ROAD MAP FOR MEMORIES FOR LIFE RESEARCH
1ST JANUARY, 2008

Comments to: info@memoriesforlife.org[1]

This proposal seeks to establish a road map of research based around one of UKCRC's Grand Challenges – Memories for Life (M4L).[2] It has arisen from a number of activities – the topic was highlighted in the OST's Cognitive Systems Programme,[3] and again in the first Grand Challenges in Computing Workshop held in Edinburgh.[4] An EPSRC Network grant was awarded in 2004 and a number of workshops and meetings have led to the framing of this proposal – culminating in wide circulation to researchers in the UK community. We believe M4L is a challenge that can harness the energies, interest and work of a wide range of scientists and researchers across a range of disciplines. It represents an exciting and important set of research questions.

INTRODUCTION

The M4L Grand Challenge has excited considerable interest and raises very substantial research challenges to computer science, psychology and allied disciplines. For the first time, computing power and application development have reached a point where it is feasible – if not necessarily desirable – to record and store all the vast quantities of information gathered during one's lifetime, and to present that information to large numbers of people. Our understanding of human-computer interfaces means that this information could also be made available to sectors of society whose IT-literacy may not be so great. For this to happen, we need to focus not only on storage and presentation, but also on making sense of the information, making it usable in real-world contexts. Merely storing data is no triumph if they cannot be used easily, or without special expertise, in the future.

In this context, there is a historic opportunity to promote and direct research that facilitates the development of digital Memories for Life. To take advantage of this opportunity, a research programme is required that:

- Supports the development of common large-scale data sets for empirical research.
- Drives empirical research into potential user requirements, to avoid development being too technology-driven. Furthermore, it is essential that the M4L work being proposed can be publicly validated and inspected, so sensitivity to human user is essential.
- Supports the development of IT infrastructure capable of supplying M4L technology to a wide user base through time and space (e.g. secure systems, long-term data formats).
- Supports a programme or programmes of "blue sky" research into novel representations, content-organising methods (e.g. ontologies, narratives, smart materials), etc.
- Supports the analysis and re-engineering of existing technologies (most obviously the WWW, Semantic Web, Web 2.0) to facilitate M4L functionality.

---

[1] This road map was edited by Kieron O'Hara and Nigel Shadbolt. The editors would like to thank the many people from the EPSRC M4L Network who contributed, including Neil Beagrie, Christopher Brewster, Hamish Cunningham, Alan Dix, Wendy Hall, Alan Newell, Ehud Reiter, Mark Sanderson and Yorick Wilks. In particular, they would like to pay tribute to the late Karen Spärck Jones, whose enthusiasm for this work was reflected in the time she devoted to it, and her many contributions, comments and constructive criticisms.

[2] http://www.memoriesforlife.org/.

[3] http://www.foresight.gov.uk/Previous_Projects/Cognitive_Systems/index.html.

[4] http://www.ukcrc.org.uk/grand_challenges/index.cfm.

Perhaps the major stimulus to the programme is the continuing increase in power of information technology. By 2025 on current trends it is anticipated that our storage devices will be around 1000 times more powerful than today – the laptop on which a document such as this would be prepared will have a terabyte of RAM and 100 terabytes of secondary storage. Information densities could be on the order of 10^{20} bytes on a sugar cube size of volume. The computing fabric will be ubiquitous with high bandwidth wireless networks. We could have witnessed a move to service oriented computing or data centric environments where your data effectively surrounds you calling on applications as needed (see for example the development of USB U3 "smart drives").

This power will provide the capacity for new kinds of computing applications, whose realisation will call for sustained and adventurous research in areas as varied as multimedia encoding, retrieval, story generation, associative data mining, viewpoint reconstruction, ontology development, temporal representation, privacy enhancing technology, ethnography, sociology and psychology. Technical realisability will have to be balanced with our psychosocial needs; issues such as information overload, durable formats for long-term data storage, identity management and digital divides will also loom large. But whilst there is commercial activity in this area (cf for example MyLifeBits from Microsoft – Gemmell et al 2002) there is no general programme of research that seeks to explore the fundamental challenges raised in this field and to examine the potential applications that might arise.

The UK is well placed to launch an internationally competitive strand of research in this area. Projects such as the AKT[5] and EQUATOR[6] IRCs have performed work in this area and also researched many of the underlying techniques and methods that would be required. The DTI OST Cognitive Systems Programme looked at the challenge of M4L from a cross disciplinary viewpoint (Morris et al 2006), and of course GCC adopted M4L as one of its six challenge topics. Furthermore, as new initiatives are being put into place in Web Science, the study of the Web and weblike structures (Berners-Lee et al 2006a, 2006b),[7] there are clearly synergies with M4L; the development of the Web is likely to determine many of the more popular and well-used technologies for storing, retrieving and sharing memories for the wider (lay) population.

APPLICATIONS FOR M4L

To emphasise the practical benefit of the proposed programme, it is worth briefly reviewing some application opportunities that can be foreseen. M4L systems are likely to pervade all aspects of our lives in the future, but a number of applications are particularly interesting and worth supporting because they relate to areas of general concern or interest.

- **Healthcare and well being**. There are several potential applications, particularly with respect to extending people's quality of life. Demographic trends indicate that older people will provide an important technology market for the future, and there are potential roles for cognitive prostheses, tools to make sense of large quantities of information collected over a lifetime, or therapeutic tools. Interface issues also loom large in this space.
- **Personal life**. While many personal applications of M4L technology will more naturally be dealt with in the commercial sector, there are still many open research issues relating to, for instance, organising multimedia collections, or making sense of a series or block of unfamiliar files, where academic work could make an impact.
- **Education**. M4L is not intended to duplicate e-learning work, but there are many areas where memory prostheses will have educational potential. For instance, problems with

[5] http://www.aktors.org/akt/.
[6] http://www.equator.ac.uk/.
[7] http://webscience.org/.

working memory have been implicated in dyslexia (McLoughlin et al 2002, Mortimore 2003).
- **Entertainment**. M4L will also have important value in the wealth-creating sphere. Usable tools to manage the annotation or delivery of large quantities of archive entertainment content would have clear application to the future of television, for instance.
- **Science**. M4L technology would apply perfectly well to corporate or group memories as well as those of an individual, and an obvious application area of that sort is the corporate memory of a scientific research group, which is generally now spread over websites and individual holdings of papers, slides, notebooks, datasets etc.

