

“Memories for life”

Managing information over a human lifetime*

Andrew Fitzgibbon (awf@robots.ox.ac.uk)

Ehud Reiter (ereiter@csd.abdn.ac.uk)

May 22, 2003

People are capturing and storing an ever-increasing amount of information about themselves, including emails, web browsing histories, digital images, and audio recordings. This tsunami of data presents numerous challenges to computer science, including: how to physically store such “digital memories” over decades; how to protect privacy, especially when data such as photos may involve more than one person; how to extract useful knowledge from this rich library of information; how to use this knowledge effectively, for example in knowledge-based systems; and how to effectively present memories and knowledge to different kinds of users. The unifying grand challenge is to manage this data, these digital memories, for the benefit of human life and for a lifetime.

For example, it is now possible to store every digital photograph one takes. A near-term challenge is to reliably organise and search an “infinite photo album”. Searching of textual information is well understood, but indexing of images and audio remains an open problem. This challenge will be forced upon us in this decade; however the optimal course of action is far from obvious, and the efforts of researchers from across all computing disciplines will be needed to ensure a successful outcome. In the longer-term, we might extract and indeed create not just individual photographs, but connected stories, such as “my son learns to swim.” This requires semantic analysis of images and other memory data to understand how they are connected, and which are most important to the story. The second strand of the challenge is to learn these semantic rules from the memories themselves.

The wider challenge is to extract knowledge from the data, and use this knowledge to build more intelligent tools. For example, we have for decades been able to build medical diagnosis systems that outperform most human doctors, *if* they have access to sufficient data about the patient. Unfortunately the data needed is almost never available, because it includes information that currently must be entered manually, such as the patient’s visual appearance, behaviour, environment, and history. If this information could be obtained automatically by analysing stored memories, it would revolutionise medical informatics and indeed medicine.

The creation, analysis, and usage of very large data sets is currently a hot topic in many areas of computer science, and recent advances are part of why we can hope

*This document is based on the UKCRC (UK Computing Research Committee’s) Grand Challenges in Computing workshop. It attempts to integrate ideas put forward by many people in submissions to the workshop, in discussions at the workshop itself, and in contributions to a web-based discussion forum after the workshop. For more information on UKCRC and its Grand Challenges initiative, see <http://www.ukcrc.org.uk>

to meet this challenge over the next two decades. For example, it is now possible to learn to identify individual objects in a photograph, a task which was previously thought to require human-like intelligence. Given several thousand photos that contain Ringo (some of which may also contain other people, such as Paul and John), and labels that indicate who is in each photo, we can use machine learning techniques to automatically construct a “Ringo-detector” that reliably identifies areas containing “Ringo” in new images. Learning techniques require large amounts of data, and this will be a natural byproduct of the management of digital memories. The digital archive of even one person in the year 2019 is likely to consist of petabytes of linked images, documents and audio; the potential for extracting useful knowledge from this archive is stupendous, and only limited by our imagination.

Digital memories clearly offer tremendous potential for science and technology. We must also ensure that they help society by widening access to information technology, so that everyone, not just well-educated people with no disabilities in rich countries, could benefit from the information revolution. The challenge is to develop detailed models of an individual’s abilities, skills, and preference by analysing his or her digital memories; and to use these models to optimise computer systems for individuals. For example, a short-term challenge could be to develop a model of a user’s literacy level by analysing examples of what he or she reads and writes, and linguistically simplify web pages based on this model; this would help the 20% of the UK population with poor literacy. A longer-term challenge might be presenting a story extracted from memories in different modalities according to ability and preference; for example, as an oral narrative in the user’s native language, or as a purely visual virtual reality reconstruction for people such as aphasics who have problems understanding language. Limited examples of such systems can be built now; the challenge is in mining the wealth of information latent in digital memories so that fully competent systems could be in use in fifteen years.

From a scientific perspective, this proposal is a challenge for many areas of computing research. The above examples have alluded to some of the specific scientific challenges in artificial intelligence, information retrieval, and human-computer interaction, but there are many challenges for other areas of computer science as well. For example, we will need to develop techniques for storing large amounts of complex data over decades and indeed centuries, in a manner that is robust to changes in hardware, operating systems, and indexing strategies. The computer and programs which operate on the data will change frequently over a human lifetime, but the data must outlast the systems which analyse it. Questions will be asked of the data which were not predicted when the data was indexed, so the indexing strategies must change over time. Security research must face the challenge of protecting information over decades, in a way that is robust to advances in computational power or mathematical knowledge, but without imposing untenable constraints on the user’s activity; and also the challenge of rigorously proving to a sceptical public that their memories are secure from hackers, amoral companies, and “Big Brother” governments. These are just a few of the scientific challenges of Memories for Life, more are listed below.

