Four Notes on 'Memories for Life'

Karen Sparck Jones
Computer Laboratory, University of Cambridge
William Gates Building, JJ Thomson Avenue, Cambridge CB3 OFD, UK
sparckjones@cl.cam.ac.uk
15 August 2005

These four notes comment on, and develop issues arising from, the document by A. Fitzgibbon and E. Reiter, ' "Memories for Life" : Managing information over a human lifetime', 22 May 2003, (via *www.memoriesforlife.org*), that became the starting paper for the UK Memories for Life Grand Challenge activity.

The notes deal respectively with
(1) interpretations of "memories for life" and their different implications;
(2) questions about data needs to support research on memories for life;
(3) salient features of a particular dataset, being developed by the European AMI Project, that could support some memories for life research requirements;and
(4) the consequences of forgetting, as well as remembering, in relation to views on memories for life and research data desiderata.

Comments on Grand Challenge Document (GCD): 'Memories for Life': managing information over a human lifetime

Karen Sparck Jones
Computer Kaboratory, University of Cambridge
January 31 2004
(reformatted, with minor edits, March 16 2005)

Summary

The GCD *assumes* that issues about data representation and manipulation are common to a wide range of computing applications, and so can be taken to justify labelling this data "memory for life". The correct view of a GC is that it is to *demonstrate* that this common, generic capability exists, so that substantiating it computationally is an advance for computer science as opposed to a miscellanous bunch of application developments.

There seem to me to be several threads in the GCD which are not adequately distinguished, and that while there are technical issues that are common eg managing large databases, handling multimedia material, there are significant differences between the threads which are not sufficiently clearly recognised by the set of exemplars. There is a particular problem in the ambiguity and hence confusion triggered by the use of the rather sexy word "memory". Thus it is not obvious that there are concepts and technologies to be developed that carry across all the exemplars.

I am deliberately starting with what 'memory for life' might mean before considering before considering the implications of the way(s) it is used in the GCD.
[Note: I use "data" to stand for "data"/"information"/"knowledge" as differences between these are not relevant here.]

Two very basic distinctions about memory are a useful start:
A: 'internal' vs 'external', and
B: 'for me' vs 'for others'.
A refers to whatever is 'in' my mind as opposed to how this externally manifested or expressed (contrast my mental image of Aunt Maisie with my expression of this in some communication medium). Internal is for me alone, but external can be for me (My Diary) or for others (my snowstorm story for my pub mates).

Unless we are talking about wiring up people's brains we're not into internal for me (or at least, leave it to the neuroscientists etc).

However we still have external for me or for others.

There is a tacit assumption throughout the paper that the role of automation is to enhance or enlarge a person's memory in some way. We are not very interested, or not primarily interested, just in manipulating whatever expressions I choose to produce as external versions of whatever is in my (known/aware) memory.

Interpretations of 'Memory for Life' :

Considering now the possible informal interpretations of 'Memories for life' (MfL): it can refer to all of the following:

1) MfL SuperMe

Consider a body of data which is intended to be an electronic enhancement of *my* memory. The presumption is that there is a mass of data items originating from me - as for example I might take a photograph, visit a store, give a talk and so on - that could go into an electronic deposit memory that was far more reliable than my actual human memory and which could thus be invoked to amplify what I actually remember. Call this interpretation of MfL the *'SuperMe'* one.

There are several assumptions here including:

a) that the form in which the data items is recorded is like that in my own memory: but (to put the point crudely) a snapshot is not a mental image;

b) that the way in which items in the electronic memory are organised and invoked is like that used in my own memory;

c) that the reason why I don't remember such things myself is simply human frailty and that I would like to have a much larger memory necause, after all, all these data items are part of my experience.

The model underlying this view of MfL is of a personal, individual, memory prosthesis in a rather strong, cognitivist sense - like the sort of prosthesis which is wired into my nervous system.

However there are clearly serious problems with all of assumptions a) b) and c). More fundamentally, as my initial remarks suggest, I don't think that we should be in this ball game at all: this not a version of MfL for *computer science*, however much computationally modelling it might be of interest to cogno/neuro/psycho scientists.

But there further versions of MfL. First,

2) MfL Deposit

Consider a data repository of items that I associate with myself, whether as original producer or as receiver and 'adopter', that I have stashed away at the time they were produced/adopted. An item might be a sketch, or a note, or a recorded bit of music. The repository could be as simple as a big box into which things are thrown or some neat filing cabinet. There is no presumption that I can at any time say what all the items are of the top of my head, but the items are both in principle and in practice accessible in some way or other, less or more efficiently. Call this view of MfL the *'Deposit'* one.

The assumptions here are:

a) that all the data items mean something to me, in some very primitive sense (that withered carnation is the buttonhole I wore ...), which is not the same as what they might look like to others (just a withered carnation);

b) thus the way that I might characterise the data items expressively for myself is not necessarily how they might be expressed for others;

there also appears, though more tacitly, to be the assumption:

c) that the way I store and organise (to whatever degree I do this) the data items has some useful relation to my internal thought processes in relation to my internal memory (I think I've got something about Joey in the box): this relation between internal real memory and external 'pseudo-memory' may of course be very complex [think Wittgenstein -'ah yes, this must be the picture of Joey'].

Both MfL 1 and MfL 2 can evolve over my lifetime, but by definition are of no direct use to anyone else. They only become of use to anyone else, in the same way as they are of use to me by what might be called revelation: expression in some form for external consumption, whether by design or accident. The difference is between my knowing this is a letter from my Mother to me and showing someone the letter and saying it's from my Mother to me (and also someone having the letter without any personal context for me).

SuperMe disappears with me; Deposit may survive but not with the same flavour. Thus the essential feature of both MfL 1 and MfL 2 is that in themselves they are not for anyone else but me. However we may also have another version of MfL:

3) MfL Persona

Consider a data declaration by me for consumption by others. This might be deemed my public memory, which may be modified over time. It is not clear whether this ought to be labelled a "memory' but since the term is used in the GC document to refer to a collection of items of interest to a person and pertaining to that person, and persisting over time, it ought at least to be considered. Call this version of MfL the *'Persona'* one.

The assumptions here are:

a) that I have some way of connecting this external memory with, or generating it from, MfLs 1 and 2;

b) that, though for public consumption, it is something over which I have some control, since it contains what I choose to put there.

This interpretation of MfL may not appear very convincing, but it is helpful because it contrasts with another, important version, namely:

4) MfL Assembly

Consider an assembly of data items *about* me, probably including but not necessarily restricted to those that I have 'authorised' (as in 3) or of which I am aware. This may seem a peculiar use of "memory" but is illustrated by examples like a long-lived doctor's record about a patient which may contain all kinds of items whose presence is justified by connection with the patient but which the patient may not have provided (a comment by another medical specialist), or been shown and hence 'adopted' even in the most minimal sense not presupposing much understanding of the item (Dr X: your unit count is low. KSJ: what are these units? Dr X: I've no idea - but there should be less than 20).

Assemblies may be 'friendly' or 'unfriendly' but that is a separate matter: even a friendly assembly could contain eg pictures of me taken by others that I have never seen and would suppress if I saw how awful they were.