RESEARCH ISSUES

We have already noted that whilst there has been progress in some areas there are many research issues which need to be addressed – in a more application-oriented and integrated way – in order to realise the ambitions sketched above. The research issues listed here are not exhaustive but they are each topics that have been discussed at some length in the M4L Network workshops that preceded this proposal. We begin by reviewing some basic issues that need to be addressed in the research; next come wider-ranging strategic issues; finally, issues that need to be resolved in the research context are explored. **We do not claim that these research issues exhaust the space.**

Core Scientific Research Issues

Many issues are fundamental to the functionality of M4L technologies. The development of some M4L systems will result in large-scale, very complex, dynamically self-modifying systems, which will provide real challenges for theoretical computer science. Powerful concepts involving general representations, inference, interaction, security, privacy, trust, stochastic behaviour and adaptivity come into play. No single model can accommodate these; we need to develop a series of models, and to understand how each model is realised or implemented in models elsewhere.

Further research is needed to devise representations, languages, processes, tools and techniques to realise theoretical approaches to M4L in applications. Application scenarios will reveal further challenges to fundamental theory. When we build our applications can we guarantee their behaviour? Do we know if networks of memory components have particular topological properties?

In this section we review some of the issues that follow from such theoretical challenges, such as representation; memory types, content and organising principles; and access, storage and retrieval issues.

Memory Representation

A core set of research questions centres on the appropriate types of representation for machine-based representations of memory. There is a long tradition of using computational models to suggest possible representation methods to psychologists, but in this case the flow of inspiration might well be reversed. We need new ways of understanding how memory is organised across modalities, if and how it exploits sparse encoding methods, what types of cue will aid retrieval – if indeed retrieval is the best way to conceive of the problem. There is substantial evidence that memories are reconstituted in the very act of recall and may in important respects be changed to meet the current context.

Some particular niche areas of research into representation are likely to be more than usually prominent. For instance, life representations, whether individual or group, inevitably have a time dimension, since lives are defined by time. The representation of temporal information has become a specialist area in AI and knowledge representation (cf. Allen 1983, Manaf & Beikzadeh 2007) but has been linked directly to Internet representation since the thesis of Setzer (2001) and to canonical concepts of time annotation (cf. e.g. Pustejovsky et al 2003, and papers in Mani et al 2005) and their role in organising the web. Since then there have been proposals for massive time stamping of Web content (CRONOPATH Project,[8] Catizone et al 2006) so as to use this to build coherent life patterns from web trawls. This would disambiguate lives with the same name, a standard Web search problem. Future M4L research will interact directly with such initiatives and their successors.

Another type of representational strategy is to combine content from more than one source seamlessly, over some kind of well-understood organising scheme (such as time-streams, or alternatively spatial maps). The scheme helps make sense of the disparate content: such arrangements, called *mashups* in the parlance of Web 2.0, have proved powerful and popular among many classes of user.

The crucial issue in M4L is not gathering or storing whole life data but developing techniques for handling and organising it in a perspicuous manner. A development in the natural language processing area in recent years has been a renewed interest in large factual knowledge bases derived by some form of pattern matching from corpus data, ideally the whole web (Kilgarriff & Grefenstette 2003). This enterprise would only be tractable if done on a huge scale in a shallow but reliable manner, and it forms an alternative data driven approach to the classic AI approach such as CYC (Lenat 1995). One obvious application of such technologies is within M4L as the derivation of all relevant facts from a whole-life-as-corpus, so as to seek an initial basis for coherence and structure in that life narrative.

It is also to be hoped that in the near future, genuinely interdisciplinary research involving information scientists, computer scientists, cognitive psychologists and cognitive neuroscientists will be able to advance the field of artificial memory representation in such a way as to determine synergistic types of representation, understandable by humans while also being machine readable (cf. O'Hara et al 2006).

Types of Memory

Memories for life can have different forms. They share a common focus on the individual, i.e. are memories belonging to, of, or about an individual. However the forms range from the hardwired prosthesis for the brain intimately connected with, and working like, neural activity in the brain to the external database centred on information bearing on the individual but operating in a more powerful, i.e. flexible, knowledge-based way, than conventional databases. In social science, the distinction between individualised and collective views of memory has not often been clearly articulated (Olick 1999), and this distinction carries over into the technological sphere.

The medium term research focus is likely to be the important, and challenging, middle of the range where memory content, organisation and operations are highly personalised and where the automated memory for life offers significant support to, and enhancement for, the individual's biological memory.

Memories for life will be characterised by vast quantities of material, derived from a wide range of perceptual and other inputs, held, used, and modified over the long term. The inputs to, internal products of, and outputs from memory will be extremely heterogeneous, taking many forms.

[8] http://nlp.shef.ac.uk/cronopath/.

Individual memory items will have variable certainty, and memory as a whole will have variable connectivity and consistency. Two major research challenges for any M4L instantiation as a system will be first, to provide appropriately rich and robust forms of content representation for memory and, second, to provide modes of inference that will make it possible to place items appropriately, to connect them, and to retrieve them, and to manipulate and modify the structures that organise sets of items. These processes will have to operate both autonomously in response to new memory inputs and according to requirements in response to calls for memory outputs. Since the material to be held in memory may take many forms, an effective memory for life will have to be able to deploy, and combine, different types of process for operating on memory.

However, although we might assume a common focus on the individual, that does not mean that memories need be person- or individual-unique. For instance, applications can be envisaged that integrate individual memories into group memories, such as family memories, company or corporate memories, or specific types of relationship such as patient/carer relations (in which the role of ‘carer’ may refer to a specific person or to a role played by different people over time). Furthermore, it could also be envisaged that certain data stores held in common could realistically incorporated as memories as a result of integration with M4L technology, for instance large commercial databases (payrolls, human resource and personnel stores, credit ratings). Steps are being taken to developing accounts based on simple easy-to-gather information (Glaser 2006, Tuffield et al 2006b). Clearly, the type of memory involved will dictate particular representational types, and indeed particular solutions to other research issues such as those that follow.

Memory Content

There is a two way dynamic between people and places and things. On the one hand places and things participate in the memories of our life, we capture them in one manner or another digitally and thus our personal collection of places is a reflection of our identity. Equally places and things speak to us, remind us of past events and trigger memories. They can play an important role in narrative construction and the association of memories (Glos & Cassell 1997). From a research perspective, the issues for M4L here are ones of cataloguing and annotating appropriately on the one hand, and locating in a digital universe on the other.

