetting of foreign languages. The advantage of better initial learning is maintained over 50 years. [Bahrck and Phelps 1987]. Figure copied from [Baddeley 1999]

3.1.2 Frequency of Memory Failures

Terry performed experiments in non-workplace settings in which he had subject maintain “forgetting diaries.” Both the frequency and the type of forgetting were recorded. After collecting 751 forgetting instances from 50 subjects, his results suggested that most memory failures are a result of failing to perform some action [Terry 1988]. Eldridge performed more detailed studies in the workplace with the purpose of designing technological assistance based on common memory problems [Eldridge et al. 1992][Eldridge et al. 1994]. Through a diary study of 100 people who submitted 182 separate memory problems, she classified these memory problems into three categories (Table 1):

Problem	Frequency	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 post a letter, forgetting a lunch 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

Table 1: Eldridge’s classification of common memory problems in the workplace [Eldridge et al. 1992]

Under Schacter’s taxonomy, retrospective memories are directly analogous to transience. Furthermore, what Eldridge labels “prospective memory” and “action slips” are both forms of what Schacter calls “absent-mindedness.”

Retrospective memory problems are the most common and out of those remembering facts and past events form the bulk of the set. This suggests that reminding people of past events could help alleviate this problem.

Both Terry's and Eldridge's work confirm that most memory failures are a result of forgetting to do something instead of forgetting facts. However, Eldridge's work suggests greater fact forgetting occurs in the workplace than in Terry's broad everyday situations. For a memory prosthesis designer, this is fortunate since Eldridge claims that it is easier to design technology to support retrospective memory failures than it is for prospective failures or action slips.

3.1.3 Long-term Memory Studies

Wagenaar and Linton performed separate, influential six-year diary studies in which they recorded salient experiences every day into a written journal [Wagenaar 1986][Linton 1982]. Wagenaar wrote down the time, place, who was with him, and a brief statement about the event. At the end of the recording period, an assistant tested his recollection of randomly selected episodes. The significance of these works rests in their magnitude and application to real-world memories (granted the retrievals were performed in the lab).

Figure 3 shows the retention curves from Wagenaar's experiments. For controls, he did separate 1-year experiments before and after the 4-year main period of the experiment. This illustrates a sharp decay over the first year and then a steady decay afterward. Figure 4 shows retention curves as a function of retrieval cues and suggests more retrieval cues leads to better recall.

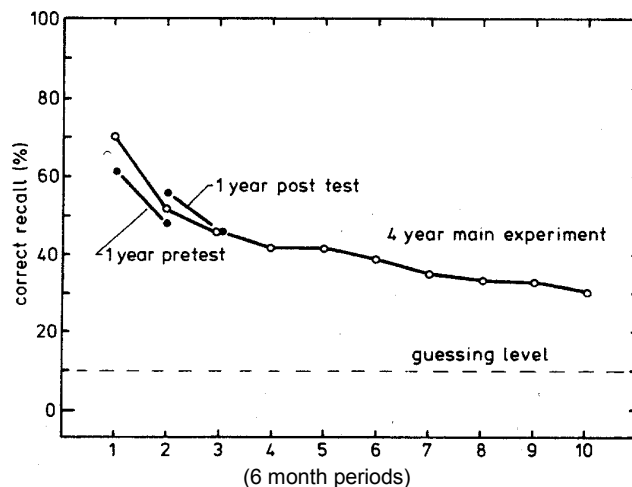


Figure 3: Retention curves obtained in three periods of experimentation [Wagenaar 1986]




Figure 4: Retention curves as a function of the number of retrieval cues. The curve marked “critical details” represents the recall of critical details after the presentation of all four retrieval cues. The curve marked “retention judgment” represents the number of events not judged to be completely forgotten [Wagenaar 1986]

It should be noted that there are differences between Ebbinghaus’ forgetting curve, which showed an exponential decay, and Wagenaar’s and Linton’s results, which both showed linear decay at 5% per year. Linton suggested the difference was due to the salience of the material. Ebbinghaus used nonsensical syllables with his experiments whereas Linton’s and Wagenaar’s results were based on real-world salient experiences.

3.1.4 Landmark Events and Context Dependency

Wagenaar’s work suggests that more triggers help elucidate memories. Loftus performed experiments demonstrating that some memories are better than others. *Landmark events* are significant episodes in one’s life (birth of a child, graduation). They can also be events shared by many (JFK shot, Space Shuttle Challenger exploded, Sept. 11 terrorist attacks). Loftus performed a questionnaire study in which subjects were asked about crimes committed against them relative to a landmark event that happened 6 months earlier (the Mt. St. Helens volcanic eruption in northwest United States). Her work suggests that subjects can more accurately determine the relative timing of past events when put in the context of a landmark event as opposed to exact dates. There are implications for accuracy of witness statements in court testimony. The problem is called “forward telescoping”: people tend to report past experiences occurring more recently than they actually did. [Loftus and Marburger 1983].

Loftus’ work specifically looked at transient memory problems, but one can speculate about the impact landmarks may have on absent-mindedness and blocking. There is evidence suggesting that *context dependency*, or the ability to remember better when the original context of the desired memory is reconstructed, has an impact on the

retrievability of a memory [Gooden and Baddeley 1980]. Intuitively, context dependency suggests one can retrieve a memory better when “returning to the scene of the crime.” This phenomenon is most pronounced under physical contexts, but there is evidence suggesting it can be mood based also [Blaney 1986].

Like Loftus’ work, Gooden and Baddeley’s study was also limited to transient memory. Both absent-mindedness and blocking describe instances in which situation specific memories are not readily available. Unfortunately, there is no research specifically looking at the impact of landmark events and context dependency on these two classes of memory failures.