Call this version of MfL *'Assembly'*.

The assumptions here are:

a) items in the memory can be connected with me as 'myself' but perhaps only with a good deal of gear crashing. However it is not obviously necessary for Assembly to work for all purposes that such connections need to be made;

b) because this form of memory brings together data about an individual supplied by others, but is also designed for use by others, it can support high mutual comprehensibility through the nature of the representation(s) used for the various data.

Unlike 1 and 2, versions 3 and 4 may of course persist after the person disappears.

Also, as this suggests, Assembly leads to a further version of MfL which has some similarity to 2, Deposit, but also a crucial difference. This is the idea of collective, or social memory, ie

5) MfL Collective

Consider a body of data items associated with a number of different people, with some

sharing or connectivity between the people but no presumption that all of the items apply to (or are already known by) all of the people. Libraries are an oldfashioned and conventional form of such memory, The Web a more modern form. Label this version of MfL *'Collective'*.

This version of MfL differs from 2, Deposit, in that while both might be deemed extensions of 'initial' memory, an individual need not, in fact will not, be aware of all it contains, and there is no presumption that the data is characterised or organised in a way fitted to any specific person, though it is clearly of no value unless an individual can establish some connection with the data. The motivation for Collective in the electronic age is that the data it contains is more conveniently reached, for extraction and incorporation in any individual's MfL 2 or indeed MfL 4 than if the person had to go foraging hither and yon among other places and people to find whether they know anything about some matter of interest.

Thus this version of MfL is based on the assumption:

a) that data created by others is accessible, ie interpretable, by others; and this applies in principle not only to data created by other humans but to data supplied by eg sensors.

The foregoing may look like a lot of MfLs. But they are different in many ways, both technically and politically. Thus for example, from the technical point of view,

1,2 and 3, SuperMe, Deposit and Persona, require that I can recognise and deploy the means of data representation, for all of the data in the base. This does not apply to 4 and 5, Assembly and Collective,

1,2 and 3 do not presuppose, and are indeed probably exclude, a single common form of data representation.

The same points apply to the operations on representations.

From the political point of view,

1 and 2, SuperMe and Deposit, involve a completely different notion, perhaps individually different notions, of privacy from 3 and 4, Persona and Assembly, while 5, Collective might have no privacy constraints.

The various MfLs have some common properties, but these are general rather than specific. Thus we may suppose that any of the MfLs contain data of varying intensional and extensional status, data that is inconsistent, and data that is heterogenous in representation. However because of the technical variations just mentioned, it does not follow that techniques developed to deal with the manifestation of these general properties in some particular form of MfL will carry over to others. For instance multimedia technology in relation to 2 and 5 respectively.

Implications for the Grand Challenge :

The explicit assumption of the GC Document is that while there are different research threads, there is a single 'essential' challenge that will be differently mined by various exemplars:

"each exemplar has access to all the data stored about an individual, or to a subset pertinent to some aspect of a person's life", and the exemplars are linked by themes, namely (that there is) "the deep, persistent model of the user", adaptive "sensory interaction between the user and the computer", "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", and "adaptation of representations to tasks whose specifications continually evolve" beyond the initial form of representation and operation.

My discussion of MfLs implies that there is a fundamental problem about the notions of each exemplar having access to all the data stored about an individual (or a pertinent subset), constituting or supporting a deep persisting model of the user. I interpret the document's statements to mean that there is *just one* ground database for an individual human user, that is the same underpinning for all of the exemplars, even if they exploit only part of it or use it in different ways. But what is this single body of 'all the data' stored about an individual? Setting aside 1, the most obvious candidate is 2, Deposit. However the presumption that any exemplar has access to all the data rather suggests 3, Persona or 4, Assembly.

It may be that we can do a lot of interesting and worthwhole things of a GC sort by assuming that we have a sort of detached MfL 2 which, by being sufficiently detached would be more like an MfL 3 and indeed, if sufficiently open, could merge into a 4 without difficulty. The differences between MfLs 2,3 and 4 on this basis would not be ones at the core, namely the nature of the data representations and operations, but much more 'managerial', according to whether there was a requirement a person should be responsible for, or own, the data, and through variations in the definition and needs for privacy.

The GCD simply takes it for granted that there can be a single individual database, that this can uncontentiously be labelled that individual's Memory for Life, that it will have data representations of entirely familiar external technological forms ie sound recordings, digital images, character strings, and hence that access to data items in memory is not a problem in principle for a computational system though, as in the case of image retrieval, it may be a practical one.

This might be deemed a shallow - though *not* trivial - view of memory. However while it largely collapses my versions (except 1 and, probably 5) it is not clear that doing this eliminates the real complexity in some of the exemplars, and hence makes achieving them easier or more satisfactory (though that of course depends on the actual degree of simplification permitted in the task).

Reviewing the listed exemplars in the light of *my* analysis of MfL, I see them as follows:

A) Multimedia searching: while a prerequisite or enabling technology, *in some form or

other* for all the MfL versions, is in no sense in itself a Grand Challenge. Whether it is or not depends on the form of representation appropriate to the various MfLs.

B) Electronic GP: appears to refer to MfL 3, Persona, as elicited from 2, Deposit, but also to 4, Assembly. The implications of manipulating these distinct MfLs needs clarity, or alternatively a proper case making for ignoring these distinctions.

C) Stories: also appears to refer to 3, Persona (or several different Personas), elicited from 2. What happens if the distinction is removed?

D) Simple Pages: refers to 3, Persona, ie is fairly obviously public.

E) Newpaper: refers to the process of adding to Deposit, 2. How will we know whether it works? This implies use of 3, Persona.

F) Intelligent Maths: refers to Persona.

G) Elderly Aid: this supposes direct access to Deposit, 2, though as the system is not the same as the person this contradicts my requirement for 3, Persona, as a mediator.

H) Virtual Memories: assumes access to Deposit, 2, and is therefore subject to the same criticism as Elderly Aid.

Another reaction to the GCD is to say, let's assume we can ignore distinctions between MfLs 2,3 and 4 (and set 1 and 5 aside altogether): to what extent, on this basis, do the exemplars constitute Grand Challenges in the sense of having computers able to do the type of task of which each is a concrete exemplar.

Getting practically worthwhile systems for some of the exemplar tasks at least, is clearly a big deal. But is what is being sought any different in kind from lots of AI, AI-ish, or AI-related, research already underway? The presumption is that the exemplars all require non-trivial stores of varied data and a battery of operations over these. But unless there is a better case made for the proposition that each of the exemplars will require a step improvement in system abilities to handle complex data *of the same general kind*, there is no single Grand Challenge. It is evident that the various exemplars would require the ability to manage complex data, though variously in fact and especially with respect to people's intensionality. But the issue is precisely whether the same data types and operations are needed across the board. The GCD starts by assuming that they are. My purpose here has been to propose a closer look at what 'MfL' really implies, and in particular to suggest that instead of assuming a common MfL type of apparatus across diverse tasks, differing only in local detail and not in essentials, the whole point of the GC is the other way round. It is to say that we hypothesise that there is such a generic apparatus so that, if we can understand it in the sense if being able to substantiate it through exploring various exemplars, this would constitute significant progress in computer science because it could enormously enhance the general capabilites of computers.