The existence of extensive gazetteers combined with unique identifiers make places uniquely identifiable which may be of immense use in the management of digital photography (Tuffield et al 2006a); collecting and even re-constructing the sequence of places visited on a journey (using the GPS record in some arbitrary digital device). How can this be integrated in a useful manner so as to facilitate the digital lifetime of the future? Will each object having its own RFID make it easier to find one’s lost keys?

Content-Organising Principles

The content and representations of memories will have to be organised to facilitate retrieval. Several content-organising principles are likely to feature in M4L research, but for the purposes of this road map we might focus on three: narratives, ontologies and annotations.

Narrative is central to the human ability to convey knowledge, retain memory and make sense of disparate events (Schank & Abelson 1975, Schank 1995, Brunner, 1990, 1991; Dautenhahn 2003). A research challenge for M4L is to address the automated transformation of digital information (images, texts, emails, formal and informal documents, calendars, bank transactions, medical records, etc.) into coherent narrative. This might have a number of applications. It may help meet the needs of some elderly people or those not otherwise deeply engaged in the digital world. It may help make some sense of human life for future generations. It may be used to piece together a person's activities over a period (for example, suspected criminals or terrorists, missing persons). With the growth of our personal digital footprint, our tendency to depend on prosthetic aids to the organisation of both present and past experience will increase.

Narrative is dynamic. We tell different stories at different times and to different audiences. Thus we can see different narratives generated from the same basic stock of information depending on the audience and the intended purpose. The construction of narrative is not authored necessarily by a single individual. It may be the product of a family or other social group (Barret 1989). In this sense it is connected to the social networks that are currently the source of much interest.

With the growth of ontologies as a vehicle for knowledge representation and integration, M4L must engage with the appropriate transformation of these knowledge structures in order to be able to manage the data generated over a lifetime. There are a number of challenges here:

- Ontologies are static knowledge structures not designed to represent the dynamics of a lifetime. New categories and knowledge structures need to be researched and developed for this end. As mentioned above, the provision of dynamic and diachronic knowledge representation structures is a major challenge with which M4L should engage (Schärfe 2004). Facilitating the use and maintenance of ontologies remains an important research issue, with a great deal of work already in progress, but most of this takes place in the context of science, academic research or commerce and business. Methods for ontology use either by non-experts, or by systems designed for non-experts, in private settings are thin on the ground.
- If ontologies are going to provide the knowledge backbone to the narrative of a digital life, then methodologies need to be developed to seamlessly annotate across a variety of media (cf. the X-Media project[9]). This involves the integration of multiple sources of information, and investigations are being made into how this might be done inexpensively (Tuffield et al 2006a).
- The M4L context raises particular challenges for the mapping and maintenance of ontologies. As fundamentally linguistic structures, the meaning associated with particular categories will change, so how will these structures be maintained over time? How will one map from one version to another without placing an unnecessary load on the end-user? Or will the imposition of ontologies on a memory for life ossify the currently fluid dynamics of memory and interpretation?

There is an interesting contradiction in that ontologies are supposed to be *social* constructs and yet in applying ontology-like structures to a personal life we imply the acceptability of personal ontologies (e.g. Kim et al. 2001). When an individual's digital lifetime becomes formalised is it necessarily in terms that are socially communicable? Or will they retain their individuality? This may be closely related to the privacy issues raised elsewhere in this document.

Furthermore, when sufficient content is assembled, and the audience sharing the information is large, local decisions by users to tag content with metadata can result in the creation of interesting emergent structures which organise the content according to the use made of it by users. Such structures (which include 'folksonomies' – emergent concept structures that share some features with ontologies) are also the object of study in a number of research programmes (cf. Berners-Lee et al 2006a, Steels 2006). Folksonomies and similar emergent structures have been developed under the general umbrella of 'social software' or 'Web 2.0'. Another type of organising principle

[9] http://www.x-media-project.org/X-Media%20Project/X-Media%20Project%20Home.html.

characteristic of Web 2.0 has been mentioned above – the use of underlying and well-understood conceptual schemes such as timelines or spatial maps to create mashups.

Annotating content is important to enable search and retrieval, especially when mediated through ontologies. Annotation is a relatively labour-intensive practice, and so automation is looked to to provide a means for increased efficiency. There are various methods to decrease the annotation overhead in some contexts. Routine actions of scholars can be used to generate annotations (referencing, note-taking, citations, indexing, updating, reviewing). When modelled as annotation, machines can help us view documents as living. Formal and authoritative annotation by scholars is most useful when accompanied by a machine-interpretable semantics or ontology. Just as XML has transformed the interchange of structured data by providing a universal syntax, the richer schema languages developed in semantic web research can facilitate the provision of higher-level information services.

Web 2.0 applications support fairly pain-free tagging, and structure emerges from collaborative tagging. Flickr is a photo tagging and publishing service which "solves" the image analysis problem by soliciting annotation from users which then serve as search terms for retrieval purposes. The resultant folksonomy constitutes annotation as behaviour mining (or DIY indexing) and is a significant current application of annotation. More complex models of annotation are also proving successful, for example in companies specialising in extracting specific and detailed information from the web. When combined with structured data sources this annotation can meet a variety of emerging needs for enhanced control, security and access. The service provided by security startup Garlik[10] is to answer the question "who's saying what about me online?" and therefore to give individuals more power over their personal information.

But semantically-enabled annotation has not spread rapidly, partly because without automation the task is too great for most people to bother with, partly because there are relatively few DIY cases (Flickr, del.ic.ious and Garlik, among others, have managed to identify common needs, but, as Tablan et al 2006) put it, in most areas we don't have enough folk for a folksonomy), and partly because, when annotating to provide material for uncertain future uses of data, it is hard to be specific enough for any real-world task (Tablan et al 2006). Improving both the quality and quantity of annotations

Memory Entry, Storage, Access, Retrieval

It is hard to predict quite how users will want to access and visualise a lifetime's data. For paper data, including photos, for many the dominant paradigm is box and sort, but for the larger volumes of digital data the box-in-the-virtual-attic solution is not tenable. Some requirements we can guess: going back to a certain time; finding things about a particular person, object or place; building some sort of overview of all or some period of your life; but others will only emerge over time. Existing projects have started to wrestle with these issues, for example, the LifeLines interface used to review case histories (Plaisant et al 1996) and MyLifeBits project (Gemmell et al 2002) visualising Gordon Bell's digitised archive (interestingly both of these are heavily timeline-based). Photologging, and indeed blogging itself have emphasised the need for interfaces to share memories (see also the proceedings of the workshop on Designing for Collective Remembering at CHI 2006[11]). Flickr is of particular note because, as well as allowing more directed browsing and sharing, its associative links through tagging and its montage views support serendipitous discovery.