The implication for the design of a prosthesis is to try to capture as many landmark events and as much context as possible. This remains a challenge as it is still not clear what constitutes a “landmark event” and different people will have different landmarks. One way alluded to earlier would be to capture news headlines associated with a given day. Also, it is unclear if a computational prosthetic’s reconstruction of a past event (i.e., graphical and audio) will have a similar effect as the full physical experience of returning to the actual context.

3.2 Retrieval

In addition to discussing the nature of forgetfulness, it’s worth discussing the same about how people remember or retrieve memories.

The act of retrieval mandates that the memory exists and is accessible and people store more than they can retrieve at a given moment. For example, category cues help retrieval and the act of retrieval helps in the learning of the material [Tulving and Pearlstone 1966]. Also, for a given retrieval session, while there is variability in the specific information remembered, the total amount of information remembered remains constant [Tulving 1967]. Finally, Brown and McNeill studied “tip of the tongue forgetting” (a form of memory blocking) in which the person knows that they know the forgotten memory but cannot seem to bring it forward [Brown and McNeill 1966].

It is general belief that long-term memories are encoded semantically. Craik and Lockhart demonstrated the “depth of processing” theory that suggests the more processing done during encoding leaves a stronger memory trace for future [Craik and Lockhart 1972].

3.3 Prior Work on Prostheses

This section describes previous projects that attempted to aid human memory. Most of these projects tried to address transience. Forget-me-not and the Remembrance Agent are both wearable devices that were put into practice in real-world situations, but were limited to text or iconic representations of past events. Memoirs and Lifestreams are prototypes designed to allow users to navigate computer data using temporal interface with memory triggers. Piles and Spatial Hypertext attempt to do the same, but using a physical metaphor. The Audio Notebook and Filochat explore the idea of navigating past audio experiences, but were limited to note taking situations. WearCam attempted to do the same but with a visual history in real-world situations. Finally, ScanMail is a more recent tool that allows indexing and retrieval of audio in the voicemail domain.

3.3.1 Forget-me-not

The “Forget-me-not” project designed one of the first portable human-memory prostheses: the ParcTab. It is a hand-held computer designed to work within the Rank Xerox Research Center in Cambridge, England. The environment is outfitted with sensors that allows the ParcTab to record a limited set of activities including personal location, encounters with others, workstation activity, file exchange, printing, and telephone calls. The device can be used as a portable episodic memory retrieval tool by allowing the user to search and browse their history of activity to try to recall past events. The interface (Figure 5) shows a short history of past significant events with iconic representations of people, location and data exchanges [Lamming et al. 1994].

The ParcTab is designed to address retrospective (transience) and prospective (absent-mindedness) memory failures. For retrospective memory failures, the ParcTab passively records information and allows the user to review past events when a forgetting incident occurs. To assist with prospective failures, it is contextually sensitive to the environment and proactively reminds users of situations-specific information they may need to know or be reminded.



Figure 5: The ParcTab [Lamming et al. 1994]

3.3.2 Remembrance Agent

The Remembrance Agent [Rhodes 2000] is a wearable computer that allows the user to write notes during conversations in a minimally intrusive way (Figure 6). A one-handed keyboard and a heads-up display allow the user to keep focused on their conversation partner while still writing notes. When the user is typing, the RA's search engine proactively identifies past notes pertinent to the current situation (determined by keyword matching). In this way memories and notes of past episodes can be elucidated through conventional information retrieval techniques. Rhodes collected several years worth of notes and cited anecdotal evidence suggesting the RA helped mainly with transience but also helped with absent-mindedness, misattribution, suggestibility, and bias to a lesser degree [Rhodes pers. comm. 2001].

all such events in the context of all other events. Hence, if a landmark event occurred proximal to a less salient document manipulation event, the user presumably would be able to localize that event quickly. Both of these systems specifically address transience, but the detailed histories they collected can also help with any of the first six memory sins.

3.3.4 Piles and Spatial Hypertext

MEMOIRS and Lifestreams use time as a navigation tool. Other projects attempt to use physical location of information as a memory trigger. The Spatial Hypertext project suggests that people better remember spatial arrangements of information instead of navigation paths [Shipman et al. 1995]. As mentioned in the background section, contextual dependency gives credence to this approach.

Piles [Rose et al. 1993] uses a metaphor of organizing information into stacks (similar to piles of paper on a desk) instead of hierarchical folders. Piles tend to have less structure than folders or directories and do not force users to commit to premature partitioning and organization of their documents. Users can arrange these piles in physical positions on the computer screen. To help users understand their partitioning schemes, Piles examines the contents of all documents in a given pile, determines similarities, and presents these to the user when they browse the pile (Figure 8).




Figure 8: An example “Pile” with information about the exposed document and the context surrounding it [Rose et al. 1993]

3.3.5 Audio Notebook, Filochat

Audio Notebook [Stifelman 1997] (Figure 9) and Filochat [Whittaker et al. 1994] capture audio during note-taking situations (e.g., a lecture or a meeting). Both systems allow users to make handwritten notes on a digital notepad while simultaneously audio recording. The audio is indexed based on the handwritten

stroke occurring at the same moment. Selecting a section of handwritten text retrieves the corresponding audio segment recorded at the time of the handwritten stroke. Both the audio recording and written record serve as memory triggers to the past event. Given the context under which both projects anticipate use, both are optimally assistive to transient memory failure.