It follows from this that, as my MfL discussion suggests, the exemplars need more analysis to establish whether they do genuinely require the same sort of capability. (I don't regard

simply saying that they've got to be able to handle multi-media' as establishing this. From this point of view it is especially important to consider how systems for some of them could be effectively evaluated: for example, it is far from clear how Newpaper, as described, could be *objectively* evaluated. Further, the GCD's thrust, as the exemplars very clearly show, is human oriented and applications oriented. But, just as with the development and use of the notion of 'agent', it may be that the idea of MfL has something to contribute to the definition of computer system capabilities, regardless of their initial motivation in people and human-oriented applications, and that thinking about this possibility may be important for a better understanding of the idea of "memory for life" (for instance, is a system's internal log a MfL (or at any rate a potential MfL?).

In my view, the problem about the GC view of MfL is that it's all really about managing a lot of miscellaneous stuff adhering to a person that's accessible to the system, both technically and formally. Calling it creating or using a person's Memory for Life is a misnomer.

Memories for Life (MfL): Some questions about data

Karen Sparck Jones
Computer Laboratory, University of Cambridge
March 15 2005

Summary

This note examines the challenge of trying to find common material as a shared resource for projects working on memories for life. It considers the implications of the different interpretations that can be given to 'Memory for Life', as presented in an earlier note, and their consequences not only for the form of the data but, even more importantly, for the needs that data serves.

Background

Community research programmes in other areas have shown how valuable common, i.e. shared data can be. Such data has been particularly important when used for evaluating task performance, e.g. for establishing the absolute or comparative merits of different document retrieval systems. But common data can also be helpful when used simply as a shared platform for study and investigation, e.g. for exploring discourse structure as characterised by different structure models.

I will use

- *evaluation data* to refer to material which consists of both input material and desired output material
  - for example, for document retrieval, the evaluation data consists of both input documents and search requests and output data in the form of statements of those documents that are relevant to a request and thus ought to be retrieved;
- *working data* to refer to material which consists only of input material
  - for example for retrieval, working data could consists of document sets, or request sets, allowing the researcher to investigate potentially pertinent properties of documents and requests;
- *study data* to cover both forms of data as adopted for research; and
- *data* simply to refer to anything that may be considered or adopted as study data.

The key points about these distinctions are that

- data does not necessarily make sensible study data;

- working data is not evaluation data.

Having common study data is not just a rather vague matter of encouraging a sense of community. There are much more important reasons for it. One reason is improving data quality and spreading data costs. Thus for evaluation data in particular, costs can be high. In retrieval it requires relevance assessment for many documents in relation to many requests. This can be done on a larger scale and in a more organised way, allowing bigger, more solid experiments for the same overall provision cost as an bunch of small and hence less useful data sets. The other reason is encouraging comparisons between different research approaches. This is especially important for task evaluation, but is also helpful for exploratory investigation.

The problem for MfL is thus not really whether to have common data, but what common data to have.

It is far from obvious that there can be any question yet, given the general state of MfL research and goals for an MfL grand challenge programme, of providing evaluation data, i.e. data that represents what the memory processes to be modelled should deliver. For example, using retrieval as an analogy, if you don't know what retrieval is and hence what a retrieval system should do (beyond the totally vacuous notion that it should retrieve), you can't say what it ought to retrieve.

The issue for MfL, now, is thus whether it would be useful to have common working data, and what this should be like. So, by analogy, if you have got to the stage of thinking about document retrieval, it can help to get an idea of what documents are like, even if you don't know too precisely what you might want in retrieval: thus whether document sets consist of one-line slides, or alternatively of books, might help you to think about what retrieval might involve.

The specific problems with working data for MfL are whether you are assuming

- you have some memory already, i.e. you have some material which is taken to be given as in memory;
- you have a tabula rasa as initial memory, and some material which is to be taken into and form memory; or
- you have a combination of some material already there, and some more coming in to be absorbed.

The first and third of these commit you to some view of the nature of memory, i.e. what things are in it and how they are organised for use (aka how experience is represented and exploited). The second does not so commit you - memory organisation and use are up for investigation.

However, when it comes to developing common working data, different views of MfL imply not merely distinct but incompatible working data.

In a previous note (1) I identified 5 different interpretations of 'Memory for Life'. These have different implications for MfL study data, and more especially, working data. This note is therefore intended to stimulate discussion about what is required, would be useful, and might be available, as working data for research under each interpretation, and thus also about the extent to which there could be working data that could be shared across interpretations. Certainly, it is not clear that the kinds of data that have so far been suggested would be appropriate for different interpretations, in particular because the assumptions about memory on which they are based have not been spelled out.

For instance, to take a random illustration, having the images produced by a buttonhole camera taking pictures of the scene in front of the wearer every 30 seconds would not meet all MfL interpretations or research interests.

The next section examines the common study data implications of the 5 MfL interpretations. More specifically, on the assumption that the MfL enterprise is oriented at basic research and that it is premature to look yet for common evaluation data, I will concentrate on what having common working data involves. For convenience I will abbreviate common working data to *common data*, where by common I mean shared between research groups, and by data I mean at least *stuff*, but at least possibly more than that. As noted, the MfL enterprise is concerned both with stuff already in memory and stuff to go into memory, and hence with the form of the conversion between external stuff, *e-stuff*, and internal stuff, *i-stuff*. A crucial point about whether one can have common data is precisely whether one can do just fine by starting with no more than pure stuff, especially pure e-stuff, i.e. a bit/pixel/whatever stream, or you have to have something that already has some status as data, i.e. this is an image, this is me saying something in English.

MfL interpretations and data implications

1. MfL SuperMe

This is a prosthetic electronic enhancement of my memory in a strong sense. This version of MfL implies that memory content is represented and organised in the same way as my own wetware memory is, or at least is sufficiently close to this for there to be seamless connectivity.

While I accept that humans are all alike and that their basic memory forms and processes will be essentially the same, it does not follow that my individual electronic memory is going to be the same as anyone else's. So even if we knew what the nature of SuperMe memory is (and this is where we have to interact with the neuroscientists and are presumably computationally modelling their theories), I do not see scope for common data because this would be preempting what we want to find out. The whole point of research on this interpretation of MfL is to find out how stuff becomes data and is managed and exploited as data. It is not even clear how far it would be meaningful to work with *common stuff*, and specifically common i-stuff

It is not evident that we can be more than illustrative and indicative about the sort of thing we think a SuperMe memory might contain: e.g. a bit of neural net encoding some

'view' of a plate of breakfast cereal, linked to some blobs that are deemed to embody generic concepts of oatflakes or niceness. But of course the whole point of the enterprise is to find out what this memory content is, so it cannot be preempted by being given as common data. Further, since the name of the game is how e-stuff is converted to memory data, unless one is in the bootstrapping from tabula rasa game and simulating the newborn, there has to be some existing memory to which new e-stuff is to be related. It then follows that what happens to some new e-stuff depends on what some *individual* memory already has in it.