[10] https://www.garlik.com/index.php.
[11] http://www.comp.lancs.ac.uk/~corina/CHI06Workshop/accepted%20papers.html.

There are many research challenges in this area. First is to try to envisage what patterns of access will be required (Whittaker 2006). Whilst we will not be able to predict all the ways in which lifelong data will be accessed and viewed, some understanding is essential for other research areas, for example, in order to gather appropriate test data sets or to choose appropriate storage and indexing mechanisms.

Data reduction is another challenge as the data sets grow, especially if the record includes automatically captured data using devices such as Microsoft's SenseCam[12] (Hodges et al 2006). Reviewing has to happen in faster than real time and this requires new forms of automatic salience detection, key framing, etc. Automatic capture itself requires interfaces to control what is or is not required. The use of explicitly stored events to index implicitly stored data is one technique that we are beginning to see in this area. Any sort of temporal indexing also has in some way to bridge the differences between objective calendar time and coordinate geography and users' own subjective understandings. This is perhaps even more important in serendipitous interfaces where without understanding unfortunate connections might be drawn: for example, showing images of a football match and funeral alongside one another.

As well as directly accessing points in time, people or things, we can envisage that users will want to create narratives for themselves and others, create or view overviews to make sense of the bigger picture of their lives. To some extent this is like a TV editor creating a programme from disparate film shots. However, the volume of data we are addressing is potentially bigger, less directed in its creation and furthermore not manipulated by specially trained professionals. Existing systems that use combinations of projection and gestures to organise and sort give some inkling of the kinds of interfaces that may be developed, but are still operating on comparably tiny data sets. For those posthumously reviewing a person's life data, this will be even more complex as it involves understanding the user's evolving ontology and what we might call id-onomies (personal keywords, relationships etc.). The EU project COMPANIONS,[13] whose director Yorick Wilks was a prominent member of the EPSRC M4L network, is addressing these issues directly.

Furthermore, making sense of a large corpus of data about ourselves in the future will be rendered even more problematic by the likelihood that such data will be highly distributed. Even the data in one's own possession will be held in many places, on PCs, laptops (perhaps one several owned by different family members), the hard drives of DVD recorders and digital cameras, on computer discs, USB memory sticks and their technological successors, as email attachments, stored on media rented from major firms such as Google, and no doubt many more. The challenges of searching and retrieving multimedia archives in these circumstances, which perhaps have developed over many years and include several different storage formats some of which may not be easily supported by current machines, will in themselves be great. But of course the distribution is likely to go beyond data repositories owned and curated by oneself, and much data will be held by others – in friends' private stores, in businesses' secure and protected repositories, in publicly available data, such as blogs or newspaper reports about one, cultural repositories and so on. There are many challenges here to do with basic functions, such as assembling such a distributed collection of heterogeneous data, as well as the querying and retrieval processes appropriate to the various interests of users – private citizens concerned with their own data, communities or organisations who wish to construst some kind of corporate history, governments and police forces

[12] http://research.microsoft.com/sendev/projects/sensecam/.
[13] http://www.companions-project.org/.

who may have to determine some pattern of behaviour, or historians or journalists who need to create a central narrative of record.

The interaction between personal digital collections and research repositories in the cultural sector is currently being explored in the Digital Lives project, under PI and member of the M4L network Neil Beagrie,[14] which brings together expert curators and practitioners in digital preservation, digital manuscripts, literary collections, web-archiving, history of science, and oral history. These materials form a substantial part of the collections of national and regional libraries and archives which are widely used by academic staff and students, and outcomes are therefore expected to be of significant interest within the Arts and Humanities and libraries, archives, and information sector. It will also be of potential interest to researchers exploring applications of digital memory in other areas such as health and ageing populations and for individuals who wish to manage their own personal digital collections for family history or other purposes.

Combinations of capture technology and video processing also offer potential for entirely new ways of reflecting and remembering – for instance, being able to view time lapse views of one's home or self and so see changes over slow time (weeks, months, years) of which we are often unaware, or being able to find oneself in other people's images and recordings and so complete gaps in one's records. One example in this area is work at Georgia Tech. keeping short buffers of implicit captures that are only permanently stored if the user explicitly requests – however, this only operates over short timescales (e.g. Truong et al 2001).

Devices and modes of interaction are also deeply problematic especially for 'on the fly' interactions with life data while in the midst of lived life. For example, existing cyborg experiments at Georgia Tech and MIT using heads-up displays have shown evolving ways of dealing with the fact that I am talking with you, but also (invisibly) reading a web page, or engaged in instant messaging (cf. e.g. Starner 1999). This is both a hardware and a software challenge, but more fundamentally addresses issues of identity and interpersonal communication (cf. e.g. Zylinska 2002, Gouveia & Biocca 2002). For example, in one cyborg-related arts installation where these interactions were 'pushed' further than normal, the cyborg's post-event interview data showed unusual, perhaps confused ideas between 'me' and 'it' (Zylinska & Hall 2002).

One particular issue of interest with the storage of memories is the longevity of storage – what are the relative merits of indefinite storage (hardware permitting), and culling? Forgetting is an ambiguous and troubling phenomenon. It has one obvious disadvantage: the loss of (access to) information that may be potentially useful later on. It has one obvious advantage: the decrease in size of information search spaces. Human forgetting has this dual aspect: we do forget 'useless' information, such as the colour of the bank clerk's tie, but equally we forget 'important' information such as where we put the car keys down.

When we use 'forgetting' as a metaphor in the computational world, the analogy is not quite exact, in that we have control over the information that gets lost (the issue of accidental loss of information through hardware failures, which may be analogous to human forgetting in other ways, is not the focus here). Can we isolate information that is not needed any more, or that is inaccurate, or that has ceased to be up to date? Many information repositories (the World Wide Web being an obvious one) retain such information, by default. A useful goal would be 'learning to forget', the process of organised forgetting, and pulling the salient forward. This should not be impossible,

[14] http://www.bl.uk/digital-lives/index.html.

given that the scientific endeavour generally, especially the data-heavy paradigm of e-science, has much in common with the creation and testing of explanatory hypotheses out of a mass of data.