Figure 9: The Audio Notebook Prototype [Stifelman 1997]

3.3.6 WearCam

One of WearCam's functions was to act as a "visual memory prosthetic" With the goal of alleviating "visual amnesia." The system included a wearable camera, a facial recognition system, and the ability to provide "flashbacks" or visual replays of recent images. These replays allowed the wearer to better encode their experiences by being reminded of previously experienced faces. Figure 10 shows images from WearCam's facial recognition system. [Mann 1996]




Figure 10: Images from WearCam. Left image showing template and right image showing information about match. Images copied from [Mann 1996]

3.3.7 ScanMail

Scanmail (Figure 11) provides a graphical user interface similar to email, but for voicemail [Bacchiani et al. 2001]. The system also includes speaker identification technology and uses large vocabulary automatic speech recognition technology to attempt to transcribe the voicemail. ScanMail's design is based on a set of studies that suggest frequent voicemail users tend to use such systems for tracking action items and other pending tasks. This suggests that in practice it could be used to alleviate absent-mindedness. Presently, no formal studies have been conducted to conclusively prove this point.

ScanMail also represents one of the only present systems designed to use large vocabulary automatic speech recognition (ASR) for browsing, searching, and retrieving past events – in this case, voicemail. Section 3.4 will discuss the use of ASR for memory prosthesis in more detail. However, the studies of ScanMail are inconclusive of the value of ASR in that system.

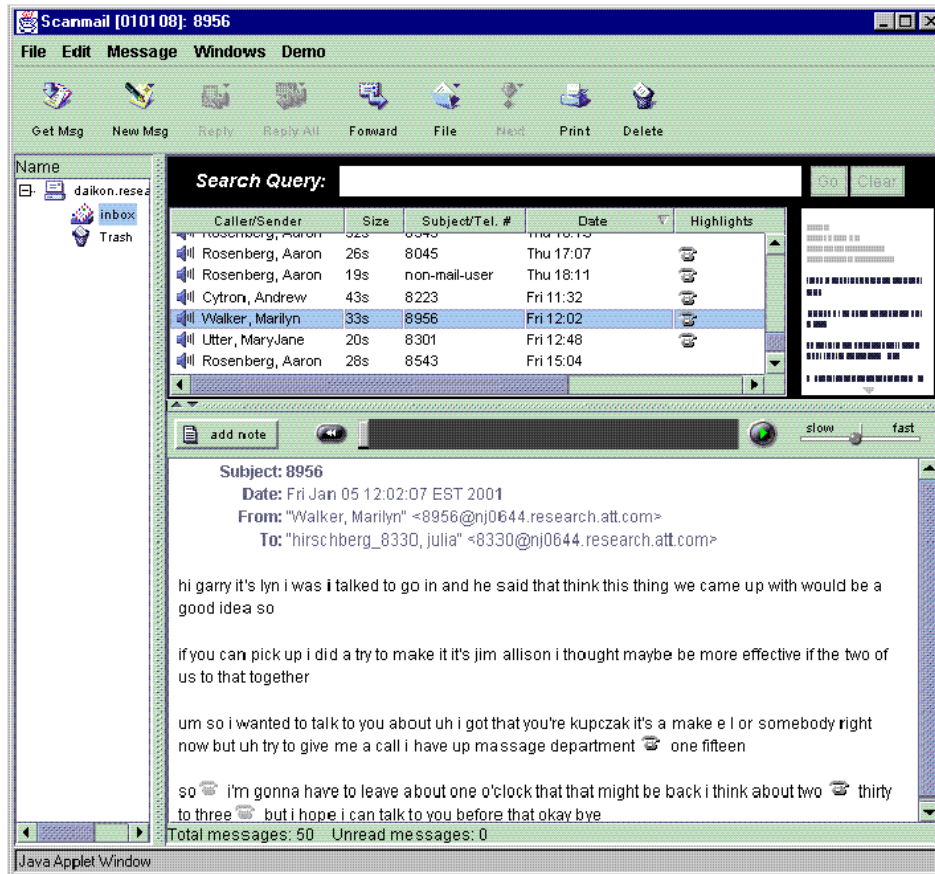


Figure 11: ScanMail interface [Bacchiani et al. 2001]

3.4 Speech Recognition

Though not a memory topic, a brief treatment of automatic speech recognition is worth covering as it has pertinence to the design of the present memory prosthetic. Specifically, can a large vocabulary automatic speech recognizer help in content searches of conversation recordings? The Spoken Document Retrieval Track of the Text Retrieval Conference (TREC SDR) examines the capabilities of various text search methods on the output of imperfect speech recognizers. Results from the latest TREC SDR indicate that keyword searching can be performed effectively on speech recognized broadcast news audio sources. Furthermore and surprisingly, there is almost no degradation in performance even when word error rates reach 25%, significant degradation does not begin until 60% WER, and one could argue that some effectiveness remains even up to 80% (Figure 12) [Witbrook 1999] [Witbrook 1997] [Garofolo et al. 2000].