Thus the most that might be achieved through commonality is agreement that different teams could learn from working with the same micro world, let's say having muesli breakfasts and, perhaps, from taking the same neutral e-stuff, let's say some sensor data of plates of cereal on checked tablecloths as visually seen and tactically felt. Even getting some agreement on the scoping description of the common micro world, and on some neutral e-stuff is fraught with challenge, both from the technical and political point of view ('who cares about breakfast muesli, I'm into motion memory'; 'my system doesn't do tactile input and I'm not interested in it' ...).

I do not know whether even such limited commonality goals would be helpful, or could be achieved. Even if they could, it is evident that many interested in MfL hare focused on the other interpretations of MfL. But I believe that it could be helpful to those working on MfL SuperMe to examine whether some shared e-stuff might be helpful.

2. MfL Deposit

This is an external electronic addition to my memory which might contain (representations of) photographs I took, notes I made, recordings of me, blah. In this version of MfL there is no reason to assume that the objects, the way they are represented and organised, etc, have to be like the things, structures or processes in my head. All that's required is that they are available to and usable by my conscious cognitive processes in the same way as perfectly normal things like e.g. books are available to me to be read and the information in them used.

This version of MfL is not concerned just with unanalysed stuff for a human's internal mental processes. MfL Deposit presupposes some representational and processing scheme for what it contains. Thus while we might take e.g. raw video camera images as e-stuff for MfL Deposit, the name of the game is to interpretively process this in some way so as to retain whatever is deemed worthy of retention about it. But we might also take something more like e-data, i.e. material that has at least some recognised semantic aspect, e.g. text in a language I know, perhaps e.g. diagrams whose complete semantic basis I already know.

It may be that what is taken as input to this version of MfL is nearer the pure stuff end of the spectrum than the characterised data end. But either way the key presumption is that whatever the input is it is accepted, or selected, because it is pertinent to the particular individual involved.

There is nothing whatever to stop someone assembling some mass of study data of this sort and doing whatever they want with it as motivated by what they think would be useful to

themselves and in the process developing notions, tools, etc that apply elsewhere. For example developing some snappy way of finding similar images and some methods for customising this; or similarly with text. But this is the sort of thing lots of people are working on already e.g. Microsoft's 'Stuff Ive seen'. There may indeed be significant technical challenges, e.g. summarising videos to suit my interests, and potential general utility in the memory management concepts and tools to which such work leads. Some might argue that, though there is lots more work to do in this area, there has been sufficient already that a little light evaluation data, not merely study data, might not come amiss. However, supposing this is not the case and the present need is for study data, what does this imply?

The issue is whether some individual's data - and this version of MfL cannot work without being focused on an individual's data - could possibly also be *common* study data. For example, if Gordon Bell, who is recording everything for himself, or Wendy Hall's photographs from her time as BCS President, are made available, this material cannot possible be of interest to other researchers in the same way as to Gordon Bell or Wendy, and thus is highly unlikely to stimulate the sort of research in personal memory management with which the MfL enterprise is concerned. Such data could be used to investigate some types of problem, e.g. integration across data types, or image analysis, say, but it is difficult to see how someone for whom the material has no personal connection or resonance or interest could come to address the specific issues of *personal* memory management.

It may of course be worth investigating whether some body of material could be obtained or created which would have enough of the sorts of properties that one supposes are likely to apply to any individual's Deposit memory for it to be worth thinking about having some common data of this sort. There is still the question of how far one operates in (unrealistic) tabula rasa mode, or takes some material to populate an initial state of Deposit memory (implying some assumptions about the form and organisation of internal memory). There is also the question of whether whatever material taken as new input is chosen so as to be unanalysed and uncharacterised stuff, e.g. a recording stream of everything in person P's auditory environment over a week (phew), or already emodies some selectivity as to pertinence (e.g. a tape of X's lecture), or is more fully explicitly coded as some organised data set (e.g. a set of emails with explicit headers).

So the discussion here is whether it is worth taking some bunch of material that passes the minimum necessary requirement of all being associated with some *one* person's environment and interests (otherwise one is simply not meeting the requirement that one is dealing with individual memory), as a playground even if the material has no specific resonance for the researchers involved, only a generic similarity to the kind of thing they each might be interested in if they could have an effective Deposit memory of their own.

However even this requires careful thought about the range of input types, their distribution, their relations, their complexity etc: there is no likely utility to be got from having a motley assembly consisting of a few R&B recordings, some photographs of King's Parade, some distribution list emails about security research, and the texts of some Shakespeare plays, especially if these are not all associated with any one person, but even if they are. On the other hand a large collection of travel snapshots all taken by one person does not have the type heterogeneity that one wants.

To make a start on common data under this heading thus requires both hard initial analysis and definition of what is wanted, assessment of candidate material with respect to where it (necessarily) goes on the stuff-data spectrum, and whether it is material that can in fact be both properly and legally used by any random MfL researcher. Closed common data is no real research resource.

3. MfL Persona

MfL Persona is a body of material that presents me for public consumption. It is material about me that I have chosen to make available for others to use. The notion is that it is drawn from MfL 1 or MfL 2; and it could take various forms depending on assumptions about how this 'memory' of me is accepted as input by other people. As a crude illustration one could imagine offering one's Persona as a relational database, though the notion also allows presentation as an autobiographical text. Clearly there are some assumptions behind the form that MfL Persona takes about the ability of others to interpret and use what is offered.

Of course at the cognitive modelling level everything is much more complex, opaque, and unconscious than such a simple example as presenting myself as a database implies: what goes on in choosing and reformulating whatever is in my own memory for this public purpose. But even with more crude or surface-level versions of MfL 1 or 2, especially MfL 2, Deposit, there are research challenges about the conversion process and the reasons for choosing some material for Persona and putting it in a certain presentational light.

In practice it might be all more explicit, de novo, memory creation. For instance I might say I am going to make a new Persona and I'm going to choose to put into it whatever I select from any comptationally accessible memory there may be around for me: or, more realistically, I might simply decide to create this computational memory from my actual non-computational memory. So I might declare I like romantic novels and sentimental films, or I might declare I like 'Anna Karenina' and 'Madame Bovary' and 'Gone with the wind' and 'Some other soggy film' and have the Persona development software do a bit of autonomous processing (if it has been designed to do this sort of inference) to say (therefore) Karen likes romantic novels and sentimental films. Depending on the nature of the Persona development software I have more or less control over Persona and hence how I actually look to other people.

Whether MfL Persona is deemed to be connected with MfL 1 or MfL 2, it shares the intrinsic problems they have about common data. But the type of issue that common data for MfL Deposit raises apply even to the 'shallowest' version of MfL Persona. Suppose I simply decide to create a conventional database expressing what I (as a human) choose to present as me, indicating e.g. that I like novel N. The same limits of value apply to this as to any common data instantiation of MfL Deposit, namely that it all, by definition, is going to look different to other people pretending to be me than it actually does for me myself.