Memories for Life also raises public policy issues which must be addressed, in particular about control and access rights. For example, should courts or the police have the right to access memories that are relevant to a legal case or criminal investigation? Should people who are included in another person’s memories (in a digital photograph, for example) have any control over how these memories are used? Should aggregate information from memories be made available for medical and other kinds of scientific research? Such issues must be resolved in a way that is satisfactory to the community,

in order for Memories for Life to reach its scientific and technological potential.

Meeting the challenge requires the convergence of a number of disparate threads of scientific research. Acceptance of the agenda among the wider scientific community and the public must be as important a consideration as scientific viability. The challenge will be met when the majority of people can efficiently manage their information stream, and when all of us can benefit from our digital memories.

Exemplars of the Challenge

This proposal develops this programme using illustrative “exemplars”: specialist systems or tools which encapsulate the essence of the challenge. Each exemplar has access to all data stored about an individual, or to a subset pertinent to some aspect of the individual’s life, such as their personal or professional activities. Themes which link the exemplars include:

- The deep, persistent model of the user which is inherent in the digital memories, but which will be differently mined by each challenge;
- Sensory interaction between the user and computer which adjusts to the abilities of each user, including visual, aural and haptic interactions, and which allows all people access to digital information;
- Extraction of deep structure from the repository of memories, first to index the information, and then to present new views of the knowledge embedded therein;
- Adaptation of the representations to allow tasks whose specifications continually evolve, and for which the appropriate algorithms and data structures can not be known at the time the representation is first designed.

The exemplars are presented as addressing specific tasks, however each implies several scientific advances which have wider scope than the narrow domain of the exemplar. The challenges are labelled with an approximate time to completion, but many might be expected to remain research topics well into the 21st century.

Multimedia searching (5-year challenge): Search for images or audio by presenting examples, rather than text. Current technology uses textual annotations of non-text data, which are expensive to produce and can never be complete: an image of Mardi Gras may not be labelled “people”; the description of the “Mona Lisa” might run to more than a thousand words.

Electronic GP (10-year challenge): Analyse stored memories to create a model of a person’s activities, life style, behaviour and health. Use this information to give advice when the person has health concerns; as noted above, the biggest problem with current computerised medical diagnosis systems is lack of information. Advice could include diagnosis, referral to specialists, and suggested lifestyle changes; it should be presented in a manner that is relevant to the person’s situation and knowledge.

Stories from a Life (15-year challenge): Analyse stored memories and re-present them as “stories”. These may use a different modality from the original memory (eg, textual story from visual memory), and should be tailored for different audiences (eg, grandchildren vs. police) and contexts. Stories are created either automatically or under the guidance of the individual. Other people’s “memories” may be accessed when allowed and desirable. Many older people in particular might find this very valuable.

Personal Simplified Web Pages (5-year challenge): Acquire a model of an individual’s linguistic competence from the stored memories, especially emails and audio

data. Use this model to linguistically simplify Web pages, for example replacing words the individual may not know with paraphrases using simpler words.

“Newpaper” (10-year challenge): Develop “smart electronic paper” that lets anyone (even people without formal IT education) write down thoughts, scribbles, drawings, or whatever, and have these incorporated into the person’s digital memories. In addition to “clean sheet” information, newspaper also lets people display and annotate existing documents, including their other memories. Newspaper can be used anywhere (bus, bed, bath), so people always have access to their memories

Intelligent Mathematics Tutor for Children (15-year challenge): Analyse stored memories of all of a child’s mathematical attempts, including both formal schoolwork and real-life usage of mathematics. Create a model of the child’s mathematical knowledge, and use this to drive an intelligent tutoring system. Personalise examples and feedback from the tutoring system based on what the child is interested in and on what he has recently done.

Aid for Elderly with Short-Term Memory Problems (5-year challenge): Analyse stored memories, including visual data, to create a schedule of the person’s typical day. Monitor his or her activities, and help the person through the day by prompting. Also alert the person or his/her carer if there is a significant deviation from the normal schedule which is likely to be a consequence of short-term memory problems.

Virtual Memories (10-year challenge): Create a virtual world that represents an incident from a person’s life, using stored memories of that incident. Reconcile and integrate memories of different modalities (eg, video and emails), and interpolate as necessary in space and time to fill in gaps. For example, a 3D birthday party, or an action replay of one’s greatest sporting moment.