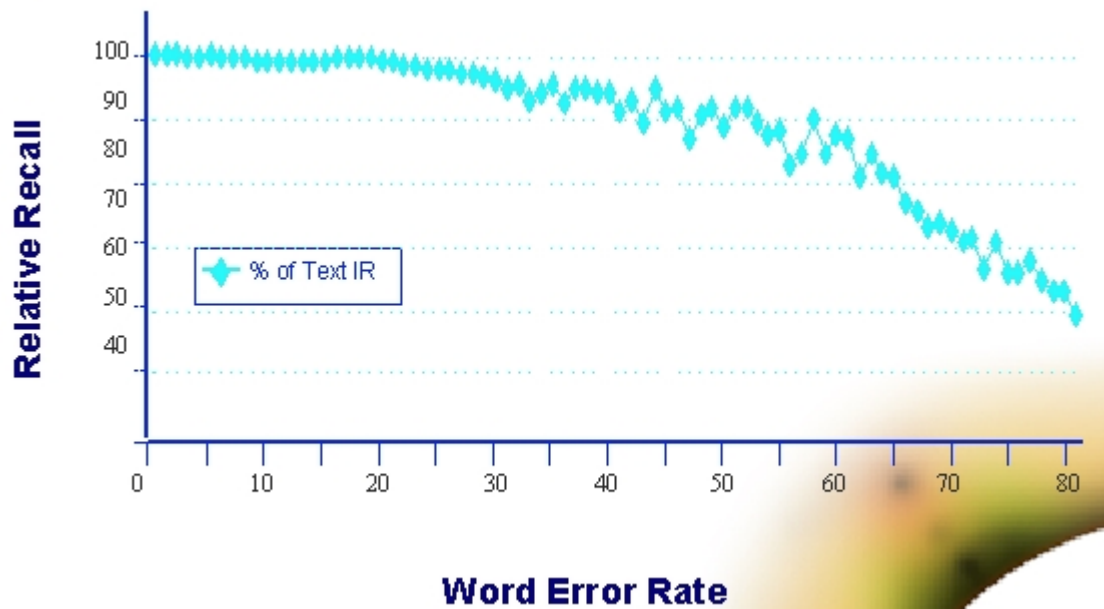


Figure 12: Effects of word error rate on recall. Figure copied from [Witbrook 1999]. The relative recall dimension shown here corresponds to mean average precision (MAP); 100% relative recall is approximately 0.4 MAP.

These results may suggest optimism on the use of speech recognition, but high word error rates are expected under the relatively noisy, multi-speaker, poorly microphoned and generally unconstrained settings of everyday conversations. Minimally, these results do not eliminate speech recognition as a possible tool. Conversely, an optimistic perspective suggests that even poor mean average precision would still be useful since a former participant in the conversation performs the retrieval. The goal is to remember, either by finding the relevant piece of information or triggering the memory of it. In effect, the combination of an imperfect speech recognition transcript and one's own recollection of a past event can help elucidate forgotten memories.

In an earlier section, ScanMail was described as useful for absent-mindedness because it is a common work practice to use voicemail for action item tracking. Those experiments

were also inconclusive with regards to the impact of speech recognition and one can hypothesize that speech recognition output is not useful for action-item tracking or absent-mindedness. However, one can hypothesize that memory prosthetics will be used in practice for transient memory failure. If so, users reviewing imperfect speech recognition for this purpose may be a viable option for content-based memory searches.

3.5 Memory and Biometrics

Studies of biometrics and memory recall tend to fall into two categories: the memorability of the fear response in animals [LeDoux 1994] and the role of affect on the formation of flashbulb memories [Finkenaur et al. 1998]. 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. Very few and sometimes only one learning instance need occur and the result is a conditioned response reaction that is often permanent or takes a long time to unlearn. 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., a 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 our intuitive notion that a memory of a fearful experience can elucidate a biometric and emotional reaction without actually reenacting the experience.

Flashbulb memories refer to memories of events of such significance that observers can recall their circumstances (who, what, where, when, how the felt, etc.) with abnormally high detail, clarity, and durability. Many people claim to experience this phenomenon with news events of high significance (e.g., JFK assassination, Space Shuttle Challenger disaster, September 11 terrorist attacks). There are several theories on how these memories are formed and why they are more durable and clear than other memories. Some suggest alternate memory mechanisms for encoding memories under intense affective state, while others suggest the simple increased instances of rehearsal and repetition would naturally occur with such poignant events. Though the theories differ in the mechanisms involved, they are consistent in that all suggest the relationship between elevated affective state and memory.

Both fear response studies and flashbulb memory studies suggest a physiological and/or emotional response to external stimuli have an effect on memorability. Unfortunately, studies have been limited to these phenomenons and there is little evidence to suggest relationships to physiological states and memorability of other classes of memories or under less extreme situations.

The implications for a memory prosthesis are that physiological state does matter when encoding certain memories. However, empirical evidence to date is restricted to laboratory animals and anecdotal evidence from those who experience flashbulb memories. A memory prosthetic that continually captures biometric information would

help determine if there is any relationship between memories under less extreme situations and physiological state.

4. Research Plan

The proposed research includes a set of experiments designed to explore the nature of memory recall using audio and biometric cues from past events. These experiments break away from the confines of laboratory tests and focus on real-world situations. Previous research on real-world long-term memory recall was based on physical contextual cues and not audio or biometric data [Linton 1982][Wagenaar 1986].

The first study, “conference room prosthesis,” attempts to determine whether there are particular parts of conversations that can more optimally trigger a past memory than others. Intuitively, one might posit that the more interesting parts of a conversation will best serve to trigger a memory of the entire conversation. The experiments described below investigate this intuition, attempt to identify the nature of these “optimal memory triggers,” and using only audio and biometric analysis, determine if these parts can be computationally isolated.

This research also explores the role of memory prosthetics in the unconstrained situation of everyday life. A set of experiments, titled “wearable memory prosthesis” below, will be conducted with subjects wearing a portable memory prosthetic. In addition to participating in constructed, quantitative memory tests, subjects will have the ability to retrieve arbitrary memories using a suite of interfaces. (A sub-goal of the investigation is to determine what is the right interface for memory retrieval and which interface modality works better for each kind of memory failure.) Finally, the research will examine the possibility of enhancing memory by actively improving the memory encoding process.

All of these studies intrinsically examine transience since the experimental procedures test subjects’ recollection of past events and facts. However, just like Rhodes’ RA, the wearable memory prosthesis has the potential to improve upon other memory problems (except, of course, persistence). Section 4.2.2 (Encoding Assistance) will discuss the possibility of reducing absent-mindedness. Given access to verbatim recordings, subjects are expected to experience improvement with memory distortions. However, it is not clear how much or how often this will occur since Eldridge’s studies suggest distortions are not commonplace in the work setting.

4.1 Conference Room Prosthesis

Experiments will consist of recording discussions during group meetings over an extended period of time (several months). After varying intervals of time subsequent to a meeting, subjects will be tested on their ability to remember specific information or occurrences of these events, based on selective triggering of the event.