I noted in (1) that MfL Persona did not loom large in views of what MfL is all about; but it is a perfectly possible interpretation, and might have attractions for what might be called socio-neuroscience. However my main reason for making it explicit was as a clarificatory

contrast with MfL 4, which follows.

4. MfL Assembly

This version of memory contains material associated with me but it is not, unlike MfL 2, restricted in content to what I have supplied, authorised or, indeed, even know about. It is slightly odd to refer to it as *my* memory, since it is really memory *of* me as held by others. However, past discussions of MfL have assumed that this is a form of MfL that we should address. Note that it is not necessarily an agreed consensual or unitary memory (in the sense in which a biography is agreed or unitary memory). It is material held together by referring to me, the same one person (or at any rate what is believed to be the same one person).

There are many obvious examples of MfL 4, Assembly memory, for example everything about me as a 'medical subject' or as a credit agency's 'credit subject'. Such Assembly memory may be quite heterogeneous in its nature and form of component organisation, for example my medical Assembly memory might consist of a mass of items in distinct databases, linked by my medical patient number. It may also be the case that not all of the entire database is comprehensible to, or even open to, any one user.

The key fact about this version of MfL is that its organisation and processes (overall or per component) have to be understood by other people and, given that there may be many and varied such people, this memory organisation and processing has to be explicit and stylised, and may therefore have no close connection with internal human memory processes: think, for example, of relational databases.

Research on MfL Assembly has obvious intellectual and practical attractions through the actual or potential scale and heterogeneity of the materials. But it raises major questions about common data.

The first question is about the cognitive systems relevance of working with such material, given that what it is like is dominated by all of the (rational) requirements for ordinary multiparty usability. But the important question, here, is whether it would be possible to get common data of this sort. Some might be created (assuming some agreed data representation model, or at least agreed form of i-stuff). But this would generate the same plausibility and coherence challenges that afflict MfL Deposit. It also raises the problem that would arise in a much more serious form if we simply adopted some existing data.

This is that if we want to work with data about a person that has already been assembled for some particular purpose, what reason is there to suppose that this could be made *common* data. Noone is going to let a bunch of researchers loose on all of someone's medical, educational or financial records. The requirements for memory research require individual data (even if technically anonymised), not population data.

Specific research teams may well be able to get access to such material under agreed, and careful, constraints. But having some material which is community-wide common data is a quite different matter, especially where the community is not only the current, but the future one: retrieval research has conclusively shown that common data gains in value with

use in ever more experiments over time. It is hard to imagine that data that is truly and distinctively personal could be obtained to serve as common research data. Attempts to finesse this problem by making up e.g. simulated credit records, or to neutralise the material e.g. by replacing all the medical concepts by pop music ones, do not offer easy or even possible ways out.

5. MfL Collective, and Collective subsets

In (1) I used Collective memory to refer to information about multiple people, of interest to different people, without any presumption that all of the material refers to, or is known by, any one person. The motivation for the definition was that the material in such Collective memory is a potential extension of an individual's memory, as long as some connection between the individual and the Collective memory can be established: the analogy was with a library.

Insofar as there is at least one Collective memory out there already, namely the Web, there is no difficulty in principle about finding some common data. However there are some practical problems about using the whole vast Web, and of selecting 'random' subsets of a more manageable kind.

It may nevertheless be useful to consider particular subsets of some Collective like the Web as follows:

1) a subset of material potentially related to an individual, though not already brought to the individual's attention: there are obvious examples e.g. books like what you are known to like. Such material would be candidate e-stuff or e-data for MfL 2, Deposit, in particular.

2) a subset consisting of that material which happens to refer to, mention, etc some one person which could, as a technical matter, be brought together in some way but has not been brought together for some particular purpose. Such material might be candidate e-stuff for some MfL 4, Assembly. It may not form a coherent or seamless whole in itself, and material in it could, indeed, be damaged by removal from its social memory context. There are also questions about the propriety of creating such ad hoc aggregations, even from public sources (see (2)), as well as the general legal ones of 'taking over' data.

The research issues here are thus those of motivating some choice of Collective memory definition, as well as well as practical ones of dataset management. Thus this version of MfL draws attention to the particular issue of research data 'freezing': in document retrieval, test collections are usually fixed and do not change over time because this makes comparison with previous work impossible. The potential need for similar freezing for organised MfL research has also to be addressed.

The query problem

It is evident that getting common data, even only to the extent analogous to having a common set of documents for retrieval research, for example some videos or email files, is far more

complex than might at first appear.

But the vital question about common data is: what about the queries?

The lesson, emphasised and re-emphasised, in document retrieval is that what matters is the user's query. The character of the documents is important, but the way these are handled (indexed, stored, searched, matched, scored) by the system depends crucially on the characteristics of user queries. Overall system performance also depends on the user's judgements about document content relevance to their information need, but the way documents are characterised and organised has to relate to what queries are like.

The same applies to memory: having material in memory, per se, is of no or little interest. What matters is how it is used in response to new stimuli, and the way it is returned in response depends on how it is characterised, which is a function not only of what has been input before but of how stored material has been invoked.

All of this is well known. But it is essential when considering study data, and looking for common study data. It is a particular problem, further, in any attempt to work with surrogate personal data, especially when this means working with surrogate users of actual personal data, but also when working with surrogate data. As noted under MfL 2, Deposit, there is a problem that while X's data may be significant for X, it doesn't follow it is significant in the same way for Y. Similarly if we have 'X's' data, i.e. some deprivatised version of X's data, there is no reason to suppose it will be of interest to anyone in the way that the original data is of interest to at least one person.

This problem becomes less and less of an issue as we move from MfL 1 to MfL 5. But in doing this we also, as noted, become progressively less cognitively personal about memory. However for any version of MfL, and for any candidate common study data, we have to ask the question: what would somebody want to use this material for, for real, or even for simulated real. Even if we are not in a position to think about evaluation data, but only about working data, we will not get far without a large mass of instances, concrete examples, of how people want (not merely might) use it.

The point here is not about being able to answer the questions using whatever actual memory model we have; nor is it about thinking outside the box of our current technological limitations. It is much more about having a lot of convincing, detailed, examples of what we want to be able to get from a memory.

It is not enough to handwave: 'Wouldn't it be nice to be able to invoke the image of that beautiful sunset I saw in Wyoming last year'? What is the evidence that this is the kind of question I actually ask of my own memory or, as memory is progressively distanced from my own person, that I would actually ask of, say, some MfL 3, Assembly, that happens to be to hand. One might indeed build some system to do this, but is the question real, with whatever form of reality the different MfLs involve? The danger is that, just as in all the talks about the hypothetical wonderfully wired house, the facilities that are offered to illustrate this possibility turn out to be limited and banal.

To avoid this, the proposal might be made that we should just get a big mass of MfL *N*, and see what we can do with this. This is the same as wiring up everything in some house.

We would learn something from wiring up some house. But even to do that profitably, we at least need to think carefully about whether the house we are trying to wire up is actually Ellen MacArthur's yacht, a family igloo, the White House, or Cottenham village. Even if we don't choose some existing house, but a notional one, it will still embody some view of what a house is, for example, an executive model in new development in a favoured suburb. And we won't get all that far with any of these without asking what are we trying to achieve by wiring this particular choice of house?