Complex research issues include how to discover forgettable information, and how to ensure that other parts of the repository do not depend on them. For instance, if the assertion of a proposition in a knowledge base was originally made on the basis of another now-forgotten proposition, should the former be kept or also forgotten? How do we retain completeness of a model, and indeed how best to adjust models after forgetting? How should we measure 'forgettableness'? The possibility remains that we can use forgetting in a constructive way with artificial data stores, but automated forgetting looks a very hard research challenge – although there is a wealth of psychological work to draw on here for inspiration.

Research Strategy Issues

As well as research issues such as explored in the previous section, there are a number of strategic issues which will have an effect on the development of any M4L programme. In this section, we look at investigative strategies and M4L requirements; the capture and design of corpora to support research; and issues to do with system and device tools.

Example Applications

To promote all these areas, an important imperative is to develop measures to foster the building of example applications across a range of domains which could be seen as 'Memories for Life'. In a sparsely-populated area, it is hard to generalise; in particular, it is hard to develop datasets, or investigate user requirements in a general sense without there being a number of example applications on which we can draw. The experience of building and reviewing such applications will improve our understanding of the types of issues involved, perhaps suggesting some generalisations of the types of data resource, user requirements and supporting infrastructures that a broadly-based M4L programme would demand.

To that end, an M4L research programme should deliberately include in its specification encouragement to build and share data resources, to perform a serious requirements analysis and so on. Workshops to compare experiences in investigating user requirements, data collection and so on would be welcome in this context. Furthermore, example applications should certainly be structured so as to help focus on some of the specific research issues outlined in the previous section.

Investigative Strategies

Applications will need to be focussed on what people want rather than what technology can provide. The challenge to Computer Science is how to work with other disciplines both to ascertain what is required and how to implement systems which are usable by a wide range of users. Whittaker (2006) sets out four dimensions that affect the nature of information sharing and the technology needed to support it. There is (1) the goal of sharing (social? business?); (2) the relation between the information producer and consumer (family members? community members? or maybe there is little common ground to exploit); (3) the nature of the materials to be shared (photographs? personal and potentially sensitive diaries?); and (4) the situational characteristics of the act of sharing (synchronous or asynchronous? co-located or not? home territory or in a neutral environment?).

We can see an example of the research challenges raised here in the context of applications in the health care field. Due to the enormous diversity particularly of the older population (not only in physical sensory and cognitive functioning but also in culture, literacy, and knowledge and experience of information technology), the requirements gathering task is a major longitudinal activity in its own right. There is a requirement to determine what is needed and wanted by users. Is this factual information, actual memories, evocative memories, memories necessary for social support? What are the most appropriate ways to convey this information – pictures, sound, text, smells? How do these requirements vary with the intellectual and sensory capability of the user, their culture, age, socio economic status, and current environment (e.g. at home, in residential care, in palliative care)?

To this extent, developments in the science of the Web will be an important parallel strand of research (Berners-Lee et al 2006a, 2006b). The aim of Web Science is to merge the analysis of the Web's properties with an understanding of the requirements for Web development and a specification of the important invariants of the Web experience, to produce a more informed specification for the synthesis (i.e. the engineering) of the Web. M4L will be an important part of user experience, and therefore an important input for Web Science here.

A further major research challenge is how to enable a user to obtain maximum advantage from access to this data. People's memories are usually structured and presented to other people as narratives (see the earlier research topic on this theme). Recalling a memory is often a creative process rather than a simple recollection. These are neither "verbatim" memories, nor a bare record of events. They are interpretations of past events, fitted into the threads of a person's life, with important themes motives and storylines. Memories are not treated as isolated events, they are fitted into the pattern of a person's life – which often involves enhancing certain aspects of the memory and removing others. Thus in a real sense a person's memory is more myth than reality. We need to consider how M4L systems can be fitted into such a structure, and enhance users' abilities to create appropriate stories of their lives.

We need to develop ways of empowering individuals to manage their memories and to transmit them to others in the way they want. The overall question is how we ensure that a Memories for Life system truly assists users to make sense of the past, and convey this to others. To address this challenge, Computer Scientists need to work with story-tellers, writers, script-writers, film makers, to exploit their knowledge in order to develop systems, suitable for a wide user population, which enable memories to be turned into appealing narratives effectively and efficiently. If we take a wider perspective, there will also need to be research jointly performed with ethnographers, sociologists, market researchers and specialists into cognition, aging, group interactions etc.

Capture and Design of Research Corpora

Research advances more quickly where there is synergy generated by multiple teams working in the same area. This synergy can be effectively encouraged by the provision of shared working data.

We can expect some M4L projects to focus on memory acquisition and learning, especially when starting with no, or little, prior memory content. However, much of the research will start from some existing body of material, taking this as the priming content of the initial memory state to be further organised or reorganised according to the theories of memory operation being studied and the external functions the memory is to serve. One simple example could be a collection of digitised images, perhaps with some primitive content labelling or metadata.

Different research projects will naturally have their particular concerns. This need not, and often should not, imply that all projects find their own *starting datasets*. Past research experience has repeatedly shown that the business of gathering and establishing starting datasets is costly and, when done by individual projects mainly interested in its subsequent uses, likely to be done in a minimal manner with poor long-term outcomes.

Numerous US and also European research programmes, especially in the speech, language, and image areas, have clearly shown that common datasets, to be shared by multiple teams, have been extremely effective as promoters of research. Such datasets, designed to satisfy carefully specified, and mutually agreed, criteria, and built on a scale justified by the range of uses to which they are to be put, have served two vital purposes. They have forced individual projects to work with proper, i.e. objective and solid, data; and they have stimulated research exchange, and hence advance, through the comparisons they allow across the approaches different teams adopt to common problems (for example, in some major US programmes, to text retrieval and to question answering). Conversely, the inability to work with common data sets carries a strong implication that particular methodologies will not work outside carefully crafted environments (Shadbolt et al 1999).

Simply sharing a common body of memory input data is valuable. But there is even more value when the input data is partnered by so-called answer data, i.e. independently ratified outputs that some call on memory in some context ought to produce. In the earlier stages of a research programme we cannot necessarily expect extensive and detailed specification of memory tasks to be undertaken and thus of the outputs they ought to obtain, so we propose two phases of working/test data for an M4L programme: an initial phase designed to identify and provide one, or a small number, of datasets to be used as inputs for studying memory organisation and manipulation, and a second phase designed to specify and provide the evaluation data for some agreed memory-using tasks. It would be sensible in the first phase to consider whether candidate study data could eventually support second phase tasks, but this should not be taken as an overriding condition, since with the progress of research new study data would be needed which could also be the base for evaluation data.