Research groups in the Media Lab will be requested to voluntarily participate in these memory experiments. Subject will be drawn from groups who typically use designated meeting rooms equipped with specialized recording equipment. Data collection includes

audio recording a meeting and biometric recording of selected participants. MIT's "Committee on the Use of Humans as Experimental Subjects" has already approved the conference-room research protocol described below.

4.1.1 Identifying Optimal Triggers

The recorded audio and biometrics of meetings will be analyzed to determine which audio clips might serve as the best memory triggers. The determination of "best" will be evaluated empirically. Experimental subjects will be tested on their memory of a past meeting based on a variety of cues. These cues may include but are not limited to:

- (1) A textual transcript of a segment of the meeting;
- (2) Computer synthesized speech of a phrase from the meeting;
- (3) Verbatim audio clips from the meeting.

Hypothesis: A verbatim audio clip serves as a better memory trigger of a past event than either computer synthesized speech or a textual transcript.

The next logical step is to determine whether particular phrases, sounds, or audio corresponding to biometric states trigger memories better than others. To evaluate this, a variety of computational and human detectors will be used to select segments of audio from a long conversation. For example, a computational detector might pick out high-pitched utterances, another might select audio clips corresponding to high galvanic skin response, and a third might select phrases just prior to loud noises. The StartleCam and LAFcam projects illustrated methods for biometric startle response and acoustic laughter identification respectively [Healey and Picard, 1998][Lockerd and Mueller, 2002]. Such established techniques for speech and biometric analysis will be used to design these computational detectors, as it is not the concentration of this work to build a better audio or biometric scene-analysis system. Humans will also be used as detectors. For example, they will be asked to listen to an entire conversation and pick out parts that they thought would be particularly good at helping others remember the conversation. Candidate audio segments will be played to subjects to determine if one particular type of audio clip serves to trigger memories better than another. Selection of segments will be done using all of the aforementioned techniques, but only the audio corresponding to the cues will be played to the subjects.

Hypothesis: Some phrases, acoustic, and biometric cues are better than others at triggering memories. However, not all cues will be detectable using current acoustic, speech, or biometric analyses.

Hypothesis: Human-selected memory triggers will be better than computationally-selected ones. However, it is unclear which audio or biometrically determined memory triggers will serve as good triggers.

In practice, it is unlikely to expect a human to be available to select the best nuggets of every past conversation. Hence, one goal of this work is to identify differences between the best overall triggers and the most reliable computationally-determined triggers.

A mismatch is expected. For example, suppose a specific word or phrase (e.g., “rosebud”) is found to be a reliable memory trigger, yet a speech recognizer cannot consistently identify this word. In this instance, we have an excellent memory trigger, but no reliable way of computationally detecting it. Conversely, suppose we have an excellent heart-rate sensor. But, the results indicate no correlation between heart rate and memory recall. In this case, we have a robust sensor that provides no help in identifying memory triggers in conversations.

4.1.2 Group Memory

In addition to performing tests on an individual’s recollection of events in previous meetings, this research examines how memory prosthetics can help with organizational memory.

Hypothesis: Strengthening many individual’s memories strengthens a group’s collective memory.

The previous experiments describe tests of individuals’ memories of past group experiences. These same tests can be administered to a set of co-present people. The group can collectively attempt to recall the past episode using both the given cues and each other’s separate recollection of the event. This experimental design is similar to the above experiments and will not be described in any further detail.

Hypothesis: Groups are faster and more accurate in recollection of past, shared experiences than any of the individuals in isolation.

There are implications for design rationale. Design-rationale systems seek to capture the reasoning for past decision points. Typically, researchers study engineering groups seeking to build an artifact and later wish to review why certain engineering decisions were made. Another commonly studied group are managers. Design rationale systems strive to assist them with recalling the reasoning for a particular organizational decision (e.g., why did we choose this vendor?, why did we hire person X?, who was responsible for authorizing Y?, etc.). In these settings, the end result may be captured in some formal documentation or meeting notes, but it is also often the case that the decision-making *process* was never captured. Memory prosthetics may serve to capture some aspects of design rationale since all informal discussions are captured and can be used collectively by groups to recall the decision making process.

4.2 The Wearable Memory Prosthesis

The conference-room memory tests described above serve as both a control and a test of memory recall under the constrained situation of periodic meetings. However, given the vision of ubiquitous memory prosthetics, this research also intends to explore the impact

of daily worn memory prosthesis both in its ability to help people remember and how its presence affects conversations.

The wearable prosthesis is a small device (the size of a personal digital assistant), which will be carried by a set of experimental subjects (approximately 5). The device will be equipped to audio-record conversations (assuming permission is granted by all speakers in the conversation as per Massachusetts state law) and collect biometric data about the wearer as the conversation proceeds. Each subject will be asked to wear the device for approximately one to two months (during working hours only). The subjects will be drawn from volunteers in the Media Lab community. Given the time and effort commitment for experimental subjects, it may not be possible to secure a large number of participants or ensure diversity among the participants.

There are three key differences between the conference-room study and the wearable study. First, the amount of data recorded will increase from roughly 1 hour-per-week to roughly 4 hours-per-day. Second, in the wearable setting, the types of conversations will include the broad range of spontaneous topics a person converses every day. In the conference room, the conversations will likely be focused on research-specific topics pertinent to that group. Finally, in the conference-room setting, there will be no mechanism for participants to retrieve memories at arbitrary times of forgetfulness. The participants using the wearable prosthesis will be provided with such mechanisms to retrieve past memories whenever they wish. By providing retrieval mechanism this research will investigate how a memory prosthetics will be put into use in unconstrained everyday life and what types of memory failures occurred such that the subject wished and did not wish to use the prosthesis.