Conclusion

Common data for research can be of immense value. But for MfL, it can't come easily.

References

(1) K. Sparck Jones, 'Comments on Grand Challenge Document (GCD): 'Memories for Life': managing information over a human lifetime', Computer Laboratory, University of Cambridge, January 31, 2004.

(2) K. Sparck Jones, 'Privacy: what's different now?', *Interdisciplinary Science Reviews*, 28 (4), 2003, 287-292, via http://www/cl/cam/ac/uk/ ksj

Memories for Life (MfL):
A note about the AMI Project data

Karen Sparck Jones
Computer Laboratory, University of Cambridge
April 29 2005

Background on MfL data

This note considers the potential utility of the AMI Project data as study materials for MfL. It assumes the analyses of my previous notes one what an MfL might be (ref 1) and on forms of data pertinent to work on MfL (ref 2). I will consider the data in relation to all the MfL interpretations (or types) I discussed earlier, though its pertinence to *MfL Deposit* and *MfL Assembly* is perhaps most important. I will also assume, following my earlier discussion of data, that we want at least what I have called *working data*, material that one can explore for modelling and processing, since it is not yet clear what functionally-related *evaluation data* might be required for MfL research and hence how it might be characterised.

The Augmented Multi-party Interaction (AMI) Project - see
http://www.amiproject.org -
is a large European project, aimed at enhancing collaborative working. To support its research it is gathering data in the form of recorded meetings, the *raw data*, to which many levels of annotation are being applied, the *labelled data*. The data has been carefully considered, and will be publicly available.

The *general* reasons for examining potential MfL uses for existing or currently being built data are the obvious ones of cost saving, especially for types of material that are intrinsically expensive, and of neutrality: if material produced by others is of a generically suitable kind for MfL purposes, using it avoids the danger of bias and over specification. At least, since what MfL is supposed to be like is sufficiently vague, there is every advantage in doing some initial exploration which some reasonable-looking data as may be to hand. For MfL purposes in particular, there is much to gain from community data sharing. It is useful not to be lumbered, at the beginning, with material confined to particular individuals that may either be of no interest to other people or raise confidentiality issues, or with material that presupposes particular foci of attention or approaches to memory modelling.

The following gives the salient facts about the AMI data. There is substantially more information available. The AMI Project is very positive about other projects using their data. It will be released under licence similar to the creative commons one at
http://creativecommons.org/licenses/by-nc-sa/2.0/
ie will be freely available for research purposes, with the condition that if further annotations are added to the data these are made similarly available. The intention is that most of the data/annotations are made available 6 months after gathering/making to allow AMI people their own publication time, so some is already available. If there is substantial interest within the MfL community in using the data, it would probably be helpful all round to develop some more organised communication with the Project to avoid too much fragmented, ad hoc and

hence inefficient interaction.

The AMI Project (AMIP) data

The AMIP data is already being collected. It covers several data sets. The most important subset is based on a consistent generic *scenario*. For present purposes the scenario data is of primary interest, so the description which follows, and the label 'AMIP data' refers only to this scenario data. The Project uses a Hub and Spokes model for its various data sets. The scenario data constitutes about two thirds of its 100-hour Hub data.

The scenario data is all in English, though not necessarily native-speaker. The raw data consists of many sets of four successive one-hour meetings about designing a consumer product. Each series of four meetings involves four people with different roles, and the series of meetings follows four stages in the product design process. The presumption behind the data gathering and labelling is that team members could use what would in effect be a limited memory for a project life (LMfaPL) to update or (re)brief themselves, check points, motivate follow-up etc. The long-term goal of the overall project is to show that such multimodal memory support leads to more effective working together. However quite apart from the possible value for the MfL enterprise of having such a LMFaPL as a sort of pilot - with scope for evaluating ideas about memory and its use for some external purpose, the data being collected has much potential as material for exploring general MfL issues, and is considered here from this point of view.

The raw data consists of multiple data streams. There are are four participant-focused cameras, each looking towards one participant's face, an overhead fisheye camera and other scene cameras. Each participant has a head microphone and there are other microphones covering the room. There is a whiteboard capture stream, a slide presentation capture stream, and note capture streams for each participant's electronic pen. The participants also have laptops. (The project envisages between-meeting activity eg email or document collection, and any such additional data could be added to the data streams.) All of the data streams are compatible time-stamped and so can be synchronised.

From the point of view of general MfL interests, the one lack in the raw data is 'from-participant' camera views, and thus also any direct connection between participant X's view of the room and what they individually say. The individual-centred view of what is going on would have to be taken as adequately represented by one of the other cameras. (This might however be entirely adequate for many exploratory purposes.) There is also no direct recording of the individual's own auditory input signal: the presumption, again adequate for many purposes, is that the individual hears what the overall microphones record (and this can be correlated with who said it through the individual lapel microphone data).

The labelled data covers many layers and types of labelling. For present purposes this can be divided into *lower level* and *higher level*. Lower level includes such things as the stream source identification eg camera X, microphone Y, time marking and perhaps also location and speaker identification as this is tied to the recording layout. For convenience lower level annotation might be taken to include automatic speech transcription (ASR). Higher level labelling includes linguistic annotation on the speech eg part of speech tagging, conversational

turn taking, semantic annotation eg topic marking, pragmatic annotation on both sound and visual streams eg gesture information, and task annotation eg decisions made. The higher level annotation is typically hand annotation, which is also useful as a gold standard when automatic annotation is developed.

All of this data capture and annotation relies on an elaborate storage infrastructure and set of management tools. The basic data capture is well underway. Annotation is also in progress, more advanced for lower levels and less for higher. However much material could be available for MfL researchers now or soon.

From the MfL point of view the two aspects of this data are

1) how it relates to the different MfL types, and
2) whether it can be used to study both memory *input* is handled and how memory *content* is managed.

The second distinction is pertinent to some MfL research because it allows for work specifically focused on how 'naked' perceptual data might be categorised or given meaning through the creation of some interpretive scheme, as opposed to work that relates new input to some existing (though possibly evolving) interpretive scheme. The potential value of the AMIP data for MfL research is that it appears to have the potential to serve multiple needs with the additional advantage that these can all be connected through the common initial input, the recordings themselves, as well as through any particular forms of annotation that are pertinent to more than one MfL need. The data origin in the meetings scenario and AMI project goal is also useful for those who are interested in databases designed to serve applications project needs.

The AMIP data's relevance to MfL types and uses can be sumarised as follows.

MfL SuperMe: as the point just made implies, there is some scope for simple modelling of a personal cognitive extension.

MfL Deposit: the AMIP data could be very well used to support research on MfL Deposit, ie on how some individual might find it helpful to organise their personal experience. This basically depends only the relatively modest leap of imagination needed to consider oneself as a meeting participant who happens to have done this X by this time T. Deposit research also depends on at least some existing memory content to which new content is assimilated. However the AMIP background of scenarios etc, and the concepts used in the annotations (especially the higher-level ones) could be used as the substructure to develop some existing semantic memory at least for a world in which meetings figure (eg a suitable ontology). Similarly, the AMIP background could help to develop some motivations for and characterisations of actions that are needed to study memory use. The AMIP data, with its extended time feature, is especially valuable for studying both memory input and memory content.