One early task for the SPP will be the investigation of existing datasets, for example the well-grounded and rich dataset formed by the European Augmented Multi-Party Interaction Project (AMI),[15] that could serve, ready made, as starting test material for M4L research (Sparck Jones 2005). It is important, in dataset formation, to address, right from the start, issues of rights, privacy, community use, etc: experience has shown that research gains very largely from the existence and persistence of data sets as public research goods. It is clear that for materials intended to be strongly tied to individuals, very careful protocols and practices for proper data use will have to be established at the beginning of the programme. The public good view of data would also rule out some kinds of data as too subject to protection rights. The AMI data has the advantage that it is explicitly intended to be publicly available for research.

Notwithstanding the existence of datasets, it should be a goal of an M4L programme to create and release its own datasets, and projects set up under the M4L umbrella should include the aim of creating and supplying datasets of their own.

System and Device Tools

A very real problem for M4L is that a person's memories are and will be stored in a diverse range of formats and databases held in different physical and network locations, access to part of it being often controlled by commercial interests. Ensuring seamless, unfettered, fast access to all of a

---

[15] http://www.amiproject.org/.

user's M4L presents challenges to researchers working in data formats, conversation of formats, federated information access, as well as information retrieval and information extraction.

M4L presents many challenges to microelectronic system and device designers. We will are likely to exploit the results from SPPs such as WINES[16] with its emphasis on increasingly capable ubiquitous systems. Such systems will need to be agile in their use of the electromagnetic spectrum. They will need to be low power and capable of very long term deployment. Low power systems may call for very different microprocessor architectures. They will be required to keep pace with our need for increasingly cheap storage on more powerful processors and across all media.

These advances may derive from research in new electronic substrates or else in systems that integrate intelligent software into the very core of devices. It is likely that the results of existing programmes such as WINES or Web Science will be of direct relevance to our programme. But it is also possible that genuine hardware projects could be submitted in support of the M4L programme.

In the past data capacity has often been a limiting factor in both digital and physical storage. However, the relentless force of Moore's law has meant that, even with increased volumes of data being generated, capacity has not been a problem. In the MyLifeBits project storing large amounts of ongoing data about Gordon Bell is estimated at about 1Gb/month. Even full AV recording has been estimated at only around 30Gb/month with current compression technology. These are large numbers, but manageable at current storage capacities, and over time even more so. Indeed, on current trends, by the time a baby born now is 70 years old, full lifetime AV recording will be storable in a cube of silicon the size of a grain of dust.

Storage is not a problem – however, indexing, accessing, interfacing and probably transmitting this information do pose significant challenges. For example, a potential 'query' might be "I've met the person in front of me before: when and where?" which would involve an image match across perhaps 30 years of previous meetings. Alternatively, data mining of lifetime health data might include looking for correlations of wellbeing detected by voice and physiological measurements against diet and environmental conditions but taking into account significant life events. Whilst for each of these examples one can find existing related work (in image retrieval and conventional data mining), the special aspects of lifetime data storage change the assumptions. In some case volume may make things easier, for example, face image retrieval might actually be easier on full life AV recording as it will be possible to use more exact matching rather than attempting to correct for viewpoint, in the same way that Google does not use stemming. However, in other cases, the volume and heterogeneity of data is clearly a major challenge, rivalled only by some high security applications.

The long timescales between storage and use are also major challenges. This requires at very least ways of dealing with data formats and evolving personal ontologies. Physically maintaining long term data is problematic – it is easy to lose data due to a disk crash, and CDs and DVDs have surprisingly short lifetimes. Compared to paper data, most people's current digital archives are short lived. Central storage, for example the Mac iDisk and similar rumoured initiatives from Google, may offer more secure storage in the short term, but in the longer term companies tend to have shorter lifetimes than people. Furthermore, even today personal data is stored on devices such as phones and video recorders, not just PCs, and increasing ubiquity of computing systems (cf. the UbiComp Grand Challenge[17]) mean that there will be greater distribution of at least primary points

---

[16] http://www.epsrc.ac.uk/CallsForProposals/Archive/WINESIII.htm.
[17] http://www-dse.doc.ic.ac.uk/Projects/UbiNet/GC/index.html.

of data. So issues of data integrity and data integration amongst even personally ‘owned’ devices will be an increasingly salient issue.

This multiplicity of devices also highlights problems of synchronisation and versioning, although the latter is also a problem for single repositories. On the one hand how do we recognise the ‘same’ data item (file, image, person, place) on different repositories by content, context or metadata, while on the other how do we manage to link and interact with these, seeing the ‘most current version’ or the version at a particular time (for faces as well as documents)?

Where data is stored over multiple repositories these issues multiply still further. At the time of the first Data Protection Act, it was estimated that an individual would need to pay access fees for 10,000 government data bases to obtain all information held by them. For other reasons governments are aiming to integrate many of these, but this should not be regarded as an easy task – even when ultimately controlled by the same body. When we include commercial sources as well as government, even finding where data is stored will be difficult and issues of managing different ontologies more problematic still.

Wider Issues

Finally, we will briefly examine some of the wider issues about the context in which M4L research will be carried out. The proposed programme in M4L raises many issues of societal concern. The Network has already received significant contact from the media. M4L research is bound to raise issues of privacy, personal empowerment, limits of use by other agencies. It will be important that projects are sensitive to and aware of these issues.

Security, Privacy, Trust

Information security is of course a major issue in all large-scale or sensitive IT applications where data is being stored, and there are flourishing research programmes into how to keep information secure, how to engineer security, how to engineer privacy, how to map provenance, how to publish policies and interrogate information against them, how to manage rights to information, and so on (O’Hara & Shadbolt 2008). These general research programmes will be of great importance for M4L, as many envisaged M4L applications specifically involve storing large quantities of personal or sensitive information attached to individuals, who in particular may not have the legal, financial or institutional resources to pursue a claim for tort.