It is expected that prosthesis wearers will accrue anecdotal evidence on social aspects of memory prosthesis usage. For example: What level of comfort do conversational partners have with wearers of memory prosthetics? Do people shy away from a person wearing a memory prosthesis; do they come forward more readily? Does the existence of the prosthesis affect the conversation or what people are willing to say? Is there a novelty effect of the prosthesis?

Hypothesis: There will be a novelty effect, which will become less pronounced over time. Furthermore, some conversation partners will shy away from and cease to converse with prosthetic wearers unless they disable the device, others will be more forthcoming.

It is currently unclear how much time will pass when conversational partner become accustomed to a colleague wearing a memory prosthesis and repeatedly asking for permission to record all conversations. Beyond that, it is unclear how prostheses will affect conversations and either personal or professional relationships.

4.2.1 Retrieving Memories

There is little point in recording vast amounts of conversations and memories without the ability to retrieve a pertinent memory in a timely fashion when needed. Hence, an important aspect of this research is to evaluate the best ways for prosthetic wearers to trigger memories and retrieve recordings when a forgetting incident occurs.

Ideally, a memory prosthetic could detect when the wearer is experiencing a memory failure and would proactively provide assistance. Someday in the future, one can envision connecting a prosthetic to the neurological pathways in the brain responsible for memory recollection. However, this avenue is well beyond the scope of this research and our understanding of the brain. Until that day, conventional interfaces between humans and computers must suffice.

For the purpose of this research, three different interfaces will be evaluated to determine which one works best for different categories of memory failure. The first interface will use the full graphical and computational power of a typical personal computer. Software will be designed to allow users to browse and search memories. This will include a graphical representation of the audio recording, indications of salient events in the conversation (as indicated by users explicitly marking salient events at the time of the original recording, or by post processing analysis of the audio and biometric recording indicating a possible salient event). Finally, an errorful speech-recognition transcript of the conversation will be provided. The exact visual interface is not defined yet and is part of the ongoing research. However, this interface will be based on previous implementations of systems such as the “Intelligent Ear” [Schmandt 1981] and ScanMail [Bacchiani et al. 2001]. This interface will run on a conventional personal computer or more likely for our subjects, a laptop which they can carry around throughout the day (a common practice among Media Lab personnel).

The second interface explores the possibility of a speech-only interface. Here, there is no visual interface whatsoever. Only spoken commands can be given via a command and control speech recognition engine. The prosthetic’s responses will be audio only. Since subjects are already carrying a PDA with them to capture recordings, this will simply be an additional feature of the system that allows them to access memories.

The final interface will also be implemented on a PDA. It will be a blend of the first two: a search mechanism implemented on the PDA that can use speech and pen for input, and graphics and audio for output. This interface will be limited by the lesser computational and graphic abilities of a PDA but with the benefit from the portability of the small device.

Hypothesis: This PDA interface is anticipated to be the preferred mechanism for retrieval due to the increased convenience of a small device versus a laptop computer. It is also anticipated that subjects will be faster at searching and browsing using a graphical interface compared to a speech-only interface. But, it is expected that users will be less effective with both browsing and searching with

the limited graphics PDA interface compared to the full graphics laptop interface due to the lower bandwidth in which the computer and human can communicate. Nevertheless, it is expected that more categories of memories failures will be helped and there will be more successful retrievals using the PDA interface versus either alternate.

The interfaces detailed above describe tools that can be used when a prosthetic wearer is actively seeking a forgotten memory. These would be specifically useful for transient memory failures and possibly with blocking failures. Furthermore, they could also be used to actively find concrete evidence to remedy all three memory distortion phenomenon. However, such interfaces will have no use for absent-mindedness as one will not remember to look at a memory prosthesis if they have forgotten that they need to remember in their present circumstance.

4.2.2 Encoding Assistance

Can something be done about absent-mindedness? After all, it constitutes about half of all memory failures in the workplace [Eldridge et al. 1992][Eldridge et al. 1994]. This section describes how a memory prosthetic can be used to help strengthen one's biological memory soon after the memory is encoded. The hope is that strengthening memory soon after an event will mitigate the need for future prosthetic assistance. Ideally, this could help with all of the first six memory sins.

Current theory suggests humans encode memories in summary form. 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 given during the initial encoding process. Stored memories are encoded in gist with details filled in as more time, effort, and attention are spent on a particular event. The more details one encodes, the more association paths will there be to the memory. This suggests that memory prosthetics can aid in future memory recall by improving the encoding process.

As described earlier, the wearable prosthesis will be capable of identifying possible salient events in a conversation. The retrieval section above described how these could be used to help with active searches of forgotten memories. However, these salient events can also be presented to the wearer moments after a conversation or perhaps later the same day. Ebbinghaus' forgetting curve (Figure 1) suggests we lose over 50% of our memory of an event after the first hour and 66% after the first day.

Hypothesis: A summary of salient parts of an event within that first hour or the first day will improve future recollection of that event. Furthermore, the memory improvements will help all three types of forgetting and three types of distortion.

This encoding assistance need not be limited to the first 24 hours of a memory. Linton's results show that a single test of a previous memory, even after years, has significant impact on its subsequent retrievability [Linton 1975]. This suggests a memory prosthesis

that occasionally elucidated past relevant events (even if only once every few years) would help improve retention of such events.

To investigate this, subjects will be provided with summaries of some of their conversations at varying times after the conversation (either within an hour or within 24 hours). These summaries will be either human or computer generated in a manner similar to the experiments described in the “conference room” section above. At a later time (perhaps days, weeks, or months), subjects’ recollection of these conversations will be tested. One set of questions will pertain to those conversations in which the subject received summary information and another set of questions will pertain to conversations in which they did not receive summary information. These experiments are intended to show the differential effect of immediate encoding assistance on memory versus no encoding assistance.

Despite its potential benefits, prosthetic encoding assistance could interfere with the normal, biological selection of memorable events. Baddeley and Hitch showed that 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 [Baddeley and Hitch 1977]. By imposing a daily summary review, this may have the adverse effect of artificially amplifying the strength of prosthesis selected salient memories in lieu of an individual’s biological selection process.