MfL Persona: the AMIP data is less obviously pertinent to MfL Persona, but could nevertheless be used to support some research on it, for example by providing a record of past experience that could be exploited to project a new model of an individual: as a simplistic example, information about what happened at the meeting could be taken as a base for producing a revised record. Thus the transcript of participant P's actual contributions could

be edited to produce a follow up record that might be eg more decisive. The AMIP data would support both input and content aspects of Persona.

MfL Assembly: the AMIP data is implicitly assembly data already, since it contains material generated by one participant that bears on another, for example spoken comments by X on what Y has argued. It could clearly be taken as a base for work on more systematic or larger development representing multiple party input to a base on a single person. Again, both input and content aspects would be covered.

MfL Collective: the AMIP data is also, naturally, a base for working on collective MfL.

General comments

The AMIP data does not seem much in quantity or rich in content variety. But in fact 100 hours of multi-stream is quite a lot of material for detailed study; and the detailed semantic content (linguistic or visual) of meetings even under scenario constraints is large; at the same time it avoids the problem that many large data collections have, namely of a lack of sufficient internal connectivity to test categorisation, searching, restructuring etc. The general aims behind and organisation of the AMIP data make it well suited to many MfL purposes. In particular, the multi-level annotation can be used to give a helping hand to get a memory instantiation off the ground, without having to develop definitions of relevant entities, properties or relationship from square one. However for those who want to focus on that, it is possible to ignore the higher-level, or all, annotation completely.

(1) K. Sparck Jones, 'Comments on Grand Challenge Document (GCD): 'Memories for Life': managing information over a human lifetime', Computer Laboratory, University of Cambridge, January 31, 2004.

(2) K. Sparck Jones, 'Memories for Life (MfL): some questions about data', Computer Laboratory, University of Cambridge, March 15, 2005.

Memories for Life (MfL): Forgetting vs deleting, forms of memory, and research data

Karen Sparck Jones
Computer Laboratory, University of Cambridge
July 5 2005

Abstract

In this note I argue that 'we can keep it all' assumptions about computational memory, so there will be no role for forgetting, or need to delete anything, are at least simplistic and may be fundamentally unsound. Much depends on the type of memory for life envisaged. Considering forgetting, and deleting, thus has implications for ideas about the kind of working data it would be useful to have for MfL research.

Memory capacity

In relation to cognitive science what follows is amateur and simplistic. But I believe the key points about computational memory hold.

It has been noted more than once in MfL discussions that human memory processes involve forgetting as well as remembering, and also that forgetting can have a positive role. It is not just a negative human failing. Thus forgetting has positive value in removing redundant or unwanted clutter and hence helping to make concepts and concept relations that matter more salient. Nevertheless the basic view is that given the vast and continual flow of input, even after some initial perceptual filtering, the limits of memory capacity mean that some stuff has to be dumped even if it might have turned out handy later. Thus loss of memory content is independent of failure to retrieve content that is actually there, though there may be a connection between the two since content that never gets fingered may be more likely to be abandoned.

In most discussions about computational memory, the explicit or implicit assumption is that ever-growing computing power means there is no competition for memory capacity. Ever more content (stuff or information) can be held in memory, which is an unequivocally Good Thing. The only issues are how to organise the content to support retrieval and how retrieval works. Indeed it is not only that there is no pressure on space to force content ejection. Setting aside technological questions of preservation, which are at another level, we also get a different take on the implications of operational failures to retrieve: these need not lead to permanent loss of content. Content can stay around for ever and maybe later, in new favourable circumstances, be picked up. Even in human memory, failures to retrieve on particular occasions need not imply immediate permanent loss. The tacit assumption in the computational case seems to be that there need never be any loss, however little something is fingered.

Thus the simple computational model is not just of a house with ordinary rooms with frequent activity and an attic of moderate size from which never-used lumber is removed at

intervals, but of a house with an infinite attic from which nothing is ever removed, on the general principle that it might come in handy some day.

The infinite attic model may seem attractive. But it raises serious issues, specifically in relation to the different definitions of what a Memory for Life might be I distinguished in an earlier note (1), namely MfL1 - Superme, MfL2 - Deposit, MfL3 - Persona, MfL4 - Assembly, MfL5 - Collective.

Forgetting leads straight to these issues. But as there are vocabulary traps on the way, I shall set up some terms for talking as follows.

Some terms

We already habitually use some words, notably 'memory' itself, in both people and machine contexts. We also apply machine metaphors to humans eg 'compute' for 'think'. But we still tend to use one lot of words for people, and another for computers, in relation to memory. Thus, for example,

> for people we have *remember* vs *forget*
> for machines we have *keep* vs *delete*.

The angle here is whether stuff is preserved in memory or not. 'Remember' is actually ambiguous between whether something is being preserved so it can be recalled, and the process of recall. Using 'remember1' and 'remember2' is tiresome, so I shall reserve 'remember' for 'remember1' and use *recall* for 'remember2'. For machines we correspondingly have *retrieve*.

There are lots of other pertinent words about, for instance 'store', 'file', 'record', usually referring to machines, which are ambiguous about permanence, 'archive' which suggests permanence but is ambiguous between mere dumping and considered conservation, 'purge', 'recover', and so on. 'File' is a fine problem term, suggesting keeping as opposed to not keeping, but we still delete files.

Now, what are we actually doing when we go along with the computational MfL as infinite attic, and assume we need never again suffer the embarrassment or inconvenience of forgetting anything? More mundanely, will we never need to delete anything?

Forgetful machines?

For humans, remembering, forgetting and recalling are unconscious (being conscious about them is a second-order business). For machines, keeping, deleting, and retrieving are (so to speak) conscious. This is the case even if we simulate unconscious processes by such things as exploiting machine learning to develop personal interest profiles for browsers: these are currently independent, external programs with the same logical black box status for us as other software.

The fact that we do not (at present) know how human memory processes work is irrelevant. Even if we did, it is not clear that this would enable us to consciously manipulate our own,

especially in fine-grain routine operation.

Suppose, therefore, that we now ask what happens if we amplify our human memory in the various MfL ways and ask, specifically, what effect this has on forgetting. If one of the major benefits of computational enhancement is to abolish forgetting, how will this actually work?

MfL1 - Superme

This notion of computational memory is as a seamless extension of our existing mind, just increasing our current mental capacity but operating in the same way. This implies that insofar as forgetting has a positive role in concentrating memory content on what matters to the human involved, then forgetting will extend to the computational add-on. The add-on would not be a genuine prosthesis if it did not, because it would imply the human and computational memories were not fully integrated. The implication is that if forgetting is needed to make human experience manageable, the extent of computational amplification will not be great. Science fiction aside, we might think of making people's memories and hence mental powers somewhat better, but we cannot think of orders of magnitude capacity enlargement. Equally, the implication is also that all the processing would be as unconscious as it is now.