Answers to Questions

“Is it driven by curiosity about the foundations, applications or limits of basic Science?”

In considering how digital memories can be used to build systems which embody human-like knowledge, it might be thought that the grand challenge is to understand how information becomes knowledge. However, such goals are too grand. The challenge is driven by the need to overcome the limitations inherent in our current technologies for managing enormous, heterogeneous, and continually expanding information repositories. Overcoming these limits will certainly require basic research in many areas of computer science, for example:

- *Database systems*: We need to store many different types of data (text, audio, visual, log files) over a very long period of time (a lifetime and beyond). How can we do this in a manner which easily adapts to new hardware and software, which easily allows new types of information to be integrated when technology advances, and which easily allows new types of questions to be asked as society changes?
- *Security*: How can we protect people’s privacy, especially when one person’s “memories” contain information about someone else? What control should people have over information about them in other people’s memories, and how can this control be implemented? How can we prove both to the scientific community and to the general public that memories are secure from attackers?
- *Operating systems*: A person’s memories will contain petabytes of data and last for decades. How should this data be distributed across physical filestores, in a way that maximises accessibility and reliability (including reliability in the face of “once in a century” disasters)? Peer-to-peer networking is likely to be an important contributor to the system.
- *Artificial intelligence*: How can we interpret audio and visual data, with a minimal amount of annotation and guidance from the person? How can we learn useful generalisations from the interpreted data, and how can we represent and reason with these generalisations?
- *User modelling*: How can we represent people’s knowledge, experiences, beliefs, emotions, intents, abilities, and so forth in a coherent and unified fashion? How should we update such models when people change?
- *Human-computer interaction*: What is the best way from an HCI perspective for annotating and searching memories? How will new sensor technologies such as haptic interfaces be integrated? How should we adapt interfaces, web pages, and documentation so that they are well matched to the information about people that can be extracted from their memories (see user modelling above). In particular, how can we support people with disabilities or skill impairments, or people in stressful environments?
- *Graphics and virtual reality*: Given such a rich source of knowledge about the world, how can we use it to build virtual models of the world that integrate all the different types of sensory information in a coherent and consistent way?

“Does it promise a revolutionary shift in the accepted paradigm of thinking or practice?”

The challenge depends on a paradigm shift: from the concept of a “computer” for life, to a “memory” for life: an information repository which is conceptually separate from the computers which manage that repository. In turn, we must stop thinking of information in terms of disjoint data types (images, audio recordings, text files, web pages), and instead think holistically of information as giving different perspectives on people, events, and the world.

“Will its promotion as a Grand Challenge contribute to the progress of Science?”

Memories for Life will help science as a whole in many ways. Firstly, the data collected will itself be an invaluable resource for the cognitive sciences. To take one small example, detailed long-term data about individuals would be a tremendous resource for longitudinal studies about child development, disease progression, and so forth,

Secondly, better techniques for managing and analysing large data sets would be very helpful to many fields of science, ranging from genomics to economics. We live in an era where the amount of scientific data is exploding exponentially; managing and analysing this data is a challenge for all of science.

Finally, an argument could be made that the organisation of information is a fundamental constituent of scientific enquiry, and that a challenge to automate that organisation is essentially a challenge to understand the mechanisms which underly human thought. While this is not a goal of this challenge, it again indicates the importance of information management to all scientific endeavour. More practically, if some problems which currently require human intelligence are found to be amenable to computation, the implications for neuroscience and psychology will be significant.

“Does it have the enthusiastic support of established scientific communities? ”

The challenge speaks to over 30% of the submissions to the UKCRC Grand Challenges workshop and provides a provocative position on the nature of computer science. The specific scientific challenges we outline above are extensions of key goals in many areas of Computer Science.

“Does it appeal to the imagination of the general public?”

We hope the individual exemplars will excite people, because they can see the benefit of things like multimedia searching or intelligent tutoring. On a broader scale, our memories are what define us, and we believe that with careful presentation, the general public would be excited and inspired by our vision of computing with digital memories. Certainly the amount of media attention received by Microsoft’s MyLifeBits project¹ suggests that the media consider this topic to be of widespread interest.

People will certainly be concerned about the extent to which information about them is captured, and about who owns and has access rights to this information. Such issues are already being widely debated in the popular press. This is a public policy issue, not a scientific one, but people may be hostile to this challenge unless they feel confident that they have sufficient control over their memories.

¹<http://research.microsoft.com/barc/mediapresence/mylifebits.aspx>

“Does it avoid duplicating evolutionary development of commercial products?”