Additional experiments will investigate the possible negative effect of a daily summary review session. In these experiments, the test subjects will be given summaries of some (not all) of their recorded conversations. The control subjects will not be given any summaries. Both will then be tested on their ability to remember details about conversations in which they did not receive summaries.

Hypothesis: The introduction of a summary reviewing activity soon after an event will not significantly reduce the memorability of events in which no summary was provided.

The reason for this hypothesis goes back to Baddeley and Hitch’s work. Additional activity between event and retrieval will make the event less memorable. However it is expected that the act of reviewing a summary will not be an additional activity. Rather, it will substitute some other activity. Hence, the overall activity remains constant, as should memorability.

5. Timeline

February 2002 – December 2002:

Recruit subjects, collect data, and perform memory tests in conference-room experiments

February 2002 – March 2002:

Complete development of wearable memory prosthesis

March 2002 – June 2002:

Initial round of wearable memory prosthesis experiments focusing on transient memory failure and debugging the system. Tests limited to retrieval using laptop and PDA interfaces.

June 2002 – September 2002:

Second round of wearable prosthesis experiments. Test transience as well as experiments designed to evaluate encoding assistance. Deploy speech-only retrieval interface for prosthesis wearers. Continue only quantitative memory tests on round 1 subjects (i.e., they no longer wear prosthesis)

September 2002 – December 2002:

Third round of wearable prosthesis experiments. Same as round 2, except new subjects, software debugging, and appropriate modifications to experimental protocol. Continue quantitative memory tests on both round 1 and 2 subjects.

6. Resources

Conference-room recording equipment. This includes a reasonable quality omni-directional microphone, microphone amplifier, and computer equipment to do recording.

Enough portable digital assistants and accessories for development team and all subjects. Presently, we are developing on the iPaq PDA. Biometric sensors for each prosthesis. Currently I anticipate needing 6 total configurations (5 for subjects, 1 for development).

Server and large amount of hard disk space to hold all of the audio data securely.

Appropriate software for speech analysis and software development. Presently, all tools I use are either free or site licensed to the Media Lab.

UROPs to help with development.

7. References:

[Bacchiani et al. 2001] Bacchiani, M., Hirschberg, J., Rosenberg, A., Whittaker, S., Hindle, D., Isenhour, P., Jones, M., Stark, L., and Zamchick, G. (2001). SCANMail: Audio Navigation in the Voicemail Domain. In *Proceedings of the Workshop on Human Language Technology*.

[Baddeley 1999] Essentials of Human Memory. Baddeley, A.D. (1999) Psychology Press Ltd. pp. 112.

[Baddeley and Hitch 1977] Recency Re-examined. Baddeley, A.D., & Hitch, G.J. In S. Dornic (Ed.), *Attention and Performance* (pp. 647-667). (1977). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

[Bahrck et al 1975] Fifty years of memory for names and faces: A cross-sectional approach. Bahrck, H.P., Bahrck, P.O., & Wittlinger, R.P. (1975). In *Journal of Experimental Psychology: General*, **104**, 54-75.

[Bahrck and Phelps 1987] Retention of Spanish vocabulary over eight years. Bahrck, H.P. & Phelps, E. (1987). In *Journal of Experimental Psychology: General*, **104**, 54-75.

[Blaney 1986] Affect and Memory: A review. Blaney, P.H. (1986). *Psychological Bulletin*, **99**, 229-246.

[Brown and McNeill 1966] Brown, R. & McNeill, D. (1966) The "tip of the tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior*, **5**, 325-337.

[Conway et al 1991] On the very long-term retention of knowledge. Conway, M.A., Cohen, G. & Stanhope, N.M. (1991). In *Journal of Experimental Psychology: General*, **120**, 395-409.

[Craik and Lockhart 1972] "Levels of processing: A framework for memory research." Craik, F.I.M., & Lockhart, R.S. (1972) in *Journal of Verbal Learning and Verbal Behavior*, **11**, 671-684.

[Ebbinghaus 1885] "Memory: A contribution to Experimental Psychology." Trans. H. A. Ruber and C. E. Bussenius. 1964. New York: Dover. Original work published 1885.

[Eldridge et al 1992] "Memory Problems at Work: Their Range, Frequency, and Severity". Margery Eldridge, Abigail Sellen, Debra Bekerian. Technical Report EPC-1992-129. Rank Xerox Research Centre.

[Eldridge et al 1994] "Autobiographical Memory and Daily Schemas at Work". Margery Eldridge, Philip Barnard, Debra Bekerian. *Memory*. Vol. 2(1), 1994, pp. 51-74.

[Fertig et al. 1996] "Lifestreams: an alternative to the desktop metaphor"; Scott Fertig, Eric Freeman and David Gelernter; *Proceedings of the CHI '96 conference companion on Human factors in computing systems: common ground*, 1996, Pages 410 - 411.

[Finkenauer et al. 1998] "Flashbulb memories and the underlying mechanisms of their formation: Towards an emotional-integrative model" Finkenauer, C., Luminet, O., Gisle, L., El-Ahmadi, A., Van Der Linden, M., and Philippot, P. *Memory & Cognition*. 1998 **26**(3), pp516-531.

[Garofolo et al. 2000] "Spoken Document Retrieval Track" slides. J. Garofolo, J. Lard, E. Voorhees (2000) in *Proceedings of The Ninth Text REtrieval Conference (TREC 9)*. Slides available only on-line at http://trec.nist.gov/pubs/trec9/t9_proceedings.html

[McKenna and Glendon, 1985] "Occupational first aid training: Decan in cardiopulmonary resuscitation (CPR) skills." McKenna, S.P. and Glendon, A.I. *Journal of Occupational Psychology*, **58**, 109-117. 1985.