MfL2 - Deposit

MfL Deposit is not a seamless extension of my open mind. It is nevertheless an extension for my exclusive use, so the presumption is that the material in computational memory is stored and organised in some useful way that is also connectible with my own memory. We can therefore expect to have linking or clustering or foregrounding or whatever. But there seems to be no reason to forget anything or rather, since MfL is now an arena for explicit computational memory operations, to delete anything.

However if we consider the need for foregrounding etc more carefully, this seems to imply that even if we do not delete anything, material may become so backgrounded it is thoroughly deactivated, and there is thus a question about how it might be reactivated.

MfL3 - Persona

MfL Persona is me for public consumption, so there is no need for any very strong connection between the way this memory is organised and used and how my own 'real' memory is organised and used.

This MfL type almost suggests a continuous refashioning or versioning, implying unused material gets older and older. However, in contrast to the previous case, it is not evident that there is any requirement to retain older versions since they will not be directly usable in new, changed circumstances. There is a natural reason to forget them and even to delete them. Certainly, if they are retained indefinitely, we have a question about how earlier material is recovered other than by some simplistic time-tagged means. For example, what kind of semantic evolutionary-chaining over the material needs to be available?

MfL4 - Assembly

MfL Assembly is yet further from me myself. This suggests that the notion of forgetting is not appropriate and the issue is the role of deletion. Is it necessary to keep entire older versions (it may of course be necessary to keep particular old content, recognised as still pertinent to current needs)? Again, will older versions be deleted, or do we want to preserve history? If we do, is there any need to do this in an elaborate semantic way, or will time-stamping be good enough?

Certainly, keeping older versions would appear to be rather differently motivated from retention in the earlier cases (apart, of course, from the shared need in all cases for vulgar crash recovery). The need would rather appear to be of the public evidence sort, eg to demonstrate that some fact was in the database and has disappeared, for perhaps improper reasons.

But more importantly, this memory type has other users than myself: so who controls any forgetting or deleting?

MfL5 - Collective

MfL Collective has much the same characteristics as MfL4, but with added complexity since the memory is not that of a single person.

The crucial problem

This review shows that forgetting, or deleting, for memories over time does not have the same flavour for the different memory types. But it may not suggest, other perhaps than in relation to MfL1, that forgetting or deleting are *necessary* for computational memory. The whole attraction of computational memory is that capacity seems to remove the need for irrecoverable loss.

The further tacit assumption is that everything in memory will be relatively accessible, since an enlarged content store is the name of the computational game. Souping things up with foregrounding or salience mechanisms, dynamic associative linking and so forth, is viewed as supplying essential but convenience-oriented overlays that do not block out underlying material and that can be connected with the 'substrate'. This will be a challenging business for the large volumes of material involved, but is not taken as logically incompatible with effective ways of organising content for present contexts and needs.

Thus having doubts about the compatibility of necessary processes that focus and reinforce some memory content while never losing other content (including structural relations, not just individual items), may seem like a lack of faith in the power of computing. The growth of the Web suggests quite the reverse: we can still find our way through ever more billions of pages.

But this ignores two key considerations. The first is the ways in which the computational processes and memory illustrated by, for example, Google's search engine and page set compare with our own. The second is the effect of the sheer volume of input to be dealt with, and in particular the huge volume of similar inputs.

Web engine search

Web engine search is *extremely* crude by comparison with whatever humans must do. It is also based on explicit concept characterisation of a particular form, namely through words (even image searching is primarily via linguistic keys). The large search support embodied in URLs and links also reflects conscious decisions about labelling and packaging items and about cross-referencing them. Language is a thoroughly good thing, and we rely on it, but it does not follow that the ways in which alphanumeric 'semi-linguistic' URL labels and simple word keys are used for the Web have much to do with the the kind of memory MfL is aiming at, or at least with the more challenging cognitive-end types. The comparison with the Web is also misleading because it implies similarities in *granularity*. But there is no evidence that web-page granularity (even allowing for great variation) has much to do with experience and memory granularity.

Current work on the Semantic Web does not suggest anything radically different is envisaged; or at least, insofar as Semantic Web effort is devoted to ontologies, it has built in notions of atomic concepts, relations and so forth that it is the business of MfL research to explore, not take for granted.

Data volumes

Humans experience enormous streams of very similar unit inputs - using the analogy of video image frames or speech stream sound slices, there will be hundreds (thousands ?) of these per minute. Human processing selects/merges/compresses/filters these at perceptual input to avoid getting swamped, but there are still enormous numbers reaching short-term memory and very many entering longer-term memory. The kinds of recording currently familiar as computational data do not begin to reflect the volume of input experience humans have to cope with. For example sensors are limited to specific signal types, monitoring cameras have low resolution, conventional video has been edited, music comes in prestructured composed form, and so on. Even those playing with high-class 'on all the time taking pictures from my lapel' cameras are not in fact taking pictures with anything like the frequency that human vision receives images.

MfL inputs

The 'more is better' assumption characteristic of much MfL discussion does not really face up to the realities of true raw data volumes as experienced by humans. These volume realities imply that there will in fact be far more need, in building MfLs, for serious and sophisticated forgetting, or deletion, mechanisms than has hitherto been allowed, if not also some non-trivial input filtering. Moreover, even though current automatic data recording is much less dense and repetitive than human input, we can expect its density and repetitiveness to increase quite rapidly in line with technological recording developments, with a corresponding need to address forgetting/deletion.

Thus there are major research issues to tackle in exploring experience and memory granularity and sameness as the basis not just for necessary input filtering but, much more importantly, as the basis for necessary later forgetting and deletion, before we can make any

significant progress with MfL. That is, before we can make progress with something worthy of the description 'memory for life', as genuinely more than just a bigger database of an essentially conventional kind.

MfL research data

Addressing these issues in turn depends on having appropriate data for research. Initially, as discussed in (2), this means *working data* that allows us to explore data properties and ways of responding to these, even if we are not yet in a position to specify *evaluation data* that will allow us to test our ideas about MfL in a rigorous way.

The foregoing implies that if we are to tackle what 'memory for life' really means, we cannot sell the pass by starting with working data that has already been significantly processed so that it embodies explicit granularity and structure ideas, for example data from a lapel camera that only takes a picture when certain previously-specified type of environment change occur, or news video, since this is highly edited, or audio recordings of prepared talks, or email files.

It may indeed be useful for some purposes, and is probably compatible with the least 'cognitive' and most immediately technological views of memory for life, to work with data that is far from naked and raw. Even large bodies of organisational documents of many different types are far from neat and tidy. But these are far from presenting the crucial challenges of dealing with really raw stuff. It is therefore vital for the MfL endeavour that any community research data also includes some bodies of material that is sufficiently raw to present at least some challenge at the most basic 'is this any sort of experience item in any way?' level.

(1) Karen Sparck Jones, 'Comments on Grand Challenge Document (GCD): "Memories for Life": managing information over a human lifetime', January 2004/March 2005.

(2) Karen Sparck Jones, 'Memories for Life (MfL): some questions about data', March 2005.