In the M4L context, however, there are special aspects of questions of security and trust. For instance, part of the M4L programme is the provision of memories for *life*. The data that an individual may store could be digital records generated many years, possibly decades, ago. M4L therefore needs to generate formats that are highly flexible and/or long-lasting, designed with security in mind. Flexibility and security may trade off against each other. The answer may well lie with policy-aware computing – recent work on this in the context of the WWW shows how it is possible to combine a rule-based policy system with theorem-proving in an open, distributed environment with heterogeneous users to facilitate a scalable access management system (Weitzner et al 2005). The need to keep security on a system with information encoded in older formats, possibly no longer actively supported, is an open-ended and wide-ranging problem.

Rights and access management are further complicated in this space by the likelihood that information about an individual may well be stored, possibly implicitly, in repositories owned by others. Tracing information about oneself is likely to remain an important subtask in the preservation of privacy. The security of the other’s system is not sufficient here: holding excessive, intrusive or inaccurate information behind secure firewalls is not an answer to the privacy problem.

Although there are a number of systems for promoting security, such as PKI and cryptographic variants, such systems have proved to be hard for individuals to use and understand. The wide user base that M4L hopes to achieve will require suitable interfaces and protocols for individual security management – a complex and ongoing research issue (Adams & Lloyd 2002).

The personal management of privacy and security also creates an interface issue. As noted, it is already hard for people to manage security settings on files, internet sites etc. With larger and larger volumes of captured data how do we make it easy for users to mark data as available to oneself, one's friends, one's family, doctor etc, implicitly as well as explicitly? The interactions between technologies here are likely to problematic. For example, automatic URL completion in web browsers is effectively using captured data (visit history) and implicitly revealing that to anyone who is watching as you enter a URL. This is compounded when the data involves the privacy of others, and not just oneself. It is likely too that privacy is likely to be time-related: "keep secret from person X until after their birthday or until after I am dead."

Human Factors and Social Context

The human factors of any Memories for Life must be accessible and inclusive. To be genuinely inclusive, they must address the challenges presented by the multiple minor physical, sensory and cognitive impairments found in older people together with the effects of major impairments which are found in almost 50% of people over 65. In addition to the traditional interface and interaction design challenges there is the challenge of developing systems which encourage creativity and the iterative development of narratives.

The aim is to do work that can be publicly inspected and validated. Hence the programme must be sensitive to the requirements of human user context. The aim is to be as user-driven as possible, and to avoid being too driven by the no doubt immense possibilities of the technology. The human factors that will be taken into account will of course depend on the application area under consideration.

INVOLVEMENT OF RESEARCH COUNCILS

Much of what we have described as M4L most naturally falls under the auspices of the EPSRC. But M4L offers particular opportunities for collaboration with other research councils. Certainly research supported by the ESRC is germane – from psychology through sociology and ethnography. The AHRC could also offer up interesting collaborations relating to literary research, narrative structure, or the role of cultural and artistic content in supporting various shared memories (cf. Zylinska 2002).

The biology and neuroscience of memory will inform the M4L programme – and as in the past the requirement to engineer such M4L systems could be of real interest to medical and neuroscience researchers. Thus both the MRC and BBSRC might be potential supporters of certain projects part funded by EPSRC. For further discussion of the interdisciplinary possibilities of M4L research, see (O'Hara et al 2006).

The Department of Health could be a potential supporter given the particular application possibilities offered in areas such as memory support. Certainly the DTI and OST who initiated the Cognitive Systems Programme could be interested in joint funding possibilities. As noted above, there are a number of application areas where the benefits of M4L technology seem particularly clear, and these application areas (such as healthcare or education) might well bring a number of interests together to support particular M4L development programmes.

NATURE AND SCALE OF ACTIVITY

A rolling programme of calls or solicitations could be structured according to some of the principles, research issues and opportunities suggested in this paper. There are six elements of the M4L research programme set out in the opening section. Their implicit relation to each other can be given shape by the way they could make an appearance in three calls.

- Development of example systems (First call/solicitation). These systems would in particular be encouraged to create and make available data sets, and to investigate user requirements.
- Development of data sets (First call/solicitation).
- Investigating user requirements (First and second calls/solicitations).
- Development of IT infrastructure (Second and third calls/solicitations).
- Representation issues (Second and third calls/solicitations).
- Web Science and study of existing technologies (First, second and third calls/solicitations).

As we come to understand user requirements to a greater extent, particular calls could be focused on particular application areas, or alternatively parallel calls, funded at least partially by alternative funding sources, could be sent out for appropriate application areas. The success, or otherwise, of example systems in the first call would help determine which application areas were most promising.

Whilst the EPSRC already has a delivery plan we believe that M4L offers great scope for focused activity in support of the general research objectives of the ICT programme. It is also of significant interest to new programmes of work such as Digital Economies; in the future we will see products and services supporting our digitally deposited memories and experiences. Such M4L resources will have social as well as economic ramifications. We believe it is important to develop a sustained programme of research in this area.

REFERENCES

Adams, C. & Lloyd, S. (2002). Understanding PKI: Concepts, Standards & Deployment Considerations, New York: Addison-Wesley

Allen, J.F. (1983). Maintaining Knowledge About Temporal Intervals. Comm. ACM 26(11), 832-843.

Barrett, E. (1989). The Society of Text: Hypertext, Hypermedia and the Social Construction of Knowledge. MIT Press

Berners-Lee, T., Hall, W., Hendler, J.A., O'Hara, K., Shadbolt, N. & Weitzner, D.J. (2006a). A Framework for Web Science, *Foundations and Trends in Web Science*, 1(1), 1-147, http://www.nowpublishers.com/product.aspx?product=WEB&doi=1800000001.

Berners-Lee, T., Hall, W., Hendler, J.A., Shadbolt, N. & Weitzner, D.J. (2006b). Creating a Science of the Web. Science 313, 769-771, http://eprints.ecs.soton.ac.uk/12615/.

Brunner, J. (1990). Acts of Meaning. Harvard University Press.

Brunner, J. (1991). The Narrative Construction of Reality. Critical Inquiry, 1, 1-21

Catizone, R, Dalli, A & Wilks, Y (2006). Evaluating Automatically Generated Timelines from the Web. *LREC*. http://www.sdjt.si/bib/lrec06/pdf/702_pdf.pdf.

Dautenhahn, K. (2003) Stories of Lemurs and Robots: The social origin of story-telling, in Mateas, M. & Sengers, P. (eds.) Narrative Intelligence. John Benjamins.