Certainly there is commercial interest in better ways of organising and searching personal archives of photographs and emails. Many commercial enterprises are built around search technology, and have an interest in extending the types of data which can be searched, or in moving the technology onto our local disks. Although most commercial R&D is focused on short-term incremental improvements to search technology, Microsoft in particular is also working on longer-term research in this area, in the MyLifeBits project mentioned above.

We welcome the interest of Microsoft and other companies in this area, but we believe that the extent of the challenge means that the path to a successful implementation of any exemplar will involve many false starts, and will require many minds working in parallel to achieve results. A single commercial organisation would find it difficult to take the risk that many person-years of work might lead to an unsaleable result. The development of universally trusted security protocols may require that no central authority controls privacy, and may depend on open standards and open research and development models which conflict with commercial objectives.

“When was it first proposed as a challenge? Why has it been so difficult so far? Why is it now expected to be feasible in a ten to fifteen year timescale? ”

An argument could be made that this was first proposed by Vannevar Bush in the 1940s. But it is only now that technology (disk capacity, sensors, processor speed) permits the acquisition and storage of large diverse collections of digital memories. We are just at the point in time where an individual could stream video to disk at ISDN rates and never be limited by disk capacity. Thus now is certainly the time to start managing these information collections. Confidence in the feasibility of the programme is reinforced by recent advances in areas such as search technology, computer vision and graphics, natural language processing, among others. Recent successes in machine learning and statistics have indicated that many of the problems that were previously considered hard enough to require cognition are in fact soluble in a purely data-driven manner.

“What are the most likely reasons for failure?”

Outright failure is unlikely—information management will always improve. It may prove, however, that some tasks cannot be solved without human cognition, and cannot be usefully automated. On the other hand, current progress in machine learning offers hope that this set of tasks is smaller than previously thought. On a larger scale, a significant risk is perhaps public hostility if the privacy implications are not carefully addressed.

“Is there a clear criterion for the success or failure of the project after fifteen years?”

Success criteria can be defined for the individual exemplars. For example, an Electronic GP could be evaluated by comparing the quality of its advice to that given by a human doctor; and Virtual Memories could be evaluated by asking people to rate their fidelity and internal consistency. It is more difficult to evaluate the challenge as a whole, since it is primarily a research direction and framework. However, certainly one measure of success is the degree to which algorithms and representations are shared and reused by the different exemplars.

“What kind of long-term benefits to science, industry, or society may be expected from the project even if it is only partial successful?”

Each of the above exemplars and scientific challenges should produce useful science and technology even if only partially successful. For example, even a partially successful “Personal Simplified Web pages” should give us a much better scientific understanding of how language use varies among individuals, and how web pages should be written to be accessible to as many people as possible.

“Does it have international scope?”

This challenge is of interest to researchers around the world. Indeed, DARPA in the US currently has a “Lifelog” programme² which is similar to Memories for Life in many ways, although focused more on the next 5 years than the next 20 years. Memories for Life is a general challenge about the acquisition, storage, and use of data about individuals; there is nothing in it that is specific to the UK or any other country. Research, wherever conducted, which makes a significant advance will be internationally applauded. There is good scope for involving developing as well as developed countries in aspects of this challenge, for example user models that incorporate cultural information about people.

“What calls does it make for collaboration of research teams with diverse skills?”

This challenge is particularly inclusive in that it naturally offers a place for many aspects of computer science. It is not narrowly focused on one subarea of computing, we believe that most computing science researchers will be able to contribute to the challenge.

“How can it be promoted by competition between teams with diverse approaches?”

The number of approaches which exist for the comparison and summarisation of non-text data is already large. Many blind alleys will be encountered in the search for efficient and reliable information management techniques. No single team can explore all the possibilities, but any individual advance will assist all researchers working on the challenge. The separation into exemplars makes the work naturally divisible into independent research efforts, and thus naturally allows collaboration and competition.

“What are the first steps?”

One obvious next step is to organise workshop(s) on Memories for Life, to bring interested researchers (international as well as UK) together and further develop the challenge. Beyond this, several contributors to our web discussion suggested that we consider creating an example Memories for Life corpus, which researchers could use to develop and test algorithms, applications, and so forth. Certainly corpora have been extremely valuable in many other areas of Computing research. Creating a corpus for Memories for Life would be expensive, time-consuming, and perhaps beyond the means of most individual research groups; but if such a corpus was created, individual groups could use it to investigate the research issues mentioned above.

²http://www.darpa.mil/ipto/Solicitations/PIP_03-30.html