[Gooden and Baddeley 1980] "When does context influence recognition memory?" Gooden, D., & Baddeley, A.D. (1980) in *British Journal of Psychology*, **71**, 99-104.

- [Healey and Picard, 1998] Healey, J. and Picard, R.W. StartleCam: A Cybernetic Wearable Camera. Proceedings of the Second International Symposium on Wearable Computing, Pittsburgh, PA, October 19-20, (1998).
- [Lamming et al. 1994] "The Design of a Human Memory Prosthesis" Mik Lamming, Peter Brown, Kathleen Carter, Margery Eldridge, Mike Flynn, Gifford Louie, Ptere Robinson, Abigail Sellen. in *The Computer Journal*. Vol 37(3), 1994, pp. 153-163.
- [Lansdale and Edmonds 1992] M. Lansdale, E. Edmonds. "Using memory for events in the design of personal filing systems" in *International Journal of Man-Machine Studies*. v. 36, 1992, pp. 97-126.
- [LeDoux 1994] "Emotion, Memory, and the Brain" LeDoux, J.E., *Scientific American*. June 1994. pp. 50-57.
- [Linton 1975] "Memory for real-world events" in D.A. Norman & D.E. Rumelhart (eds.), *Explorations in cognition* (pp. Chapter 14). Linton. M. (1975) San Francisco: Freeman.
- [Linton 1982] "Transformations of memory in everyday life" in U. Neisser, Ed. *Memory Observed: Remembering in Natural Contexts*. San Francisco. W.H. Freeman.
- [Lockerd and Mueller, 2002] LafCam: Leveraging Affective Feedback Camcorder.] A. Lockerd and F. Mueller. To appear in Proceedings of CHI 2002.
- [Loftus and Marburger 1983] "Since the eruption of Mt. St.Helens, has anyone beaten you up? Improving the accuracy of retrosptcie reorts with landmark events." Elizabeth F. Loftus and Wesley Marburger. *Memory and Cognition*. Vol 11(2), pp. 114-120. 1983.
- [Mann 1996] "Wearable Tetherless Computer-Mediated Reality: WearCam as a wearable face-recognizer, and other applications for the disabled." S. Mann. In *Proceedings of AAAI Fall Symposium on Developing Assistive Technology for People with Disabilities*, 11 November 1996, MIT
- [Rhodes 2000] "Just-In-Time Information Retrieval." Bradley J. Rhodes. Ph.D. Dissertation, MIT Media Lab, May 2000.
- [Rose et al. 1993] "Content awareness in a file system interface: implementing the pile metaphor for organizing information;" Daniel E. Rose, Richard Mander, Tim Oren, Dulce B. Poncéleon, Gitt Salomon and Yin Yin Wong; *Proceedings of the sixteenth annual international ACM SIGIR conference on Research and Development in Information Retrieval*, 1993, Pages 260 - 269.
- [Schacter 1999] "The Seven Sins of Memory: Insights from Psychology and Cognitive Neuroscience" by Daniel L. Schacter. March 1999, *American Psychologist* Vol. 54, No. 3. pp. 182-203
- [Schmandt 1981] The Intelligent Ear: A Interface to Digital Audio. *Proceedings, IEEE International on Cybernetics and Society*, IEEE, Atlanta, GA, 1981.
- [Shipman et al 1995] "Finding and using implicit structure in human-organized spatial layouts of information"; Frank M. Shipman, Catherine C. Marshall and Thomas P. Moran; *Conference proceedings on Human factors in computing systems*, 1995, pp 346-353.
- [Stifelman 1997] Stifelman, L.J. "The Audio Notebook: Paper and Pen Interaction with Structured Speech." Ph.D. dissertation. Massachusetts Institute of Technology. September, 1997.
- [Terry 1988] Everyday forgetting: Data from a diary study. Terry, W.S. (1988) *Psychological Reports*. 62, 299-303.

[Tulving 1967] Tulving, E. (1967). The effects of presentation and recall on material in free-recall learning. *Journal of Verbal Learning and Verbal Behavior*, **6**, 175-184.

[Tulving and Pearlstone 1966] Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, **5**, 381-391.

[Wagenaar 1986] "My Memory: A study of Autobiographical Memory over Six Years" Willem A. Wagenaar. *Cognitive Psychology*. 18, pp. 225-252. 1986.

[Warrington and Sanders 1971] The fate of old memories. Warrington, E.K. & Sanders, H.I. (1971). In *Quarterly Journal of Experimental Psychology*, **23**, 432-42.

[Whittaker et al 1994] "FILOCHAT: handwritten notes provide access to recorded conversations;" Steve Whittaker, Patrick Hyland and Myrtle Wiley; Conference proceedings on Human factors in computing systems: 1994, pp 271 - 277

[Witbrock 1997] Michael J. Witbrock and Alexander G. Hauptmann. *Speech Recognition and Information Retrieval: Experiments in Retrieving Spoken Documents*. In Proceedings of the DARPA Speech Recognition Workshop, 1997.

[Witbrock 1999] Michael Witbrock, <http://www.infonortics.com/searchengines/boston1999/witbrock/index.htm>, Lycos 1999.

[Wilding and Rugg 1997] An event-related potential study of memory for words spoken aloud or heard. Wilding, E.L., Rugg, M.D. (1997). *Neuropsychologia*. **35**(9), 1185-1195.

8. Author's Bio

Sunil Vemuri, a rabid Yankees fan, is an IBM Fellow in the MIT Media Lab's electronic publishing group. Sunil received his B.S. in cybernetics (UCLA) in 1992 and M.S. in computer science (Stanford) in 1994. From 1991-1996, he was a member of Apple Computer's Advanced Technology Group where he worked on search engines, organization memory, and technology to support learning.