Gemmell, J., Bell, G., Lueder, R., Drucker, S. & Wong, C. (2002). MyLifeBits: Fulfilling the Memex Vision. ACM Multimedia 02, http://research.microsoft.com/~jgemmell/pubs/MyLifeBitsMM02.pdf.
Glaser, H. (2006). Semantic Squirrels. ISWC 2006, http://eprints.ecs.soton.ac.uk/13486/.
Glos, J.W. & Cassell, J (1997). Rosebud: a Place for Interaction Between Memory, Story, and Self" 2nd International Conference on Cognitive Technology (CT '97).
Gouveia, F.R. & Biocca, F. (2002). Proceedings of the 5th International Workshop on Presence, http://www.temple.edu/ispr/prev_conferences/proceedings/2002/Final%20papers/Presence2002-all%20papers.pdf.
Hodges, S., Williams, L., Berry. E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N. & Wood, K. (2006). SenseCam: A Retrospective Memory Aid. Proceedings of UBICOMP 2006, 177-193, http://research.microsoft.com/~shodges/papers/SenseCam%20Ubicomp%202006%20(camera-ready).PDF.
Kilgarrif, A. & Grefenstette, G. (2003). 'Introduction to the special issue on the Web as corpus', *Computational Linguistics*, 29(3), 333-348, http://www.kilgarriff.co.uk/Publications/2003-KilgGrefenstette-WACIntro.pdf.
Kim, S., Hall, W. & Keane, A. (2001). Using Document Structures for Personal Ontologies and User Modeling, Lecture Notes in Computer Science, Volume 2109, 240.
Lenat, D.B. (1995). CYC: a large-scale investment in knowledge infrastructure. Communications of the ACM, 38(11).
Manaf, N.A.A. & Beikzadeh, M.R. (2007). Crisp-Fuzzy Representation of Allen's Temporal Logic. Proc. 25th IASTED, 174-179.
Mani, I., Pustejovsky, J. & Gaizauskas, R. (2005). The Language of Time: A Reader. Oxford University Press.
McLoughlin, D., Leather, C. & Stringer, D. (2002). The Adult Dyslexic: Interventions and Outcomes, London: Whurr.
Morris, R., Tarassenko, L. & Kenward, M. (eds.) (2006). Cognitive Systems: Information Processing Meets Brain Science. Elsevier.
Mortimore, T. (2003). Dyslexia and Learning Style: A Practitioner's Handbook, London: Whurr.
O'Hara, K., Morris, R., Shadbolt, N., Hitch, G.J., Hall, W. & Beagrie, N. (2006). Memories for Life: A Review of the Science and Technology Journal of the Royal Society Interface, 8.
O'Hara, K. & Shadbolt, N. (2008). The Spy in the Coffee Machine: The End of Privacy as we Know It. Oneworld.
Olick, J.K. (1999). Collective Memory: The Two Cultures. Sociological Theory 17(3), 333-348.
Plaisant, C., Milash, B., Rose, A., Widoff, S. & Schneiderman, B. (1996). LifeLines: Visualizing Personal Histories. Proceedings SIGCHI, 221-227.
Pustejovsky, J., Hanks, P., Sauri, R., See A., Gaizauskas, R., Setzer, A., Radev, D., Sundheim, B., Day, D., Ferro, L. & Lazo, M. (2003). The TIMEBANK Corpus. Proceedings of Corpus Linguistics, 647-656.
Schank, R.C. (1995). Tell Me A Story: Narrative and Intelligence. Northwestern University Press.
Schank, R.C. & Abelson, R.P. (1975). Scripts, Plans and Knowledge. In P.N. Johnson-Laird & P.C. Wason (eds.) Thinking: Readings in Cognitive Science.
Schärfe, H. (2004). Narrative Ontologies. In Cao, C.G. & Sui Y.F. (eds.) Knowledge Economy Meets Science and Technology – KEST2004, Tsinghua University Press: 19-26.
Setzer, A. (2001). Temporal Information in Newswire Articles: An Annotation Scheme and Corpus Study. Doctoral dissertation, University of Sheffield.
Shadbolt, N., O'Hara, K. & Crow, L. (1999). The Experimental Evaluation of Knowledge Acquisition Techniques and Methods: History, Problems and New Directions. International Journal of Human-Computer Studies 51(4), 729-755.

Spärck Jones, K. (2005). Memories for Life (MfL): A Note About the AMI Project Data. EPSRC M4L Network.

Starner, T.E. (1999). Wearable Computing and Context Awareness. Doctoral dissertation, MIT, http://www.cc.gatech.edu/ccg/publications/starner-phd/.

Steels, L. (2006). Semiotic Dynamics for Embodied Agents, IEEE Intelligent Systems.

Tablan, V., Peters, W., Maynard, D., Aswani, N., Bontcheva, K. & Cunningham, H. (2006). From Sumer to the Internet: Recent and Future Research in Semantic Annotation for Cultural Heritage. Unpublished paper, University of Sheffield.

Truong, K.N., Abowd, G.D. & Brotherton, J.A. (2006). Who, What, Where, When, How: Design Issues of Capture and Access Applications. Proceedings of UBICOMP 2001, Springer, 209-224.

Tuffield, M.M., Harris, S., Dupplaw, D., Chakravarthy, A., Brewster, C., Gibbins, N., O'Hara, K., Ciravegna, F., Sleeman, D., Wilks, Y. & Shadbolt, N. (2006a). Image Annotation with Photocopain. SWAMM 06, http://eprints.ecs.soton.ac.uk/12004/.

Tuffield, M.M., Loizou, A., Dupplaw, D., Dasmahapatra, S., Lewis, P.H., Millard, D.E. & Shadbolt, N.R. (2006b). The Semantic Logger: Supporting Service Building From Personal Context. CARPE Workshop, http://eprints.ecs.soton.ac.uk/12899/.

Weitzner, D.J., Hendler, J., Berners-Lee, T. & Connolly, D. (2005). Creating a Policy-Aware Web: Discretionary, Rule-Based Access for the World Wide Web. In E. Ferrari & B. Thuraisingham (eds.), Web and Information Security, Hershey PA: Idea Group Inc.

Whittaker, S. (2006). Design Considerations for Collective Remembering. CHI 2006 Workshop on Designing for Collective Remembering, http://www.comp.lancs.ac.uk/~corina/CHI06Workshop/Papers/Whittaker.pdf.

Zylinska, J. (ed.) (2002). The Cyborg Experiments: The Extensions of the Body in the Media Age. Continuum International Publishing.

Zylinska, J. & Hall, G. (2002). Probings: An Interview With Stelarc. In Zylinska 2002, 114-130.